

On the Ability to Inhibit Thought and Action: A Theory of an Act of Control

Gordon D. Logan

University of British Columbia, Vancouver,
British Columbia, Canada

William B. Cowan

National Research Council of Canada, Ottawa,
Ontario, Canada

Many situations require people to stop or change their current thoughts and actions. We present a theory of the inhibition of thought and action to account for people's performance in such situations. The theory proposes that a control signal, such as an external stop signal or an error during performance, starts a *stopping process* that races against the processes underlying ongoing thought and action. If the stopping process wins, thought and action are inhibited; if the ongoing process wins, thought and action run on to completion. We develop the theory formally to account for many aspects of performance in situations with explicit stop signals, and we apply it to several sets of data. We discuss the relation between response inhibition and other acts of control in motor performance and in cognition, and we consider how our theory relates to current thinking about attentional control and automaticity.

Thought and action are important to survival primarily because they can be controlled; that is, they can be directed toward the fulfillment of a person's goals. Control has been a central issue in the study of motor behavior since the turn of the century (e.g., Sherrington, 1906; Woodworth, 1899; see Gallistel, 1980, for a review), and it has been important in psychology since K. J. W. Craik's seminal papers in 1947 and 1948. Students of motor behavior have not forgotten the importance of control and have developed sophisticated theories that integrate behavioral and physiological data (e.g., Feldman, 1981; Kelso & Holt, 1980; Navas & Stark, 1968; Robinson, 1973; Young & Stark, 1963). However, psychologists have strayed from the path somewhat over the years.

Craik's papers, which described the human performer as an engineering system, provided a framework in which to study tracking tasks and stimulated interest in the (possibly intermittent) nature of the control system in such tasks. This approach kindled interest in the psychological refractory period (e.g., Hick,

1949; Vince, 1948), which led to the formulation of single-channel theory (Davis, 1957; Welford, 1952). In the hands of Broadbent (1958) and others, single-channel theory was extended to deal with many diverse phenomena of attention, and dominated theories of attention for nearly 20 years. The extended single-channel theory attracted the interest of cognitive psychologists who dealt primarily with tasks other than tracking, and, in their hands, control became less important than did other issues such as memory (Norman, 1968), expectancy (LaBerge, 1973), selectivity (Treisman, 1969), and time sharing (Posner & Boies, 1971). Single-channel theory was replaced by capacity theory (Kahneman, 1973) and multiple-resource theory (Navon & Gopher, 1979), and little attention was paid to problems of control (but see Broadbent, 1977; Reason & Myceilska, 1982; Shallice, 1972; more generally, see Gallistel, 1980; Kimble & Perlmutter, 1970; Miller, Galanter, & Pribram, 1960; Powers, 1978).

Recently, cognitive psychologists have become interested in control once more, in the guise of research on automaticity and skill (e.g., Anderson, 1982; Hasher & Zacks, 1979; Logan, 1978; Posner, 1978; Shiffrin & Schneider, 1977), but the studies bear little resemblance to the early fruits of Craik's seminal thinking and even less resemblance to studies of motor behavior. Whereas the earlier studies in Craik's

This research was supported by Grant U0053 from the Natural Sciences and Engineering Research Council of Canada to Gordon D. Logan.

Requests for reprints should be sent to Gordon D. Logan, who is now at the Department of Psychological Sciences, Purdue University, West Lafayette, Indiana 47907.

tradition and the studies of motor behavior focused on how a process is controlled and on what variables are adjusted to achieve control, modern studies of automaticity are often concerned with determining whether a process is controlled and what classes of processes might be automatic.

In this article, we attempt to bridge the gap between studies of motor control and studies of cognitive control by presenting a framework within which control can be understood from either perspective and by developing a formal model of a simple act of control that is relevant to both domains. We believe that control can be understood in terms of the interaction between an *executive* system that forms intentions and issues commands to realize the intentions, and a *subordinate* system that interprets the commands and carries them out. The behavior of the executive in this interaction can be viewed as a series of *acts of control*, which can be analyzed behaviorally.

There are probably many different acts of control, each designed to solve a particular problem. The act of control that we chose to study is the ability to inhibit thoughts and actions when they are no longer relevant to the current goals. Inhibition is not likely to be the only act of control in the executive's repertoire, but it is important in motor control (e.g., stopping ourselves from batting a baseball pitched outside the strike zone) and in cognitive control as well (e.g., suppressing an irrelevant thought or an inappropriate remark). So it seemed to be a good place to begin bridging the gap between the two domains. Our investigation focused on three issues that concern both domains.

1. *Measuring the difficulty of control.* It is important to determine whether one situation is easier or harder to control than another and to determine when two situations are equally difficult to control. The theory provides a rationale and methods for comparing the ability to control in different subjects, strategies, conditions, or tasks.

2. *Measuring the latency of control.* An act of control, like any other act, must take time. The theory provides methods for measuring the latency of control even when the act of control is not directly observable.

3. *Measuring the ballistic component of the process being controlled.* Some components of

a process may be ballistic in the sense that they cannot be inhibited once they begin; they must run on to completion before they stop. The theory suggests ways to detect the presence of ballistic processes and to measure their duration.

The plan of this article is to introduce the *stop-signal procedure* as a paradigm for studying control, and then to present the theory and its predictions. The theory is shown to account for the basic findings in the stop-signal paradigm. It is developed further to suggest ways to measure the difficulty and the latency of control and to estimate the duration of the ballistic component. Deeper predictions that test the basic assumptions of the theory are then developed and tested. Finally, we suggest ways of extracting more information from limited data by theory-based extrapolation. The theory is then generalized to other tasks and paradigms, and common themes running through the literature on control are discussed in terms of the theory.

The Stop-Signal Paradigm

The theory focuses on the stop-signal paradigm, in which subjects are given a primary task to perform and, on occasion, a stop signal is presented that tells them not to respond on that trial. Many investigators have studied the stop-signal paradigm for various reasons. Students of the psychological refractory period have found that responses to signals to stop a movement or to change some of its parameters may not be subject to the same refractory delays as are responses to other types of signals (e.g., Bartlett, Eason, & White, 1961; Brebner, 1968; Gottsdanker, 1966; Harrison, 1960; Slater-Hammel, 1960; Vince & Welford, 1967; Welford, 1952), a finding which has been important in shaping current theories of attention. Some investigators have used stop signals to learn about catch trials in reaction time tasks (e.g., Grice, Hunt, Kushner, & Morrow, 1974; Grice, Hunt, Kushner, & Nullmeyer, 1976; LaBerge, 1971; Ollman, 1973), and still others have been interested in stopping for its own sake (e.g., Ladefoged, Silverstein, & Pappun, 1973; Lappin & Eriksen, 1966; Logan, 1981, 1982, 1983; Logan, Cowan, & Davis, 1984). Although a fair amount is known in each of these domains, there is no general the-

ory of the acts of control they require, and it is not obvious how the specific treatments would generalize beyond their own limited domains. This article is intended to provide such a theory.

The major dependent variable in the stop-signal paradigm is the probability of inhibiting the response to the primary task or, equivalently, the probability of responding to the primary task given that a signal has occurred. We are also concerned with the latency of responses when no stop signal is presented (*no-signal* reaction times), with the latency of responses that escape inhibition (*signal-respond* reaction times), and with the latency of the internal response to the stop signal (*stop-signal* reaction times).

We illustrate the stop-signal paradigm and describe its major findings by describing an experiment we performed to obtain detailed data for analysis in later sections of this article. Some of the major findings can be seen in the data directly; others cannot be seen without a model or theory to guide interpretation. We present results until no more can be elucidated without the model, and, at that point, we introduce the model, apply it to the basic findings, and develop it to reveal the less perspicuous findings that are nonetheless important.

Method

Subjects. Three adult subjects performed the task for six 1-hour sessions each. All of the subjects had extensive experience with the stopping task prior to the experiment.

Apparatus and stimuli. The stimuli for the choice reaction time task were the letters A–X, printed in capitals by illuminating about 20 points in a 5×7 matrix on a computer-controlled cathode-ray tube. Viewed at a distance of 60 cm, each letter subtended $.43^\circ \times .57^\circ$ of visual angle.

Each trial began with a fixation point illuminated for 500 ms. It was then extinguished and followed by the letter for that trial, which was exposed for 500 ms. After the letter, the screen was blank for a 2-s intertrial interval. Subjects responded to the letter by pressing one of two telegraph keys with the index and middle fingers of their right hands.

The stop signal was a 500-ms 900-Hz tone presented through a speaker behind the cathode-ray tube at a comfortable listening level. When a stop signal occurred, it occurred at one of 10 delays, 50, 100, 150, 200, 250, 300, 350, 400, 450, or 500 ms after the onset of the letter for that trial.

Procedure. The choice task involved classifying single letters drawn from a set of four. Two of the letters were assigned to one response and two were assigned to the other response. A different set of four letters was used each session to keep reaction time relatively constant. Two

subjects had the first four letters of the alphabet on the first session, the second four on the second session, and so on (i.e., ABCD, EFGH, IJKL, etc.), and one subject had the opposite (i.e., UVWX, QRST, MNOP, etc.).

Each letter was presented 200 times in a session, for a total of 800 trials. The stop signal occurred on 25% of the trials (i.e., 200 times per session), equally often at each delay (i.e., 20 times per delay per session), and equally often with each letter at each delay (i.e., 5 times per letter per delay per session). Within these constraints, the order of letters, stop signals, and stop-signal delays was random. A separate random order was prepared for each subject each session.

The first session was considered practice, and data from it were not analyzed. Data for each subject were collapsed across the last five sessions and entered into the analyses described below.

Results and Discussion

Inhibition functions. The probability of responding when a stop signal occurred is presented for each subject as a function of delay in Figure 1. There are 10 points for each subject, and each point is based on 100 observations.

The probability of responding given a stop signal increased with stop-signal delay for each subject. This finding is typical of the literature: Similar results have been found in studies of eye movements (Lisberger, Fuchs, King, & Evinger, 1975), typewriting (Logan, 1982), simple reaction time (Lappin & Eriksen, 1966; Ollman, 1973; Logan, Cowan, & Davis, 1984), choice reaction time (Logan, 1981, 1983; Logan et al., 1984), arm movements (Henry & Harrison, 1961), and anticipatory responses (Slater-Hammel, 1960). We are concerned with the form of the inhibition function and how it depends on parameters of the primary-task reaction time distribution and the stop-signal reaction time distribution.

Primary-task reaction time. The mean and standard deviation of the reaction times and the error rate from the trials on which no stop signal was presented are displayed in Table 1. Each datum in the table is based on 3,000 observations.

Subjects differed in primary-task reaction time, and the differences were reflected in the inhibition functions: At a given delay, the probability of responding given a stop signal was generally higher for subjects with faster primary-task reaction times. This is also typical of findings in the stop-signal paradigm. In general, the probability of responding given




Figure 1. The probability of responding given a stop signal— $P(\text{respond}|\text{signal})$ —as a function of stop-signal delay for subjects J. M., G. L., and J. C.

a stop signal is higher the faster the primary-task reaction time whether reaction time varies due to differences between individuals (Lisberger et al., 1975), differences between strategies (Lappin & Eriksen, 1966; Logan, 1981; Ollman, 1973), differences between conditions (Logan, 1981, 1982, 1983), or differences between tasks (Logan et al., 1984). The model shows that much of the differences between inhibition functions can be accounted for by differences in primary-task reaction time.

Signal-respond reaction times. Reaction times from trials on which a signal occurred but subjects responded anyway are presented in Table 2. In each subject's data, signal-re-

spond reaction times were faster than no-signal reaction times, and signal-respond reaction times increased with stop-signal delay. Both of these findings are consistent with the literature, whenever signal-respond reaction times have been reported (Lappin & Eriksen, 1966; Lisberger et al., 1975; Logan, 1981, 1983; Logan et al., 1984). These effects, too, follow from the model.

The results presented so far are all that can be extracted from the data without a model or theory to guide interpretation. We now turn to the model and apply it to the data.

The Horse-Race Model

The model accounts for response inhibition in the stop-signal paradigm in terms of a "horse race" between two sets of processes, one that generates a response for the primary task and one that responds to the stop signal: If the primary-task process finishes before the stop-signal process, the response is executed; if the stop-signal process finishes before the primary-task process, the response is inhibited. To model this situation, the finishing times of the primary-task process and the stopping process

Table 1
Mean Reaction Time, Standard Deviation, and Proportion of Errors From Trials on Which No Stop Signal Occurred

Measure	Subject		
	J. M.	G. L.	J. C.
Mean RT	482	452	378
SD of RT	110	50	58
P (error)	.01	.04	.02

Table 2
Mean Reaction Times From Trials on Which Stop Signal Occurred but the Subject Responded (Signal-Respond Trials)

Condition	Stop-signal delay										Mean
	50	100	150	200	250	300	350	400	450	500	
J. M.											
Observed	617	—	365	356	371	397	415	444	441	455	434
Predicted	315	—	305	315	361	382	420	442	464	461	434
G. L.											
Observed	391	438	428	378	423	447	440	457	454	451	445
Predicted	312	325	340	378	418	438	449	449	450	452	441
J. C.											
Observed	337	351	362	349	372	384	375	371	380	390	374
Predicted	293	302	315	347	359	372	376	378	378	378	366

Note. Mean = mean weighted by frequency of occurrence.

are assumed to be independent random variables, and this complicates the situation. Response inhibition is stochastic, reflecting the probability that the stopping process finishes before the primary-task process. The model developed below spells out the implications of these relatively simple assumptions, showing how they account for factors affecting the probability of inhibition and how they can be exploited to estimate parameters of the stopping process and the ballistic component of the primary-task process.

The basic predictions of the model can be seen clearly in the situations depicted in Figure 2. The top panel shows a hypothetical distribution of primary-task reaction times, the onset of the stimulus for the primary task, the onset of the stop signal, and the point in time at which the internal response to the stop signal occurs. For the present, to simplify the presentation, we talk as if the latency of the response to the stop signal does not vary and as if there were no ballistic phase. Later, we show that the picture does not change substantially when these factors are included in the model.

According to the horse-race model, response inhibition depends on the relative finishing times of the stopping process and the primary-task process. This is represented in the top panel of Figure 2 by the vertical line extending upward from the point at which the response

to the stop signal occurs. On the left side of the line, the response to the primary task is faster than the response to the stop signal, and the subject responds. The area to the left of the line represents the probability of responding to the primary task given a stop signal. On the right side of the line, the response to the stop signal is faster than the response to the primary task, and the subject inhibits his or her response. The area to the right of the line represents the probability of inhibition. Most of the variance in the probability of inhibition can be accounted for in terms of relative finishing time, as we see below.

Inhibition Functions

The model accounts for the inhibition functions by asserting that stop-signal delay biases the relative finishing time of the stopping process and the primary-task process, handicapping the race in favor of one process or the other. The effect can be seen in Figure 2 by comparing the top panel with the second panel. In the second panel, the stop signal has been delayed, relative to the top panel, so the response to the stop signal cuts off more of the primary-task reaction time distribution, reducing the probability of inhibition and increasing the probability of responding given a signal. It is easy to see that if the stop signal

occurs early enough, the response is always inhibited, and that if the stop signal occurs late enough, the response is always executed. The points in between trace out the inhibition function.

To lay a groundwork for the later analyses, let us formalize this account. Let T be a random variable representing the finishing time of the primary-task process. T is distributed by the distribution $f(t)$, so that




Figure 2. Graphic representation of the assumptions and predictions of the horse-race model, indicating how the probability of inhibiting a response— $P(\text{inhibit})$ —and the probability of responding given a stop signal— $P(\text{respond})$ —depend on the distribution of primary-task reaction times, stop-signal reaction time, and stop-signal delay.

$$1 = \int_{-\infty}^{\infty} f(t)dt.^1$$

It has a mean \bar{T} and variance σ^2 given by

$$\bar{T} = \int_{-\infty}^{\infty} tf(t)dt \quad (1)$$

and

$$\sigma^2 = \int_{-\infty}^{\infty} (t - \bar{T})^2 f(t)dt.$$

Now, let t_d be the stop-signal delay and let t_s be the time between the stop signal and the inhibition of the response (i.e., stop-signal reaction time). Here, t_s is taken to have a fixed value that is independent of t_d (but see Appendix). A response occurs if the random variable T is less than the sum $t_d + t_s$. Thus, $P_r(t_d)$, the probability that a response occurs given a stop signal, depends on the stop-signal delay, t_d , in the following manner:

$$P_r(t_d) = \int_{-\infty}^{t_s+t_d} f(t)dt. \quad (2)$$

Note that varying t_d produces an inhibition function corresponding to the cumulative distribution of $f(t)$.

Primary-Task Reaction Time

The model also accounts for the effects of primary-task reaction time in terms of biasing the relative finishing times of the stopping process and the response-generation process. This can be seen in Figure 2 by comparing the top panel with the third panel. In the third panel, primary-task reaction time is delayed relative to the top panel, so a stop signal at the same delay cuts off less of the distribution in the third panel than in the top panel, increasing the probability of inhibition and decreasing the probability of responding given a stop signal. It is easy to see that for any given delay, the probability of inhibition is higher and the probability of responding given a signal is lower the longer the primary-task reaction time.

Signal-Respond Reaction Time

The model also predicts that signal-respond reaction times will generally be faster than no-signal reaction times, and will increase with stop-signal delay. These effects can be seen in

Figure 2. The mean signal-respond reaction time represents the mean of the portion of the primary-task reaction time distribution that is to the left of the line representing the internal response to the stop signal. It is necessarily faster than the mean of the no-signal reaction times because it excludes the longer tail of the primary task distribution, whereas the no-signal mean includes the whole distribution. The increase in signal-respond reaction time with delay occurs because the internal response to the stop signal cuts off more of the primary-task reaction time distribution as delay increases, including progressively longer reaction times in the calculation of the mean.

More formally, the mean signal-respond reaction time, \bar{T}_r , is

$$\bar{T}_r = \frac{1}{P_r(t_d)} \int_{-\infty}^{t_s+t_d} tf(t)dt. \quad (3)$$

It is necessarily shorter than the mean no-signal reaction time,

$$\begin{aligned} \bar{T} &= P_r(t_d)\bar{T}_r + \int_{t_s+t_d}^{\infty} tf(t)dt \\ &> P_r(t_d)\bar{T}_r + \bar{T}_r \int_{t_s+t_d}^{\infty} f(t)dt \\ &> \bar{T}_r \end{aligned}$$

because $t_s + t_d$ is always greater than \bar{T}_r . Moreover, as t_d increases, so will \bar{T}_r , approaching the mean no-signal reaction time in the limit.

The expression for the mean signal-respond reaction time represents a prediction from the model that can be tested against our data. Given the observed no-signal reaction time distribution and the probability of responding given a signal, we can estimate the mean signal-respond reaction time for any given delay by rank ordering the no-signal reaction times and taking the mean of the fastest n of them, where n is the number of reaction times in the empirical distribution multiplied by the probability of responding given a signal at that delay.

We calculated the predictions at each delay for each subject. They are presented in Table

¹ Note that the lower limit on this integral—as with all other integrals of reaction time distributions—is taken to be negative infinity. This was done to make the equations less complicated. The negative reaction times it allows may either be taken as resulting from anticipatory responses or be excluded by stipulating that $f(t) = 0$ for $t < 0$.

2 with the observed values. Visual inspection suggests a close match for subjects J. M. and G. L. for delays with more than 10 reaction times going into the mean (delays of 250 ms or more for J. M. and 200 ms or more for G. L.). The last column in the table shows the means across delay, weighted by the number of observations at each delay. Again, J. M. and G. L. show a close match between prediction and observation. For subject J. C., the predicted reaction times underestimated the observed ones at all but two delays, and the difference remained, although somewhat attenuated, in the means weighted by the number of observations at each delay. The underestimation could be due to variance in the stop-signal reaction times, or to a tendency to pull responses when the signal occurred (cf. Rabbit, 1978), resulting in reaction times that are less forceful and slower than they should be. Note, however, that the discrepancy between prediction and observation is not large. Taken together with the good agreement in subjects J. M. and G. L., the results are encouraging.

Stop-Signal Reaction Time

The model also suggests a way to estimate the reaction time of the internal response to the stop signal at any given delay. All that we need to know is the probability of responding given a stop signal and the distribution of no-signal reaction times. We can then integrate the no-signal reaction time distribution from negative infinity until the integral equals the probability of responding given a stop signal. The value of the abscissa at that point represents the point in time at which the internal response to the stop signal occurred (i.e., $t_s + t_d$). Subtracting out stop-signal delay, t_d , yields stop-signal reaction time, t_s . In Figure 2, this amounts to moving a vertical line across the primary-task reaction time distribution until the area to the left of the line equals the probability of responding given a signal, then reading off the value of the time axis, and then subtracting out the stop-signal delay.

More formally, we can estimate stop-signal reaction time as $r(t_d) - t_d$, where $r(t_d)$ is defined as the point on the cumulative no-signal reaction time distribution that is equal to $P_r(t_d)$. That is

$$P_r(t_d) = \int_{-\infty}^{r(t_d)} f(t) dt. \quad (4)$$

This estimate assumes that stop-signal reaction time is a constant. Later we treat it as a random variable and show that the estimate given here is a good first approximation to the mean of the random variable (see Appendix). Accepting the arguments in the Appendix, we can use Equation 4 to calculate estimates from the data. The procedure is relatively simple: No-signal reaction times are rank ordered, and then the time at which the response to the stop signal occurred is set equal to the n th reaction time, where n is the number of reaction times in the empirical no-signal distribution multiplied by the probability of responding given a signal at that delay. We then subtract out stop-signal delay and are left with the estimate of stop-signal reaction time.

Stop-signal reaction times were estimated in this manner at each delay for each subject. They are displayed in Table 3. The estimates ranged from 150 ms to 300 ms, with a mean close to 200 ms. These values are about what might be expected for simple reaction time to a tone presented by itself (Woodworth & Schlosberg, 1954), and are much faster than simple reaction time to a tone presented concurrently with another task (Logan, Zbrodoff, & Fostey, 1983; Posner & Boies, 1971). Together, these contrasts suggest that there was no interference between the stopping process and the primary-task process; stop-signal reaction time was more like single-task reaction time than like dual-task reaction time. This is important because it is consistent with the assumption of the model that the stopping process and the primary-task process are independent.

It is also worth noting that the stop-signal reaction times observed here were consistent with estimates from a wide variety of tasks (see Table 6 on p. 316), including eye movements (Lisberger et al., 1975), typewriting (Logan, 1982), speech (Ladefoged et al., 1973), and step tracking (Vince, 1948; Vince & Welford, 1967). The different tasks must tap a very general control process.

Note that the estimated stop-signal reaction times tend to decrease as stop-signal delay increases. The decrease is not consistent with the current version of the model, which assumes that t_s is constant. However, it need not be inconsistent with the general assumption that stopping processes and primary-task processes are independent.

Table 3
Reaction Time to Stop Signal for Each Subject

Subject	Stop-signal delay										Mean
	50	100	150	200	250	300	350	400	450	500	
J. M.	286	—	173	136	153	135	148	151	216	154	172
G. L.	278	246	213	207	209	200	209	159	130	—	206
J. C.	258	222	193	202	182	190	189	—	—	—	205

The decrease in stop-signal reaction time with delay may be a consequence of variability in stop-signal reaction time: If stop-signal reaction time had a constant mean and a nonzero variance, then different portions of its distribution would produce response inhibition at different delays. At the short delays, nearly all stop-signal reaction times would be faster than the primary-task reaction times, whereas at the long delays, only the fastest stop-signal reaction times would be faster than the primary-task reaction times. As a result, the stop-signal reaction times that produced response inhibition would be longer at the short delays than at the long delays, just as we observed. However, this does not prove that the mean stop-signal reaction time remained constant.² In later sections, we estimate limits on the variance of the stop-signal reaction times and we relate the delay effects in the stop-signal reaction times to delay effects in the probability of inhibition and the signal-respond reaction times. We find evidence that stop-signal reaction times may have substantial variance, and we show that, within limits, the delay effects conform to the relation specified by the model.

Inhibition Functions and Primary-Task Reaction Time

The model predicts that the effects of stop-signal delay and of primary-task reaction time can compensate for each other. This can be seen in Figure 2 by comparing the top panel with the fourth panel. In the fourth panel, primary-task reaction time is delayed relative to the top panel, but the stop signal has been delayed by a corresponding amount so that the response to the stop signal cuts off the same proportion of the distribution in the two panels, leaving the probability of inhibition and the probability of responding given a stop signal unchanged.

In general, the probability of responding given a signal should not change despite changes in primary-task reaction time, stop-signal delay, and so on, as long as the response to the stop signal occurs at the same point on the primary-task reaction time distribution in the different conditions. According to the current version of the model, the point on the distribution at which the response to the stop signal occurs depends on four factors: stop-signal delay, t_d ; stop-signal reaction time, t_s ; and the mean and standard deviation of the primary-task reaction times, \bar{T} and σ . That point can be expressed as a Z score:

$$ZRFT = \frac{\bar{T} - t_d - t_s}{\sigma}$$

It should be possible to account for differences in inhibition functions due to variation in parameters of the primary-task reaction time distribution by plotting the probability of responding given a signal as a function of ZRFT. The data from the experiment reported here are plotted as a function of ZRFT in Figure 3, using the method described above to estimate t_s . The estimates of t_s were averaged across delay, and the average value was used

² The decrease in stop-signal reaction time with delay may result from the fact that the stop signal "aged" or became more likely to occur as delay increased. Subjects may have responded faster to the more probable stop signals, producing the observed decrease with delay. However, the aging effect was slight. The probability that a signal would occur, given that it had not occurred already, was .025, .026, .026, .027, .028, .029, .029, .030, .031, and .032 for the 50, 100, 150, 200, 250, 300, 350, 400, 450, and 500 ms delays, respectively. The difference in probability between the most likely signal (at the 500 ms delay) and the least likely signal (at the 50 ms delay) was .007, a difference that is far too small to affect reaction time. We feel confident that the decrease in stop-signal reaction time with delay was not a consequence of aging stop-signal delays.

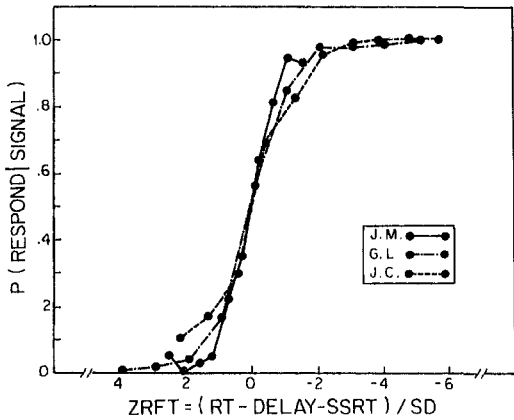


Figure 3. The probability of responding given a stop signal— $P(\text{respond}|\text{signal})$ —as a function of the relative finishing times of the primary-task process and the stopping process, expressed as a Z score, for subjects J. M., G. L., and J. C. (RT = primary-task reaction time; SSRT = stop-signal reaction time; SD = standard deviation of primary-task reaction time.)

to calculate each value of $ZRFT$. Clearly, most of the differences in the inhibition functions for the three subjects were absorbed by $ZRFT$.

The model predicts that inhibition functions should be aligned when plotted against $ZRFT$, but in practice, approximations to $ZRFT$ have been sufficient. In studies of choice reaction

time (Logan, 1981, 1983), eye movements (Lisberger et al., 1975), and typewriting (Logan, 1982), primary-task reaction time minus stop-signal delay (RT - delay) has been sufficient to bring into alignment inhibition functions from different subjects or conditions. The eye-movement and typing data are presented in Figures 4 and 5, respectively, to illustrate how effectively RT - delay accounts for differences in the inhibition functions.

Logan et al. (1984) plotted inhibition functions from simple and choice reaction time tasks against various transformations of delay to see what was necessary to bring the functions into alignment. The data from their 8 practiced subjects are plotted in Figure 6. The figure shows that there were large differences between tasks when the inhibition functions were plotted against delay, and that most of the differences were removed by plotting the functions against RT - delay. Some of the points from the simple task remained outside the range of the other points (i.e., five points from the simple task are to the right of the main cluster of points), suggesting that no-signal reaction time does not account for all of the differences between inhibition functions. The simple task and the choice task differed in variance as well as mean no-signal reaction

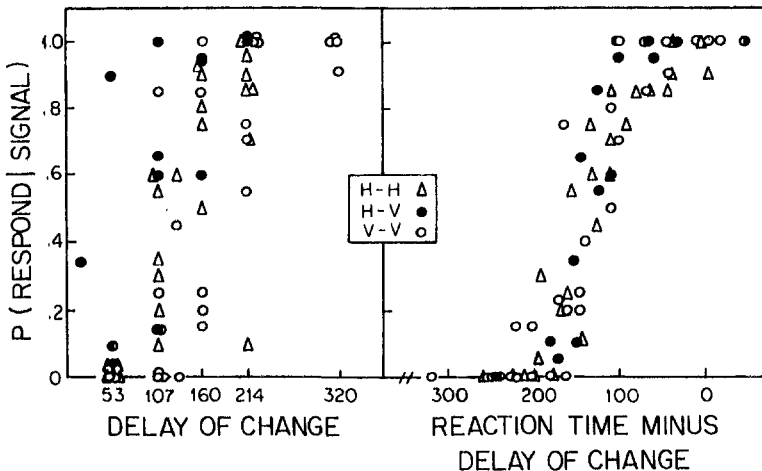


Figure 4. Probability of a saccadic response to the first of two successive fixation stimuli— $P(\text{respond}|\text{signal})$ —in a saccadic step-tracking task, expressed as a function of the delay of the second stimulus in the left panel and as a function of primary-task reaction time minus delay in the right panel. (H-H = horizontal movement followed by a horizontal movement; H-V = horizontal followed by vertical; V-V = vertical followed by vertical. From "Effect of Mean Reaction Time on Saccadic Responses to Two-Step Stimuli With Horizontal and Vertical Components" by S. G. Lisberger, A. F. Fuchs, W. M. King, and L. C. Evinger, *Vision Research*, 1975, 15, 1021-1025.)

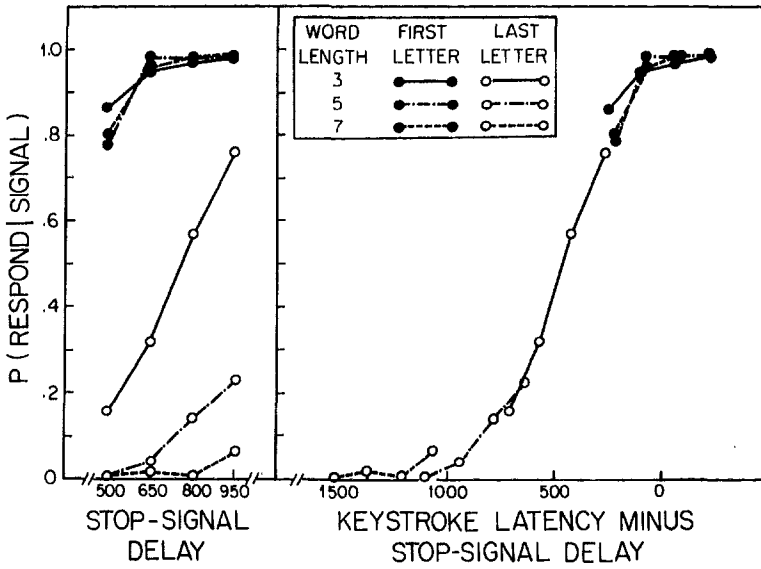


Figure 5. Probability of typing the first and last letters of a word given a stop signal— $P(\text{respond}|\text{signal})$ —as a function of stop-signal delay (left panel) and as a function of keystroke latency minus stop-signal delay (right panel). (Word length, 3, 5, and 7 letters is the parameter. From “On the Ability to Inhibit Complex Movements: A Stop-Signal Study of Typewriting” by G. D. Logan, *Journal of Experimental Psychology: Human Perception and Performance*, 1982, 8, 778–792. Copyright 1984 by the American Psychological Association, Inc.)

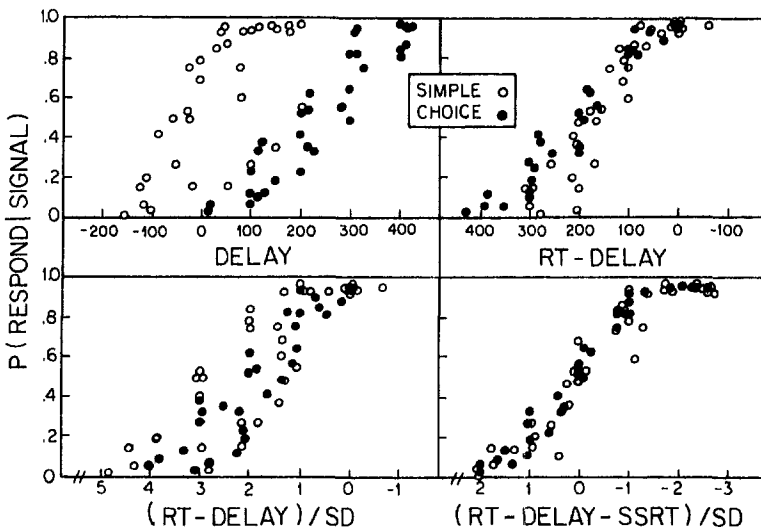


Figure 6. Probability of responding to a choice or simple reaction time task given a stop signal— $P(\text{respond}|\text{signal})$ —as a function of various transformations of stop-signal delay for eight practiced subjects ($RT = \text{primary-task reaction time}$; $SSRT = \text{stop-signal reaction time}$; $SD = \text{standard deviation of primary-task reaction time}$). From “On the Ability to Inhibit Simple and Choice Reaction Time Responses: A Model and a Method” by G. D. Logan, W. B. Cowan, and K. A. Davis, *Journal of Experimental Psychology: Human Perception and Performance*, 1984, 10, 276–291. Copyright 1984 by the American Psychological Association, Inc.)

time, and the discrepant points reflect the smaller variance in the simple task. Logan et al. (1984) tried to take variance into account by plotting the inhibition functions against $(RT - \text{delay})/SD$, but as Figure 6 shows, this increased the scatter of the points. The scatter increased because $(RT - \text{delay})/SD$ does not represent the relative finishing times of the two processes; it represents the interval between the beginning of the stopping task and the end of the reaction time task, and the model does not predict that the functions will be aligned when plotted against that interval. The model suggests that *ZRFT* is the best way to account for differences in variance because it represents relative finishing time, and Logan et al. (1984) found the best alignment when they plotted inhibition functions against *ZRFT* (see Figure 6). However, the alignment was not much better than the alignment produced by plotting the functions against $RT - \text{delay}$, which suggests that differences in the mean no-signal reaction time are more important than are differences in its variance.

Thus, it seems that *ZRFT* is always sufficient to bring into alignment inhibition functions from different subjects, tasks, conditions, or strategies, but it is not always necessary. When the subjects, tasks, and so on produce primary-task reaction times with similar variances, and when they do not differ much in stop-signal reaction time, $RT - \text{delay}$ is all that is necessary to bring the inhibition functions into alignment.

The idea that the model can be used to account for differences in inhibition functions (i.e., by plotting them as a function of *RFT* and *ZRFT*) is very important. Inhibition functions reflect the ability to control, and comparisons between functions allow us to compare the ability to control in different individuals, strategies, tasks, and conditions. If different inhibition functions can be brought into alignment by plotting them against *ZRFT*, then we can conclude that the situations that gave rise to them are controlled equivalently because the same control process applies to them (i.e., the horse-race model). We can then ask whether the parameters of the model are the same for the different functions. If stop-signal reaction time is the same, we can conclude that the different conditions are equally difficult to control; if stop-signal

reaction time differs between conditions, then the condition with the longer stop-signal reaction time is the more difficult one to control. It may not, in some cases, be possible to bring different inhibition functions into alignment by plotting them against *ZRFT*. In those cases, different control processes apply. Typically, one inhibition function follows the predictions of the horse-race model, and the other departs from them substantially. The departure may indicate the presence of a ballistic process (see Logan, 1982; also see the General Discussion section).

Given the importance of using the model to account for differences in inhibition functions, it is worthwhile formalizing the logic necessary to do so: A different subject, task, or condition has a different distribution of primary-task reaction times, $f'(t)$, and a different stopping time, t'_s . The probability that a response occurs, $P'_r(t_d)$, is

$$P'_r(t_d) = \int_{-\infty}^{t'_s+t_d} f'(t)dt.$$

Various transformations of t_d are possible. For example, suppose we plot P_r against $y = t_d + t_s$, and P'_r against $y' = t_d + t'_s$. Then

$$P_r(y) = \int_{-\infty}^y f(t)dt,$$

and

$$P'_r(y') = \int_{-\infty}^{y'} f'(t)dt.$$

The two graphs, P_r versus y and P'_r versus y' , are the same if and only if the reaction time distributions $f(t)$ and $f'(t)$ are the same.

More likely, the reaction time distributions will not be the same. Intuitively, we may split the difference into difference between the means, difference between the variances, and other differences in shape. If the distributions differ only in their means, then the plot of P_r against $y - \bar{T}$ and the plot of P'_r against $y' - \bar{T}'$ will be the same. The fact that inhibition functions for different subjects and conditions in eye movements (Lisberger et al., 1975) and for different conditions in typewriting (Logan, 1982) and choice reaction time (Logan, 1981) line up when plotted as a function of $y - \bar{T}$, suggests that the different conditions and subjects differed only in mean primary-task reaction time.

Two distributions may also differ in variance. If the distributions differ in mean and variance but are otherwise the same shape, then the plot of P_r against $z = (y - \bar{T})/\sigma$ and the plot of P'_r against $z' = (y' - \bar{T}')/\sigma'$ are the same. (The independent variable, *ZRFT*, discussed above, is an estimate from the data of the variables z or z' .) More precise intersubject comparisons for choice reaction time (Figure 3) and intertask comparisons between simple and choice reaction time (Logan et al., 1984; see the present Figure 6) require this transformation to have the same function for the probability of responding given a stop signal, indicating that intersubject or intertask functions have the same shape, although the mean and variance of primary-task reaction time varies between subjects and tasks.

Independence of Processes

The model assumes that the stopping process and the primary-task process are independent random variables. The assumption of independence is important because it simplifies the formal development of the model. It should not be accepted lightly, however, because in most dual-task situations the concurrent processes are not independent (e.g., Kantowitz, 1974; Welford, 1952). Here we consider the implications of the independence assumption and those of some alternatives to independence.

Independence of processes predicts that no-signal reaction times and stop-signal reaction times should not be different from single-task controls. These predictions have not been tested explicitly, but nevertheless, the no-signal reaction times observed in stop-signal studies are about as fast as reaction times in comparable single-task studies and stop-signal reaction times are about as fast as simple reaction times to similar stimuli (see e.g., Woodworth & Schlosberg, 1954). The independence assumption also predicts that signal-respond reaction times should be faster than no-signal reaction times and that they should increase with stop-signal delay. As noted above, both of these predictions have been confirmed many times. On the balance, then, the data are consistent with the independence assumption. To accept the assumption with confidence, however, we need to show that the data are inconsistent with alternatives to independence.

The first major alternative to independence is the possibility that the stopping process and the primary-task process race against each other but share limited resources, so that the two processes cannot function at the same time without interfering with one another. The pattern of interference depends on the subject's strategy for allocating resources. We consider two possibilities: First, subjects may allocate resources dynamically, dividing them among the processes that demand them, changing the allocation as new processes enter the competition. In the stop-signal task, all the resources may be allocated to the primary task at the beginning of a trial (or for the duration of a no-signal trial), but if a stop signal occurs, some resources may be taken away from the primary task and allocated to the stopping task. Under this strategy, the stopping process would be operating with less resources than it needed, so stop-signal reaction times should be elevated substantially. The primary-task process would be operating with less resources on stop-signal trials than on no-signal trials, so signal-respond reaction times should be substantially longer than no-signal reaction times. Also, the sooner the stop signal occurred, the sooner the primary task would have to operate with reduced resources, and so signal-respond reaction times should be longer the earlier the delay (i.e., signal-respond reaction times should decrease with stop-signal delay). All of these predictions have been disconfirmed by our data and by previous data (Lappin & Eriksen, 1966; Logan, 1981, 1983; Logan et al. 1984), so this alternative to independence can be safely ruled out.

Second, subjects may fix the allocation of resources to the two processes at the beginning of each trial and maintain it rigidly throughout the trial, independent of stimulus events. Under this strategy, the primary-task process would be operating with less resources than it needed on no-signal trials as well as on stop-signal trials, so there should be no difference between no-signal reaction times and signal-respond reaction times. Both should be elevated substantially above single-task controls, however. The stopping process would also be operating with less resources than it needed, which would elevate stop-signal reaction times. The comparisons between stop-signal studies and other single-task studies, though indirect, tend to rule out this possibility. Thus, resource

sharing can be ruled out as an alternative to independence; the data are accounted for better by the horse-race model, predicated on the assumption of independent processes.

Another alternative to independence is the possibility that subjects make the processes nonindependent by the strategies they choose to employ. For example, they may choose to respond or not to respond before each trial, independent of stimulus events. Under this strategy, the probability of responding given a signal would be the same as the probability of responding given no signal, and it would not vary with stop-signal delay. These predictions have been disconfirmed in every stop-signal study in the literature; subjects respond on virtually every no-signal trial and respond less often on stop-signal trials, and stop-signal delay has substantial effects on the probability of responding given a signal.

Subjects may adopt a more subtle version of this strategy that would be harder to detect: They may choose to respond regardless of stimulus events on a certain proportion of the trials and follow the horse-race model on the remaining trials. The proportion of trials on which subjects choose to respond independent of the stop signal would be the same on stop-signal trials as on no-signal trials because subjects would have no way of knowing in advance whether a stop signal was forthcoming. Thus, reaction times characteristic of responding independent of the signal would contribute equally to no-signal and signal-respond reaction times, diluting the characteristic effects of the horse race without changing their pattern. It would be difficult to reject this strategy on the basis of signal-respond and no-signal reaction times. However, the inhibition functions would be different under this strategy than under a pure horse race: Under this strategy, the probability of responding given a signal should never be as low as zero, regardless of how early stop signals are presented. If there are enough early delays to define the lower asymptote of the inhibition function, it would not be zero as the horse-race model predicts, but rather, it would equal the proportion of trials on which subjects chose to respond regardless of the stop signal. The regular inhibition function would be compressed, ranging from that value to 1.0. A hypothetical inhibition function that would be produced by this strategy is illustrated in Figure 7.

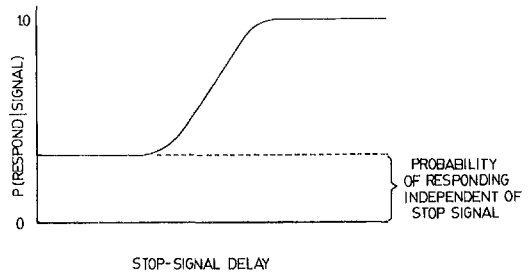


Figure 7. A hypothetical inhibition function for subjects who choose to respond on a certain proportion of trials regardless of whether there is a stop signal. (The lower asymptote of the inhibition function reflects the proportion of trials on which subjects choose to respond regardless of the signal.)

The data are generally inconsistent with this prediction. The lower asymptote of the inhibition function, when it has been measured, is usually close to zero, as the horse-race model predicts (see the present Figure 1; also see Logan, 1982; Slater-Hammel, 1960). The only study in which the strategy might have occurred was reported by Logan (1983, Experiment 6). In that study, mean no-signal reaction time was 1,126 ms and the mean stop-signal delay was 400 ms, yet subjects responded on 40% of the stop-signal trials. It seems likely that they could have responded less often. The inhibition function reached its lower asymptote with the probability of responding given a signal equal to .29, suggesting that subjects ignored the signal on 29% of stop-signal trials.

In general, however, the data are inconsistent with the alternatives to the assumption that the stopping process and the primary-task process are independent, and they are consistent with the predictions of the horse-race model, which are predicated on the assumption of independence. It seems reasonable to conclude, then, that the assumption is warranted.

The Complete Model

In this section, we present a deeper analysis of the horse-race model, exploring the consequences of nonzero variance in stop-signal reaction time and including the mean and variance of the duration of the ballistic phase in the analysis. We develop and test predictions that are not immediately obvious in the approximations to the complete model we pre-

sented earlier, and we provide methods for estimating stop-signal reaction time from inhibition functions and estimating complete inhibition functions from partial functions obtained empirically.

When stopping time and the ballistic process are included as random variables, the formal model becomes much more complicated.

1. Response is now made up of two processes: an early process that can be inhibited, and a subsequent ballistic process that cannot be inhibited once it has been started. Both processes are assumed to be stochastic. The early process is represented by a random variable T_1 , distributed by $f_1(t)$, and the ballistic process is represented by a random variable T_2 , distributed by $f_2(t)$. The complete response process has mean reaction time given by

$$\bar{T} = \bar{T}_1 + \bar{T}_2 \tag{5}$$

and variance given by

$$\sigma^2 = \sigma_1^2 + \sigma_2^2. \tag{6}$$

2. The stopping process, which inhibits the response, is taken to be stochastic. It is represented by a random variable T_s , distributed by $f_s(t)$.

3. The response is inhibited if the stopping process terminates before the early process, that is, if

$$T_s + t_d < T_1.$$

Now, using these assumptions, we can deduce several conclusions.

Inhibition Functions Revisited

Now we are prepared to analyze inhibition functions more deeply in terms of the complete model. First, we define the inhibition function for the complete model, then we consider how it is related to signal-respond reaction times and stop-signal reaction times, and finally, we treat it as a cumulative distribution and use its moments to estimate parameters of the stopping process and the ballistic phase of the primary-task process.

The inhibition function is defined by calculating the probability that a response occurs, $P_r(t_d)$, as a function of stop-signal delay, t_d .

$$P_r(t_d) = \int_{t < u+t_d} f_1(t) f_s(u) dt du$$

$$\begin{aligned} &= \int_{-\infty}^{\infty} f_1(t) \int_{t-t_d}^{\infty} f_s(u) du dt \\ &= \int_{-\infty}^{\infty} f_s(u) \int_{-\infty}^{u+t_d} f_1(t) dt du. \tag{7} \end{aligned}$$

Relations between inhibition functions, signal-respond reaction times, and stop-signal reaction times. According to the model, the probability of inhibition, no-signal reaction time, and stop-signal reaction time all depend on stop-signal delay, t_d . We have already seen how the probability of inhibition depends on t_d . Now let us consider the others. First, consider the signal-respond trials. Two quantities can be measured: \bar{T}_r , which is the mean signal-respond reaction time we defined earlier in Equation 3, and $\langle T_r \rangle$, which is the mean reaction time for all trials on which a stop signal occurred, taking $RT = 0$ for those trials when no response occurs. Thus, $\langle T_r \rangle$ is the sum of the reaction times divided by the number of trials on which a stop signal occurred at a given delay, t_d (i.e., $P_r(t_d) \bar{T}_r = \langle T_r \rangle$). Let us examine $\langle T_r \rangle$. Because the ballistic process occurs on each trial where a response occurs,

$$\begin{aligned} \langle T_r \rangle &= \langle T_{1r} \rangle + \langle T_{2r} \rangle \\ &= \langle T_{1r} \rangle + P_r(t_d) \bar{T}_2. \tag{8} \end{aligned}$$

We can calculate $\langle T_{1r} \rangle$ from the distributions f_1 and f_s

$$\begin{aligned} \langle T_{1r} \rangle &= \int_{t < u+t_d} t f_1(t) f_s(u) du dt \\ &= \int_{-\infty}^{\infty} f_s(u) \int_{-\infty}^{u+t_d} t f_1(t) dt du. \tag{9} \end{aligned}$$

Next, consider the trials on which stopping occurs. Analogous to \bar{T}_r and $\langle T_r \rangle$, there are two quantities \bar{T}_s and $\langle T_s \rangle$ with

$$\langle T_s \rangle = P_s(t_d) \bar{T}_s.$$

We can calculate $\langle T_s \rangle$ from the distributions f_1 and f_s :

$$\begin{aligned} \langle T_s \rangle &= \int_{t > u+t_d} u f_s(u) f_1(t) dt du \\ &= \int_{-\infty}^{\infty} f_1(t) \int_{-\infty}^{t-t_d} u f_s(u) du dt. \tag{10} \end{aligned}$$

Both $\langle T_{1r} \rangle$ and $\langle T_s \rangle$ depend on stop-signal

delay, t_d , so that we may differentiate them with respect to t_d :

$$\begin{aligned} \frac{d\langle T_{1r} \rangle}{dt_d} &= \int_{-\infty}^{\infty} f_s(u)(u + t_d)f_1(u + t_d)du \\ \frac{d\langle T_s \rangle}{dt_d} &= - \int_{-\infty}^{\infty} f_1(t)(t - t_d)f_s(t - t_d)dt \\ &= - \int_{-\infty}^{\infty} f_s(u)(u)f_1(u + t_d)du. \end{aligned} \quad (11)$$

Adding the two derivatives, we get

$$\begin{aligned} \frac{d\langle T_{1r} \rangle}{dt_d} + \frac{d\langle T_s \rangle}{dt_d} &= t_d \int_{-\infty}^{\infty} f_s(u)f_1(u + t_d)du \\ &= t_d \frac{dP_r(t_d)}{dt_d}, \end{aligned} \quad (12)$$

where the last step follows from differentiating Equation 7. This identity cannot be checked in its present form because neither $\langle T_{1r} \rangle$ nor $\langle T_s \rangle$ is observable. However, $\langle T_r \rangle$ is. Define for the random variable T_s a quantity analogous to T_r . That is,

$$T'_s = T_s + T_2,$$

which augments the stopping time by the same ballistic process that occurs with a response. The reason for this definition will become clear below where methods of estimating $\langle T'_s \rangle$ are discussed.

$$\langle T'_s \rangle = \langle T_s \rangle + P_s(t_d)\bar{T}_2$$

$$\langle T_r \rangle + \langle T'_s \rangle = \langle T_{1r} \rangle + \langle T_s \rangle + \bar{T}_2 \quad (13)$$

because $P_r(t_d) + P_s(t_d) = 1$. \bar{T}_2 is assumed not to depend on stop-signal delay in any significant way, so that

$$\frac{d\langle T_r \rangle}{dt_d} + \frac{d\langle T'_s \rangle}{dt_d} = t_d \frac{dP_r(t_d)}{dt_d}. \quad (14)$$

This expression is important because it predicts a relation between inhibition functions, signal-respond reaction times, and stop-signal reaction times. We cannot test the prediction as it is stated now because we do not know the form of the distribution functions that are being differentiated. However, we can test a special case of the relation in the following manner, without knowing the form of the functions.

Looking at the data, find a range of stop-signal delays $t_{d0} < t_d < t_{d1}$ for which the probability of responding given a signal increases linearly. No theoretical significance is attached to the linearity; it is just a phenomenological description of the data. Over this range we can regress $P_r(t_d)$ against t_d , finding constants A and B such that

$$P_r(t_d) = At_d + B \quad t_{d0} < t_d < t_{d1}.$$

There is also a boundary condition

$$P_r(t_{d0}) = At_{d0} + B. \quad (15)$$

We can substitute this expression into Equation 14 and integrate from t_{d0} to t_d :

$$\begin{aligned} \int_{t_{d0}}^{t_d} \frac{d}{dt_d} [\langle T_r \rangle + \langle T'_s \rangle] dt_d &= \int_{t_{d0}}^{t_d} t_d A dt_d \\ \langle T_r \rangle + \langle T'_s \rangle - [\langle T_r \rangle_0 + \langle T'_s \rangle_0] &= \frac{1}{2} A t_d^2 - \frac{1}{2} A t_{d0}^2 \\ \langle T_r \rangle + \langle T'_s \rangle &= \frac{1}{2} A t_d^2 - \frac{1}{2} A t_{d0}^2 \\ &\quad + [\langle T_r \rangle_0 + \langle T'_s \rangle_0]. \end{aligned} \quad (16)$$

Thus, within the interval $t_{d0} < t_d < t_{d1}$, the sum $\langle T_r \rangle + \langle T'_s \rangle$, represented as quadratic in t_d , has a leading term which is exactly one half the linear term in $P_r(t_d)$. As with the linear form of $P_r(t_d)$, no theoretical significance should be ascribed to the quadratic form of this equation. It is, once again, a phenomenological description of the data. The argument given here runs, in essence: if you describe one set of data, $P_r(t_d)$, by a given form, a linear function of t_d , and if the horse race holds, then when you describe another set of data, $\langle T_r \rangle + \langle T'_s \rangle$, in a given form, a second order function of t_d , then certain quantitative relations must hold between the two forms. We show, for our data on choice reaction time, that these quantitative relations hold.

We estimate $\langle T'_s \rangle$ using $P_r(t_d)$, the probability of responding when the stop-signal delay is t_d , and $P(T < t_0)$, the probability that the reaction time, on trials where no stop signal is given, is less than t_0 . We measure $\langle T'_s \rangle$ using the method that assumes that the variance in the stopping process is zero (see Equation 4), which is equivalent to the lowest order approximation to the estimate derived when stopping time is allowed to vary (see Appendix), minus the lowest order correction. We

estimate A from $P_r(t_d)$, Equation 15, and from $\langle T_r \rangle + \langle T'_s \rangle$, Equation 16, using the estimate described above.

We fit the functions to the data of each subject individually. For subject J. M., the inhibition function was relatively linear from the 200-ms delay to the 450-ms delay. Regressing the probability of responding against delay in that range yielded a slope of .00372. Regressing the quantity $P_s(t_d)$ (stop-signal reaction time) + $P_r(t_d)$ (signal-respond reaction time) against delay squared in the same range yielded a slope of .00180. The ratio of the former slope to the latter was 2.07, compared to the ratio of 2.0 predicted by the model.

For subject G. L., the inhibition function was relatively linear from the 150-ms delay to the 350-ms delay. Regressing the probability of responding against delay in that range yielded a slope of .00512. Regressing the quantity $P_s(t_d)$ (stop-signal reaction time) + $P_r(t_d)$ (signal-respond reaction time) against delay squared in the same range yielded a slope of .00241. The ratio of the former slope to the latter was 2.12, compared to the ratio of 2.0 predicted by the model.

For subject J. C., the inhibition function was relatively linear from the 100-ms delay to the 300-ms delay. Regressing the probability of responding against delay in that range yielded a slope of .00422. Regressing the quantity $P_s(t_d)$ (stop-signal reaction time) + $P_r(t_d)$ (signal-respond reaction time) against delay squared in the same range yielded a slope of .0018. The ratio of the former slope to the latter was 2.34, compared to the ratio of 2.0 predicted by the model.

For each subject, the 95% confidence limits for the slope of the linear function relating the probability of responding to t_d contained the value corresponding to twice the slope of the quadratic function relating $\langle T'_s \rangle + \langle T_r \rangle$ to t_d . Thus, the observed inhibition functions, signal-respond reaction times, and stop-signal reaction times were related in a way that was not significantly different from the relation predicted by the model.

Inhibition functions as cumulative distributions. Finally, we use the model to estimate some parameters of the latency distributions of stopping and ballistic processes. To do so, consider the probability that a response will occur given a stop signal at delay t_d . We have

written it as $P_r(t_d)$, and we can see that it has the following properties:

1. As $t_d \rightarrow -\infty$, $P_r(t_d) \rightarrow 0$. That is, if the stop signal occurs early enough, no response occurs.

2. As $t_d \rightarrow +\infty$, $P_r(t_d) \rightarrow 1$. That is, if the stop signal occurs late enough, there is always a response.

3. $0 \leq P_r(t_d) \leq 1$.

4. If $t'_d > t_d$, then $P_r(t'_d) \geq P_r(t_d)$; $P_r(t_d)$ is monotonic, increasing with t_d . That is, the later the stop signal occurs, the greater the probability of a response.

These properties indicate that we can treat $P_r(t_d)$ as the cumulative distribution of a random variable T_d . T_d will be distributed by the function $\phi(t_d)$, defined by

$$\phi(t_d) = \frac{dP_r(t_d)}{dt}$$

$\phi(t_d)$ may be calculated directly from experimental data. T_d has a mean and variance calculated from $\phi(t_d)$ in the usual way:

$$\bar{T}_d = \int_{-\infty}^{\infty} t\phi(t)dt$$

$$\sigma_d^2 = \int_{-\infty}^{\infty} (t - \bar{T}_d)^2\phi(t)dt. \quad (17)$$

We will want to compare these quantities to properties of the reaction time distribution for trials on which no stop signal occurs. For these responses, which consist of an early plus a ballistic component, the cumulative distribution, $P(T < t) = F(t)$, is given by

$$F(t) = \int_{t'+t'' < t} f_1(t')f_2(t'') dt' dt''$$

The corresponding density function $f(t)$ is

$$f(t) = \int_{-\infty}^{\infty} f_1(t')f_2(t - t') dt'$$

Now, we wish to compare the moments of $f(t)$ with those of $\phi(t)$. They are all easily calculated from experimental data.

First, calculate $\phi(t)$. Differentiating Equation 7, we get

$$\phi(t) = \int_{-\infty}^{\infty} f_s(u)f_i(u + t)du. \quad (18)$$

From Equation 18 we get the mean

$$\begin{aligned} \bar{T}_d &= \int_{-\infty}^{\infty} t \int_{-\infty}^{\infty} f_s(u)f_1(u+t)du dt \\ &= \bar{T}_1 - \bar{T}_s \end{aligned}$$

and the variance

$$\begin{aligned} \sigma_d^2 &= \int_{-\infty}^{\infty} (t - \bar{T}_d)^2 \int_{-\infty}^{\infty} f_s(u)f_1(u+t)du dt \\ &= \sigma_1^2 + \sigma_s^2. \end{aligned}$$

By comparison,

$$\begin{aligned} \bar{T} &= \int_{-\infty}^{\infty} t \int_{-\infty}^{\infty} f_1(t')f_2(t-t')dt' dt \\ &= \bar{T}_1 + \bar{T}_2, \end{aligned}$$

and

$$\begin{aligned} \sigma^2 &= \int_{-\infty}^{\infty} (t - \bar{T})^2 \int_{-\infty}^{\infty} f_1(t')f_2(t-t')dt' dt \\ &= \sigma_1^2 + \sigma_2^2. \end{aligned}$$

Consider the quantity $\bar{T} - \bar{T}_d = \bar{T}_2 + \bar{T}_s$. The horse-race model predicts it to be greater than zero because both \bar{T}_2 and \bar{T}_s are greater than zero. It also provides an upper limit for both \bar{T}_2 and \bar{T}_s . That is,

$$0 \leq \bar{T}_2 \leq \bar{T} - \bar{T}_d$$

and

$$0 \leq \bar{T}_s \leq \bar{T} - \bar{T}_d.$$

We estimated the mean of the T_d distribution for each of our subjects, and used the estimate to calculate $\bar{T}_2 + \bar{T}_s$. The estimates of $\bar{T}_2 + \bar{T}_s$ are presented in Table 4 with other estimates of stop-signal reaction time. Note

Table 4
Estimates of Stop-Signal Reaction Time Plus the Duration of the Ballistic Component from the Primary-Task Reaction Time Distribution and from the Mean and Median of the T_d Distribution

Measure	Subject		
	J. M.	G. L.	J. C.
RT (distribution)	172	206	205
T_d Mean	164	182	180
Median (slope)	158	206	199
Median (interpolation)	156	210	202

Table 5

Estimates of the Standard Deviation of T_d Distribution

Measure	Subject		
	J. M.	G. L.	J. C.
Full function	75	62	77
T_d Median (slope)	107	78	95
Median (interpolation)	69	51	51

that the values calculated from the T_d distribution are somewhat smaller than the values calculated previously from the primary-task reaction time distribution for all three subjects, but they are still relatively close.

Consider the quantity $\sigma_d^2 - \sigma^2 = \sigma_s^2 - \sigma_2^2$. Its sign shows which of the stopping and ballistic processes has the smaller variance. Furthermore, because both variances must be greater than zero, its absolute value represents a lower limit for the magnitude of the larger variance.

We estimated the variance of the T_d distribution for each subject. The square roots of the values (i.e., standard deviations) are presented in Table 5 together with other estimates of the standard deviation of T_d (see below). For subjects G. L. and J. C., the T_d distribution was more variable than the primary-task reaction time distribution, suggesting that the variance of their stop-signal reaction times was larger than the variance of their ballistic components. The lower limits on the standard deviation of their stop-signal reaction times were 37 ms and 51 ms, respectively. For subject J. M., the primary-task reaction times were more variable than the T_d distribution, suggesting that the variance of the ballistic component was larger than the variance of the stop-signal reaction times. The lower limit on the standard deviation of her ballistic process was 80 ms.

Thus, consideration of the two distributions yields a prediction and various limits on parameters of the stopping and ballistic distributions. The limits may be regarded as estimates, so that the numbers so derived should be compared with expected magnitudes to provide a further test of the reasonability of this analysis. Calculation of moments could

easily be pushed to moments higher than the second. It is unlikely, however, that data of quality high enough to justify the effort will be available.

Parameterizing the Model

The theoretical development of the model did not assume that the various distributions had any particular form (i.e., normal, exponential, etc.). This makes the model very general. The moments of the different processes can be calculated, whatever form the distributions may take. However, the analysis requires very high-quality data. The moments must be calculated from the inhibition functions, treated as cumulative distributions, and few studies can provide sufficient data to do so. To estimate the distribution $\phi(t)$, some stop signals must occur so early that subjects never respond and some must occur so late that subjects always respond. Moreover, there must be many delays in between these extremes to capture the shape of the distribution with any fidelity. Only one study in the literature has used an appropriate range of delays (Slater-Hammel, 1960).

These considerations make it difficult to apply the model to data and extract all of the information that is potentially available. One way around this difficulty is to assume particular forms for the various distributions. Then the assumed distributions can be fitted to the data and the parameters of the fitted distributions can be used to calculate the parameters of the model. Under some assumptions, it may not even be necessary to fit the assumed distributions to data; characteristics of particular distributions can be exploited to calculate parameters of the model directly from the data.

Consider, for example, the distribution $\phi(t)$. If we assume it to be symmetrical (e.g., normal, t), then the mean and the median are identical. Under the assumption that the empirical function relating the probability of inhibition to stop-signal delay is the cumulative form of $\phi(t)$, we can estimate the median (and hence the mean) of $\phi(t)$ by finding the stop-signal delay at which subjects inhibit their responses 50% of the time. This value can then be entered into the equations for the model to calculate the latency of the stopping process and the

duration of the ballistic component (i.e., $\bar{T}_s + \bar{T}_2$). Note that the parametric analysis does not require a complete inhibition function; the only requirement is that the empirical inhibition function passes through the point at which 50% of responses are inhibited. Nearly all of the studies in the literature provide data for which this is true (see Table 6, rows 1, 3, 8, 9, 10, 17, 18, and 19).

To illustrate, we estimated the median of the T_d distribution for each subject by calculating it from the regression equation relating the probability of responding given a signal to delay and by interpolating from the two points on the inhibition function that bounded it. These values were then subtracted from the mean primary-task reaction time to estimate $\bar{T}_2 + \bar{T}_s$. The estimates are presented above in Table 4 along with estimates from the primary-task reaction time distribution and the mean of the T_d distribution calculated previously.

First, note that the values were about the same whether the median was estimated from the regression equation, which used several points on the inhibition function, or from interpolation, which used only two points. This suggests that the median can be estimated reasonably from only two points on the inhibition function, and nearly all of the studies in the literature provide data sufficient to do at least this. Second, note that the estimates from the median are relatively close to the estimates from the primary-task and T_d distributions. For subjects G. L. and J. C., the estimates from the medians were closer to the estimates from the primary-task distribution than to estimates from the T_d distribution; for subject J. M., the estimates from the median were closer to the estimates from the T_d distribution.

If we assume a particular symmetrical distribution (e.g., normal), then we can calculate the variance from the empirical function relating the probability of inhibition to stop-signal delay. In symmetrical distributions, the variance is proportional to the value of the ordinate at the median of the distribution or, equivalently, to the value of the slope of the cumulative distribution at its median. For the normal distribution,

$$\text{slope at median} = \frac{1}{\sqrt{2\pi} \sigma}$$

and for the rectangular distribution,

$$\text{slope at median} = \frac{1}{2\sqrt{3} \sigma}.$$

These formulas can be used to calculate estimates of the variance of the T_d distribution from estimates of the slope of the inhibition function around the point where 50% of the responses are inhibited. Of course, one must know the form of the distribution to select the appropriate equation or one must be willing to assume a particular form. In either case, it is possible to estimate the variance of the T_d distribution without a complete inhibition function; all that is required is a function that passes through the point at which 50% of the responses are inhibited.

We assumed that the T_d distribution was normal and estimated its variance in two ways. First, we calculated it from the slope of the regression equation relating the probability of responding given a signal to delay to illustrate an estimate derived from several points on the inhibition function. Then, we estimated the variance by calculating the slope of the inhibition function between the two points that bounded the median. The square root of the values for each subject (i.e., standard deviations) are presented above in Table 5 along with estimates calculated previously from the whole T_d distribution. For all three subjects, the three estimates were relatively similar. The estimates from the regression equation led to the same conclusions as did the estimates from the whole T_d distribution. For G. L. and J. C., stop-signal reaction times were more variable than the ballistic process, but for J. M., the ballistic process was more variable. However, the estimates based on interpolation led to different conclusions for subject J. C. The ballistic phase seemed more variable than the stop-signal reaction times, whereas the opposite was true with the other estimates.

The parametric analysis described here may be appropriate in most of the domains that have been studied in the literature. Most of the inhibition functions plotted in this article seem normal, or at least, reasonably symmetrical. Thus, it may be possible to apply the model in all its detail to rather scanty data, as long as the tasks are sufficiently similar to the tasks studied previously. When exploring

new domains where little is known in advance about the forms the various distributions are likely to take, the full-scale, nonparametric analysis should be applied (also see Logan et al., 1984).

General Discussion

A major justification for developing a model of performance in the stop-signal paradigm was to answer three general questions about control. In this section we review the answers that the model provides, and then consider ways in which the model might be generalized so that it can be applied to situations other than the stop-signal paradigm. After that, we conclude with a discussion of how the model fits into current thinking on the nature of control, attention, automaticity, and skill.

Issues in the Study of Control

1. *Measuring the difficulty of control.* The model suggests a two-step procedure for measuring the difficulty of control. First, we compare inhibition functions and attempt to remove differences between them by plotting them as a function of some transformation of delay (e.g., *ZRFT*) that takes into account differences in the mean and variance of primary-task reaction time and differences in stop-signal reaction time. Functions that cannot be brought into alignment by these transformations involve different control processes, and at least one of them may contain a substantial ballistic component (see below).

If the different functions can be brought into alignment, then we take the second step and compare the parameters of the model. According to the model, stop-signal reaction time is the critical parameter. If conditions, tasks, and so forth, differ in stop-signal reaction time, the one with the larger value is the more difficult to control. Estimates of stop-signal reaction time could differ because of underlying differences in the latency of the response to the stop signal, t_s , or differences in the duration of the ballistic component, t_2 . In both cases, it seems appropriate to consider the condition with the longer estimated stop-signal reaction time the more difficult to control.

Studies that compare inhibition functions suggest that there are few differences between

subjects (Lisberger et al., 1975), strategies (Logan, 1981), tasks (Logan et al., 1984), and conditions (Logan, 1981, 1982, 1983) that cannot be removed by transforming delay in the ways specified by the model. The ability to inhibit action seems not to vary much between subjects, tasks, strategies, or conditions, suggesting that performance is often controlled rather closely.

We expect that most inhibition functions can be brought into alignment with the various transformations. Stop-signal reaction time, which enters into the calculation of *ZRFT*, is a free parameter that can be set to minimize the difference between inhibition functions. If we calculated a different value of stop-signal reaction time for each delay and used it to calculate *ZRFT* for that delay, we could bring together functions that differed substantially in shape, trivializing the comparison of inhibition functions. However, we have avoided this trivialization in practice by using the same estimate of stop-signal reaction time to calculate *ZRFT* for each delay (see above; also see Logan et al., 1984). Also, the problem does not exist in situations in which differences in mean reaction time are enough to account for differences in inhibition functions (e.g., Lisberger et al., 1975; Logan, 1981, 1982, 1983). Thus, comparisons of inhibition functions can be informative.

Estimates of the latencies of various acts of control were taken from 19 studies in the literature and presented in Table 6. Studies of motor control (rows 1–8), reaction time (rows 9–13), and higher cognitive processes (rows 14–19) all yield estimates around 200 ms. Some of the deviations can be accounted for in terms of different criteria for deciding whether a response has stopped (rows 1, 2, 14, 15, and 16), and others may reflect different resource-allocation strategies that place less importance on stopping quickly (rows 17, 18, and 19). Possibly, some tasks take longer to control than do others or some acts of control may take longer to execute than do others; these are open questions of future research. For now, it is interesting that such diverse activities can be brought under control in roughly the same amount of time. This general finding corroborates the analysis of inhibition functions in suggesting that subjects typically have very close control over their actions. Some im-

plications of this conclusion are discussed below.

2. *Measuring the latency of control.* Estimates of stop-signal reaction time provide a measure of the latency of control. Stop-signal reaction time can be estimated by two different methods, one calculated from the probability of inhibition and the primary-task reaction time distribution, and one calculated from the mean of the inhibition function and the mean primary-task reaction time. The two methods yield similar estimates whether they are applied to the same data (see Table 4) or to different data (compare rows 11, 12, and 13 of Table 6 with rows 1, 3, 8, 9, 10, 17, 18, and 19). Both methods have advantages and disadvantages.

The former method allows us to calculate stop-signal reaction time at each delay and thereby assess it as a function of delay. The observed effects of delay can be compared with those predicted by the model to test its assumptions (see *Relations between inhibition functions, signal-respond reaction time, and stop-signal reaction time*, above), and the delay effects can provide information about ballistic processes in the primary task (see below). However, the method assumes that there is no variance in stop-signal reaction time, which is almost certain to be false. We proved that it was false in 2 of our 3 subjects (see Tables 3 and 5; also see Logan et al., 1984). We also showed that the assumption is not very important: We developed expressions for stop-signal reaction time in the full model, and found that the lowest order approximation was identical to the estimate that assumes no variance in stop-signal reaction time (see Appendix).

The second method, using the mean of the inhibition function and the mean primary-task reaction time, was developed from the full model. Its main advantage is that it can be calculated from rather scanty data (see rows 1, 3, 10, 17, 18, and 19 in Table 6). However, it yields but one estimate per inhibition function; it cannot be used to assess the effects of delay.

In principle, there is no reason why both methods cannot be used on the same data to produce converging estimates. Note that in both methods, the estimate of stop-signal reaction time includes the latency of the internal

Table 6
Estimates of Latency of Various Acts of Control (in ms) From Studies in the Literature

Study	Task	Stop-signal RT	Method of estimation
1. Henry & Harrison, 1961	Stopping arm movement	69 ^a	RT - Median of T_d ; based on time to onset of deceleration
2. Megaw, 1972	Error correction	64	Time to onset of deceleration
	Manual step tracking	92	Time to reversal of movement
3. Vince, 1948	Manual step tracking	241 ^a	RT - Median of T_d ; based on amplitude
4. Vince, in Welford, 1952	Manual step tracking	363 ^b	Time until movement stops
5. Vince & Welford, 1967	Manual step tracking	Stopping: 347 Speeding: 343 Slowing: 429	Not described Not described Not described
6. Hick, 1949	Changing response force	304	Time until movement begins
7. Newell & Houk, 1983	Changing constant force	Increasing: 152 Decreasing: 184	Time until compensatory movement begins Time until compensatory movement begins
8. Lisberger, Fuchs, King, & Evinger, 1975	Saccadic step tracking	134	RT - Median of T_d
9. Slater-Hammel, 1960	Stopping transit reaction; Anticipation	166	RT - Median of T_d
10. Poulton, 1950	1: Stopping a complex response	250-450	Not described adequately
	2: Stopping simple RT on catch trials	208 ^a	RT - Median of T_d
11. Lappin & Eriksen, 1966	Stopping simple RT	224	RT distribution
12. Logan, Cowan, & Davis, 1984	Stopping simple RT	205	RT distribution
	Stopping choice RT	222	RT distribution
13. Logan, 1981	Stopping choice RT	1: 231 2: 212	RT distribution RT: distribution
14. Ladefoged, Silverstein, & Papcun, 1973	Stopping speech	≈ 200	Time spent speaking after signal
15. Long, 1976	Error detection in typing	449 ^a	Typing rate × letters typed after error
16. Rabbitt, 1978	Error detection in typing	1: 166 ^a 2: 198 ^a	Typing rate × number of letters typed after error Typing rate × number of letters typed after error
17. Logan, 1982	Stopping typing	1: 261 ^c 454 ^a 2: 244 ^c 368 ^a	Time spent typing after signal RT - Median of T_d Time spent typing after signal RT - Median of T_d

Table 6 (continued)

Study	Task	Stop-signal RT	Method of estimation
18. Logan, 1983	Stopping thought	1: Category decision 326 ^a	RT - Median of T_d
		Rhyme decision 360 ^a	RT - Median of T_d
		2: Category decision 455 ^a	RT - Median of T_d
		Rhyme decision 423 ^a	RT - Median of T_d
19. Zbrodoff, Logan, & Barber, 1984	Stopping simple arithmetic	Vocal responses 399 ^a	RT - Median of T_d
		Manual responses 391 ^a	RT - Median of T_d

Note. The estimates were calculated by the original investigators and published in their reports unless otherwise noted.

^a Calculated from data in published reports.

^b Reported in Welford (1952) Table 2.

^c Calculated from data not presented in published report.

response to the stop signal and the duration of the ballistic process. We now turn to procedures for separating these components.

3. *Measuring the ballistic component of the process being controlled.* The estimates of stop-signal reaction time described above place upper limits on the duration of the ballistic component, t_2 , but they are confounded with the latency of the internal response to the stop signal, t_s . In practice, the confound may not be much of a problem because situations that lengthen t_s should be very different from situations that lengthen t_2 . Consider, for example, Logan et al.'s (1984) experiments on simple and choice reaction time in which stop-signal reaction time was often found to be faster with the simple task than with the choice task (see Table 6, row 12). The difference could mean that t_s was longer in the choice task, as if the primary task and the stopping task competed for resources. Alternatively, it could mean that t_2 was longer in the choice task, as if the choice task had a more substantial ballistic component. The former seems plausible and is consistent with existing data (e.g., Logan, 1980), whereas the latter seems implausible. More generally, it should be possible to unravel the confound by manipulating parameters of the experiment that affect t_2 and t_s differently.

We can also assess the importance of the ballistic component by examining the effects of delay on stop-signal reaction time and by examining inhibition functions. If some component of a task must run on to completion before stopping, stop signals presented before the component finishes should not be effective.

This means that the probability of inhibiting responses based on the component should be close to zero and should not increase until delay is so long that the response to the stop signal occurs after the component has finished. Thus, flat inhibition functions suggest the presence of a lengthy ballistic component (see Logan, 1982). If stop-signal reaction time can be measured (e.g., if the task is continuous, like typing), then the estimates should decrease linearly with a slope of -1 as delay increases. This follows because the stopping process would have to wait until the ballistic component finished, and the waiting time depends on delay. Note that this prediction is similar to the predictions for the psychological refractory period made by single-channel theory (Welford, 1952), although the response to the stop signal is refractory for different reasons.

Two instances of ballistic processing have been discovered using this method. Logan (1982) found that skilled typists tended to type out the word *the* and the following space before stopping, and they tended to complete verbs with familiar endings (i.e., *ed* and *es*). Otherwise, typists seemed able to stop typing at any point within or between words. Ladefoged et al. (1973) found evidence that people begin speaking ballistically. They found that subjects responded to stop signals presented during speech in about 200 ms regardless of where in the utterance the signal occurred, but subjects responded to signals presented before speech began more slowly and the latency decreased with a slope of -1 until speech began.

However, the findings of Ladefoged et al.

(1973) may have been an artifact of their method of measuring stop-signal reaction time. With their procedure, subjects had to begin speaking in order to register a stop-signal reaction time, and their data from signals before speech began may reflect those trials on which the response to the stop signal was so slow that speech could not be inhibited before it began. Trials on which the response to the stop signal was faster than the onset of speech would not contribute to their estimates of stop-signal reaction time because the speaking response would be inhibited and no reaction time could be recorded. Ideally, stop-signal reaction time could be estimated from the inhibition data, using the techniques described here, and the values could be compared with the estimates from the offset of voicing. Unfortunately, Ladefoged et al. (1973) did not present data on the probability of inhibiting speech before it began, so this possibility cannot be checked.

Zbrodoff, Logan, and Barber (1984) presented some data that bear on the issue: They presented subjects with problems in simple arithmetic, and had subjects respond by pressing keys or by speaking. Stop signals were presented 100, 300, 500, or 700 ms after the onset of the problem—well before the response was expected to occur. If Ladefoged et al. (1973) interpreted their data correctly, it should have been more difficult to inhibit the vocal response before it began than to inhibit the key press before it began. However, it was not. Stop-signal reaction time, estimated from the inhibition functions in the manner described above, was 399 ms for vocal responses versus 391 ms for key-press responses (see Table 6, row 19). Thus, the conclusions of Ladefoged et al. (1973) that speech cannot be inhibited before it begins should be regarded with caution.

Expanding the Focus of the Model

Our original intention in developing the model presented here was to describe in general the way people take control of their thoughts and actions. To develop a formal model, we focused on the stop-signal paradigm, which involves stopping discrete voluntary actions. We now wish to suggest ways in which the model might be generalized to

other situations that require different acts of control, to fulfill our original intention.

Continuous versus discrete action. The model developed so far is addressed to discrete tasks with discrete responses. It would seem to be relatively straightforward to generalize the model to continuous responses. All responses that involve moving parts of the body are continuous because physical principles dictate that the body must move continuously. The responses in a reaction time task, such as the one studied here, are discrete only insofar as a discrete criterion is imposed on the movement, such as the closure of a telegraph key. The responses seem discrete largely because we are used to thinking of them as discrete. On the other hand, properties of seemingly continuous responses, such as arm movements, are often studied by imposing discrete criteria, such as a change in direction or acceleration (e.g., Gottsdanker, 1966; Henry & Harrison, 1961; Newell & Houk, 1983; but see e.g., Viviani & Terzuolo, 1973). Thus, the distinction between discrete and continuous responses seems arbitrary and artificial. There is no reason why all of the properties of the model developed here cannot be applied directly to tasks with so-called continuous responses.

A potential problem is that the parameters of the model may depend on the discrete criterion applied to the continuous movement (see Table 6, rows 1, 2, and 17). For example, Henry and Harrison (1961) had subjects reverse an arm movement on signal, and found that when the signal occurred 100 ms after the stimulus to begin the movement, subjects began decelerating before completing the movement on 58% of the trials, but they reversed the movement before completing it on only 3% of the trials. Similarly, Megaw (1972) observed subjects correcting errors in a step-tracking task by reversing the direction of movement, and found that subjects began decelerating some 28 ms before they reversed direction.

In situations like these, it would be interesting to plot inhibition functions for several different criteria. If the different criteria reflect different portions of the same process, it should be possible to account for the differences between the inhibition functions in terms of pa-

rameters of the reaction time distribution and parameters of the stop-signal reaction time distribution for each component. It is beyond the scope of this article to generalize the formal model to deal with these complications, but the basic idea seems worth pursuing.

Continuous tasks. Tasks themselves may be either discrete or continuous. The distinction between discrete tasks and continuous tasks may be fuzzy, but extreme cases are not difficult to distinguish. Tasks such as choice reaction time that involve a single stimulus and a single response are clearly discrete, whereas tasks such as typing that involve a chain of stimuli and a chain of responses are clearly continuous. The model can easily be generalized from the discrete tasks it was developed with to continuous tasks, by imposing discrete criteria on the continuous tasks. For example, Logan (1982) generated inhibition functions for typewriting by calculating the probability of inhibiting particular keystrokes (the first and the last in a word) as a function of stop-signal delay (see the present Figure 5).

Continuous tasks are interesting because they offer a way of estimating stop-signal reaction time that does not depend on inhibition functions: The chain of responses stops some time after the stop signal, and the time between the onset of the stop signal and the occurrence of the last response to be emitted can be used as an estimate of stop-signal reaction time. Ladefoged et al. (1973) used this procedure to estimate stop-signal reaction time in stopping speech,³ and Logan (1982) used the procedure to estimate stop-signal reaction time in stopping typewriting (see Table 6, rows 14 and 17; cf. rows 15 and 16). Estimates made in this way are interesting in relation to estimates made from inhibition functions. Whereas values calculated from inhibition functions may overestimate stop-signal reaction time and the latency of the ballistic phase of the primary task (see Appendix), values calculated from the latency of the last response after a stop signal may underestimate them. This follows because subjects could inhibit their action any time between the last response after the stop signal and the expected occurrence of the next response. Alternatively, in tasks with complex responses, the component of the response the experimenter is observing

may not be the last component the subject executes before stopping (cf. Henry & Harrison, 1961). Thus, the time between the stop signal and the last component of the response provides a lower limit on stop-signal reaction time and the duration of the ballistic phase of the primary task.

Despite the possibility that they may underestimate values calculated from inhibition functions, estimates from the last response after the stop signal are important. They are clearly in the same neighborhood as estimates from inhibition functions, and they should respond in the same way to the same experimental manipulations. Ideally, they could be used in conjunction with estimates from inhibition functions to test the same hypotheses and confirm the conclusions.

Thought versus action. The model has focused on stopping action because action is relatively easy to observe. In principle, the model can be applied to thought as well. Thoughts, like actions, can be directed toward goals, and like actions, thoughts sometimes need to be inhibited. The model assumes only that two processes race until the first one finishes; it does not require that the process being stopped generates an overt response.

A major problem in extending the model to the control of thought is that thoughts are generally not observable. Cognitive psychologists typically make thoughts observable by setting a task in which some response is made contingent on thought and observing the latency and accuracy of the response. This would be difficult in the stop-signal paradigm when subjects inhibit their responses. However, thoughts have other consequences besides generating overt responses, and those consequences can be observed whether or not a response occurs. People generally remember what they have thought about, and their

³ Since this article went to press, Levelt (1983) published an important article on error correction in spontaneous speech. His Figure 4 contains a distribution of the number of syllables uttered before stopping to correct an error, which could be used to estimate an analog of stop-signal reaction time in speech. Unfortunately, Levelt did not report an estimate of the rate of speech, which is necessary to convert the distribution into an estimate of stop-signal reaction time.

thoughts can be made observable through their memories.

Logan (1983) did just that. He presented subjects with pairs of words and had them make category and rhyme judgments about them. Occasionally, he presented a stop signal and told subjects to inhibit their response to the word pair when they heard the signal. Logan was interested in whether subjects stopped thinking about the word pairs when they stopped their responses to them. To find out, he tested subjects' ability to recognize words whose responses they had inhibited as well as words that received a complete response. If thoughts were inhibited with their corresponding actions, memory performance should improve as stop-signal delay increased, and words whose responses were inhibited should be remembered less well than words that received a complete response. However, if thoughts went on to completion when the actions were inhibited, memory performance should not depend on stop-signal delay, and should be about the same whether or not subjects inhibited their responses.

The data, displayed in the present Figure 8, generally supported the hypothesis that thoughts went on to completion. Memory performance was independent of delay when subjects stopped their responses on signal (the top two panels of Figure 8) and when subjects had to respond overtly to the tone as well as stop their responses to the word pair (the middle two panels of Figure 8). Apparently, thoughts went on to completion when the corresponding action was inhibited, even when attention was diverted to a new task. The only way Logan was able to inhibit thought was to replace the word pair with another when the stop signal occurred. Then, memory performance improved as delay increased. The increase was about the same when subjects made a new judgment about the new word pair (bottom left panel of Figure 8) as when they simply stopped their response to the first pair (bottom right panel of Figure 8), suggesting that it was the change in the display that inhibited thought.

All thoughts do not run on to completion ballistically. Using a test of recognition memory, Zbrodoff et al. (1984) found evidence that the thoughts underlying simple and complex arithmetic can be inhibited by interrupting

the overt actions that are associated with them. It is not clear why arithmetic should be different from the category and rhyme judgments Logan (1983) studied. Perhaps arithmetic is not as well practiced and so it is less automatic. In any case, an important goal for future research is to determine what kinds of thoughts are ballistic and what factors make thoughts ballistic.

The investigations need not rely on recognition memory entirely; other tests of memory can be used to converge on conclusions drawn from studies of recognition memory. For example, Logan (1984) used repetition priming as a test of memory and found results that agreed in some respects with previous studies of recognition memory (Logan, 1983) and disagreed in other respects.

These studies suggest that the horse-race model can be extended to processes that are not directly observable, which is an important development. Of course, it would be premature to apply the formal model to these domains, but nevertheless, the studies suggest that qualitative analyses can be quite informative.

Broader Implications

The application of the horse-race model to the stop-signal paradigm and others (see Table 6) suggests that subjects have very close control over their actions. Typically, they can stop the current course of action in about 200 ms. The analysis of inhibition functions showed that subjects could stop a response up to the point of executing it; the probability of inhibition depended on the relative finishing times of the stopping task and the primary task, not on their relative starting times (also see Logan, 1981, 1982, 1983; Logan et al., 1984). Together, these conclusions have implications for the nature of attentional control and the nature of automaticity, and they suggest further investigations of acts of control.

Attentional control. In general, the facts about stopping and changing are consistent with a *late-selection* theory of attention (Duncan, 1980; Norman, 1968). Late-selection theories propose that attention has its selective, controlling influence at later stages of processing, after stimuli have had complete access to the semantic system. Attention is necessary primarily to control access to the motor sys-

tem, which it does by gating the relevant stimulus information. The late locus of control would account for the relatively short ballistic component observed in stop-signal studies; the ballistic component should take no more time than a response takes to execute after the command to execute it has been given (cf. Cher-

nikoff & Taylor, 1952). The late locus of control would also account for the fact that response inhibition depends on relative finishing time rather than on relative starting time; control is exerted at the point that stimulus information reaches the motor system, regardless of how long it took to get there.

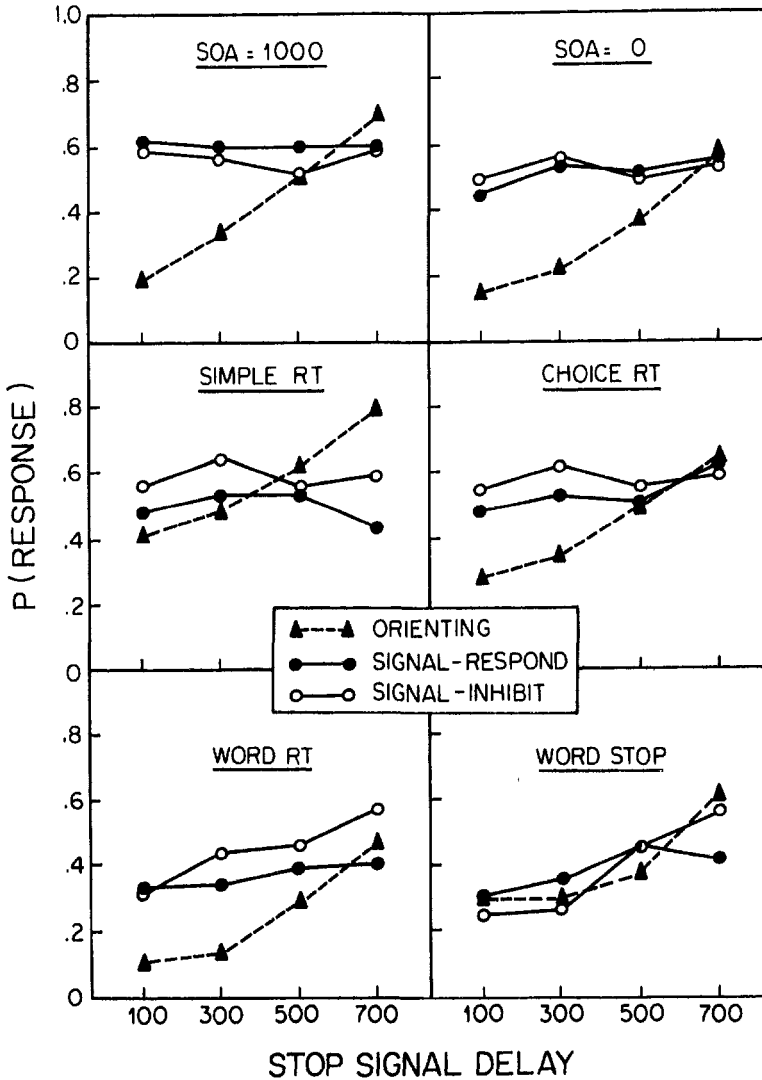


Figure 8. Inhibition functions for thought and action plotted against stop-signal delay. (Orienting = probability of responding given a stop signal in the orienting task; signal-respond = probability of recognizing a word from the orienting task that occurred with a stop signal and received a response; signal-inhibit = probability of recognizing a word from the orienting task that occurred with a stop signal and did not receive a response. From "On the Ability to Inhibit Simple Thoughts and Actions: I. Stop-Signal Studies of Decision and Memory" by G. D. Logan, *Journal of Experimental Psychology: Learning, Memory and Cognition*, 1983, 9, 585-606, Figure C1. Copyright 1983 by the American Psychological Association, Inc. Reprinted by permission.)

The facts are also consistent with a *hierarchical* theory of attention (Broadbent, 1977; Reason & Mycielska, 1982). Hierarchical theories propose that attention has its selective, controlling influence by acting as an executive that gives orders to subordinate systems. The subordinate systems interpret their orders and carry them out on their own until they finish or are countermanded by the executive. A hierarchical system could respond quickly to a signal to stop by countermanding the orders to the subordinate systems. The subordinate systems, having no orders to go on, would grind to a halt relatively quickly. In principle, control could be exerted at any point from stimulus reception to response; the motor system would be the last possible point at which a response could be inhibited, so inhibition functions would depend on relative finishing times.

The fact that stop-signal reaction times did not show the refractoriness typical of many double-stimulation paradigms (Kantowitz, 1974; Welford, 1952) is easier to reconcile with hierarchical theories than with late-selection theories. The data suggest that the system is refractory to some kinds of signals but not to others. Stop signals and signals to modify some parameters of a response appear to have privileged access to the system, whereas signals that require an overt response that is sufficiently different from the response to the first signal appear to have to wait for the system to deal with them (see Brebner, 1968; Gottsdanker, 1966; Megaw, 1972, 1974; Vince, 1948; Vince & Welford, 1967). Given that the same stimulus (e.g., a tone) can be used as a signal to respond or to stop responding, how does the system know what kind of stimulus it is dealing with? In hierarchical theories, the executive system monitors the consequences of its commands as well as issues commands to subordinates. The (executive) system that decides whether a signal requires privileged access is different from the (subordinate) system that the signal accesses. In late-selection theories, there is generally one limited capacity system for monitoring, issuing, and executing commands. The system that decides whether a signal requires privileged access is the same system that the signal accesses; the two processes must be time shared, which should produce refractory delays. Thus, hierarchical the-

ories account for the absence of refractory effects in stop-signal data better than late-selection theories do.

The evidence from stop-signal studies may support hierarchical theories of attention, but it does not rule out late-selection theories entirely. The treatment here considered *modal models* that were abstracted from several different theories. It is entirely possible that members of each class of theory can be found that mimic the predictions of at least one member of the other class. For example, a late-selection theory could simulate a hierarchical system by devoting a fixed portion of its limited capacity to monitoring and issuing executive commands. The remaining capacity could be allocated to the various signals in accord with commands from the executive portion. If the portions allocated to the executive and subordinate really were fixed, the two systems would not time share, just as the executive and the subordinates in a hierarchical theory do not time share (cf. McLeod, 1977). Thus, we cannot decide for once and for all which theory is better.⁴ However, we can treat the two classes of theory as frameworks for organizing data and for asking theoretical and empirical questions. From that perspective, hierarchical theories seem better than do nonhierarchical theories because they make it easier to talk about control.

Automaticity and control. Automaticity is a natural concept. It is part of the nontechnical vocabulary that lay people use to describe thoughts and actions that begin and run on to completion relatively autonomously, with little conscious effort. Automaticity is also a scientific concept, part of experimental psychology since the last century (e.g., James, 1890). Modern attempts to understand automaticity often do so by drawing a contrast between automatic and "controlled" process-

⁴ The goals to be attained by performing may force a hierarchical structure on the system whether or not the system is hierarchical to begin with. Some goals will be subordinate to others in achieving the main goal, and the processing that is directed toward fulfilling the subordinate goals may be considered subordinate to the processing that is directed toward fulfilling higher level goals. It may be difficult to separate the hierarchical goal structure from the hierarchical processing structure, but in either case, it would be appropriate to describe performance as hierarchical.

ing, claiming that automatic processing is fast, parallel, and relatively immune to dual-task interference, whereas "controlled" processing is slow, serial, and subject to severe dual-task interference (e.g., Hasher & Zacks, 1979; Logan, 1978; Posner & Snyder, 1975; Shiffrin & Schneider, 1977). Some even claim that automatic processes cannot be inhibited once they begin, whereas controlled processes can be stopped easily (Hasher & Zacks, 1979; Shiffrin & Dumais, 1981).⁵ These modern theories carry the implication that automaticity and control are opposites, which is contrary to evidence from stop-signal studies and evidence from studies of everyday errors.

Stop-signal studies show that highly skilled actions such as typing and speaking are controlled very closely (Ladefoged et al., 1973; Levelt, 1983; Logan, 1982; Long, 1976; Rabbitt, 1978), yet skilled actions are automatic by most conventional criteria (i.e., they are fast, their components run in parallel wherever possible, and they are not subject to severe dual-task interference). The time required to stop typing or speaking is about the same as the time required to stop much simpler actions (see Table 6). Apparently, the claim that automatic processes are hard to inhibit is wrong.

Studies of everyday errors also suggest that automatic reactions may be controlled. Reason and Myceilska (1982) report a case in which a person absentmindedly put coffee into a pot intended for tea. The erroneous behavior was automatic by some criteria (e.g., it was clearly done without intention; cf. Posner & Snyder, 1975), but it was nevertheless controlled. The person opened a jar and poured coffee from it into the pot. The movements required to open the coffee jar were different from the movements required to open the tea jar, so the behavior was controlled in the sense that it adapted to constraints imposed by the jar. Furthermore, the coffee was poured into the pot, not spilled haphazardly about the kitchen. Thus, automatic processes are usually not uncontrolled; the same behavior may be described as automatic and controlled.

The relation between control and automaticity can be understood by examining more closely what is meant by control. In general terms, control means directing thought and action toward a goal (Craik, 1947, 1948; Miller, Galanter, & Pribram, 1960; Powers,

1978). Thought and action that are directed toward a goal are controlled by definition, regardless of their speed, susceptibility to dual-task interference, and so on. Automatic processes are controlled insofar as they are engaged to bring about a goal (e.g., we control the automatic process of typing in order to write a letter). Current theories of automaticity seem to address a different meaning of control that has to do with the amount of work done by an executive process. The amount of processing required to accomplish a goal depends on the difference between the current state of affairs and the goal state. Some of the work is done by the executive and some is done by the subordinates. According to current theories of automaticity, work done by the executive is controlled and work done by subordinates is automatic. From the present perspective, this is not correct. All goal-directed activity is controlled, whether it is carried out by the executive or by the subordinates. It may be appropriate to call work done by subordinates automatic, but it need not be uncontrolled. The distinction between executive processing and subordinate processing is important, and it will certainly spawn more research. However, it is not a distinction between controlled and uncontrolled processing.

The perspective adopted here suggests a new direction for research on automaticity. We can ask what aspects of the subject's interaction with the environment are being controlled in a given task, and we can ask what the subject does to bring those aspects under control. We can then determine how much of the work is done by the executive and how much is done by the subordinates, and we can ask whether different divisions of labor would be more effective. Current studies of automaticity often ask the last two questions but rarely ask the first two; early studies in Craik's (1947, 1948) tradition and modern studies of movement control often ask the first two questions and rarely ask the last two. It seems important to ask all four questions together. The first two

⁵ Wickens (1939) and Kimble and Perlmutter (1970) have used inhibitability as a defining characteristic of voluntary processes, arguing that voluntary processes can be inhibited, whereas involuntary processes cannot. We have no quarrel with their definition, but we do not believe that automatic processes are necessarily involuntary.

provide a context for the last two, making them more meaningful than they would be alone, and the last two elaborate the answers given to first two.

Other acts of control. In this article, we have focused on inhibition of thought and action as acts of the human control system. We do not mean to imply that inhibition is the only act of control in the human repertoire; only that it is relatively simple and it is useful in controlling the motor system and the cognitive system. How it works together with other acts of control is an important question for future research.

The role of inhibition as a control strategy is perhaps clearer in motor skills, which tend to be continuous, than in cognitive skills, which often involve concatenation of discrete components. Motor skills may often be controlled by setting a reference level or a goal that subordinate systems try to maintain despite perturbations; inhibition may be used only when things get out of hand (i.e., when the subordinate systems cannot compensate for the perturbations). Alternatively, inhibition may be employed early in the acquisition of a motor skill, such as a novice skier deliberately falling down to prevent disaster on a run designed for experts. It may be replaced by more sophisticated, less drastic control strategies as skill develops, such as traversing to control speed on a steep run. Again, further research is needed to evaluate these possibilities.

Very little is known about the control of cognitive processes. Possibly, the control-system-engineering approach that has been so successful in the motor domain may not be as useful in analyzing cognitive skills. In cognitive skills, achieving a goal is often more of a problem than maintaining a goal state in the face of perturbations. Nevertheless, inhibition is an important act of cognitive control, and future research will determine the extent to which it is sufficient and the extent to which it must be supplemented by other acts of control.

References

- Anderson, J. R. (1982). Acquisition of cognitive skill. *Psychological Review*, 89, 369-406.
- Bartlett, N. R., Eason, R. G., & White, C. T. (1961). Latency of ocular fixation upon the second of two successive stimuli. *Perceptual and Motor Skills*, 13, 259-268.
- Brebner, J. (1968). Continuing and reversing the direction of responding movements: Some exceptions to the so-called "psychological refractory period." *Journal of Experimental Psychology*, 78, 120-127.
- Broadbent, D. (1958). *Perception and communication*. London: Pergamon.
- Broadbent, D. (1977). Levels, hierarchies, and the locus of control. *Quarterly Journal of Experimental Psychology*, 29, 181-201.
- Chernikoff, R., & Taylor, F. V. (1952). Reaction time to kinesthetic stimulation resulting from sudden arm displacement. *Journal of Experimental Psychology*, 43, 1-8.
- Craik, K. J. W. (1947). Theory of the human operator in control systems: I. The operator as an engineering system. *British Journal of Psychology*, 38, 56-61.
- Craik, K. J. W. (1948). Theory of the human operator in control systems: II. Man as an element in a control system. *British Journal of Psychology*, 38, 142-148.
- Davis, R. (1957). The human operator as a single channel information system. *Quarterly Journal of Experimental Psychology*, 9, 119-129.
- Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychological Review*, 87, 272-300.
- Feldman, A. G. (1981). The composition of central programs subserving horizontal eye movements in man. *Biological Cybernetics*, 42, 107-116.
- Gallistel, C. R. (1980). *The organization of action: A new synthesis*. Hillsdale, NJ: Erlbaum.
- Gottsdanker, R. (1966). The effect of superceding signals. *Quarterly Journal of Experimental Psychology*, 18, 236-249.
- Grice, G. R., Hunt, R. L., Kushner, B. A., & Morrow, C. (1974). Stimulus intensity, catch trial effects, and the speed-accuracy tradeoff in reaction time: A variable criterion theory interpretation. *Memory and Cognition*, 2, 758-770.
- Grice, G. R., Hunt, R. L., Kushner, B. A., & Nullmeyer, R. (1976). Associative processes and strategies in disjunctive reaction time. *Memory and Cognition*, 4, 433-445.
- Harrison, J. S. (1960). Psychological refractoriness and the latency time of two consecutive motor responses. *Research Quarterly*, 31, 590-600.
- Hasher, L., & Zacks, R. T. (1979). Automatic and effortful processes in memory. *Journal of Experimental Psychology: General*, 108, 356-388.
- Henry, F. M., & Harrison, J. S. (1961). Refractoriness of a fast movement. *Perceptual and Motor Skills*, 13, 351-354.
- Hick, W. E. (1949). Reaction time for the amendment of a response. *Quarterly Journal of Experimental Psychology*, 1, 175-179.
- James, W. (1890). *Principles of psychology*. New York: Holt.
- Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice-Hall.
- Kantowitz, B. H. (1974). Double stimulation. In B. H. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition* (pp. 83-131). Hillsdale, NJ: Erlbaum.
- Kelso, J. A. S., & Holt, K. G. (1980). Exploring a vibratory systems analysis of human movement production. *Journal of Neurophysiology*, 43, 1183-1196.

- Kimble, G. A., & Perlmutter, L. C. (1970). The problem of volition. *Psychological Review*, 77, 361-384.
- LaBerge, D. (1971). On the processing of simple visual and auditory stimuli at distinct levels. *Perception and Psychophysics*, 9, 331-334.
- LaBerge, D. (1973). Identification of two components of the time to switch attention: A test of a serial and a parallel model of attention. In S. Kornblum (Ed.), *Attention and performance IV* (pp. 71-85). New York: Academic Press.
- Ladefoged, P., Silverstein, R., & Papcun, G. (1973). Interruption of speech. *Journal of the Acoustical Society of America*, 54, 1105-1108.
- 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.
- Levitt, W. J. M. (1983). Monitoring and self-repair in speech. *Cognition*, 14, 41-104.
- 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. (1978). Attention in character classification tasks: Evidence for the automaticity of component stages. *Journal of Experimental Psychology: General*, 107, 32-63.
- Logan, G. D. (1980). Short-term memory demands of reaction time tasks that differ in complexity. *Journal of Experimental Psychology: Human Perception and Performance*, 6, 375-389.
- Logan, G. D. (1981). Attention, automaticity, and the ability to stop a speeded choice response. In J. Long & A. D. Baddeley (Eds.), *Attention and performance IX*. (pp. 205-222) Hillsdale, NJ: 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. (1984). On the ability to inhibit simple thoughts and actions: II. Stop-signal studies of repetition priming. Submitted for publication.
- Logan, G. D., Cowan, W. B., & Davis, K. A. (1984). On the ability to inhibit simple and choice reaction time responses: A model and a method. *Journal of Experimental Psychology: Human Perception and Performance*, 10, 276-291.
- Logan, G. D., Zbrodoff, N. J., & Fostey, A. R. (1983). Costs and benefits of strategy construction in a speeded discrimination task. *Memory and Cognition*, 11, 485-493.
- Long, J. (1976). Visual feedback and skilled keying: Differential effects of masking the printed copy and the keyboard. *Ergonomics*, 19, 93-110.
- Megaw, E. D. (1972). Directional errors and their correction in a discrete tracking task. *Ergonomics*, 15, 633-643.
- Megaw, E. D. (1974). Possible modification to a rapid on-going programmed manual response. *Brain Research*, 71, 425-441.
- McLeod, P. (1977). A dual task response modality effect: Support for multiprocessor models of attention. *Quarterly Journal of Experimental Psychology*, 29, 651-667.
- Miller, G. A., Galanter, E., & Pribram, K. (1960). *Plans and the structure of behavior*. New York: Holt.
- Navas, F., & Stark, L. (1968). Sampling or intermittency in hand control systems dynamics. *Biophysical Journal*, 8, 252-302.
- Navon, D., & Gopher, D. (1979). On the economy of the human processing system. *Psychological Review*, 86, 214-255.
- Newell, K. M., & Houk, J. C. (1983). Speed and accuracy of compensatory responses to limb disturbances. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 58-74.
- Norman, D. A. (1968). Toward a theory of memory and attention. *Psychological Review*, 75, 522-536.
- Ollman, R. T. (1973). Simple reactions with random countermanding of the "go" signal. In S. Kornblum (Ed.), *Attention and performance IV* (pp. 571-581). New York: Academic Press.
- Posner, M. I. (1978). *Chronometric explorations of mind*. Hillsdale, NJ: Erlbaum.
- Posner, M. I., & Boies, S. J. (1971). Components of attention. *Psychological Review*, 78, 391-408.
- Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola Symposium* (pp. 55-85). Hillsdale, NJ: Erlbaum.
- Poulton, E. C. (1950). Perceptual anticipation and reaction time. *Quarterly Journal of Experimental Psychology*, 2, 99-112.
- Powers, W. T. (1978). Quantitative analysis of purposive systems: Some spadework at the foundations of scientific psychology. *Psychological Review*, 85, 417-435.
- Rabbitt, P. M. A. (1978). Detection of errors by skilled typists. *Ergonomics*, 21, 945-958.
- Reason, J., & Myeilska, K. (1982). *Absent minded: The psychology of mental lapses and everyday errors*. Englewood Cliffs, NJ: Prentice-Hall.
- Robinson, D. A. (1973). Models of the saccadic eye movement control system. *Kybernetik*, 14, 71-83.
- Shallice, T. (1972). Dual functions of consciousness. *Psychological Review*, 79, 383-393.
- Sherrington, C. S. (1906). *The integrative action of the nervous system*. London: Constable.
- Shiffrin, R. M., & Dumais, S. T. (1981). The development of automatism. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 111-140). Hillsdale, NJ: Erlbaum.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84, 127-190.
- Slater-Hammel, A. T. (1960). Reliability, accuracy, and refractoriness of a transit reaction. *Research Quarterly*, 31, 217-228.
- Treisman, A. M. (1969). Strategies and models of selective attention. *Psychological Review*, 76, 282-299.
- Vince, M. A. (1948). The intermittency of control movements and the psychological refractory period. *British Journal of Psychology*, 38, 149-157.
- Vince, M. A., & Welford, A. T. (1967). Time taken to change the speed of a response. *Nature*, 213, 532-533.
- Viviani, P., & Terzuolo, C. A. (1973). Modeling of a simple motor task in man: Intentional arrest of an ongoing movement. *Kybernetik*, 14, 35-62.

Welford, A. T. (1952). The 'psychological refractory period' and the timing of high-speed performance: A review and a theory. *British Journal of Psychology*, 43, 2-19.

Wickens, D. D. (1939). A study of voluntary and involuntary finger conditioning. *Journal of Experimental Psychology*, 25, 127-140.

Woodworth, R. S. (1899). The accuracy of voluntary movement. *Psychological Review*, 3, 1-114.

Woodworth, R. S., & Schlosberg, H. (1954). *Experimental psychology* (2nd ed.). New York: Holt.

Young, L. R., & Stark, L. (1963). A discrete model for eye tracking movements. *Institute of Electrical and Electronic Engineers Transactions on Human Factors in Electronics*, 4, 38-51.

Zbrodoff, N. J., Logan, G. D., & Barber, C. Y. (1984). On the autonomy of mental processes: A case study of arithmetic. Submitted for publication.

Appendix

In this appendix, we show that the method of estimating stop-signal reaction time from the primary-task reaction time distribution and the probability of responding given a signal, which assumes that stop-signal reaction time is constant (see Equation 4), provides reasonable estimates even if stop-signal reaction time has nonzero variance. By Equation 13,

$$\langle T'_s \rangle = \langle T_s \rangle + P_s(t_d)\bar{T}_2,$$

where $\langle T'_s \rangle$ is the mean stopping time on trials where stopping occurs, for a given delay t_d , and \bar{T}_2 is the mean response time of the ballistic component. Now, evaluate $\langle T'_s \rangle$.

$$\langle T'_s \rangle = \int_{u+t_d < t} uf_s(u)du f_1(t)dt,$$

where $f_s(u)$ is the stopping time distribution and $f_1(t)$ is the distribution of the early process.

$$\begin{aligned} \langle T'_s \rangle &= \int_{-\infty}^{\infty} uf_s(u) \int_{u+t_d}^{\infty} f_1(t)dt du \\ &= \int_{-\infty}^{\infty} uf_s(u)du(1 - F_1(u + t_d)), \end{aligned} \quad (A1)$$

where $F_1(t)$ is the cumulative distribution of the early process. The stopping time distribution $f_s(u)$ is peaked around its mean, \bar{T}_s , so that we want to calculate $F_1(u + t_d)$ in that region. If $f_s(u)$ is narrowly peaked, we may take

$$F_1(u + t_d) = F_1(\bar{T}_s + t_d) + (u - \bar{T}_s)F'_1(\bar{T}_s + t_d) + \text{second derivative terms.}$$

Substituting in A1 and evaluating gives

$$\begin{aligned} \langle T'_s \rangle &= \int_{-\infty}^{\infty} uf_s(u)du(1 - F_1(\bar{T}_s + t_d) \\ &\quad - (u - \bar{T}_s)F'_1(\bar{T}_s + t_d) - \dots) \\ &= \bar{T}_s(1 - F_1(\bar{T}_s + t_d)) \\ &\quad - \sigma_s^2 F'_1(\bar{T}_s + t_d) - \dots \end{aligned} \quad (A2)$$

All of the terms in A2 are unknown, but we can calculate most of them from the probability of stopping and the reaction time distribution when no stop signal is given. From Equation 7, the probability of stopping is

$$\begin{aligned} P_s(t_d) &= 1 - P_r(t_d) \\ P_r(t_d) &= \int_{-\infty}^{\infty} f_s(u)du \int_{-\infty}^{u+t_d} f_1(t)dt \\ &= \int_{-\infty}^{\infty} f_s(u)du F_1(u + t_d). \end{aligned} \quad (A3)$$

Furthermore, the probability that the reaction time is less than t_0 when no stop signal is given is

$$\begin{aligned} P(T < t_0) &= \int_{t+t' < t_0} dt' f_1(t') dt f_2(t) \\ &= \int_{-\infty}^{\infty} f_2(t)dt \int_{-\infty}^{t_0-t} f_1(t')dt' \\ &= \int_{-\infty}^{\infty} f_2(t)dt F_1(t_0 - t). \end{aligned} \quad (A4)$$

In A3, expand $F_1(u + t_d)$ about $u = \bar{T}_s$; in A4, expand $F_1(t_0 - t)$ about $t = \bar{T}_2$. The rationale for this step is the expectation that most of the weight of $f_s(u)$ is near \bar{T}_s , so that the value of $P_s(t_d)$ is most sensitive to fluctuations of $F_1(u + t_d)$ in that region. It is assumed that $F_1(u + t_d)$ varies relatively smoothly across the region where $f_s(u)$ is significantly nonzero. Similar considerations apply to A4. Performing the substitutions, we get

$$\begin{aligned} P_r(t_d) &= \int_{-\infty}^{\infty} f_s(u)du(F_1(\bar{T}_s + t_d) \\ &\quad + (u - \bar{T}_s)F'_1(\bar{T}_s + t_d) + \dots) \\ &= F_1(\bar{T}_s + t_d) + \text{second derivative terms} \end{aligned} \quad (A5)$$

as well as

$$\begin{aligned}
 P(T < t_0) &= \int_{-\infty}^{\infty} f_2(t) dt (F_1(t_0 - \bar{T}_2) \\
 &\quad + (\bar{T}_2 - t) F_1'(t_0 - \bar{T}_2) + \dots) \\
 &= F_1(t_0 - \bar{T}_2) + \text{second derivative terms.}
 \end{aligned}$$

Note that in both cases, the first derivative term vanishes. This is important, because it makes the estimate in Equation 4 better than it might seem at first glance.

Now, choose t_0 so that $P_s(t_d) = P(T < t_0)$, which makes t_0 depend on t_d . This is precisely the procedure followed for estimating t_s in the case where the stopping time is assumed to have a fixed value. The present calculation shows the approximation is better than it seems to be, owing to the vanishing first derivative terms. Call the function so derived $r(t_d)$ (i.e., $r(t_d) = t_0$). Then, ignoring second derivative terms,

$$F_1(\bar{T}_s + t_d) = F_1(r(t_d) - \bar{T}_2).$$

Because $F_1(t)$ is monotonic,

$$\bar{T}_s + t_d = r(t_d) - \bar{T}_2$$

or

$$\bar{T}_s + \bar{T}_2 = r(t_d) - t_d.$$

Now, substituting A5 into A2,

$$\langle T_s \rangle = \bar{T}_s P_s(t_d) - \sigma_s^2 F_1'(\bar{T}_s + t_d),$$

and using the definition of $\langle T_s' \rangle$

$$\begin{aligned}
 \langle T_s' \rangle &= (\bar{T}_s + \bar{T}_2) P_s(t_d) - \sigma_s^2 F_1'(\bar{T}_s + t_d) \\
 &= (r(t_d) - t_d) P_s(t_d) - \sigma_s^2 F_1'(\bar{T}_s + t_d).
 \end{aligned}$$

In this result, all of the components of the first term, $r(t_d)$, t_d , and $P_s(t_d)$, can be estimated directly from the data for each value of stop-signal delay. The final term, which is a correction to the lowest order approximation, can be estimated from the same data. The correction depends on σ_s^2 , which cannot be measured directly. We derived a limit on its value in the section on inhibition functions as cumulative distributions.

Received June 24, 1983

Revision received November 11, 1983 ■