



Countermanding saccades in humans

Doug P. Hanes^a, R.H.S. Carpenter^{b,*}

^a *Vanderbilt Vision Research Center, Department of Psychology, Wilson Hall, Vanderbilt University, Nashville, TN 37240, USA*

^b *The Physiological Laboratory, University of Cambridge, Cambridge CB2 3EG, UK*

Received 3 October 1997; received in revised form 6 August 1998

Abstract

We used a countermanding paradigm to investigate the relationship between conflicting cues for controlling human saccades. Subjects made a saccade to a target appearing suddenly in the periphery; but on some trials, after a delay, a stop-signal was presented that instructed subjects to inhibit the saccade. As we increased this delay, subjects increasingly failed to inhibit the movement. From measurements of this relationship, and of saccadic latency in control trials, we estimated the average time needed to inhibit the saccade (the stop-signal reaction time or SSRT). SSRTs were similar across subjects, between 125 and 145 ms, and did not vary with target luminance. We then investigated a race model in which the target initiates a response preparation signal rising linearly with a rate varying randomly from trial to trial, and racing against a similarly rising signal initiated by the cue to inhibit the saccade. The first process to cross a trigger threshold determines whether the saccade is initiated or not. In Monte Carlo simulations, this model correctly predicted the probability of successful saccade inhibition as a function of the stop-signal delay, and also the statistical distributions of saccadic latency during trials in which a stop-signal was presented but the subject failed to inhibit the saccade. These findings provide a comparison to results previously described in the monkey, and show that a simple race model with a linear rise to threshold may underlie behavioural performance in tasks of this kind. © 1999 Elsevier Science Ltd. All rights reserved.

Keywords: Reaction time; Countermanding; Saccade; Eye movement; Monte Carlo simulation

1. Introduction

Recent neurophysiological studies have begun to elucidate the neural processes that regulate saccade production in macaque monkeys (Munoz & Wurtz, 1992, 1993a,b, 1995; Hanes & Schall, 1996; Hanes, Patterson & Schall, 1998). The outcome of these processes, which arise out of a balance between gaze-holding and gaze-shifting mechanisms, is either the initiation or inhibition of saccades. The extent to which this information can be used to help understand the generation of saccades in humans depends on the similarity of monkey and human saccadic eye movements.

One recent approach used to investigate the neural processes that regulate saccade production in monkeys is the use of a *countermanding paradigm* (Hanes & Schall, 1996; Patterson & Schall, 1997; Hanes & Paré,

1998; Hanes et al., 1998), originally developed to investigate the voluntary control of action (Vince, 1948; Lappin & Eriksen, 1966; Logan & Cowan, 1984; Osman, Kornblum & Meyer, 1986, 1990; DeJong, Coles, Logan & Gratton, 1990; DeJong, Coles & Logan, 1995 reviewed by Logan, 1994). A subject's ability to control the production of movements voluntarily is evaluated in a reaction time task by infrequently presenting an imperative *stop-signal*. The subject is instructed to withhold the impending movement if the stop-signal occurs.

One important aspect of the countermanding task is the latency for cancelling rather than initiating the response. Because of the stochastic nature of the behaviour, for any particular delay before the stop-signal is presented, sometimes the inhibition will be successful and sometimes not; the shorter the delay before the stop-signal, the greater the probability of inhibition. Generally speaking, if we determine the delay for which inhibition occurs on 50% of trials, and subtract this from the mean reaction time when no stop-signal is given, we obtain an estimate of the time needed to

* Corresponding author. Tel.: +44-1223-333886; fax: +44-1223-333840.

E-mail address: rhsc1@cam.ac.uk (R.H.S. Carpenter)

cancel the movement, known as the *stop-signal reaction time*, or SSRT. Stop-signal reaction times are around 95 ms in monkeys (Hanes & Schall, 1995; Hanes et al., 1998), and one element of the current investigation was to determine whether this value was similar in humans.

A simple interpretation of this process is in terms of a model with a race between a go process and a stop process that race towards a finish line (Fig. 4) (Logan & Cowan, 1984; reviewed by Logan, 1994). The go process prepares and generates the movement, while the stop process inhibits movement initiation: whichever process finishes first determines whether a saccade will be initiated, or not. Recent work by Hanes and Schall (1996) has shown that single neurons within the frontal eye fields, an area in the prefrontal cortex that lies at the interface of visual processing and eye movement production (Schall, 1997), of macaque monkeys appears to embody the same linear rise to threshold model (LATER: Linear Approach to Threshold with Ergodic Rate) that provides a good description of human saccadic reaction times (Carpenter, 1981; Carpenter & Williams, 1995). The LATER model embodies a signal

that rises to threshold linearly at a rate that varies in a Gaussian manner from trial to trial: as a rise-to-threshold model it is similar to the diffusion model proposed by Ratcliff (1978). By using very large numbers of trials we wished to determine whether the specific architecture of the LATER race model would provide an accurate description of subjects' behaviour, in terms of predicting both the proportion of trials on which the stop-signal was obeyed as a function of the stop-signal delay, and also the statistical distribution of latencies under the various conditions.

Some of the findings presented in this report have appeared in abstract form (Hanes & Carpenter, 1997).

2. Methods

2.1. Experimental procedure

Four subjects (AC, female, age 18; DM, male, 21; DH, male, 28; RC, male, 51) performed a countermanding task while their eye movements were recorded. All experiments were carried out with the understanding and consent of each subject. The visual stimuli consisted of rectangular yellow diffuse LEDs subtending 14×23 min arc presented by means of a beam-splitter against a colour-matched background of uniform luminance of 4.5 cd m^{-2} . In high-contrast experiments the luminance of the target LEDs was 9 cd m^{-2} , giving a contrast of 200%; in low-contrast trials the luminance was adjusted to give a contrast of 20%; in all cases the contrast of the fixation LED was 200%. All trials began with the presentation of a central fixation target accompanied by a brief warning tone, followed after a random interval in the range 0.5–1.0 s by extinction of the fixation target and illumination of a peripheral target at 9° eccentricity on the horizontal meridian, either randomly to the left or right, or in some experiments, always on the left (Fig. 1). The subject had previously been instructed to fixate the central target and then to make a saccade to look at the peripheral target when it appeared: on 65% of the trials (*control* trials) this is all that happened. However, on 35% of the trials (*stop* trials) the fixation spot reappeared after a delay of δ ms, and remained on, together with the peripheral target, for the entire duration of the trial. Stop-signal delays, δ , ranged from 10–170 ms in 40 ms steps for the high contrast targets and from 50–210 ms in 40 ms steps for low contrast targets; all trials lasted 1000 ms from the appearance of the peripheral target. In some stop trials the subject obeyed the countermanding signal and no saccade was generated; in others a saccade to the target was nevertheless initiated. Experiments were performed in blocks of 200 or 300 trials, with all the values of δ used occurring within a block. Each block of trials contained only high or low contrast targets.

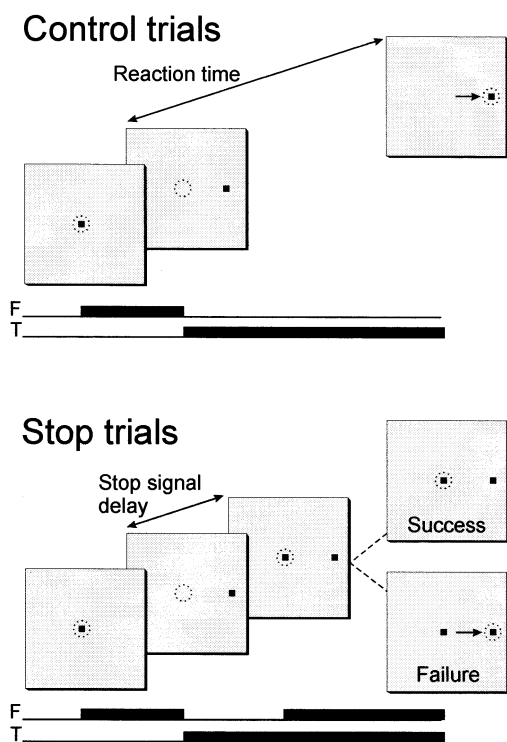


Fig. 1. Trial displays for the countermanding task. The dotted circle indicates the focus of gaze at each interval, and the arrow indicates the saccade. All trials began with the presentation of a central fixation spot. After fixation of this spot for a variable interval, it disappeared, while a visual target simultaneously appeared in the periphery. In *Control trials* the subjects then made a saccade to this target. In *Stop trials*, the fixation spot reappeared after a delay, referred to as the stop-signal delay: this instructed the subject to withhold the movement. During these stop trials subjects sometimes successfully countermanded the saccade and sometimes they did not.




Fig. 2. Latency distributions for the four subjects using high (filled symbols) and low contrast targets (open symbols), shown as reciprobbit plots (i.e. cumulatively with a reciprocal time axis and probit ordinate).

2.2. Training procedures

Subjects were initially trained to generate saccades to the peripheral targets without the introduction of stop trials. For all four subjects, examination of the distributions of saccade latencies showed small but significant differences in the saccade latency distributions for targets to the left and right of the central fixation spot. Since we required the distributions of saccade latencies for subsequent calculations, the necessity to evaluate the data for saccades to the right and left separately would have required the subjects to perform twice the number of trials. Therefore, for two of the four subjects (AC and DH) the target was always presented 9° to the left of the central fixation spot. For subject DM the target always appeared randomly on either the right or left of the central fixation spot. To check that left only and left/right target pre-

sentation conditions were comparable, data were collected from RC during blocks of trials under both conditions. There were no significant differences between these two target presentation conditions, justifying the comparison of the data from all four subjects.

Once the subjects had performed approximately 400 trials without a stop-signal, stop trials were introduced. Subjects were told that on some trials the fixation spot would reappear (i.e. the stop-signal) at a variable delay after the target was presented. They were told that some stop-signals would occur early enough that they would be able to inhibit saccade production and some would occur so late that they would not be able to do so. Subjects were instructed to try to inhibit the saccade when the stop-signal was presented, but not to be concerned if they were not able to inhibit it. Subjects were given the opportunity of resting between blocks.

2.3. Data collection and analysis

Eye movements were measured by means of an infra-red scleral oculometer (Carpenter, 1988), with a frequency response flat to 500 Hz and linear to 1% over a range of some $\pm 10^\circ$. Its output was sampled at 10 ms intervals by the PC-based saccadic analysis system SPIC (Carpenter, 1994) which controlled the presentation of the stimuli, and displayed and stored the eye movement data, detecting saccades in real time by means of a velocity and acceleration criterion. At the end of a series of blocks, the operator would go through the stored records, eliminating those with blinks or other errors. SPIC was then used to calculate and display raw and cumulated histograms and to perform statistical tests (Kolmogorov-Smirnov).

2.4. Monte Carlo simulations

SPIC was also used to perform the Monte Carlo simulations. For each simulated trial, values r_{go} and r_{stop} for the rates of linear rise of the go and stop-signals respectively were selected randomly from a pair of Gaussian populations of means μ_{go} and μ_{stop} , and standard deviations σ_{go} and σ_{stop} (Fig. 4), taking the threshold, θ , as unity. (The parameters μ , σ and θ are not independent: if all three are equally scaled, the resultant distribution is identical). Starting τ ms after the beginning of the trial, iterations were performed that corresponded to 10 ms intervals of real time; at each iteration the value x_{go} of a decision signal, initially zero, was incremented by r_{go} . In the same way, after the stop-signal delay δ , and with the same added constant delay of τ ms, the value of a second decision signal x_{stop} was incremented at each iteration by r_{stop} . If x_{stop} reached a threshold value, θ , before x_{go} did, then the trial was taken as one in which the saccade was successfully inhibited; if on the other hand x_{go} reached θ , before x_{stop} , then the saccade was taken to have escaped inhibition, and the time of reaching threshold was taken as the simulated saccadic latency. The value of τ was held constant at 60 ms for all simulations, an estimated value based upon the onset latency of visual cells in the macaque visuomotor system (e.g. Goldberg & Wurtz, 1972; Thompson, Bichot & Schall, 1997). It is worth noting that the behaviour of the model is in fact little affected by changes in τ . Randomisation was achieved by means of a standard linear congruence method with a period of over 30 000 (Abramowitz & Stegun, 1965).

3. Results

3.1. Behavioural performance

Across all subjects, we analysed a total of 6832

control and 3482 stop trials performed with the low contrast target, and 7461 control and 3711 stop trials with the high contrast target. A total of 2262, 9347, 2456, and 7421 trials were collected from subjects AC, DH, DM, and RC, respectively. From this raw data, with both the high and low contrast targets, three determinations were made: the distribution of the latencies of saccades during control trials, and of saccades that failed to be inhibited in stop trials; and of the probability of inhibiting a saccade when a stop-signal was given as a function of stop-signal delay.

3.2. Control trials

Fig. 2 shows the distribution of latencies to high contrast targets (●) and low contrast targets (○) for all subjects during control trials (i.e. when there was no stop-signal). Saccade latency distributions are plotted cumulatively on a probit scale with a reciprocal time axis. This method of plotting means that latency distributions, which normally exhibit a skewed upper tail, result in a straight line (Carpenter, 1981; Carpenter & Williams, 1995). As is evident from Fig. 2, the distribution of saccade latencies to low contrast targets is shifted to the right relative to the distribution of saccade latencies to high contrast targets. Across all subjects, the average latency L was significantly less for the high contrast targets, 227 ± 0.4 ms, than for the low, 259 ± 0.4 ms ($t = 57.47$; d.f. = 14291; $P < 0.01$): values for individual subjects are shown in Table 1.

3.3. Stop trials

In stop trials, the generation of saccades depends on the stop-signal delay, δ . With short delays the subjects successfully inhibited saccade initiation, but as δ increased subjects increasingly failed to withhold saccade initiation. A plot of the probability of responding with a saccade (i.e. of ignoring the stop-signal) for different values of δ , which may be written as $P(\delta)$, is known as the *inhibition function*. Fig. 3 shows inhibition functions for each subject with high contrast targets (●) and low contrast targets (○), and demonstrates the effect of increasing δ very clearly. For example, with the high contrast target the probability of RC generating a saccade was 0.11 for $\delta = 50$ ms but was 0.60 for $\delta = 90$ ms. Also, for a given δ , $P(\delta)$ is always greater for the high contrast target than for the low. At $\delta = 90$ ms, for instance, the probability of DH generating a saccade to the high contrast target was 0.44, but to the low contrast target was 0.18. We posit that this increase in the probability of successfully inhibiting saccades when target contrast is low is because—as demonstrated with

control trials—the underlying latency is longer to low contrast targets than to high contrast targets. If the shift in the inhibition functions is indeed due to the increase in control saccade latency with low contrast targets, then plotting the percent inhibition as a function of $(L - \delta)$, (where L is the mean saccade latency during control trials) rather than simply of δ , should

bring the inhibition functions into alignment (Logan & Cowan, 1984). The result of this transformation is illustrated in Fig. 3B. Because the abscissae for the two data sets do not correspond, one cannot test for alignment statistically, but it is evident that the prediction is essentially fulfilled. This result supports the notion that the shift in the inhibition function for low and high

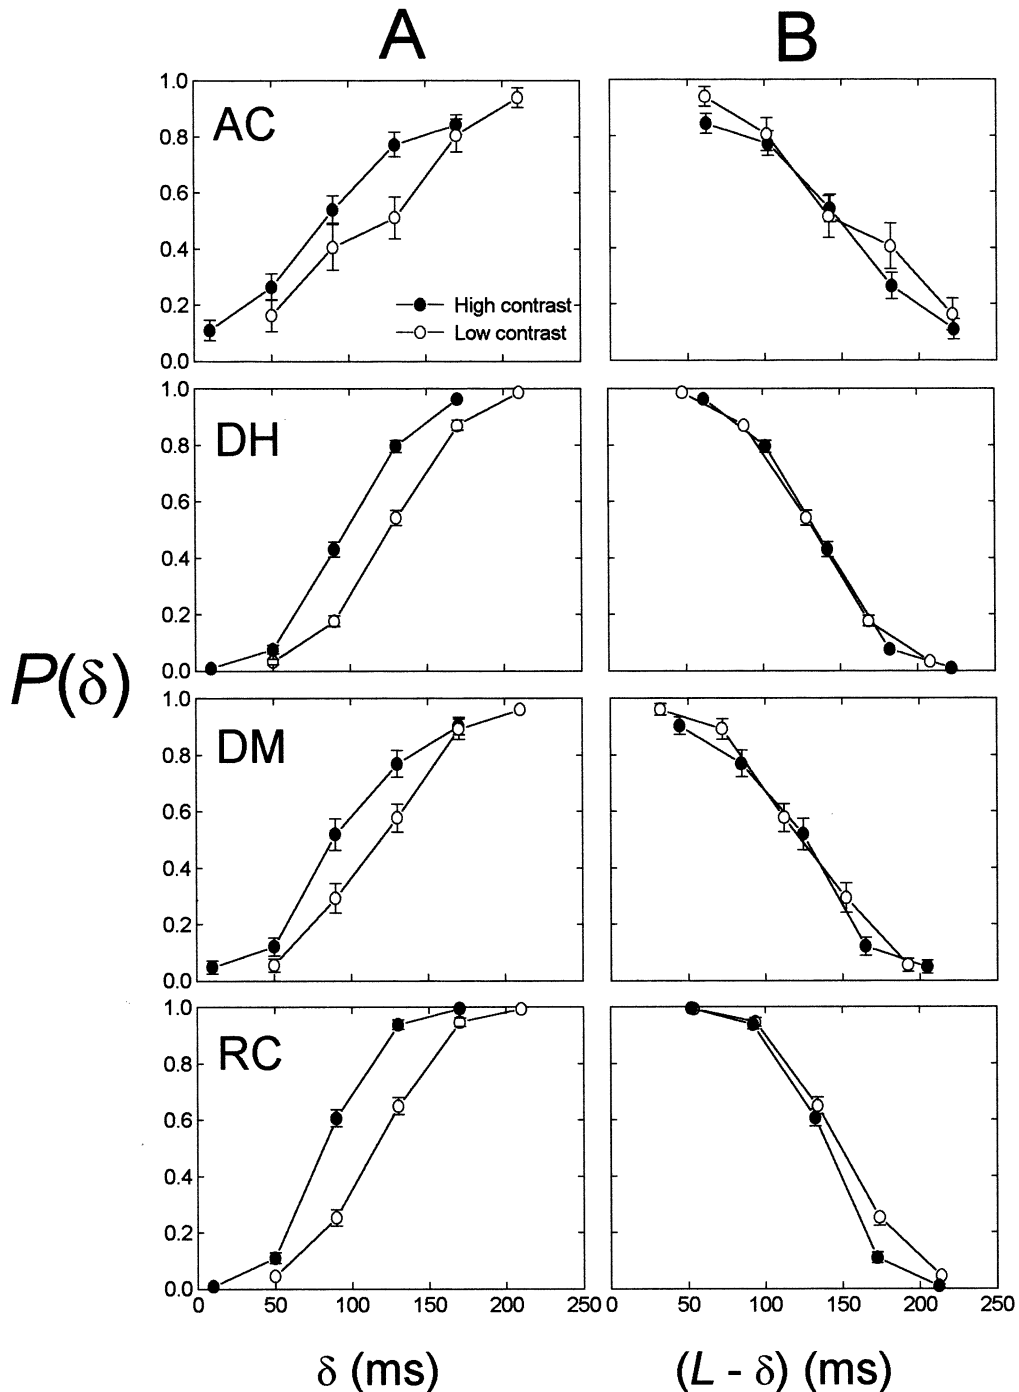


Fig. 3. Inhibition functions. Left (A), the observed probability $P(\delta)$ of each subject failing to inhibit saccade initiation is plotted as a function of the stop-signal delay, δ , for high (filled symbols) and low contrast targets (open symbols). Right (B), the same probabilities plotted as a function of $(L - \delta)$, where L is the corresponding mean saccade latency during control trials. One S.E.M. is indicated for each data point.

Table 1

Mean saccadic latencies, L , in ms to high and low contrast targets in control trials for each subject, together with the overall means across all subjects

	High contrast	Low contrast
AC	232	272
DH	232	258
DM	215	242
RC	222	264
Mean	227 ± 0.4	259 ± 0.4

contrast targets is a result of the difference in the values of control trial saccade latencies in the two cases.

One important aim of this investigation was to determine the *stop-signal reaction time* (SSRT), the length of time required to cancel the saccade being programmed and maintain fixation on the central fixation spot. The duration of this covert inhibitory process is not explicit in the behavioural data, but can be determined by the application of a race model (Logan & Cowan, 1984; reviewed by Logan, 1994), the race being between a go and a stop process (Fig. 4). The go process, initiated by the presentation of the target stimulus that increases a subject's readiness to respond, includes both the release of fixation and programming the metrics of the saccade. The distribution of movement latencies during control trials represents the outcome of the go process on its own. Previous work (Carpenter, 1981; Hanes & Schall, 1996) suggests that one possible form of the go process is a linear rise to a fixed threshold, with the rate of rise varying randomly from trial to trial. The stop process inhibits movement in response to the presentation of the stop-signal (in this case the reappearance of the fixation spot), and may equally be modelled as a linear

rise to threshold. If the go process happens to have a high rate of rise, and crosses its threshold before the stop process, the saccade is generated (Fig. 4B): If the stop process arrives first, saccade initiation is inhibited (Fig. 4A). It is worth noting that although for ease of explanation the two thresholds are shown as identical, they do not have to be, and most likely are not, at the same level. The random variability of the rates of rise of the go and stop processes is thus translated into a randomness in whether the stop-signal actually results in the inhibition of saccade initiation.

The monotonically increasing inhibition functions arise because increasing the stop-signal delay postpones the onset of the stop process, as can be seen by comparing panels A and B in Fig. 5. Here the timing of two stop trials is superimposed on the control latency distribution for subject RC, all the data being collected with the high contrast target. With small δ (Fig. 5A), the stop process more often reaches its threshold before the go process, resulting in inhibition in the majority of trials. When δ is increased (Fig. 5B), the probability that the go process will reach its threshold before the stop process is increased, and so the probability of inadvertently generating a saccade is also increased, tending to one as δ is made larger.

The inhibition function $P(\delta)$ is used in the context of the race model to estimate the stop-signal reaction time (SSRT). Two methods of estimating the SSRT were used in the current study, one based on the integration of the *control* trial saccade latency distribution and the other using the mean of the inhibition function. Detailed descriptions of these methods have appeared previously (Logan & Cowan, 1984; Logan, 1994; Hanes & Schall, 1995).




Fig. 4. Races between go and stop processes. The race model consists of a go process (solid line) and a stop process (dotted line) that are racing independently toward their respective thresholds (dashed horizontal line). The thresholds for the go and stop processes coincide only for ease of illustration. In *stop* trials, the stop process is evoked after the go process has begun. Left panel, the go and stop stimuli each trigger a signal rising linearly towards a threshold: if, as here, the stop process rises so fast that it overtakes the go process and reaches threshold first, the saccade is successfully inhibited. Right panel, if the go process reaches threshold first, the saccade fails to be countermanded.




Fig. 5. The relation between the stop-signal delay (δ), stop-signal reaction time (SSRT) and the probability of successfully countermanding a saccade in stop trials. The timing of two stop trials is superimposed on the control saccade latency distribution. Presentation of the stop-signal after a short δ (upper panel) initiates the stop process, which when it reaches its finish line (on average at a time SSRT after the stop-signal delay) successfully cancels all saccades that would otherwise have occurred after that point (light shading), but fails to cancel those occurring earlier (dark shading). If δ is increased (lower panel), the probability of failure will increase in proportion to the increased area to the left of the finish line.

The integration method assumes that the duration of the stop process is constant for a given stop-signal delay. At first sight this assumption may seem unwarranted as it is implausible that a physiological process would take a constant amount of time to execute, and indeed it will be shown later in this paper that the results are best fitted by assuming that the rate of rise of the stop process does indeed vary. However, Logan and Cowan (1984) mathematically analysed the consequences of this assumption and found that it introduced only small errors, as was confirmed by DeJong et al. (1990) using Monte Carlo simulations. By this method the SSRT is estimated for any stop-signal delay, δ , by integrating the control saccade latency distribution, beginning at the time of target presentation, until the integral equals the probability $P(\delta)$ of generating an errant saccade at that stop-signal delay. The time value at that location represents the finish line of the stop process. Thus, the time between the onset of the stop-signal and this finish line represents the *stop-signal reaction time* at this stop-signal delay. Means and standard errors of the SSRTs for each subject, measured in this way across all the values of δ , are shown in Table 2. Across all subjects the average (\pm S.E.M.) SSRTs using this method of estimation were 136 ± 4.6 and 143 ± 5.2 ms for the high and low contrast targets,

respectively. In other words, once the fixation spot reappeared it took approximately 140 ms to cancel the impending saccade.

The second method of estimating the SSRT assumes that it is a random variable. Logan and Cowan (1984) showed that the mean SSRT is equal to the difference between the mean reaction time during control trials and the mean value of the inhibition function. The mean of the inhibition function was determined by treating the inhibition function as a cumulative distribution and converting it to a probability density function. If the inhibition function ranges from a

Table 2
Stop-signal reaction times for different subjects and methods of estimation^a

	High contrast		Low contrast	
	Mean	Integration	Mean	Integration
AC	150	140 \pm 15	140	150 \pm 9
DH	135	138 \pm 10	130	142 \pm 6
DM	123	122 \pm 11	121	128 \pm 6
RC	138	142 \pm 8	145	150 \pm 6

^a One method uses the mean of the inhibition function and the other is based upon integration of the *control* trial saccade latency distribution. Times are given in ms, \pm 1 S.E.M. where appropriate.

probability of 0–1, then the mean is the difference between the probability P_i of responding at the i th stop-signal delay and the probability P_{i-1} of responding at the $(i-1)$ th stop-signal delay, multiplied by the i th stop-signal delay, summed over all stop-signal delays (Logan & Cowan, 1984):

$$\text{Mean of inhibition function} = \Sigma[(P_i - P_{i-1}) \cdot \text{SSD}_i]$$

During most conditions the actual inhibition functions had a minimum greater than 0, or a maximum of less than 1. To account for this, the mean of the inhibition function was rescaled to reflect the range of the probability of responding. This was accomplished by dividing the mean of the inhibition function by the difference between the maximum and the minimum probabilities, P_{\max} and P_{\min} , of responding:

$$\text{Mean of inhibition function} = \frac{\Sigma[(P_i - P_{i-1}) \cdot \text{SSD}_i]}{(P_{\max} - P_{\min})}$$

The average SSRTs during trials to the high and low contrast target are shown for all subjects in Table 2. Across all subjects the average (\pm S.E.M.) SSRTs using this method of estimation were 137 ± 5.5 and 134 ± 5.3 ms for the high and low contrast targets, respectively.

A central premise of the race model used to estimate the SSRT is that the go and stop processes are stochastically independent: specifically, that the finish times of both process are uncorrelated. Violation of this premise is not fatal; it only means that the estimate of the SSRT will vary as a function of stop-signal delay (Logan & Cowan, 1984; DeJong et al., 1990). The possibilities are either that the stop process is affected by the go process, or vice-versa. In the first case, SSRTs would then differ for high and low contrast targets, since the duration of the go process was different for high and low contrast targets (Fig. 2). However, across all subjects the average SSRT using both methods of estimation for the high contrast target was 136.0 ms and was 138.3 ms for the low contrast target, which were not significantly different (t -test; $P > 0.05$). This result supports the independence assumption and indicates that the go process does not influence the finish time of the stop process.

The converse possibility, that the go process is affected by the stop process, has been shown to be false by a number of previous studies (Logan & Cowan, 1984; DeJong et al., 1990; Hanes & Schall, 1995; Hanes et al., 1998). We addressed this possibility by performing a median test to determine whether reaction times during stop trials in which the saccade escaped inhibition were different from those lying within the equivalent part of the distribution of control trials (i.e. those with reaction times less than the sum of the SSRT and δ). This test was done for all SSDs in which more than 20 saccades escaped inhibition, a criterion that was

fulfilled in 32 data sets. Only two sets, both from subject RC, showed significant differences in the medians ($P < 0.05$), which may well be the result of a modification in the distribution of the longer response times, apparent in the responses for small values of δ shown in Fig. 7: a possible mechanism is discussed later. That apart, this analysis provides further evidence for independence of the finish times of the go and stop processes.

3.4. Monte Carlo simulations

Previous work using the countermanding task has posited that a race model between independent go and stop processes underlies a subject's behavioural performance. In other work it has been shown that the process that initiates a saccade can be modelled as a linear accumulator with a variable rate of growth across trials (LATER model) (Carpenter, 1981; Hanes & Schall, 1996). We therefore implemented Monte Carlo simulations to determine whether a LATER model of the go and stop processes could in fact give accurate quantitative prediction of actual human behaviour. The procedures have already been described in Section 2.

In order to compare the *distributions* of simulated data with the *distributions* of actual data between 2000 and 5000 trials in both the high and low contrast target conditions were necessary. This large number of trials could only be collected from subjects DH and RC; the data from the other subjects, while adequate for determining means and other global parameters, do not permit so critical a test of conformity to what is predicted. A total of four groups of simulations were performed, one simulating responses to high contrast targets and one to low, for both DH and RC. Each group consisted of six simulations, one with only the go process active, simulating control trials, and the other five with different delays between the go and stop processes to simulate stop trials. For the high contrast simulations, the five delays used were $\delta = 10, 50, 90, 130,$ and 170 ms; for low contrast, $50, 90, 130, 170,$ and 210 ms. Each of the 24 simulations contained 2048 trials, making a total of 49 152 simulation trials in all.

Four parameters were needed for each simulation, $\mu_{\text{go}}, \mu_{\text{stop}}, \sigma_{\text{go}}$ and σ_{stop} . Since only the go process is active in control trials, both μ_{go} and σ_{go} were estimated from the actual distribution of their latencies, using the values that minimised the Kolmogorov-Smirnov statistic between the simulation and the actual distribution of *control* latencies for that condition (i.e. high or low contrast target) and subject. Fig. 6A shows the distribution of simulated and actual saccade latency distributions for high and low contrast targets during control trials for DH and RC. The actual and simulated distributions were not significantly different for either the

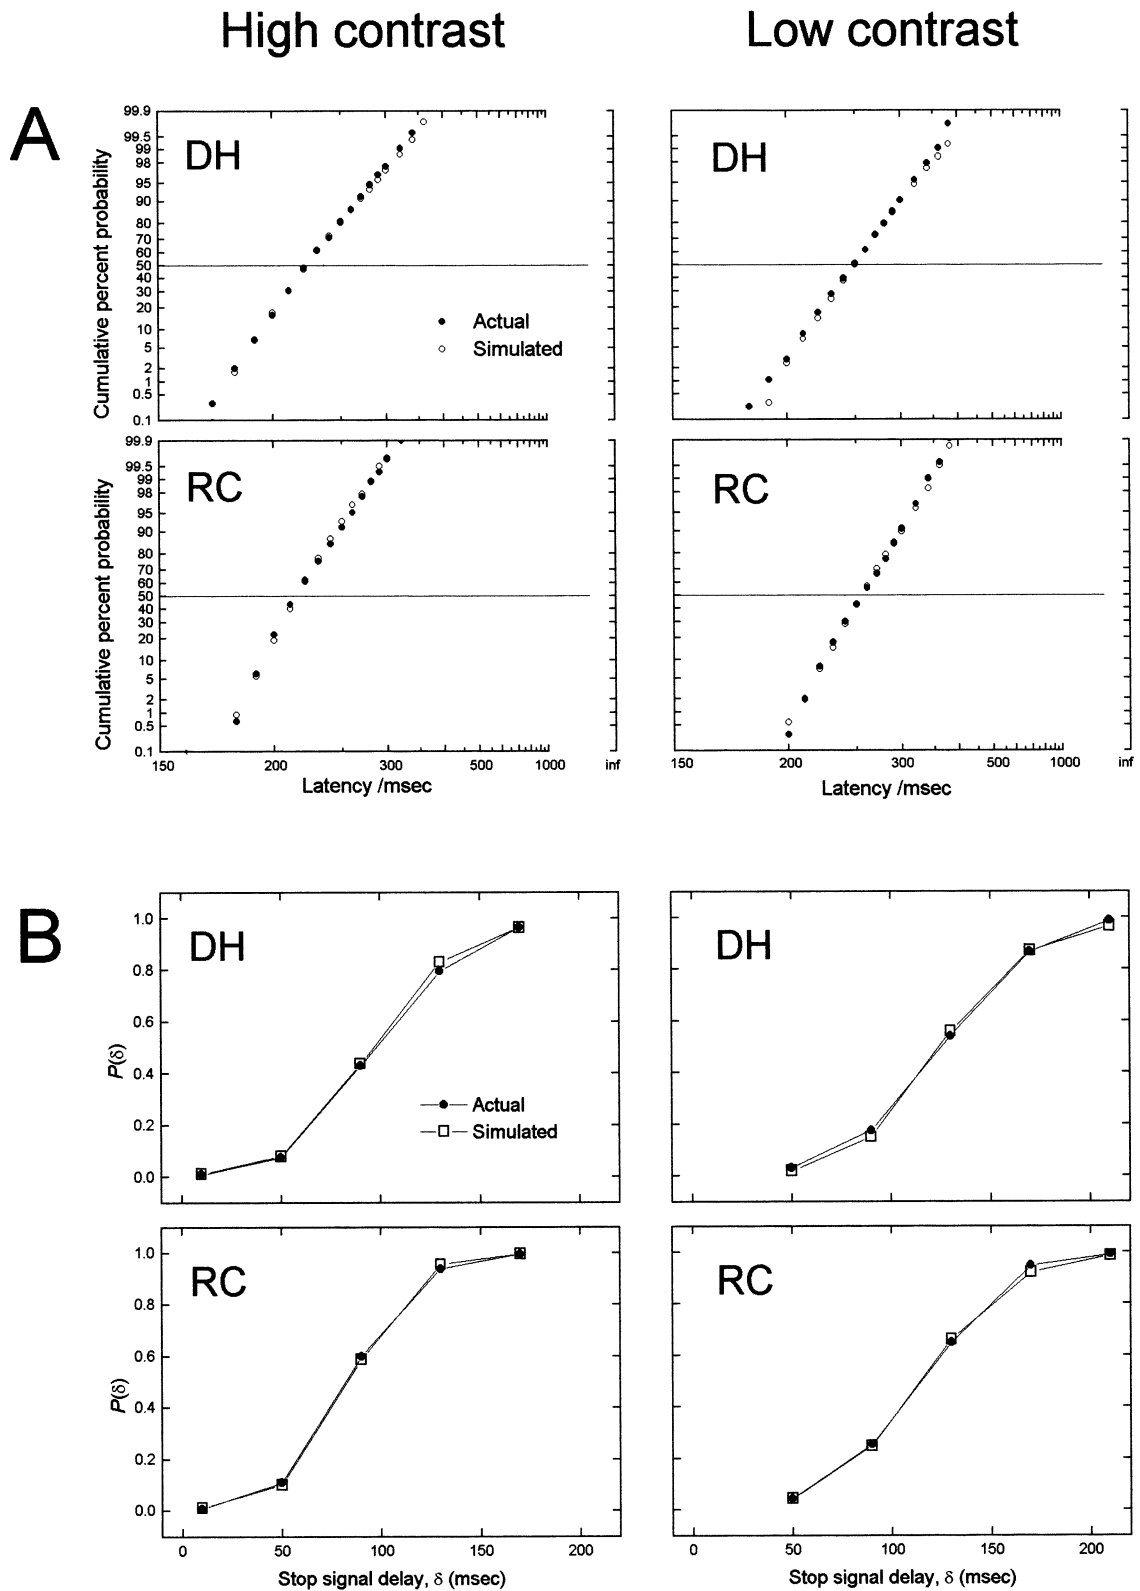


Fig. 6. Actual behaviour (filled symbols) and simulations (open symbols) for subjects DH and RC with high (left) and low (right) contrast targets. Above (A), actual and simulated latency distributions in control trials, shown as reciprob plots. Below (B), inhibition functions showing the probability of failing to inhibit saccade initiation, $P(\delta)$, as a function of stop-signal delay, δ .

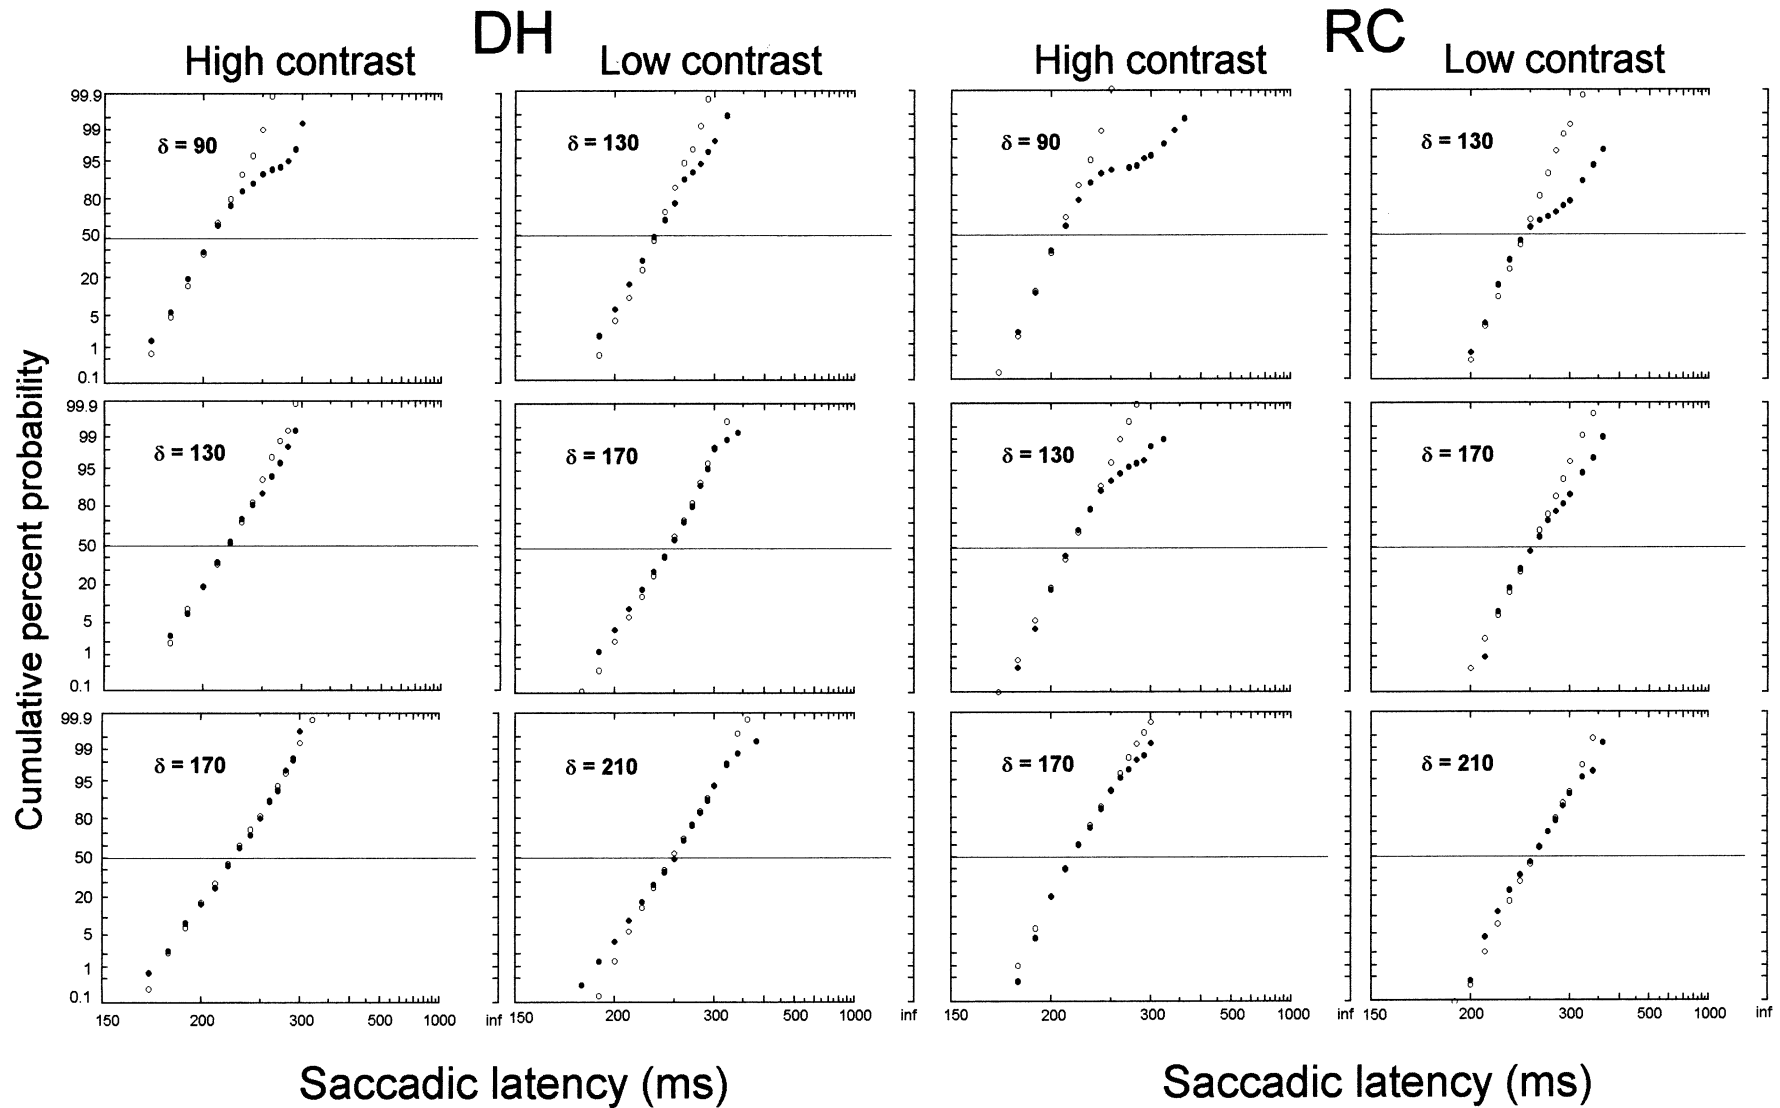


Fig. 7. Actual (filled symbols) and simulated (open symbols) latency distributions for high and low contrast targets for trials when the saccade escaped inhibition. The values of stop-signal delay, δ , in ms are shown in each case. Simulations are based on a LATER model. When either the go or stop process crosses its threshold the other process is entirely inhibited.

Table 3

Best-fit parameters for modelling distributions and inhibition functions in stop trials for subjects DH and RC, at high and low contrast^a

	High contrast		Low contrast	
	DH	RC	DH	RC
μ_{go}	6.11	6.47	5.27	5.13
μ_{stop}	13.80	13.46	14.80	11.80
σ_{go}	1.05	0.80	0.88	0.75
σ_{stop}	3.00	2.20	3.20	2.00

^a All values are in Hz, for unit threshold.

high or low contrast targets (Kolmogorov-Smirnov, $P > 0.1$).

Simulations were then run at the five stop-signal delays with these previously determined values of μ_{go} , σ_{go} and varying μ_{stop} and σ_{stop} , in order to minimise the deviation of the actual and simulated inhibition functions. The best fit values of μ_{go} , μ_{stop} , σ_{go} and σ_{stop} for both the high and low contrast targets for DH and RC are shown in Table 3.

Fig. 6B shows the actual and simulated inhibition functions for subjects DH and RC using the four best fit parameters. A binomial test showed no significant differences between any simulated and actual value of $P(\delta)$ under either the high or low contrast conditions for DH and RC ($P > 0.05$; for all but one, $P > 0.1$). In other words, the model is able satisfactorily to predict the proportion of trials in which saccades are successfully inhibited, as a function of the stop-signal delay.

An adequate model of performance in the countermanding task should be able to predict not only the inhibition function, but also the latency distribution for

those saccades that escape inhibition in stop trials. To test this more stringent prediction, we compared simulated and actual latency distributions for trials when the saccades failed to be inhibited. It is worth noting that trials for the two shortest values of δ could not be used for this purpose, since even with the thousands of trials that the experiment as a whole demanded, the number of saccades escaping inhibition at these stop-signal delays was too small for the predictions concerning the distributions to be tested. Fig. 7 shows the results of this comparison.

The actual and simulated saccade latency distributions were not significantly different (K.S. $P \pm 0.05$) in 11 of the 12 conditions tested. The exception was for the low contrast target with RC as subject and $\delta = 130$ ms. As is evident in Fig. 7, the actual distributions sometimes exhibited a peculiarly elongated long-latency tail, most noticeably at shorter values of δ , which the method of plotting tends to exaggerate even though it generally only accounted for approximately 5–10% of the entire saccade latency distribution and was statistically compatible with the simulation.

One possible explanation for the elongated tail is that when the stop process reaches its threshold, a degree of inhibition of the go process occurs that in most cases prevents it from reaching its threshold, but in a small percentage of trials is manifested as a slowing of the rate of rise of the go process which then results in a deferment of the saccade rather than its complete cancellation. Fig. 8 shows that a simple model of this kind, with suitable selection of the parameters, can mimic this peculiarity of the distribution quite well. However, there are probably many processes of this general type that are capable of modelling these delayed responses,

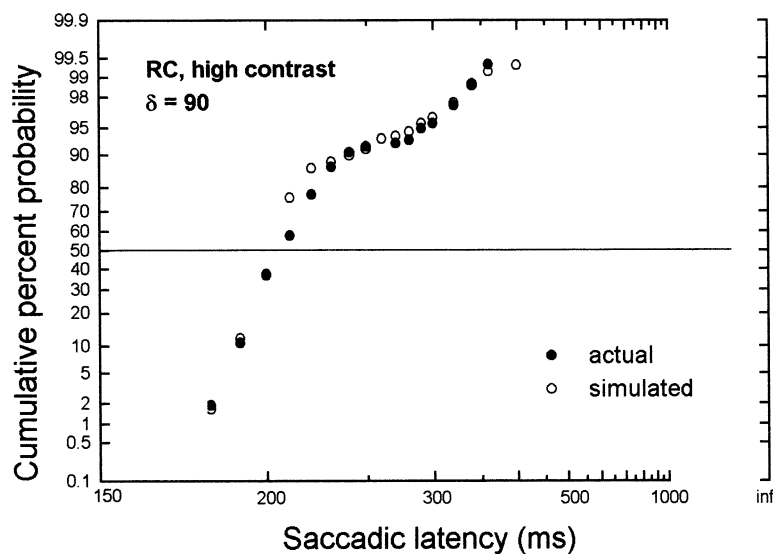


Fig. 8. Actual and simulated latency distributions for subject RC with high contrast targets, using a modified LATER model in which the go process is only partially inhibited after the stop process reaches its threshold; it is then possible to explain the extra slowing of the longer reaction times (Table 1).

whose accurate characterisations would require data sets an order of magnitude larger than what is presented here. For the moment, this aspect does not seem worth pursuing.

It is perhaps worth pointing out that of the four parameters shown in Table 3 for each condition, two (μ_{go} and σ_{go}) are pre-determined by the data from control trials, and one (σ_{stop}) turns out to be rather uncritical, in the sense that variations of the order of $\pm 20\%$ typically produce only insignificant changes in the distributions or inhibition functions. Fitting thus relies heavily on the value of μ_{stop} , and it is clear that this value is consistently more than double μ_{go} . The values of μ_{stop} should also be related to the behaviourally estimated SSRT for any particular condition, since this depends on the time taken for the stop-signal to rise to its threshold. The relationship is however not straight-forward, since the SSRT is defined with respect to the mean of the latency distribution, whereas μ_{stop} determines the median. Within this limitation, the values were nevertheless comparable except for DH with low contrast targets: here the value of μ_{stop} that best fitted the inhibition function was distinctly different from what would be predicted from the behaviourally estimated SSRT; a value compatible with the SSRT did nevertheless provide a statistically acceptable prediction of the inhibition function, but not the best fit. Overall, the results of these simulations show that the LATER race model between independent, linear go and stop processes can account not only for the inhibition functions but also for the distributions of latencies for those saccades that escape inhibition, and that the stop process rises more than twice as fast as the go process.

4. Discussion

This study represents a continuation of research aimed at understanding the processes that regulate saccade production. Previously, we have shown that rhesus monkeys can perform the countermanding task and have identified a neural correlate in the FEF of macaque of the go and stop processes that are posited to underlie a subject's behavioural performance in the this task (Hanes & Schall, 1995; Hanes et al., 1998). The extent to which this information can be used to help understand the generation of saccades in humans depends on the similarity of monkey and human saccadic eye movements.

This investigation, the first to implement a saccade version of the countermanding task in humans, yielded three main findings. First, that the latency required to inhibit the production of saccades following the presentation of a stop-signal is similar across subjects, on average 137 ms, and is approximately 40 ms longer

than in rhesus monkeys. Second, that average stop-signal reaction times do not vary significantly when the saccade latencies were altered in control trials by manipulating target contrast. Third, that a simple LATER race model between independent go and stop processes can account for both the percentages of trials in which saccades are successfully inhibited under different conditions and the distribution of their latencies when they escape the inhibition.

4.1. Relation to previous work

Several studies of manual choice and simple response time have used a race model between independent go and stop processes to explain a subject's performance in the countermanding paradigm (Logan, 1981, 1982, 1983; reviewed by Logan & Cowan, 1984; Osman et al., 1986; DeJong et al., 1990). In the current paper, we have proposed a specific type of race model in which the go and stop processes rise linearly toward a threshold; whichever processes crosses its threshold first determines if a saccade will be initiated or not. While the rates of growth of the processes do not vary within a trial they do vary across trials in a Gaussian fashion. This type of linear model, referred to as LATER, has been shown to provide a good description of human saccadic reaction times (Carpenter, 1981; Carpenter & Williams, 1995). Recently, Hanes and Schall (1996) have also shown that single neurons within the frontal eye fields of macaque monkeys appear to embody the LATER model. The Monte Carlo simulations were used to provide an additional test as to the merits of the LATER model. These simulations showed that not only can the LATER model account for the probability of successfully inhibiting saccade initiation as a function of the stop-signal delay, but it can also predict the statistical distributions of saccadic latency during trials in which a stop-signal was presented but the subject failed to inhibit the saccade.

One simplifying assumption of this race model is that the go and stop processes grow independently. Previous studies have provided evidence that is consistent with this assumption. First, the behavioural predictions based upon this model have been supported during the performance of many types of countermanding task (Logan & Cowan, 1984), including an eye movement version of the countermanding task (Hanes & Schall, 1995). Second, studies using event related potentials (DeJong et al., 1990) and single unit physiology within the frontal eye fields (Hanes et al., 1998) have also provided evidence that this assumption is valid. To test directly whether the growth of the stop process affects the growth of the go process, brain activity was compared in trials where there was a stop-signal, but it failed to inhibit the saccade, and control trials with no stop-signal. In both cases a saccade is generated to the

peripheral target, but in stop trials both the go and stop processes are racing toward their respective thresholds, whereas in control trials only the go process is active. If the stop process interfered with the rise of the go process, then the rate of growth of activity before movement initiation in stop trials should be slower than that observed before saccades in control trials. Both DeJong et al. (1990) and Hanes et al. (1998) showed that the neural activity was not different in these two cases, supporting the assumption that the stop process does not affect the go process. What happens downstream of these racing processes is of course a different matter; in the end, a 'win' by the stop process must at some point inhibit the production of a saccade by the go process. As was noted earlier, the prolongation of a very small proportion of the latencies in stop trials may well be explicable in terms of a downstream inhibitory process that on some trials slows, but does not completely abolish, the nascent saccade (Fig. 8).

Although these studies provide evidence that the stop process does not affect the rise of the go process, to show true independence one must also show the converse, that the go process does not influence the rise of the stop process. If it did, one would expect changes in the duration of go process to result in changes in the duration of the stop process. The current study has examined this question by using different target contrasts as a way of modifying the go process in isolation. The slower distribution of saccade latencies for low contrast targets than for high contrast targets implies that the go process rises more quickly for the high contrast target. Yet we have also shown that the estimated duration of the stop process, the SSRT, does not vary with target contrast. Although the estimated SSRTs were not different for high and low contrast targets, the rates of rise of the stop process used in the Monte Carlo simulations were somewhat different for the high and low contrast conditions, since these were the best-fit parameters for each subject and condition. When the rates of rise of the stop processes were constrained to be the same for the high and low contrast conditions, the results of the simulations still fit the behavioural results produced by the subjects: thus the independence assumption is not violated. Previous work using a oculomotor version of the countermanding task has shown that SSRTs are around 90 ms in monkeys (Hanes & Schall, 1995; Hanes et al., 1998). While in both previous studies using monkeys and the current study using humans the average saccade latencies during control trials were around 220–230 ms and there were only one or two possible target positions, the estimated SSRTs were approximately 40–50 ms longer in humans. The estimated SSRTs of 135 ms in the current study are similar to what has been shown before in other human studies (e.g. Lappin & Eriksen, 1966; Logan, 1982, 1983; Zbrodoff & Logan, 1986). It

seems probable that these differences in SSRTs are due to species differences. Previous studies have noted faster reaction times in general for macaque monkeys than humans (e.g. Fischer & Weber, 1993) so it is perhaps not surprising that stop-signal reaction times are also faster in monkeys. Although we hypothesise that the difference in SSRTs are due to species differences they could be due to subtle task differences. For example, the luminance and contrast of the stop signal were higher in Hanes and Schall (1995) and Hanes et al. (1998) than that used in the current study, and this may possibly lead to the differences in the SSRTs. Future work will ultimately be necessary to address this issue.

Previous investigations of gaze control have presented subjects with two target steps to probe the timecourse of saccade programming (Westheimer, 1954; Wheelless, Boynton & Cohen, 1966; Komoda, Fester, Phillips, Duckman & Young, 1973; Lisberger, Fuchs, King & Evinger, 1975; Becker & Jürgens, 1979). In fact, a comparable race model of saccade generation was formulated by Becker and Jürgens (1979), but the specific predictions and analytical procedures were not developed. In one condition of some of these double-step saccade studies, known as pulse-return, the target jumped to the peripheral location and then back to its original location. In one double-step study, Komoda et al. (1973) showed that as the delay between the target stepping away from the central fixation spot and the fixation spot reappearing was made larger, subjects increasingly failed to inhibit saccade generation to the first target step. Using the data presented in Table 1 of Komoda et al. (1973) and the assumption that the race model that we have proposed underlies behavioural performance in Komoda's task, it is possible to estimate a value analogous to stop-signal reaction time. With the use of the method of estimating the stop-signal reaction time based upon the mean of the inhibition function, we estimate that the stop-signal reaction times for Komoda's subjects to be around 120 ms. Although this value is somewhat lower than those presented in the current paper they are still reasonably similar. A double-step saccade study was also implemented by Lisberger et al. (1975). In one condition in Lisberger et al. (1975) both target steps were horizontal. The first step displaced the target to the right or left of the central fixation spot and the second step moved the target across the central fixation spot to a final position on the opposite side. It is worth noting that like Komoda et al. (1973) the initial target disappeared when the target stepped to the second location. Lisberger et al. (1975) showed that as the delay between the two target steps increased the subjects increasingly failed to inhibit saccade generation to the location of the first target step. Using the data presented in Fig. 2 of Lisberger et al. (1975) and the method of estimating the stop-signal reaction time based upon the mean of the

inhibition function we estimate that the stop-signal reaction times for Lisberger's subjects to be around 135 ms. This estimate of the duration required to cancel the movement to the first target step and program a saccade to the second target step is similar to the duration required to inhibit saccade generation shown in the current study.

Although the estimated stop-signal reaction time that we have derived from Komoda et al. (1973) and Lisberger et al. (1975) are similar to those of the current study, it is worth noting that these double step tasks and the countermanding task are significantly different. In our study, the reappearance of the fixation spot on *stop* trials represented an imperative instruction signal rather than a second target for a saccade. In fact, unlike the double-step studies, the target stimulus remained on even when the fixation spot reappeared on *stop* trials.

4.2. Future directions

With the foundation of information provided by the current and previous studies (Hanes & Schall, 1996; Hanes et al., 1998) using a saccade version of the countermanding task future work can look at the effects of systematically changing various parameters of the countermanding task. For example, does the average SSRT vary if the stop-signal is of an entirely different modality from the target, for instance an auditory tone? Does the stop-signal show the same general properties in respect of stimulus detectability as ordinary reaction times do, increasing for example when stimulus contrast is reduced? These parametric studies will provide further evidence about the relationships between gaze-holding and gaze-shifting mechanisms, which apart from being of interest in their own right can also be implemented during single unit studies in monkeys to investigate the neural signals that underlie these processes. In addition, it will be useful to determine if the countermanding task could be used as a diagnostic tool for diagnosing various neurological disorders that affect the initiation and inhibition of movements such as Parkinson's disease and Huntington's disease, and also in patients with cognitive or attentional disorders, perhaps especially in relation to frontal lobe impairment. One of the benefits of this approach from the clinical point of view is that saccades can be recorded non-invasively and automatically under computer control, generating relatively large amounts of data in a short space of time. In fact, investigators are currently testing the feasibility of using the countermanding task as a tool for monitoring levels of anesthesia (Khan, Taylor, Swart & Jones, 1998).

Acknowledgements

We thank the subjects AC and DM for their time and J. Schall for helpful comments on the manuscript. Supported in part by a Dissertation Enhancement Award from Vanderbilt University and the National Institute of Mental Health Grants F31-MH11178 to D. Hanes and RO1-MH55806 to J. Schall. D. Hanes present address: Laboratory of Sensorimotor Research, Building 49 room 2A50, National Eye Institute, NIH, Bethesda, MD 20892.

References

- Abramowitz, M., & Stegun, I. A. (1965). *Handbook of mathematical functions*. New York: Dover.
- Becker, W., & Jürgens, R. (1979). An analysis of the saccadic system by means of double step stimuli. *Vision Research*, *19*, 967–983.
- Carpenter, R. H. S. (1981). Oculomotor procrastination. In D. F. Fisher, R. A. Monty, & J. W. Senders, *Eye movements: cognition and visual perception* (pp. 237–246). Hillsdale: Erlbaum.
- Carpenter, R. H. S. (1988). *Movements of the eyes* (2nd ed.). London: Pion.
- Carpenter, R. H. S. (1994). SPIC: a PC-based system for rapid measurement of saccadic responses. *Journal of Physiology (Proceedings)*, *480*, 4P.
- Carpenter, R. H. S., & Williams, M. L. L. (1995). Neural computation of log likelihood in the control of saccadic eye movements. *Nature*, *377*, 59–62.
- DeJong, R., Coles, M. G. H., Logan, G. D., & Gratton, G. (1990). In search of the point of no return: the control of response processes. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 164–182.
- DeJong, R., Coles, M. G. H., & Logan, G. D. (1995). Strategies and mechanisms in nonselective and selective inhibitory motor control. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 498–511.
- Fischer, B., & Weber, H. (1993). Express saccades and visual attention. *Behavioral and Brain Sciences*, *16*, 588–589.
- Goldberg, M. E., & Wurtz, R. H. (1972). Activity of superior colliculus in behaving monkeys. I. Visual receptive fields of single neurons. *Journal of Neurophysiology*, *35*, 542–559.
- Hanes, D. P., & Carpenter, R. H. S. (1997). Countermanding saccades in humans. *Society of Neuroscience Abstracts*, *23*, 757.
- Hanes, D. P., & Paré, M. (1998). Neural control of saccade production studied with the countermanding paradigm: superior colliculus. *Society of Neuroscience Abstracts*, *24*, 418.
- Hanes, D. P., & Schall, J. D. (1995). Countermanding saccades in macaque. *Visual Neuroscience*, *12*, 929–937.
- Hanes, D. P., & Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science*, *274*, 427–430.
- Hanes, D. P., Patterson, W. F., & Schall, J. D. (1998). The role of the frontal eye fields in countermanding saccades: visual, movement, and fixation activity. *Journal of Neurophysiology*, *79*, 817–834.
- Khan, O. A., Taylor, S. R. J., Swart, M., & Jones, J. G. (1998). The effects of low doses of isoflurane on inhibition of saccadic eye movement. *British Journal of Anaesthesia*, *79*, 675P.
- Komoda, M. K., Festinger, L., Phillips, L. J., Duckman, R. H., & Young, R. A. (1973). Some observations concerning saccadic eye movements. *Vision Research*, *13*, 1009–1020.
- Lappin, J. S., & Eriksen, C. W. (1966). Use of a delayed signal to stop a visual reaction-time response. *Journal of Experimental Psychology*, *72*, 805–811.

- Lisberger, S. G., Fuchs, A. F., King, W. M., & Evinger, L. C. (1975). Effect of mean reaction time on saccadic responses to two-step stimuli with horizontal and vertical components. *Vision Research*, *15*, 1021–1025.
- Logan, G. D. (1981). Attention, automaticity, and the ability to stop a speeded choice response. In J. Long, & A. D. Baddeley, *Attention and performance*, vol. IX (pp. 205–222). Hillsdale: Erlbaum.
- Logan, G. D. (1982). On the ability to inhibit complex movements: a stop-signal study of typewriting. *Journal of Experimental Psychology: Human Perception and Performance*, *8*, 778–792.
- Logan, G. D. (1983). On the ability to inhibit simple thoughts and actions: I. Stop-signal studies of decision and memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *9*, 585–606.
- Logan, G. D. (1994). On the ability to inhibit thought and action: a users' guide to the stop-signal paradigm. In D. Dagenbach, & T. H. Carr, *Inhibitory processes in attention, memory and language* (pp. 189–239). Kent: Academic Press.
- Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: a theory of an act of control. *Psychological Review*, *91*, 295–327.
- Munoz, D. P., & Wurtz, R. H. (1992). Role of the rostral superior colliculus in active visual fixation and execution of express saccades. *Journal of Neurophysiology*, *67*, 1000–1002.
- Munoz, D. P., & Wurtz, R. H. (1993a). Fixation cells in monkey superior colliculus I. Characteristics of cell discharge. *Journal of Neurophysiology*, *70*, 559–575.
- Munoz, D. P., & Wurtz, R. H. (1993b). Fixation cells in monkey superior colliculus II. Reversible activation and deactivation. *Journal of Neurophysiology*, *70*, 576–589.
- Munoz, D. P., & Wurtz, R. H. (1995). Saccade-related activity in monkey superior colliculus. I. Characteristics of burst and buildup cells. *Journal of Neurophysiology*, *73*, 2313–2333.
- Osman, A., Kornblum, S., & Meyer, D. E. (1986). The point of no return in choice reaction time: controlled and ballistic stages of response preparation. *Journal of Experimental Psychology: Human Perception and Performance*, *12*, 243–258.
- Osman, A., Kornblum, S., & Meyer, D. E. (1990). Does motor programming necessitate response execution? *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 183–198.
- Patterson, W. F., & Schall, J. D. (1997). Supplementary eye field studied with the countermanding paradigm. *Society for Neuroscience Abstracts*, *23*, 474.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.
- Schall, J. D. (1997). Visuomotor areas of the frontal lobes. In K. S. Rockland, A. Peters, & J. H. Kaas, *Cerebral cortex: extrastriate cortex of primates*, vol. 14. New York: Plenum Press.
- Thompson, K. G., Bichot, N. P., & Schall, J. D. (1997). Dissociation of visual target selection from saccade programming in macaque frontal eye field. *Journal of Neurophysiology*, *77*, 1046–1050.
- Vince, M. A. (1948). The intermittency of control movements and the psychological refractory period. *British Journal of Psychology*, *38*, 149–157.
- Westheimer, G. (1954). Eye movement responses to a horizontally moving visual stimulus. *Archives of Ophthalmology*, *52*, 932–941.
- Wheless Jr., L. L., Boynton, R. M., & Cohen, G. H. (1966). Eye-movement responses to step and pulse-step stimuli. *Journal of the Optical Society of America*, *56*, 956–960.
- Zbrodoff, N. J., & Logan, G. D. (1986). On the autonomy of mental processes: a case study of arithmetic. *Journal of Experimental Psychology: General*, *115*, 118–130.