

- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992a). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, 18, 849-860.
- Raymond, J. E., Shapiro, K. L., & Arnell, K. A. (1992b). The role of post-target stimulation in producing "attentional blinks" in RSVP. *International Journal of Psychology*, 27 (3 & 4), 44.
- Reeves, A., & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review*, 93, 180-206.
- Shapiro, K. L., Arnell, K. A., & Drake, S. H. (1991). Stimulus complexity mediates target detection in visual attention search [Abstract]. *Investigative Ophthalmology and Visual Science*, 32, Supplement, 1040.
- Shapiro, K. L., Raymond, J. E., & Arnell, K. M. (in press). Attention to visual pattern information produces the attentional blink in RSVP. *Journal of Experimental Psychology: Human Perception and Performance*.
- Shapiro, K. L., Raymond, J. E., & Taylor, T. (1993). The attentional blink suppresses size and shape but not colour information. (Abstract) *Investigative Ophthalmology and Visual Science*, 34(4), 1232.
- Smith, M. C. (1967). Theories of the psychological refractory period. *Psychological Bulletin*, 67, 202-213.
- Sperling, G., & Weichselgartner, E. (1990, November). *Episodic theory of visual attention*. Paper presented at the meeting of the Psychonomic Society, New Orleans, LA.
- Taylor, D. A., & Chabot, R. J. (1978). Differential backward masking of words and letters by masks of varying orthographic structure. *Memory and Cognition*, 6, 629-635.
- Treisman, A. M., & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, 12, 97-136.
- Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, 95, 15-48.
- Ward, W., Duncan, J., & Shapiro, K. L. (1992, November). *The attentional blink does not require selection from among a nontarget stream*. Paper presented at the 33rd Annual Meeting of the Psychonomic Society, St. Louis, MO.
- Ward, R., Duncan, J., & Shapiro, K. L. (in preparation). *Target selection requiring one vs. two attributes does not affect the attentional blink*.
- Weichselgartner, E., & Sperling, G. (1987). Dynamics of automatic and controlled visual attention. *Science*, 238, 778-780.

On the Ability to Inhibit Thought and Action A Users' Guide to the Stop Signal Paradigm

Gordon D. Logan

I. INTRODUCTION

This chapter presents a users' guide to the stop signal paradigm. It explains the paradigm and the situations in the natural and artificial environment the paradigm is intended to capture. It explains how the inhibition that is apparent in the stop signal paradigm is similar to and different from the kinds of inhibition seen in other paradigms. It reviews theory and data in the experimental literature on stopping. It explains the race model, which is the commonly accepted theory of performance in the stop signal paradigm, and it explains how to apply the model to data.

II. THE STOP SIGNAL PARADIGM

A. What Is Stopping?

You are Guy Lafleur racing across the blue line with the puck on your stick. You wind up your famous slap shot . . . and the whistle blows. You are Pete Rose trying to hit a home run . . . and the pitch breaks out of the strike zone just as you begin your swing. You are Spaceman Spiff blasting aliens in the asteroid belt . . . and your teacher calls your attention back to the lesson. You are Gordon Logan giving the same lecture

for the n th time . . . and you stumble on the word "resources." What happens? Guy stops his famous slap shot in mid swing (much to the goalie's relief). Pete checks his swing before his wrists "break," avoiding a strike. Gordon stops in midword and begins again (this time with tongue out of cheek). The school child may or may not heed the teacher's call. If his diagnosis is attention deficit disorder, he is less likely to.

What do these situations have in common? The current course of thought and action is no longer appropriate. Goals change (as they did for Guy and the school child) and the world changes—the stimuli one is reacting to change (as they did for Pete). Errors introduce new goals (error correction) that take priority over current goals (as they did for Gordon). The thing to do is to stop the current course of thought and action and begin another. In this chapter, the focus is on stopping. Stopping is the first step in reorienting to new goals and new worlds, a general requirement in all kinds of cognitive control. Stopping is a clear case of executive intervention. It gives us a chance to see executive processes in action, a chance to discern their nature.

Stopping can be studied empirically in a simple laboratory analogue called the *stop signal paradigm*: People are engaged in a primary task, and occasionally, they are presented with a signal that tells them to stop their response to the primary task. The primary task is typically a visual choice reaction time task, such as discriminating X's from O's, and the stop signal is usually a tone. Other primary tasks and stop signals work just as well. The main datum is whether or not subjects withhold their response to the primary task when the stop signal occurs—do they stop? This datum, interpreted in terms of data on timing of stimuli and responses, can reveal a surprising amount about the underlying mechanism.

B. Why Is Stopping Interesting?

Stopping is interesting theoretically because it is an internally generated act of control. It is something a person does to change the current course of thought and action and bring it into line with current goals. Stopping research focuses on the role of the person in cognitive control, whereas most research focuses on the role of the environment, asking how familiar, novel, and distinctive stimuli control a person's behavior (e.g., research on involuntary attraction of attention, research on automaticity). These foci are complementary: Stopping research addresses the control process directly and research on involuntary control addresses the things that confront the control process (i.e., the demands for attention).

I think that cognitive control can be understood generally as acts of control, just like stopping. Some acts of control are more subtle and

graceful than stopping. They involve deciding to do something a little differently, like digging an edge in a little deeper to turn more sharply in skating and skiing. The parameters of action may change, but action may proceed uninterrupted. Not so with stopping. Stopping is an extreme form of control; the subject decides not to do something and action stops. The skater gives up and falls down; the skier takes a dive to stop careening down the hill. Subtle control and extreme control may differ more in degree than in kind. Both can be analyzed as acts of control that are triggered by certain circumstances and result in certain changes after some latent period.

The stop signal paradigm is advantageous because it allows a clear definition of the conditions that trigger the act of control (i.e., presentation of the stop signal) and the changes that result from executing the act (i.e., inhibition of the response). It provides a way to measure the latency of the act of control (i.e., stop signal reaction time), even though successful execution of the act produces no overt behavior (it is the absence of overt behavior that tells us that stopping was successful). Acts of control that are more subtle may be harder to observe and their latencies may be harder to measure. The disadvantage of the stop signal paradigm is that it reflects an extreme form of control that may be different in important ways from more subtle forms of control. Understanding the relation between stop signal control and other forms of cognitive control is an important goal for future research.

Stopping is interesting empirically because some important results have emerged and others are on the horizon. First, the data from young adults are consistent with the idea that a single, general mechanism underlies the ability to inhibit many kinds of thought and action. Young adults can stop a wide variety of actions very quickly. Eye movements, hand movements, key presses, squeezes, and speech can all be stopped in about 200 msec (for a review, see Logan & Cowan, 1984). The similar latencies across tasks and effector systems suggest that stopping relies on an amodal, central process.

The speed with which young adults can stop responses is impressive. It allows them very close control over their actions. The speed is even more impressive because stopping is the response to the second of two stimuli. Usually, the response to the second of two stimuli is very slow, as the burgeoning literature on dual-task interference and the psychological refractory period attests (see, e.g., Pashler, 1989; see Logan & Burkell, 1986, for an explicit comparison of stopping with dual-task performance). Somehow, stop signals seem to bypass the bottlenecks that hold back other signals (see Logan & Cowan, 1984).

Second, there are interesting changes in stopping performance over the life span that are consistent with theories of development (Bjorklund & Harnishfeger, 1990) and aging (Hasher & Zacks, 1989) that

interpret the cognitive difficulties of the young and the elderly as deficits in inhibitory processing. Younger children (Schachar & Logan, 1990a) and older adults (Kramer, Humphrey, Larish, Logan, & Strayer, 1992) are slower to stop than young adults, suggesting that the stopping process is the same over the life span but increases in speed (and therefore effectiveness) as childhood progresses to adulthood and declines in speed with senescence. There may be deficiencies earlier in childhood and later in aging. The youngest children studied by Schachar and Logan (1990a) were 7-year-olds; younger children may have qualitatively different stopping processes (or qualitatively different deficiencies). The older adults studied by Kramer et al. (1992) were all healthy and active and reported no more cognitive deficits in daily life than young adult controls. Perhaps older subjects who have significant cognitive deficits in daily life would show qualitatively different deficiencies in the stop signal paradigm.

Third, stop signal inhibition may be implicated in inhibitory psychopathology. Schachar and Logan (1990a) found that hyperactive children had trouble stopping on stop signal trials. Not only were they slower to inhibit than normal controls, they were less likely to inhibit altogether, responding more often than normals on stop signal trials. It was not the case that hyperactive children were less likely to "notice" the stop signal. Schachar and Logan (1990b) ran a dual-task experiment presenting the same stimuli they used in their stop signal experiment but requiring the children to make an overt response to the "stop signal" as well as an overt response to the primary task. Hyperactive children detected the signal as often as normal controls, and showed the same refractory effect in their reaction times. Their deficiency in stopping was not a deficiency in detecting the stop signal. Interestingly, stimulant medication (methylphenidate), which improves behavioral symptoms of hyperactive children, also improves their stopping performance (Tannock, Schachar, Carr, Chajczyk, & Logan, 1989).

III. STOP SIGNAL INHIBITION

The inhibition seen in the stop signal paradigm may be different from inhibition seen in other paradigms in important ways. Stop signal inhibition requires the subject to take a deliberate action. Stop signal inhibition represents an entire process, extending from (stop signal) stimulus to (internal inhibitory) response. Moreover, stop signal inhibition cannot be understood adequately without a formal theory of the stopping process.

These factors make stop signal inhibition different from the neurologically inspired inhibition that is seen in spreading activation and

connectionist models (e.g., Rumelhart & McClelland, 1986; see Dell & O'Seaghdha's chapter in this volume for a detailed discussion of inhibition in connectionist models). Neurological inhibition is negative activation; it is something that is subtracted from the current level of activation. The subtraction is intended to be understood neurally rather than cognitively. It happens as a result of the connectivity of the network; it does not require the subject to perform mental arithmetic. In these theories, subjects generally have little control over neurological inhibition. Inhibitory (and excitatory) connections are established by a long history of training and do not change during the course of the trial.

Stop signal inhibition is also different from what might be termed reactive inhibition. *Reactive inhibition* is inhibition that results from executing some process. It may be the most common kind of inhibition in the literature, challenging the ubiquitous "neurologically inspired" inhibition. Classical examples of reactive inhibition include response conflict (Berlyne, 1957), refractory effects in dual-task situations (Welford, 1952), and proactive inhibition in short-term memory (Wickens, 1970). Current examples include inhibition of return (Posner & Cohen, 1984), negative priming (Lowe, 1979; Neill, 1977; Tipper, 1985), top-down inhibition of distractors (Treisman & Sato, 1990), and a "center-surround" memory retrieval system (Carr & Dagenbach, 1990; Dagenbach, Carr, & Barhnardt, 1990). The idea behind reactive inhibition is that executing a process has a side effect that concurrent processes must overcome or leaves a residual effect that subsequent processes must overcome. The process that produces the inhibition may be engaged deliberately, but its inhibitory effect on concurrent and subsequent processes is not (usually) intended.

An important similarity is that all forms of inhibition work against some form of excitation. Neurologically, a tendency to decrease firing tempers a tendency to increase firing. In reactive inhibition, aftereffects of prior processing work against current-trial processing. And in the stop signal paradigm, the inhibitory process races against a "go" process. Stop signal inhibition is different from the others in that the excitatory and inhibitory processes are largely independent. The race model that is used to analyze stopping performance assumes that stop and go processes are independent, and the data for the most part are consistent with that assumption (De Jong, Coles, Logan, & Gratton, 1990; Jennings, van der Molen, Brock, & Somsen, 1992; Logan & Cowan, 1984; Osman, Kornblum, & Meyer, 1986, 1990). By contrast, neural inhibition works together in the same network as excitation, and reactive inhibition affects the same mechanisms and processes used for current-trial processing. These forms of inhibition are not independent of excitatory processes. Instead, they interact in complex ways. Nonlinearity has

important computational consequences in neural networks and connectionist models (Rumelhart & McClelland, 1986).

Relations between stop signal inhibition and other forms of inhibition have not been studied extensively. This is surprising because the experiments should be easy to do and the logic required to interpret them should be fairly straightforward. Stop signal inhibition can be compared in conditions that either require or do not require some other form of inhibition. If that other inhibition is based on the same process as stop signal inhibition, there should be interactions. Stopping may be easier when the go task requires inhibition, because the go task inhibition may prime or prepare the inhibition required for stopping. Alternatively, stopping may be harder when the go task requires inhibition if there is a limited capacity for inhibition.

Three studies have used this logic to compare stop signal inhibition with other forms of inhibition. First, Logan (1981) compared the ability to inhibit spatially compatible and spatially incompatible responses and found no difference. Apparently, the kind of inhibition required to resolve incompatible stimulus-response relationships (Kornblum, Hasbroucq, & Osman, 1990) does not interact with stop signal inhibition. Second, Kramer et al. (1992) compared the ability to inhibit responses to targets that were flanked by distractors that were associated with compatible and incompatible responses and found that subjects found it harder to inhibit responses to incompatible displays. Apparently, the kind of inhibition required to resolve response conflict (Eriksen & Eriksen, 1974) interacts with stop signal inhibition. Third, Jennings et al. (1992) compared stop signal inhibition with inhibition of the cardiac cycle and found that successful inhibition slowed heart rate. They suggested that stop signal inhibition and cardiac inhibition may depend on the same midbrain structures. These studies only scratch the surface. There is much more research to be done.

IV. BALLISTIC PROCESSES AND THE POINT OF NO RETURN

A. Discrete Tasks

A great deal of the interest in the stop signal paradigm has been focused on the possibility that some processes are *ballistic*—they cannot be stopped. Can something as automatic as Guy Lafleur's slapshot be stopped once it begins? A great deal of the modeling has been addressed toward ways of distinguishing ballistic from stoppable (or *controlled*) processes (Osman et al., 1986, 1990; see also Logan, 1981, 1982; Logan & Cowan, 1984; Zbrodoff & Logan, 1986). Early on, it became clear that no tasks were ballistic, in that virtually all responses can be inhibited if

the stop signal occurs early enough, even highly skilled responses like speaking (Ladefoged, Silverstein, & Papcun, 1973) and typing (Logan, 1982). So the focus shifted from tasks to components of tasks. Researchers asked whether components of tasks—processing stages—were ballistic or controlled (De Jong et al., 1990; Logan, 1981, 1983, 1985; Osman et al., 1986, 1990).

An important concept in research on ballistic processes is the point of no return. The *point of no return* is a cusp that divides controlled processes from ballistic ones. If the stop process finishes before the go task has reached the point of no return, the response can be inhibited. If the stop process finishes after the point of no return, the response cannot be inhibited (De Jong et al., 1990; Osman et al., 1986, 1990). Several studies have tried to find the point of no return and locate it in terms of processing stages defined psychologically or psychophysically. The studies agree in concluding that if there is a point of no return, it is pretty close to the effectors. They disagree on how close. It is possible to make a case that there are no ballistic processes—that the point of no return is “on the skin” (at the interface between the person and the manipulanda)—and it is possible to make a case that certain late processes are beyond the point of no return.

One set of studies modeled the stop signal paradigm as a race between stop and go processes, and used the additive factors logic (Sternberg, 1969) to locate the point of no return in terms of stages in the go task. The idea was to manipulate factors that affect reaction time to the go task and see if those factors affect the probability of inhibiting responses (see Logan, 1981; Osman et al., 1986, 1990). According to the race model, the stop process races against the go process. If the stop process finishes before the go process reaches the point of no return, the response can be inhibited. If the stop process finishes after the go process reaches the point of no return, the response cannot be inhibited. Factors that increase the duration of go task stages before the point of no return should increase the probability of inhibition; those factors provide subjects with more time to stop, and therefore make stopping more likely. Factors that increase the duration of go task stages after the point of no return should have no effect on the probability of inhibition. The time available for subjects to stop is not affected by increasing the duration of ballistic processes.¹

¹The predictions that were actually tested were more precise than the ones described here. The race model allows one to predict exactly how much the probability of inhibition should change given an increase in the duration of the controlled phase. Those predictions require a deeper analysis of the race model than I have presented so far. The interested reader is referred to the original papers or to subsequent sections of this chapter (see Section VIII).

Studies conducted in this fashion showed that discriminability (Logan, 1981; Osman et al., 1986), lexical status (Osman et al., 1986), stimulus-response compatibility (Logan, 1981), and response complexity (Osman et al., 1990) affect controlled stages of processing. The only factor that appeared to affect ballistic stages of processing was stimulus-response repetition (Osman et al., 1986). Repetition is a difficult factor to interpret, however, because it has many different effects on the information processing system and affects several stages (Kornblum, 1973). It may even change the way the task is organized (Fletcher & Rabbitt, 1978).

A study by De Jong et al. (1990) used psychophysiological measures to search for the point of no return. They measured responses at three different stages in the motor system: They measured lateralized readiness potentials on the scalp above the motor cortex, electromyograms from the responding muscles in the forearm, and the force with which subjects' hands squeezed the response manipulanda (zero displacement dynamometers). They argued that if any of these stages was past the point of no return, then it should not be possible to inhibit the response once it reached that stage. Thus, if central preparation were after the point of no return, it should not be possible to inhibit responses that show a lateralized readiness potential; if muscle activation were after the point of no return, it should not be possible to inhibit responses that show an electromyographic response; and if the squeeze itself were after the point of no return, then it should not be possible to stop a squeeze once it begins.

The data suggested that central preparation was before the point of no return; subjects could easily inhibit responses that produced lateralized readiness potentials. The data also suggested that muscle activation was before the point of no return; subjects could inhibit responses that produced electromyographic responses. Finally, the data suggested that the squeeze itself was before the point of no return. There were many partial squeezes on stop signal trials, squeezes in which less than the critical amount of force was exerted, as if subjects stopped in the middle of squeezing. Thus, if there are processes after the point of no return, they must be very late and very brief in duration.

Jennings et al. (1992) replicated the results of De Jong et al. (1990), finding evidence of electromyographic activity on some trials on which the final response (switch closure) was inhibited. The interval between electromyographic activity and switch closure was brief, 60–80 msec, which suggests that ballistic processes must be very brief in duration and late in the processing sequence.

The electromyographic data can be interpreted differently, given a different definition of the point of no return. De Jong et al. (1990) and Jennings et al. (1992) defined the point of no return as the point at which

subjects could not stop themselves from *completing* the motor response. Thus, De Jong et al. interpreted partial (subthreshold) squeezes as cases of successful inhibition, and Jennings et al. interpreted all trials that did not result in a key closure as cases of successful inhibition. However, Osman et al. (1986, 1990) defined the point of no return as the point at which subjects could not stop themselves from *beginning* the overt movement. From that perspective, the partial squeezes observed by De Jong et al. would be interpreted as cases of unsuccessful inhibition. From that perspective, if muscle activation is the point of no return, the motor response should always begin if there is electromyographic activity. De Jong et al. (1990) found no cases in which there was electromyographic activity and no (partial) squeezes; electromyographic activity was always accompanied by a measurable motor response. Osman et al. (1986, 1990) would interpret those data as evidence that muscle activation was the point of no return.

Logan and Cowan (1984) presented a mathematical analysis that corroborates the conclusion that ballistic processes must be very brief. They modeled the stop signal paradigm as a race between stop and go processes (see the following). In the full development of the model, they included two parts to the go process: a controlled phase and a ballistic phase. They showed that estimates of stop signal reaction time include not only the reaction time to the stop signal, but also the duration of the ballistic phase of the go task. They note that estimates of stop signal reaction time are on the order of 200 msec. This may be interpreted as an upper limit on the duration of ballistic processes. The repetition effect observed by Osman et al. (1986) may be interpreted as a lower limit: If repetition affected the ballistic stage by 25 msec (as Osman et al. estimate), then the ballistic stage had to be at least 25 msec in duration.

B. Continuous Tasks

Several investigators have studied the ability to inhibit continuous responses, such as arm movements, and found they are inhibited in much the same way as discrete responses. Stop signal reaction times are the same as, if not faster than, those observed with discrete responses (Henry & Harrison, 1961; Megaw, 1972; Vince, 1948). Continuous tasks, such as speaking (Ladefoged et al., 1973; Levelt, 1983) and typing (Logan, 1982; Long, 1976; Rabbitt, 1978), can also be stopped very quickly. Typically, subjects emit only a syllable or two, or a keystroke or two before stopping. Stop signal reaction times are on the order of 200–400 msec.

The possibility of ballistic processes in continuous tasks is interesting. One often gets the impression that very familiar continuous tasks, such as chopping onions or typing sentences, go on ballistically—on automatic pilot—once they are initiated. However, the data suggest there is very close control. Subjects may not be able to prevent themselves from making a few responses (e.g., uttering a few syllables or striking a few keys) before stopping, but this most likely reflects the latency of the response to the stop signal rather than a substantial ballistic component. The stopping process, once begun, takes time to finish, and in tasks like speaking and typing, subjects may have time to execute several responses between the onset of the stop signal and the completion of the stopping process. Much of the data can be interpreted that way. In typing, for example, the number of characters typed after a stop signal multiplied by the rate at which characters are typed yields values close to 200 msec (Logan, 1982), which is a typical stop signal reaction time for a discrete task. The only evidence of ballistic responding in a continuous task comes from Logan's (1982) study of typewriting, in which skilled typists tended to type out familiar verb endings (*es* and *ed*) and the word *the* together with the space that followed it. In a similar vein, however, Osman et al. (1990) found that subjects had a hard time inhibiting second and subsequent responses in preprogrammed sequences if they failed to inhibit the first.

How can we reconcile the close control over continuous tasks with our intuition that we perform continuous tasks ballistically? The facility with which people inhibit continuous tasks suggests they have very close control over their performance, but the control can be realized in several ways. One possibility is that the task is organized in a single feedback loop, in which the consequences of each response or two are monitored and used to base decisions on what to do next. In this view, the stopping process is an integral part of the feedback loop and stop signal reaction time represents the time it takes the feedback loop to operate. Another possibility is that the task is organized hierarchically in several feedback loops. Low-level processes may be controlled by local feedback loops whose parameters are set by higher level processes. The different feedback loops may have different latencies; lower level loops may be faster than higher level loops. The stopping process may act outside of this hierarchy, separate from its feedback loops. The stopping process might inhibit low-level loops directly or it may change the parameters that high-level loops pass to low-level loops. In either case, stop signal reaction time does not represent the duration of a feedback loop within the control system for the go task.

Hierarchical control seems more likely on other grounds (Dell, 1986; Lashley, 1951; Rumelhart & Norman, 1982). In typing, for example, there are several "spans" that reflect different levels of organization—typists

plan phrases, words, and keystrokes (Logan, 1983; Salthouse, 1986). Vallacher and Wegner (1987) pointed out that people can think of what they are doing at several different levels. I may be pressing keys, typing words, writing sentences, explaining the stop signal paradigm, or fulfilling my obligations to Dale and Tom, the editors who asked me to write this chapter. Reason and Myceilska (1982) studied everyday errors—action slips—and found they were best explained assuming hierarchical control. Many errors involve doing the wrong low-level activity in the right way (e.g., driving directly home instead of stopping for milk). Thus, the control seen in the stop signal paradigm may operate outside the go-task control system, intervening when necessary to stop action, but otherwise quiescent. This underscores the point made earlier, that stopping is an extreme and unsubtle form of control.

C. Ballistic, Automaticity, and Autonomy

In theory, it is important to distinguish between ballistic processes, automatic processes, and autonomous processes. Ballistic processes are often a defining characteristic of automatic processes; once initiated, they run on to completion whether or not the person intends them to (Posner & Snyder, 1975; Shiffrin & Schneider, 1977). However, ballistic processes are not necessarily automatic. Similarly, autonomous processes are ballistic in that once initiated, they run on to completion whether or not the person intends them to. But autonomous processes can also be initiated without intention, and that is not necessarily true of ballistic processes (Zbrodoff & Logan, 1986). Ballistic processes have to do with whether a process runs on to completion without intention, not with whether it begins without intention. In principle, a process begun intentionally can run on to completion ballistically.

The stop signal paradigm addresses ballistic processes directly. It addresses automaticity and autonomy indirectly, through their relations to ballistic processes. Thus, the stop signal paradigm may be used as part of a battery of tasks to assess automaticity and autonomy. Zbrodoff and Logan (1986) suggested using Stroop-type tasks or Eriksen-type flanker tasks to determine whether processes begin without intention, and the stopping task to determine whether a process runs on to completion without intention.

V. COMPLICATING THE STOPPING PROCESS

The stop signal paradigm, as described so far, involves a simple response to the stop signal, which is analogous to a simple reaction time

task. There is only one signal and only one response to make to it. It is like blowing a whistle to stop play in a hockey game: everything stops. There is nothing to do until "the next trial," when play begins again. Some investigators have begun to explore complications of the stopping process. Some examine the change paradigm, which models situations that require doing something else after stopping, like switching from autopilot to real thinking when a student asks a question during a lecture. Others examine selective inhibition, modeling situations in which the stop signal is not always a signal to stop (e.g., when your little brother vs. your mother discovers your hand in the cookie jar) and situations in which not everything must be stopped (e.g., trying to avoid using a sore finger, trying to avoid sensitive topics in conversation).

A. Change Paradigm

The complication of the stop signal paradigm that has been explored most often is the *change paradigm*, in which the stop signal tells the subject to change from one task to another. There are three tasks: the go task, the stop task, and the change task. Subjects perform the go task unless a stop signal occurs. The stop signal triggers the stop task and the change task. It tells subjects to inhibit their response to the go task (the stop task) and make an overt response to the stop signal (the change task). The stop signal itself is the stimulus for the change task, though it need not be. The stop signal could indicate a change in the task to be performed on the go-task stimulus, but it has not done so in the studies I know. Several studies have used simple (overt) reaction time to the stop signal as the change task (Logan, 1983, 1985; Logan & Burkell, 1986). Another presented two stopping tones and had subjects stop to both tones but press different keys for each tone in the change task (Logan, 1983). Others changed the go task stimulus and had subjects respond to the new go task stimulus (Logan, 1982, 1983, 1985).

The change paradigm is interesting because there are three processes involved, not just two. The go process and the stopping process seem to function pretty much the same in the change paradigm and the stop signal paradigm. The race model provides an accurate account of go signal and stop signal processing. Reaction times to the go signal and the stop signal may be a little slower in the change paradigm than in the stop signal paradigm, but the race model fits just as well (Logan, 1982; Logan & Burkell, 1986). The goodness of fit and the similarity of parameter values are interesting because they support the assumption that stop and go processes are independent of each other, which is a key assumption in the race model. The go and stop processes are not affected much by adding another task.

Reaction times to the change task depend on what the subject does in the stop task: Change reaction times show the usual dual-task interference or psychological refractory period if the subject fails to inhibit the response to the go task. Change reaction times are elevated on signal-respond trials, and the elevation is greater the shorter the delay. However, change reaction times show no interference at all if the subject succeeds in inhibiting the response to the go task. Change reaction times are fast on signal-inhibit trials and change very little with stop signal delay (Logan, 1985; Logan & Burkell, 1986; Schachar, Marriott, Tannock, & Logan, 1993; Tannock, Schachar, & Logan, 1993).

The change task reaction times are interesting because of their implications for theories of dual-task interference. The data show that dual-task interference depends on whether or not a response was executed. This suggests response execution as the locus of the processing bottleneck that is responsible for dual-task interference (Logan & Burkell, 1986). This conclusion contrasts with Pashler's (1984, 1989) claim that stimulus-response translation processes are the locus of the bottleneck (see also De Jong, 1993). It may be possible to reconcile the views by proposing that translation is linked to execution, such that execution is a necessary (obligatory) consequence of translation. If that were the case, then the only way to prevent execution would be to prevent translation (or to inhibit via some other, peripheral mechanism). Successful inhibition would leave the stimulus-response translation processes free to deal with the response to the change task, eliminating dual-task interference, just as we observed.

The key to this proposal is the assumption that translation and execution processes are linked, so that execution is a ballistic consequence of translation. This assumption received little support in studies that searched for the point of no return (see Section IV), where there was little evidence for ballistic processes. The specific hypothesis that translation necessitates execution was addressed by Osman et al. (1990), and they found little support for it. So it is not clear how to resolve the difference between Pashler's (1984, 1989) claim and Logan and Burkell's (1986) claim about the locus of the dual task bottleneck. The finding that inhibiting a response eliminates dual-task interference is an interesting and important challenge for future research.

The change task is interesting from a clinical perspective. Schachar et al. (1993) tested hyperactive children on the change task and found their stopping performance was poorer than that of normal controls, as Schachar and Logan (1990a) found with the stopping task. Interestingly, hyperactive children were not deficient on the change task. Their overt reaction times to the change signal were not much worse than control subjects'. They appear to have trouble stopping one task, but have no trouble switching to another. This is consistent with Schachar and

Logan's (1990b) finding that hyperactive children were no more susceptible than normal controls to dual-task interference. Tannock, Schachar, and Logan (1993) found that stimulant medication improved the stopping performance of hyperactive children in the change task, just as it did in the stopping task (Tannock et al., 1989). There was a suggestion that the dose-response curves were different: Whereas inhibition performance improved with dose in the simple stopping study, it improved and then declined as the dose increased in the more complex change study.

B. To Stop or Not to Stop

Riegler (1986) complicated the stopping task in a classical Donderian fashion. He presented two different stop signals and compared the usual simple stopping task, in which subjects stopped responses when either tone sounded, with a *stop/no-stop* task, in which subjects stopped responses when one tone sounded but not when the other tone sounded (i.e., get your hand out of the cookie jar if your mother comes, but not if your brother comes). He found substantially longer stop signal reaction times in the stop/no-stop condition, corresponding to single-task differences between simple reaction times and go/no-go reaction times to the same tones.

Riegler (1986) manipulated the difficulty of the go task independently of the difficulty of the stop task. The go task involved two or four choices. The easy conditions were much like simple stopping: stopping was fast and independent of the go task. However, performance started to break down in the most difficult combination of stop and go tasks—the four-choice stop/no-stop condition. The stop and go tasks were no longer performed independently. They interfered with each other, as if they were competing for shared resources. This result is important because it is the first suggestion that the stopping process may suffer the same limitations as other concurrent tasks.

C. Selective Stopping

Logan, Kantowitz, and Riegler (1986) complicated the stopping process by asking subjects to engage it selectively, inhibiting one response but not another. Subjects performed a two- and four-choice go task and were presented with stopping tones. In the *selective stopping* task they were required to inhibit only the key press assigned to the right index finger, as if it were a sore finger they were trying to avoid using. Key presses assigned to other fingers were to be executed whether or not a

stop signal sounded. In the simple stopping task, they stopped key presses assigned to all fingers when a stop signal sounded, just as in the usual stop signal paradigm.

Subjects were able to stop selectively. The selective stopping task produced inhibition functions that were well fit by the race model. The requirement to inhibit selectively increased stop signal reaction time substantially, and the increase was greater for the four-choice task, where one response out of four had to be inhibited, than for the two-choice task, where one response out of two had to be inhibited. By contrast, in the control task, stop signal reaction times were fast and not affected much by the number of choices in the go task. Logan et al. (1986) interpreted this interaction by proposing two modes of inhibition, a global mode for the control task that inhibited all responses, and a local mode for the selective inhibition task that focused on a single response (anticipating, perhaps, the distinction between central and peripheral inhibition mechanisms made by De Jong et al., 1990; Jennings et al., 1992). Stop signal reaction time depended on the difficulty of the discrimination that the stop process had to make in selecting the response to be inhibited. The global mode should be fast because all responses are inhibited; there is no discrimination required. The local mode should be slower because it requires discrimination, and it should be slower to discriminate the greater the number of alternatives. It should take longer to discriminate one response from four alternatives than from two alternatives, following a stop signal equivalent of Hick's law.

Logan et al. (1986) performed two experiments varying the stimuli used in the go task. One experiment used letters presented in the center of the screen that were assigned randomly to the keys. The other used X's presented in four different locations that were assigned to response keys that were spatially compatible with them. The second experiment produced much faster go task reaction times than the first, but the magnitudes and the patterns of the stop signal reaction times were the same. This is important because it rules out one interpretation of the pattern of stop signal reaction times: Subjects could postpone their stopping response until they could see which response was required in the go task. This would produce longer stop signal reaction times in selective stopping than in simple stopping because simple stopping need not be postponed in the same manner as selective stopping. It would produce longer selective stop signal reaction times in the four-choice task than in the two-choice task because subjects would have to wait longer to know which response to inhibit in the four-choice task. These predictions are in accord with the data. However, postponing stopping should also produce longer stop signal reaction times in the first experiment than in the second, corresponding to the longer go task reaction times

in the first experiment, and this prediction was not confirmed. The delay in selective stopping cannot be attributed to postponing the stop process to monitor the go process.

VI. AFTEREFFECTS OF INHIBITION

What are the consequences of inhibiting a response? What aftereffects does the inhibition process leave behind? Kramer et al. (1992) examined the immediate consequences of inhibition, asking whether inhibiting the response on trial n affected performance on trial $n + 1$. They found that go signal reaction time was slower on trials following successful inhibition than on control trials. The effect was significant for both young and old subjects, but it was stronger for old subjects than for young ones (50 msec vs. 21 msec). This effect is interesting and merits further investigation. It would be interesting to compare it with negative priming, which is also interpreted as an aftereffect of inhibition (Tipper, 1985).

The inhibition process does not appear to leave behind much in the way of longer term aftereffects. Logan (1985) compared repetition priming effects—the reduction in reaction time produced by presenting something twice—produced by stimuli whose responses had and had not been inhibited on first presentation, and found equivalent repetition priming. This corroborates Logan's (1990) conclusion from transfer experiments that memory for response execution does not play a large role in repetition priming.

Logan (1983) used the same materials (words on which subjects performed category and rhyme judgments) but tested recognition memory rather than repetition priming. He found equivalent memory for words associated with responses that had and had not been inhibited during the study phase, corroborating the results of the repetition priming experiments.

Experiments by Zbrodoff and Logan (1986) suggest that inhibition may leave some measurable aftereffects. They performed a stop signal experiment in which the go task was verification of simple arithmetic problems like $3 + 4 = 7$, and later compared memory for problems associated with responses that had and had not been inhibited. This time, memory was worse for inhibited problems, as if subjects stopped encoding the problem into memory when they stopped trying to solve it. The differences between these results and Logan's (1983) must be due to the nature of the go tasks. Arithmetic problems may be less memorable than words, particularly words subjected to semantic orienting tasks. Arithmetic verification may be less automatic than category or

rhyme judgments and therefore less likely to run on to completion ballistically.

Research on the aftereffects of inhibition has only just begun. Very little is known. Very little is even known about what to look for. This strikes me as a promising area for future research, particularly for researchers interested in relating stop signal inhibition to other forms of inhibition, which also appear as aftereffects of prior acts of control.

VII. MECHANISMS OF INHIBITION

Little is known about the mechanisms underlying the stop process, except that they take time to finish and when they finish, they stop responses. The data on young adults are consistent with the hypothesis that a single, global mechanism is responsible for inhibition in all kinds of stopping tasks; a wide variety of tasks can be stopped in about 200 msec (see Logan & Cowan, 1984, for a review). Most of the data come from simple stopping experiments, in which there is only one stop signal and it tells subjects to inhibit any and all responses. A single, global mechanism may be responsible for simple stopping performance. Selective stopping, studied by Riegler (1986) and Logan et al. (1986), may require a more complicated mechanism. A selective stopping mechanism may operate differently, depending on the nature of the selection required. More research on selective stopping will be necessary to reveal the nature of simple and selective stopping mechanisms.

Two recent psychophysiological investigations of stopping provoked speculations about the nature and number of stopping mechanisms. De Jong et al. (1990) examined event-related potentials in the electroencephalogram while subjects processed go and stop signals. Their analysis suggested two mechanisms of inhibition: a central one that operated selectively, inhibiting central preparation of the required response, and a peripheral one that operated nonselectively, inhibiting any and all responses. Two pieces of evidence suggested a central mechanism. The first was a positive deflection in the waveform of the event-related potential at frontocentral electrode sites on trials on which subjects successfully inhibited their responses (*signal-inhibit* trials) that was not present on trials on which subjects failed to inhibit their responses (*signal-respond* trials). The second piece of evidence was a difference in the amplitude of the *lateralized readiness potential*, which reflects central motor preparation. The amplitude was lower on *signal-inhibit* trials than on *signal-respond* trials, and on trials on which no stop signal was presented, as if the stopping process inhibited central motor preparation.

The peripheral mechanism was suggested when attempts to account for the observed inhibition with central mechanisms failed. Differences in both central waveforms appeared (slightly) too late to cause inhibition of muscle activation and the final response. The amplitude of the lateralized readiness potential was smaller on signal-inhibit trials, but it was still larger than the amplitude needed to trigger electromyographic activity in the responding muscles—it ought to have produced a response, but somehow the response was inhibited. Do Jong et al. (1990) argued that responses were inhibited by a peripheral inhibition mechanism that operated separately from the central processes that prepare and inhibit motor responses.

Jennings et al. (1992) monitored cardiac activity rather than cortical activity during stopping, and came to the same conclusion as De Jong et al. (1990), albeit from a different direction. Jennings et al. (1992) found that the interval between heart beats was longer on signal-inhibit trials than on signal-respond trials or trials on which no stop signal was presented. They interpreted this interaction between cardiac inhibition and stop signal inhibition as suggesting that the stop process may be controlled by the same midbrain system that controls heart rate. They argued that the midbrain system inhibited central commands for a motor response, just like De Jong et al. (1990), but this time with more direct evidence for a noncentral (noncortical) locus.

The distinction between central preparation and peripheral inhibition is consistent with current theories of motor control. Brooks (1986) and Bullock and Grossberg (1988), among others, have argued that cortical processes determine which responses are to be made and midbrain processes determine when they start and stop. Nevertheless, the evidence for central and peripheral inhibitory mechanisms is scant and depends as much on argument as on fact.

The arguments De Jong et al. (1990) made about timing depend on very small differences and required generalizations from different tasks, subjects, and paradigms, which may not be valid. The arguments about the amplitude of the lateralized readiness potential assume that the critical factor is amplitude (and not, for example, the smoothness with which the potential rises to a critical amplitude). The relation between amplitude and responding is an empirical generalization, not a theoretical premise. It may be correlated with, but conceptually different from, the critical factor that triggers the motor response.

Similarly, the arguments of Jennings et al. (1992) for a midbrain inhibition system assume that the stopping process affects cardiac control directly. The stopping process could affect cardiac control indirectly, by eliminating one of the factors that affect it. Other studies show that responding during the critical phase of the heart rate is necessary to prevent prolongation of the heart beat (Jennings & Wood, 1977). Per-

haps, stop signal inhibition prolongs heart rate by removing the response that is necessary to prevent prolongation. More research will be necessary to distinguish these alternatives. The idea of two inhibition mechanisms—cortical and midbrain—is tantalizing, but it is only the beginning. There is a great deal more to the story of the nature and the locus of stop signal inhibition processes that remains to be discovered.

VIII. THE RACE MODEL

So far, I have been able to explain most of the important results in the stop signal paradigm without getting very deeply into the race model that underlies the analysis of stopping performance. Most of the conclusions were not based on the data directly, but rather were based on the data as interpreted through the race model. In order to understand how the conclusions relate to the raw data (or to the data one might collect in one's own stopping experiment), it is necessary to understand how the race model accounts for stopping performance. The rest of this chapter is intended to provide such an understanding. Up to now, the chapter has shown the benefits gained from studying the stop signal paradigm and has suggested benefits to be gained from future studies. The rest of this chapter is the gauntlet that must be run to appreciate the intricacies of those benefits or to gain similar benefits oneself. There is no avoiding the race model.

Stop signal inhibition involves a whole process—the stopping process—working against another process—the go process. The go process is easy to understand. Much of the literature in experimental psychology is dedicated to understanding go processes, and what was learned there can be adapted to our purposes. The stop process is harder to understand because it is not directly observable. When the stop process beats the go process, there is no response to observe and no latency to measure. In order to understand stop signal inhibition, we must have a theory of the stop process. That theory must tell us when the stopping process is executed, how to measure the latency of the stopping process when it is executed, which factors affect the probability and latency of stopping. This is where the race model comes in.

The stop signal paradigm is interesting because in contrast with other paradigms, there is no controversy over the basic theory. All approaches assume that performance depends on a race between the stopping process and the go process. In some cases, the race is described informally (Lappin & Eriksen, 1966; Vince, 1948). In other cases, the race is described more formally, relating the inhibition function to the distribution of reaction times in the go task (Logan, 1981; Logan, Cowan, &

Davis, 1984). Others describe the race in terms of the distributions of finishing times of stop and go processes (Colonus, 1990; Logan & Cowan, 1984; Osman et al., 1986, 1990). Still others assume specific forms for the distributions (Logan & Burkell, 1986; Ollman, 1973). But the basic theory in each case is a race between the stop and the go processes.

The race model derives from two rules that define the conditions for the outcomes of the race: (1) If the stop process is faster than the go process, the response is inhibited. Formally, if $T_g > T_s + t_d$ then the stopping process wins and the response is inhibited. T_g and T_s are the finishing times of the go process and the stopping process, respectively, and t_d is the delay between the onset of the go signal and the onset of the stop signal. (2) If the go process is faster than the stop process, the response is executed. Formally, if $T_g < T_s + t_d$ then the go process wins and the response is executed. All theories assume that T_g and T_s are independent random variables, so the outcome of the race is stochastic rather than deterministic. On some proportion of the trials, the stop process will be faster than the go process; on the remaining proportion of the trials, the go process will be faster. That proportion depends on the distributions of finishing times for the stop and go processes and on t_d , the delay between the onsets of the stop and go stimuli.

These assumptions are sufficient to account for all of the observed behavior in the stop signal paradigm. They account for inhibition functions, the effects of go task reaction time on the inhibition functions, and the speed of responses that escape inhibition (i.e., responses that occur when $T_g < T_s + t_d$). They provide three ways to estimate parameters of the distribution of stop process finishing times (i.e., stop signal reaction time). Performance in the stop signal paradigm seems well described as a race between stop and go processes (Logan & Cowan, 1984; Osman et al., 1986).

The race model succeeds because it is general. It addresses only the finishing times of the stop and go processes, and it imposes little constraint on those finishing times. All Logan and Cowan (1984) required, for example, was that the finishing times come from some distribution. The form of the distribution did not matter (see also Colonius, 1990; cf. Ollman, 1973). More importantly, the nature of the process that gave rise to the distribution of finishing times did not matter either. The only constraint was that the stopping process and the go process should function independently, and that is possible theoretically (e.g., in multiple resource theory) and appears to be the case empirically (De Jong et al., 1990; Jennings et al., 1992; Logan & Cowan, 1984).

The generality of the race model is a weakness as well as a strength. The strength is obvious: The race model accounts for the data very well.

It makes quantitative predictions that are impressively accurate. The weakness is that the race model is not a process model. The race model does not depend on the nature of the processes that race against each other, only on their finishing times. Thus, the race model itself does not give any special insight into the nature of those processes. However, the race model may be used as a tool to investigate hypotheses about the nature of the underlying processes. Hypotheses about underlying processes are often couched in terms of reaction time and to factors that affect it. The race model provides a way to measure the reaction time of the stopping process and to determine which factors affect it. Thus, the race model is not the final answer to questions about performance in the stop signal paradigm, but it is an important and apparently necessary first step toward the final answer. A process model that explains how stopping actually occurs must be built within the constraints of the race model, filling in the details of processes that the race model treats abstractly.

A. Inhibition Functions

Inhibition functions relate the probability of inhibiting or the probability of responding to stop signal delay. They are important because stop signal delay (or some transformation of delay) is the most important independent variable in the stop signal paradigm; inhibition functions plot the effect of delay. Inhibition functions are important theoretically because they reflect the outcome of the race between the stop process and the go process. The race and the distributions of finishing times for the stop and go processes determine the shape of the function. Finally, inhibition functions are important because they are the focus of most of the data analysis in stop signal experiments.

Two questions arise immediately: Should we plot probability of inhibiting or probability of responding? And what measure of delay should those probabilities be plotted against? The probability of inhibiting is 1 minus the probability of responding, so it makes no difference formally which one is plotted. Choosing a measure of delay is more difficult. As experimenters, we have control over the delays between the onset of the go signal and the onset of the stop signal, which determine the *starting times* of the underlying processes. As theorists, we wish to control the finishing times of the underlying processes. The different measures of delay approximate control over the *finishing times*. They tell us how to set the delay between stimulus onsets—how to lag the starting times—in order to achieve a certain delay between the finishing times. They are all motivated by the race

model, so it is important to explain how the race model accounts for inhibition functions before explaining the differences between the measures.

The easiest way for me to explain how inhibition functions result from the race model is to begin with two simplifying assumptions: (1) the finishing time of the stopping process (*stop signal reaction time*) is constant, and (2) there are no ballistic processes. The assumption about stop signal reaction time makes the mathematics easier, but more importantly, it allows a graphic representation of the underlying processes that illustrates the relationships very clearly (see Figure 1). The assumption is bound to be wrong—no one would seriously propose a psychological process that took a constant amount of time to execute—but it turns out in practice that the correctness of the assumption is not very important. Logan and Cowan (1984) (mostly Cowan) analyzed the formal consequences of the assumption, and found that it introduced very small measurement errors. De Jong et al. (1990) analyzed the effects of violating the assumption with Monte Carlo simulations and found they were small if stop and go processes were independent. Logan and Cowan (1984) compared estimates of stop signal reaction time that assumed it was constant with estimates that assumed it was variable (see the following) and found good agreement between them. Measures that do not require the assumption often require better data (Colonius, 1990; Logan & Cowan, 1984).

The assumption that there are no ballistic processes is more controversial, as the earlier section on ballistic processes and the point of no return attested (see Section IV). Nevertheless, it is clear that if there are ballistic processes, they occur late in the sequence, close to or during motor execution. The clearest candidate for a point of no return is muscle activation, and this candidacy depends on defining the point of no return as the point beyond which subjects cannot stop themselves from *beginning* the movement (Osman et al., 1986, 1990). The point beyond which subjects cannot stop themselves from *completing* the movement is clearly after muscle activation (De Jong et al., 1990). If we define that point as the point of no return, then we can safely ignore ballistic processes. We adopt that definition for the remainder of this chapter, mostly because it makes the model easier to explain. Logan and Cowan (1984) develop the model completely, assuming that stop signal reaction time is a random variable and that there may be ballistic processes in the go task.

The top panel of Figure 1 depicts a time line that indicates the onset of the go stimulus (primary task stimulus) and the onset of the stop signal. The distribution of reaction times to the go stimulus is drawn above the time line. The constant finishing time of the stopping process is drawn as a vertical line that intersects the time line. It extends




FIGURE 1

Graphic representation of the assumptions and predictions of the race model, indicating how the probability of inhibiting a stimulus— $P(\text{INHIBIT})$ —and the probability of responding given a stop signal— $P(\text{RESPOND})$ —depend on the distribution of primary (go) task reaction times, stop signal reaction time, and stop signal delay. (The representation assumes that stop signal reaction time is constant.) From "On the Ability to Inhibit Thought and Action: A Theory of an Act of Control" by G. D. Logan and W. B. Cowan, 1984, *Psychological Review*, 91, 295–327, Figure 2. Copyright 1984 by the American Psychological Association. Reprinted by permission.

vertically above the time line, dividing the go process reaction time distribution into two parts. The part on the left side of the line represents go reaction times that win the race. They are faster than stop signal reaction time plus stop signal delay (i.e., $T_g < T_s + t_d$). The area under the distribution to the left of the line represents the probability of responding, given a stop signal at that delay. The part of the distribution on the right side of the line represents the go trials that lose the race to the stopping process. They are slower than stop signal reaction time plus stop signal delay (i.e., $T_g > T_s + t_d$). The area to the right of the line represents the probability of inhibiting at that stop signal delay.

The second panel of Figure 1 illustrates how inhibition functions are produced by varying the stop signal delay. It depicts the same time line as the top panel, with the go signal presented at the same time and the same distribution of go signal reaction times. However, this time, the stop signal is delayed. Consequently, the point at which the stopping process finishes is delayed. It appears farther to the right on the time line, cutting the go signal reaction time distribution in a different place. Compared to the top panel, the probability of responding is higher and the probability of inhibition is lower.

It is easy to see how this idea can be extended to produce a complete inhibition function. If the stop signal is presented early enough, the probability of responding will equal 0, and the probability of inhibition will equal 1. If the stop signal is presented late enough, the probability of responding will equal 1, and the probability of inhibition will equal 0. Points in between these extremes will produce a smooth function in which the probability of responding increases from 0 to 1 and the probability of inhibition decreases from 1 to 0 as delay increases. Formally,

$$P_r(t_d) = \int_{T_s+t_d}^{\infty} f_g(t) dt \quad (1)$$

where $P_r(t_d)$ is the probability of responding at delay t_d , T_s is the stop signal reaction time, and $f_g(t)$ is the distribution of reaction times to the go signal. The go stimulus is presented at $t = 0$. An important point to be taken from Figure 1 and Eq. (1) is that the shape of the inhibition function depends on the shape of the reaction time distribution. This relationship is emphasized when inhibition functions are plotted using the probability of responding rather than the probability of inhibiting. The "central" part of the inhibition function, where the probability of responding is around .5, depends on the mean of the reaction time distribution. The "slope" of the inhibition function depends on the variability of the reaction time distribution. The more variable the distribution, the broader and flatter the inhibition functions.

The third panel in Figure 1 illustrates the effects of mean reaction time to the go signal on the inhibition function. The stop signal is pre-

sented at the same time (i.e., the stop signal delay is the same), but the go signal reaction time distribution is delayed. Consequently, the internal response to the stop signal cuts off less of the go-signal reaction time distribution; the probability of responding is lower and the probability of inhibiting is higher.

The fourth panel of Figure 1 illustrates the major point of the race model: The inhibition function depends on the relative finishing times of the stop and go processes, not on their relative starting times. In the fourth panel, the go process reaction times have been delayed relative to the first panel, but the presentation of the stop signal has been delayed by the same amount. Consequently, the stopping process finishes at the same time, relative to the go process, as it does in the first panel. Go signal reaction time and stop signal delay are very different, but the probability of inhibition and the probability of responding given a signal are the same because the relative finishing time is the same.

The compensatory relationship between go reaction time and stop signal delay has been investigated several times. Different experimental conditions (Logan, 1981, 1982; Osman et al., 1990), tasks (Logan et al., 1984), strategies (Logan, 1981), subjects (Logan & Cowan, 1984), and subject populations (Schachar & Logan, 1990a) produce different go signal reaction time distributions and different inhibition functions, but it is generally possible to bring the functions into alignment by plotting them against some transformation of delay that takes into account the differences in go signal reaction time distributions (but see Osman et al., 1986, 1990, Experiment 2; Schachar & Logan, 1990a; Tannock et al., 1989). Simply subtracting stop signal delay from mean go signal reaction time (i.e., plotting inhibition functions against mean reaction time minus stop signal delay) is sufficient to bring most functions into alignment. When there are large differences in variability, other transformations are necessary to bring the functions into alignment (Logan et al., 1984; Schachar & Logan, 1990a). All of these transformations are intended to express the inhibition function in terms of the relative finishing times of the stop and go processes. They will be discussed more extensively later (see Section IX).

Logan and Cowan (1984) developed a complete version of the race model in which the stop process and the go process were both random variables (i.e., in which stop signal reaction time was not assumed to be constant). Their analysis is too complex to reiterate here, but the main results are straightforward. Logan and Cowan (1984) (again, mostly Cowan) treated the inhibition function as if it were a cumulative distribution (it looks like one) and related it to the distribution of stop process and go process finishing times. They showed that the mean of the

inhibition function is simply the difference between the mean go signal reaction time and the mean stop signal reaction time.² Explicitly,

$$\bar{T}_i = \bar{T}_g - \bar{T}_s \quad (2)$$

where $\bar{T}_i$ is the mean of the inhibition function and $\bar{T}_g$ and $\bar{T}_s$ are the means of the go and stop process reaction times, respectively.

The relationship in Eq. (2) is very useful. First of all, it defines the relationship between the inhibition function and go signal reaction time; the greater the go signal reaction time, the later the inhibition function. Second, it provides an easy way to estimate stop signal reaction time. The mean of the inhibition function is subtracted from the mean go signal reaction time (i.e., $\bar{T}_s = \bar{T}_g - \bar{T}_i$).

Logan and Cowan (1984) showed that the variance of the inhibition function was equal to the sum of the variances of the stop and go processes. Explicitly,

$$\sigma_i^2 = \sigma_g^2 + \sigma_s^2 \quad (3)$$

The relationship in Eq. (3) is very useful in interpreting inhibition functions. It allows us to separate the contribution of differences in go process variance from differences in stop process variance. Basically, the variance of the inhibition function determines how broad and flat the inhibition function appears. Under some conditions, the variance can be estimated quantitatively from the slope of the inhibition function (i.e., if the inhibition function is assumed to have a particular shape; see Logan & Cowan, 1984, pp. 313–314). In most conditions, the slope of the inhibition function can be interpreted qualitatively in terms of the variance of the go process and the variance of the stop process. Subjects with more variable go processes will produce flatter inhibition functions. Subjects with more variable stop processes will also produce flatter inhibition functions. It is possible to correct for differences in go process variance by transforming the stop signal delay appropriately. Residual differences in the inhibition function after these differences

²Logan and Cowan (1984) distinguished between a controlled phase and a ballistic phase in responses to the go signal. Responses could be inhibited if the stop process finished during the controlled phase; responses could not be inhibited if the stop process finished during the ballistic phase. In their analyses, the mean of the go processes was expressed in terms of both components (i.e., $\bar{T}_g = \bar{T}_c + \bar{T}_b$, where $\bar{T}_c$ and $\bar{T}_b$ are the means of the controlled and ballistic phases, respectively). However, the mean of the inhibition function was expressed only in terms of the mean of the controlled process (i.e., $\bar{T}_i = \bar{T}_c - \bar{T}_s$) and the variance of the inhibition function was expressed only in terms of the variance of the controlled process (i.e., $\sigma_i^2 = \sigma_c^2 + \sigma_s^2$). I do not distinguish between controlled and ballistic components of the go process in this chapter in order to keep the exposition clear and because the evidence suggests that ballistic processes may be negligible, brief enough to be safely ignored (De Jong et al., 1990; Osman et al., 1986, 1990; see Section IV).

have been removed, may be interpreted as differences in the variability of the stopping process. Schachar and Logan (1990a) used this technique to conclude that hyperactive children had more variable stopping processes than normal and clinical control children. The transformations will be discussed more extensively in Section IX.

B. Stop Signal Reaction Time

Probability and latency are the two most important measures of responses in experimental psychology. The inhibition function describes the probability that the response to the stop signal will win the race. But the internal response to the stop signal cannot be observed directly, so its latency must be estimated somehow. This is an important step in any stop signal study. Estimates of stop signal reaction time are the basis for most of the conclusions about the ability to inhibit thought and action. They are essential in comparing inhibitory ability between subjects and tasks, and in detecting deficient inhibition.

Three methods are available to estimate stop signal reaction time. The easiest to understand and the most commonly used is the one that assumes stop signal reaction time is constant, first proposed by Logan (1981). It is illustrated in Figure 1. Stop signal reaction time is the difference between the point at which the stop signal was presented and the point at which the stopping process finished. We know when the stop signal was presented from the experimental protocol. We have to estimate the point at which the stopping process finished from the observed distribution of go signal reaction times and the observed probability of responding given a stop signal. We do so by integrating the go signal reaction time distribution until the area under the integral equals the probability of responding. In Figure 1, this amounts to moving a vertical line across the distribution until the area to the left of the line equals the probability of responding, and then reading the value of the time axis at that point as an estimate of the finishing time of the stopping process. Then stop signal delay is subtracted out, and the remainder is the estimate of stop signal reaction time. Formally, stop signal reaction time is estimated by $r(t_d) - t_d$, where $r(t_d)$ is defined in terms of Eq. (1):

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

How do we calculate stop-signal reaction time in practice? Reaction times from go signal responses on which no stop signal occurred are collapsed into a single distribution. The reaction times are rank ordered, and the n th reaction time is selected, where n is obtained by multiplying the number of reaction times in the distribution (m) by the probability of responding at a given delay [$P_r(t_d)$; i.e., $n = m \cdot P_r(t_d)$]. The n th reaction

time estimates the time at which the stopping process finished, relative to the onset of the go signal. To estimate stop signal reaction time (the time at which the stopping process finished relative to the stop signal), stop signal delay (the interval between the onset of the go signal and the onset of the stop signal) must be subtracted from this value. This process is repeated for each stop signal delay for each subject. The results are then averaged over subjects within and sometimes across stop signal delays.

Stop signal reaction times estimated by this method are close to 200 msec in a variety of experimental tasks, conditions, and strategies in subjects in their late teens and early twenties (for a review, see Logan & Cowan, 1984). Stop signal reaction time is longer in younger children (Schachar & Logan, 1990a) and older adults (Kramer et al., 1992), following other life-span trends in reaction time. Hyperactive children show slower stop signal reaction times than normal and clinical control children (Schachar & Logan, 1990a; Schachar et al., 1993), and their stop signal reaction times are sped up by methylphenidate, a stimulant medication that relieves the behavioral symptoms of hyperactivity (Tannock et al., 1989; Tannock, Schachar, & Logan, 1993). Perhaps stop signal reaction time can reveal inhibitory deficiencies related to other forms of psychopathology and other clinical conditions.

Researchers have typically ignored factors affecting stop signal reaction time, being more concerned with factors affecting go signal reaction time (but see Logan et al., 1986; Riegler, 1986). This is unfortunate because the nature of processes are revealed by the factors that affect their latency, and we need to learn more about the nature of the stopping process.

Stop signal reaction times estimated by the method described in Eq. (4) tend to decrease somewhat as the stop signal delay increases (Logan & Burkell, 1986; Logan & Cowan, 1984; Logan et al., 1984). This reduction could reflect a refractory effect of processing the go signal, as if the stop process and the go process compete for common resources. De Jong et al. (1990) showed with Monte Carlo simulations that negative dependencies between stop and go tasks produce estimates of stop signal reaction time that decrease with delay. This is bad news for the race model; the reduction in stop signal reaction time with delay suggests that the stop task and go task are not independent, as the race model requires them to be.

Logan and Burkell (1986) tried to save the independence assumption and account for the decrease in terms of variability in stop signal reaction time. They argued that at short stop signal delays, nearly all stop signal reaction times would be fast enough to win the race, so the mean of the winners of the race would approximate the mean of the stop signal reaction time distribution. At longer delays, only the faster stop

signal reaction times would be fast enough to win the race. Responses to the go signal win the race on a significant proportion of the trials, and they screen out the slower stop signal reaction times. Thus, the mean of the stop signal reaction times that won the race would be faster than the mean of the entire distribution because the go signal responses cut off the upper tail [this logic is developed formally in Section VII(C)]. Logan and Burkell (1986) had some success fitting means of distributions with truncated upper tails to the data. Thus, decreases in stop signal reaction time with delay need not be interpreted as evidence against the independence assumption. De Jong et al. (1990) found that high variance in (simulated) stop signal reaction times produced estimated stop signal reaction times that decreased with delay, even when stop and go processes were independent, corroborating this conclusion.

The second method for estimating stop signal reaction time assumes that it is a random variable, not a constant. The method derives from Logan and Cowan's (1984) treatment of the inhibition function as distribution. It exploits the relationship between the mean of the inhibition function and the means of the stop and go processes expressed in Eq. (2): Mean stop signal reaction time equals the difference between the mean go signal reaction time and the mean of the inhibition function. Explicitly,

$$\bar{T}_s = \bar{T}_g - \bar{T}_i \quad (5)$$

The estimate in Eq. (5) is easy to calculate. All that is required is the mean of the go signal reaction times, which is easy to compute, and the mean of the inhibition function. The mean of the inhibition function can be computed in several ways. If there is a complete inhibition function, ranging from 0 to 1, then the mean is simply $\sum p_i x_i$, where p_i is the probability of responding at the i th stop signal delay minus the probability of responding at the $i-1$ th stop signal delay [i.e., $p_i = p(\text{respond})_i - p(\text{respond})_{i-1}$] and x_i is the i th stop signal delay. The mean can be calculated in a similar manner from a truncated inhibition function (i.e., one in which the smallest probability of responding given a signal is greater than 0, and the largest probability of responding given a signal is less than 1).³

It is also possible to use the median to estimate the mean of the inhibition function. This can be justified formally if the investigator is willing to assume that the inhibition function is a symmetrical distribution

³Specifically, the computation of p_i has to be rescaled to reflect the truncated range. Each value of p_i has to be divided by the difference between the maximum and the minimum probabilities of responding, that is, the probabilities of responding at the longest and shortest stop signal delays, respectively. The corrected p_i should equal $p_i / [p(\text{respond})_{\max} - p(\text{respond})_{\min}]$. In practice, it is easier to do the correction once, after summing the $p_i x_i$ values (i.e., $(\sum p_i x_i) / [p(\text{respond})_{\max} - p(\text{respond})_{\min}]$).

(such as the normal, the t , and the rectangular) so that the mean and median are guaranteed to be identical (see Logan & Cowan, 1984, pp. 313–314). Investigators not willing to make that assumption can hope that the mean and the median are not too different so that the median is a reasonable estimate of the mean. The ease of computing the median of the inhibition function may encourage tough-minded investigators to make one assumption or the other: The median is simply the point at which the probability of responding and the probability of inhibiting are .5. In principle, one stop signal delay could suffice if it were guaranteed to produce inhibition 50% of the time. More realistically, two delays are sufficient, one with the probability of responding below .5 and one with the probability above it. The median can be computed by interpolation. Logan and Cowan (1984) compared estimates based on the median to estimates based on the mean and found reasonable agreement in the three subjects they studied. Interpolation based on the two delays that surrounded the median produced estimates very close to those computed from means based on the entire inhibition function. They did not explore the possibility of extrapolating when the maximum probability of responding is less than .5 or when the minimum probability of responding is greater than .5, but estimation by extrapolation is possible in principle.

Regardless of how the mean of the inhibition function is calculated, using it to estimate stop signal reaction time is easy, and this is an important advantage. The method does not require implausible assumptions, like a constant stop signal reaction time, and this is an advantage as well. The main disadvantage is that the method provides one measure of the mean stop signal reaction time for the whole inhibition function. It does not allow separate estimates at each stop signal delay and it does not allow estimates of parameters other than the mean.

The third method of estimating stop signal reaction time was proposed by Colonius (1990; see also De Jong et al., 1990, Appendix). It involves using the observed distribution of go signal reaction times from trials on which no stop signal was presented, $f_g(t)$, and the distribution of go signal reaction times from stop signal trials on which the subjects failed to inhibit their responses (i.e., *signal-respond* reaction times), $f_r(t)$. The cumulative distribution of stop signal reaction times, $F_s(t-t_d)$, is defined as follows:

$$F_s(t-t_d) = 1 - [P_r(t_d) \cdot f_r(t)/f_g(t)] \quad (6)$$

where $P_r(t_d)$ is the probability of responding given a stop signal at delay t_d .

This method has the advantage of estimating the entire distribution of stop signal reaction times at each stop signal delay. Parameters like means and standard deviations can be calculated directly from the esti-

mated distribution. They do not have to be estimated or approximated in other ways. The disadvantage of the method is that the quality of the estimates is very sensitive to the quality of the data. The major limitation comes from signal-respond reaction times. For reasons, which follow, stop signals typically do not occur very frequently (often on 25% or less of the trials) and signal-respond trials occur even less frequently. It is difficult to gather enough observations to get stable signal-respond distributions, particularly at the tails, and the stability of the signal-respond distributions limits the stability of the estimated distributions of stop signal reaction times.

C. Signal-Respond Reaction Times

Responses that escape inhibition are observable and their latency (*signal-respond* reaction time) can be measured directly. An adequate model of the stop signal paradigm should be able to predict signal-respond reaction times and account for factors that affect them, and the race model does just that. This prediction is important because it bears on the validity of the assumptions underlying the race model. De Jong et al. (1990) and Jennings et al. (1992) used the prediction to test the assumption that stop and go processes are independent of each other (see De Jong et al., 1990; Jennings et al., 1992; Logan & Cowan, 1984).

The way that the race model accounts for signal-respond reaction times is illustrated in Figure 1. The responses that escape inhibition are those to the left of the line representing the finishing time of the stop process. Signal-respond reaction time should correspond to the mean of the reaction times in the area to the left of the line. Formally,

$$\bar{T}_r = (1/P_r(t_d)) \int_{T_s+t_d}^{\infty} t \cdot f_g(t) dt \quad (7)$$

where $\bar{T}_r$ is the mean signal-respond reaction time and the other variables are as defined in Eq. (1).

Three predictions follow from this analysis. First, signal-respond reaction time should be faster than go signal reaction time. Signal-respond reaction times come from the same distribution as go-signal reaction times, but the stopping process cuts off the upper tail. The mean of the left part of the distribution has to be smaller than the mean of the whole distribution. This prediction has been confirmed every time it is tested (De Jong et al., 1990; Jennings et al., 1992; Lappin & Eriksen, 1966; Logan, 1981; Logan & Cowan, 1984; Osman et al., 1986, 1990).

Second, signal-respond reaction times should increase as stop signal delay increases. This follows because slower go responses win the race when the stop signal is delayed. It can be seen in Figure 1: The mean of the left part of the top panel is smaller than the mean of the left part in

the second panel. This prediction has also been confirmed several times (De Jong et al., 1990; Jennings et al., 1992; Logan, 1981; Logan & Cowan, 1984; Osman et al., 1986). Sometimes it is violated with very short delays, when signal-respond reaction times are based on very few observations (Logan & Cowan, 1984), but it is generally true when there are sufficient observations.

Third, observed signal-respond reaction times should *equal* those predicted from the race model, using Eq. (2). This is a *point* prediction rather than a *difference* prediction or a *trend* prediction, so it is stronger than the first two. Nevertheless, it has been confirmed several times (De Jong et al., 1990; Jennings et al., 1992; Logan & Cowan, 1984). Again, it is confirmed more readily in the central parts of the inhibition function where signal-respond responses are more numerous.

How is the prediction tested in practice? Reaction times from go signal responses on which no stop signal occurred are collapsed into a single distribution. The reaction times are rank ordered, and the mean of the fastest n is computed, where n is the number of reaction times in the distribution (m) multiplied by the probability of responding at a given delay [$P_r(t_d)$; i.e., $n = m \cdot P_r(t_d)$]. The mean of the fastest n is the predicted value, to be compared with observed signal-respond reaction times. (This procedure is essentially the same as the procedure for calculating stop signal reaction time. The main difference is that here we take the mean of the fastest n reaction times, whereas in estimating stop signal reaction time, we took the n th reaction time and interpreted it as the finishing time for the stopping process.)

In a way, it is impressive that this prediction succeeds. It depends on the assumption that stop signal reaction time is constant, and that is unlikely to be true. That assumption also predicts an odd shape for the distribution of signal-respond reaction times; the fast tail should be typical of all reaction time distributions, but the slow tail should be truncated abruptly at the point at which the stopping process finishes. Observed signal-respond reaction time distributions do not look like this (see Lappin & Eriksen, 1966; Osman et al., 1986, 1990). The point prediction of mean signal-respond reaction time excludes the upper tail of the observed signal-respond distribution, so it is not influenced by outliers. This may be bad in principle, but in practice, it may actually be helpful. The prediction may be improved by excluding spurious outliers.

De Jong et al. (1990) and Jennings et al. (1992) took the prediction of signal-respond reaction times a step further. They tried to predict the latencies of partial responses on stop signal trials. De Jong et al. had subjects respond by squeezing a zero-displacement dynamometer, and defined partial responses as squeezes that did not reach the criterial amount of force (25% of maximum force). Jennings et al. defined partial

responses as ones that occurred when there was electromyographic activity but no completion of the button press response. The analysis involved distinguishing between three classes of responses rather than two: complete inhibition responses, partial inhibition responses, and complete go responses. The probability of responding given a stop signal corresponds to the probability of making a complete go response [i.e., $P_r(t_d)$]. The probability of inhibition [i.e., $1 - P_r(t_d)$] corresponds to the sum of the probability of complete inhibition and the probability of partial inhibition. This is illustrated in Figure 2, which depicts the race model with two finishing times for the stopping process. The first (leftmost) distinguishes between responses that are executed and responses that are inhibited partially or completely. The second (rightmost) distinguishes between partially inhibited responses and completely inhibited responses. The idea is that partially inhibited responses were faster than completely inhibited responses, almost fast enough to be executed, but not completely. The area between those lines corresponds to the probability of partial inhibition, and the mean reaction time in that area corresponds to signal-respond reaction times for the partial responses. It can be calculated using a variation of Eq. (7). De Jong et al. (1990) and Jennings et al. (1992) made such calculations and found that the predicted signal partial respond reaction times agreed well with the observed ones.

Osman et al. (1986) developed the race model a little differently and generalized these predictions to distributions of signal-respond reaction times, not just the means. According to their analysis, distributions of signal-respond reaction times and distributions of go task reaction times from trials on which no stop signal was presented (*no-signal* reaction




FIGURE 2

The race model extended to include partial responses. From R. De Jong, M. G. H. Coles, G. D. Logan, and G. Gratton (1990). In search of the point of no return: The control of response processes. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 164-182. Copyright 1990 by the American Psychological Association. Reprinted by permission.

times) should start from roughly the same minimum value, but signal-respond distributions should reach their peaks before no-signal distributions do. Moreover, signal-respond reaction times should reach their peaks sooner the earlier the stop signal was presented. In practice, this means that the cumulative distributions of reaction times should begin from about the same point and then "fan out," with the signal-respond reaction times from the earliest stop signal delay rising first and most sharply, and no-signal reaction times rising last and most gradually. They tested and confirmed this prediction in three experiments (see also Osman et al., 1990).

IX. SETTING STOP SIGNAL DELAY

Inhibition functions depend on go signal reaction time and stop signal reaction time, as well as on stop signal delay, so choosing and setting stop signal delays is a difficult matter. At some delays, subjects will inhibit all the time and at other delays they will respond all the time. The problem is to find the intermediate delays at which the probability of responding is between 0 and 1. This part of the function is the most meaningful theoretically. All of the analytic methods reviewed previously focus on it. Several methods are available, ranging from choosing delays arbitrarily to tracking several parameters of subjects', performances and setting delays contingent on the values of those parameters.

The first issues to be decided are how many delays to use and how often to present the stop signal. The data provide a clear answer: At least three or four delays should be used. In theory, at least two are required to estimate the mean and variance of the inhibition function. In practice, subjects adapt to the range of delays they experience. If there is only one delay, they will prolong their go signal reaction times in order to maximize the probability of inhibiting (Lappin & Eriksen, 1966; Logan, 1981). Prolongation can be minimized by presenting several delays, some early enough that subjects will be able to inhibit most of the time, and some late enough that subjects will usually respond when they occur (Logan, 1981). Early and late delays break the contingency between variation in go signal reaction time and the probability of inhibiting. Subjects can still increase the probability of inhibiting by prolonging reaction time, but the effects of prolongation will be small compared to the effects of early versus late stop signal delays, so subjects may not bother with prolongation. Tracking procedures, like those of Osman et al. (1986, 1990), also break the contingency between variation in go signal reaction time and the probability of inhibition. They guarantee that stop signals will be presented at different delays whose values are unpredictable to the subject.

How often should the stop signal be presented? Again, the data are clear: The greater the probability that a stop signal will occur, the greater the probability of inhibition and the slower the reaction times to the go signal. Logan (1981) found these effects presenting stop signals on 10% versus 20% of the trials. Logan and Burkell (1986) found similar (and larger) effects presenting stop signals on 20%, 50%, and 80% of the trials. Quantitative analyses of inhibition functions, described previously, showed that the increase in probability of inhibition can be accounted for completely by the increase in go signal reaction time. It appears as if subjects place greater importance on the stopping task the more often the stop signal occurs, prolonging their go signal reaction times in order to increase the probability of inhibition. This effect can be minimized by instructing subjects appropriately and presenting stop signals relatively infrequently. There is a trade-off, because the more often stop signals are presented, the more observations there are to analyze (e.g., signal-respond reaction times), but the more often stop signals are presented, the more likely the data are to be influenced by unwanted strategies. In practice, presenting stop signals on 25% of the trials is a reasonable compromise. The stop signal occurs often enough to get reasonable data, but rarely enough that subjects do not expect it and change their go signal reaction times accordingly.⁴

How should the delays be chosen? Ideally, the delays should capture the part of the inhibition function that is most informative—the part where the probability of responding changes smoothly from 0 to 1. There are several options.

A. Arbitrary Delays

Delays may be chosen arbitrarily and fixed for all conditions and all subjects. This may seem to rely a lot on luck, but some rules of thumb,

⁴Instructions are very important in the stop signal paradigm. The main difficulty is in getting subjects to give appropriate priority to the two tasks. I usually instruct subjects about the go task first, without describing the stopping task. I tell them what they are expected to do and to do it as quickly as possible without making errors. Then I introduce the stopping task, telling them about the stop signals and what to do when a stop signal occurs. I tell them that stop signal delay is varied by the experimenter, and that some stop signals will occur so early that they will always be able to stop and some will occur so late that they will rarely be able to stop. I tell them to try to stop if they can, but not to worry about failing to stop because of the way the delays are set. I tell them to perform the go task as quickly and accurately as possible. I tell them not to let the stopping task interfere with their performance on the go task. I tell them specifically not to delay their response to the go task in order to improve their chances of stopping. I have not investigated the consequences of varying these instructions, so I do not know which parts are essential. The whole set of instructions works well, however, producing reasonable inhibition functions and reasonable reaction times. Other investigators have used explicit payoff schemes to the same end (De Jong et al., 1990; Jennings et al., 1992).

which are formalized in the other methods, can guide the selection and produce a reasonable set of delays. The shortest delay should be close to zero, the longest should not be much longer than the mean go signal reaction time, and the rest should be evenly spaced in between. The top row of Figure 3 contains inhibition functions from the simple- and choice-reaction go tasks in Logan et al. (1984) that were generated choosing arbitrary delays using these rules of thumb. The left panel in the top row plots the functions according to stop signal delay; the middle and right plot them according to transformations of delay. The arbitrary delays captured most of the inhibition function between 0 and 1.

The alternative to choosing delays arbitrarily is to choose delays contingent on some aspects of subjects' performance, such as the mean go signal reaction time, or the success or failure of the last attempt at inhibition. Delays can be set dynamically, calculating performance measures and updating delays on every trial, or they can be set in advance, based on performance measures gathered in control blocks or a control session. There are several options for setting delays contingent on performance.

B. Mean Go Signal Reaction Time Minus Delay

Delays may be set relative to mean go signal reaction time. The idea here is that stop signal delay can compensate for differences in mean reaction time, which is illustrated in the top and bottom panels of Figure 1, as discussed previously. If the mean go signal reaction time for the first of two conditions, tasks, subject groups, or whatever, is x msec faster than the second, then stop signal delays for the two that differ by x msec (with the first shorter than the second) should produce equivalent probabilities of responding. This prediction has been confirmed many times by showing that disparate inhibition functions can be brought into alignment by replotting them relative to mean go signal reaction time (Logan, 1981, 1982, 1983; Logan et al., 1984; Osman et al., 1990).

The procedure by which delays are set relative to mean go signal reaction time is as follows: First some measure of mean go signal reaction time is calculated. The mean may be estimated from similar experiments with similar subjects and stimuli, or it may be calculated subject by subject. The intervals between the delays are chosen arbitrarily. As a rule of thumb, the shortest stop signal delay should be close to zero, the largest should be close to the mean, and the others should be spaced evenly in between. Schachar and Logan (1990a) set stop signal delays equal to the mean go signal reaction time minus 0, 100, 200, 300, 400, and 500 msec. The second row of Figure 3 contains inhibition functions




FIGURE 3

The effects of different methods of selecting stop signal delays on the inhibition functions on data from simple and choice reaction time tasks. From G. D. Logan, W. B. Cowan, and K. A. Davis (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, 275-291. Copyright 1984 by the American Psychological Association. Adapted by permission.

from the simple- and choice-reaction go tasks in Logan et al. (1984) that were generated by selecting delays relative to mean go signal reaction time. Mean go signal reaction time was calculated after each block, and the delays for the next block were determined by the means from the immediately previous block.⁵ The data gathered by this procedure are plotted in the middle panel of the second row of Figure 3; the left and right panels represent transformations of mean minus delay. These functions capture much of the range from 0 to 1 for both tasks.⁶

C. Relative Finishing Time as a Z Score

Inhibition functions generated by setting delay relative to mean reaction time will be aligned only if the go-signal distributions differ only in mean reaction time. If they differ substantially in variance, then functions set relative to the mean will not produce equivalent probabilities of responding. This can be seen in Figure 4, where the two panels depict two go signal reaction time distributions that differ substantially in variance. Two stop signals are presented at equivalent delays in the two conditions. Their finishing times are depicted as vertical lines intersecting the go signal reaction time distributions. The figure shows that the two stop signal delays cut off different proportions of the distributions. They cut off the upper and lower tails off the low-variability distribution, whereas they cut off the upper and lower thirds of the high-variability distribution. The inhibition function for the condition with the greater variability would be flatter than the function for the condition with the lesser variability.

One solution to this problem, proposed by Logan et al. (1984) is to plot inhibition functions in terms of a Z score that represents the relative finishing time of the stop and go processes in standard deviation units, using the standard deviation of the go signal reaction times to define the units: specifically, Z relative finishing time or

⁵Delays for the first block were set arbitrarily. Schachar and Logan (1990a) used a similar method, calculating delays relative to the mean reaction time from the previous block, but the first block did not require subjects to inhibit responses when they heard the stop signal. This allowed the experimenter to introduce the task gradually and emphasize the importance of responding quickly to the go task.

⁶There is a potential for confusion in discussing this transformation and relating it to delays defined in terms of stimulus onsets, which Schachar and Logan (1990a; mostly Logan) fell prey to in describing the procedure for setting delays (p. 712). The transformation, called *mean minus delay*, reverses the ordering of the delays. The shortest stop signal delay is the largest mean minus delay, and vice versa. It is important to keep this clear in discussing how the probability of responding (or inhibiting) changes with stop signal delay versus mean minus delay. The probability of responding increases as stop signal delay increases and decreases as mean minus delay increases.




FIGURE 4

The effects of variability of go-process reaction times on inhibition functions, shown by comparing the effects of two stop signals presented at different delays on distributions with more (top panel) and less (bottom panel) variability.

$$\text{ZRFT} = (\bar{T}_g - t_d - \bar{T}_s) / S_g \quad (8)$$

where $\bar{T}_g$, t_d , and $\bar{T}_s$ are mean go-signal reaction time, stop signal delay, and mean stop signal reaction time, respectively, as defined above, and S_g is an estimate of the standard deviation of the go-signal reaction times.

The motivation for including stop signal reaction time, $\bar{T}_s$, in Eq. (8) can be seen in Figure 1. The idea is to express the relative finishing times of the stop and go processes. The mean go signal reaction time provides some estimate of the finishing time of the go process. The finishing time of the stop process, represented by the vertical line across the go signal reaction time distribution in Figure 1, is the sum of two components: stop signal delay, t_d , and stop signal reaction time, $\bar{T}_s$. The relative finishing time is the difference between these estimates [i.e., $\bar{T}_g - (t_d + \bar{T}_s)$].

In practice, ZRFT is calculated by estimating the mean and standard deviation of go signal reaction time, estimating stop signal reaction time, and choosing arbitrary delays to produce evenly spaced Z scores.

It can be calculated post hoc to bring different inhibition functions into alignment (see Logan & Cowan, 1984; Logan et al., 1984; Schachar & Logan, 1990a), or it can be calculated dynamically, estimating each of the parameters each block. The third row of Figure 3 contains inhibition functions from the simple- and choice-reaction go tasks in Logan et al. (1984) that were generated by selecting delays in terms of ZRFT, calculating the parameters of each block, and setting the ZRFTs for the next block. Stop signal reaction time was calculated for each delay, and the average over delays was used in Eq. (8). The data collected by this method are plotted in the right panel in the bottom row of Figure 3. The other two panels contain transformations of ZRFT. These functions capture much of the range from 0 to 1 for both tasks.⁷

D. Relative to Go Signal Reaction Time Distribution

The finishing time of the stop process relative to the go process may be represented as a point on the distribution of go signal reaction times. In Figure 1, it is drawn as a vertical line across the distribution. The ZRFT method tries to estimate that point as a Z score, in order to capture differences in variability as well as differences in means. However, inhibition functions will be aligned when plotted against ZRFT only if the distribution of go signal reaction times is normal or approximately normal. If it is skewed, as reaction time distributions typically are, then ZRFT will not produce perfect alignment.

One solution to this problem, attempted by Kramer et al. (1992), is to set stop signal delays relative to the go signal reaction time distribution directly. They targeted specific percentiles of the go signal reaction time distribution and set delays so that the stopping process should finish at those percentiles. Subjects were tested in four sessions, one baseline session and three sessions with the stopping task. In the baseline session, subjects performed the go task as a single task and the 20th, 40th, 60th, and 80th percentiles of each subject's reaction time distribution

⁷There is potential for confusion with ZRFT as well, that Schachar and Logan (1990a), this time entirely Logan, fell prey to. ZRFT reverses the ordering of delays, relative to stop signal delay. ZRFT decreases as stop signal delay increases. When calculating the ZRFT values for Figure 2 (the clinical study) in Schachar and Logan (1990a), I added stop signal delay to mean minus delay instead of subtracting it. This resulted in an artifactual horizontal separation of the inhibition functions for the different groups. When ZRFT is calculated correctly and the data are replotted, the different inhibition functions are aligned horizontally. However, even with this alignment, the inhibition function was flatter for hyperactive children than for normal controls, and that is the important point to be taken from the analysis. In fact, the difference in slope was the main basis for the conclusion that hyperactive children were deficient inhibitors. This conclusion was correct even if the figure was wrong.

were calculated. Subjects also performed a simple reaction time task, responding to the stop signal tones, which were presented on the same schedule as they would be in the stopping task itself. These simple reaction times to the onsets of the tones were used as estimates of stop signal reaction times. They were calculated for each subject and subtracted from the go signal reaction times at the target percentiles to produce stop signal delays. That is,

$$t_d = x - \bar{T}_s,$$

where x is the target percentile on the cumulative distribution of go signal reaction times, $F_g(x)$. These estimates were revised each session, using the data from the previous session to set the new delays.

This technique produced nice inhibition functions. The probability of responding was close to the value of the target percentile, except for the shorter delays where subjects responded more often than predicted. The main problem with this technique is that it does not guarantee that stop delays will be equally spaced on the time axis, as the other methods do. The skew typical of reaction time distributions means that equal intervals on the (cumulative) probability axis will generally not produce equal intervals on the time axis. Equal time intervals are important primarily for some statistical analyses (such as analysis of variance); they are not important in theory.

E. Tracking Stopping Performance

All of the techniques reviewed so far track parameters of the go task. Osman et al. (1986) developed a method for setting stop signal delays that tracks subjects' stopping performance. They used three different delays—early, middle, and late—and had a different rule for changing each delay. The middle delay was increased by 50 msec every time the subject inhibited, and decreased by 50 msec every time the subject responded. On average, subjects should respond half of the time and inhibit half of the time to delays set by this rule. The early delay was decreased by 50 msec every time the subject responded, and increased by 50 msec every time the subject inhibited twice consecutively. On average, subjects should respond 29% of the time and inhibit 71% of the time with delays set by this rule. The late delay was increased by 50 msec every time the subject inhibited, and decreased by 50 msec every time the subject responded twice consecutively. On average, subjects should respond 71% of the time and inhibit 29% of the time to delays set by this rule. Osman et al. (1986) found observed probabilities of responding that were very close to the predicted values.

The observed probabilities of responding at early, middle, and late delays do not reflect the inhibition function, however. The tracking algorithm guaranteed that each subject experienced several different values of t_d for early delay, some of the same values and some different ones for the middle delay, and another overlapping range for the late delay. In order to calculate inhibition functions, Osman et al. (1986) collapsed data across early, middle, and late delays and plotted the probability of responding at each value of t_d they used in the experiment (the tracking algorithm varied t_d in 50 msec increments). This procedure produced inhibition functions that covered the range from 0 to 1, but the number of observations was not distributed evenly across values of t_d , so many of the inhibition functions were nonmonotonic. Osman et al. (1986) smoothed the inhibition function for each subject by finding the monotonic function with the best least squares fit to the observed inhibition function. The mean of the smoothed function was calculated for further analysis. The distinction between early, middle, and late delays plays no part of this analysis, so Osman et al. (1990) dropped the early and late delays in subsequent experiments and calculated inhibition functions in the same fashion. They found that the tracking algorithm for the middle delay produced enough different values of t_d to produce inhibition functions that covered the range from 0 to 1.

The main problem with this method is the difficulty of computing inhibition functions. The fact that early, middle, and late delays do not correspond to distinct values of t_d may raise problems for data analysis (for reasons described in G. Computing Inhibition Functions).

De Jong et al. (1990) used a variation of Osman's technique. Osman et al. (1990) changed delays on every stop signal trial, but that was impractical for a psychophysiological study in which event-related potentials had to be averaged across several trials to be recovered from the data. De Jong et al. needed fixed delays. They compromised by running a practice session using the Osman et al. tracking technique, calculated the average, short, middle, and long delays, and fixed those values for the test session during which psychophysiological measures were recorded. In the data from their test session, early, middle, and late delays each corresponded to single values of t_d , so they could be interpreted directly as inhibition functions, without the subsequent curve-fitting process of Osman et al. (1986, 1990). De Jong et al. got nice looking inhibition functions from this procedure.

F. Other Contingencies

Jennings et al. (1992) produced inhibition functions of a different type by holding stop signal delay constant and varying reaction time by

imposing various deadlines on go task performance. They used two stop signal delays (50 and 150 msec) and four deadlines (225, 300, 375, and 400 msec). Their analyses suggested violations of the race model with the 50-msec delay, so they focused on the 150-msec delay and compared inhibition performance across deadlines. They tested the independence assumption of the race model, for example, by comparing estimates of stop signal reaction time in the different deadlines and finding that it was nearly constant.

It is possible to present stop signals contingent on events external to the stop signal paradigm, events that are not directly involved in processing the stop signal and the go signal but are interesting nevertheless. The experiment of Jennings et al. (1992) is interesting again in this regard. Jennings et al. (1992) presented go and stop stimuli presented contingent on the phase of the subject's heartbeat, and found that responses inhibited during a sensitive period between heartbeats delayed the next heartbeat. Jennings et al. (1992) interpreted this result as suggesting that stop signal inhibition may involve the same midbrain structures as the processes that control heart rate (see also De Jong et al., 1990).

G. Computing Inhibition Functions

In theory, the different methods for choosing stop signal delays are equivalent because any one measure can be transformed into any other. In practice, however, the different methods exert different constraints on the computation and analysis of inhibition functions. Points that are equivalent under one transformation will generally not be equivalent under other transformations. If the data are grouped tightly according to values of one transformation, they may be spread across the function in some other transformation. For example, an experiment that tests all subjects at the same values of ZRFT will test different subjects at different values of stop signal delay. The actual stop signal delays will vary with the mean and standard deviation of each subject's go task reaction times and with each subject's stop signal reaction time. The tracking technique of Osman et al. (1986) is extreme in this regard. It groups observations very tightly into early, middle, and late delays, but the observations in each category are made at several different values of stop signal delay (t_d). Different subjects are unlikely to cover the same range of stop signal delays and are unlikely to have the same number of observations in the delays they happen to share.

The importance of grouping or spreading observations along the various transformations of delays depends on the data analysis techniques one wishes to employ. Standard techniques like analysis of variance

require tight grouping—several levels of the delay factor with an equal number of observations at each level. They can be applied only to the transformation by which delays are selected (i.e., the transformation in which grouping is tight). Other transformations must be analyzed by other methods.

X. DETECTING DEFICIENT INHIBITION

The main use of the stop signal paradigm is to compare the ability to inhibit in different tasks, conditions, strategies, and subject populations. A major goal in these comparisons is to detect and diagnose deficient inhibition. What can go wrong with the stopping mechanism? It can become slower and more variable, and it can fail to be triggered. Each of these effects would diminish the ability to inhibit. Slower and more variable stop signal reaction times would be more likely to lose the race with the go process, decreasing the probability of inhibition. And if the stopping process failed to be triggered altogether—if it acted as if no stop signal were presented—the probability of inhibition would decrease even further.

So far, the analysis sounds straightforward: look for decreases in the probability of inhibition. The problem with that simple analysis is that there are ways to decrease the probability of inhibition without impairing the stopping process. A faster go process, for example, would be harder to stop than a slower one simply because of its speed. So a decrease in the probability of inhibition does not necessarily imply deficient inhibition. A more complex analysis is required, one that uses the race model. The idea is to try to account for differences in the probability of inhibition—differences in inhibition functions—by testing a series of hypotheses. Differences in the inhibition functions due to differences in go processes are removed, and residual differences are interpreted in terms of the stop process. Some residual differences reflect differences in the speed of the stop process; others reflect differences in the variability of the stop process or differences in the probability of triggering the stop process.

The analysis consists of three phases. First, the inhibition functions are inspected to see if they are interpretable. In order to be interpretable, they must capture some of the middle part of the inhibition function, where the probability of responding is, say, between .1 and .9. Functions in which subjects nearly always inhibit or nearly always respond are not interpretable within the race model (although the race model can inform their interpretation; see the third phase). Inhibition functions that are flat (i.e., in which the probability of inhibition does not vary with stop signal delay) are also not interpretable within the race model.

For example, subjects may withhold a certain proportion of their responses whether or not a stop signal occurred. This would flatten their inhibition functions, adding the same number of inhibited responses to each part of the inhibition function. This strategy can be detected on go signal trials; the proportion of missed responses on go signal trials can be used to correct the inhibition function (see Tannock et al., 1989).

Flat inhibition functions are diagnostic of an inhibitory deficit. Most people can inhibit responses on request and produce reasonable inhibition functions. Someone who cannot is deficient in the ability to inhibit. Flat inhibition functions could also mean poor experimental control. The stop signals could have been presented too early or too late to capture the interesting part of the inhibition function. To tell the difference, it is helpful to have a control group with stop signal delays selected in the same manner as the group in question (i.e., so as to have the same opportunity to inhibit). If the control group's inhibition function is "reasonable" (i.e., not flat), the flat inhibition functions of the group in question reflect an inhibitory deficit. If the control group's inhibition function is also flat, the delays are inappropriate.

If the inhibition functions are reasonable, the second phase of the logic begins. The race model is fitted to the data. Inhibition functions are calculated and the transformations just described are applied in an attempt to bring the functions into alignment. If the functions can be brought into alignment, then the fit is successful. The race model applies to the data set. The different tasks, conditions, strategies, or subject populations have the same stopping process with different parameters. Differences in the parameters of the fitted model are then interpreted within the race model. If mean minus delay is sufficient to bring the functions into alignment, then the differences in the inhibition functions are due only to differences in mean go signal reaction time. The stopping process is essentially the same. It is executed with the same probability and latency in the different tasks, conditions, strategies, or subject populations. If mean minus delay brings the functions into alignment, there is no inhibitory deficit.

If mean minus delay does not bring the functions into alignment, I suggest trying mean minus delay minus stop signal reaction time (i.e., the numerator of Eq. (8): $\bar{T}_g - t_d - \bar{T}_s$) as the next transformation. If the functions can be brought into alignment with this transformation, then differences in the inhibition functions are due to differences in mean go signal reaction time and differences in mean stop signal reaction time. Differences in stop signal reaction time have to be interpreted. They could reflect a specific deficit in inhibition (as in hyperactive children) or a general slowing that affects go processes as well as stop processes (as in young children; Schachar & Logan, 1990a).

If mean minus delay minus stop signal reaction time does not bring the functions into alignment, ZRFT should be tried. If ZRFT brings the functions into alignment, then the tasks, conditions, strategies, or subject populations differ in the standard deviation of go signal reaction time. There are generally differences in mean go signal reaction time and stop signal reaction time as well when ZRFT brings the functions into alignment. They need to be interpreted if they appear. Note, however, that differences in the mean and standard deviation of go signal reaction time should be interpreted differently from differences in stop signal reaction time. They reflect differences in the go task. Differences in stop signal reaction time may indicate a deficient stopping process.

If the race model cannot be fitted to the data successfully—if the inhibition functions cannot be brought into alignment by some transformation—then the third phase of the logic begins. The failed alignment means that the race model does not apply to one of the tasks, conditions, strategies, or subject populations, and it is the experimenter's job to find out which one it is. In many cases, one of the tasks, conditions, strategies, or subject populations is considered "normal" or "standard" and the other is considered as a possible deviation. Children with no psychiatric problems are considered normal in studies of psychopathology. College-age adults provide standards against which others are compared in studies of life span development.

Often, failed alignment means that one inhibition function has a flatter slope than the other even after ZRFT correction. The difference in slope reflects a deficiency in the stop process. It can be interpreted in two ways. First, it could indicate greater variability in the stop process. The slope of the inhibition function depends on the variability of the stop process as well as variability in the go process, and the ZRFT transformation only corrects for differences in go process variability [see Eq. (3)]. Subjects with more variable stop processes would produce flatter inhibition functions even after ZRFT correction. Second, a difference in slope could mean that the inhibition process was executed or triggered less often. Failing to trigger the inhibition process on a constant proportion of the trials would produce a shallower inhibition function. The lowest probability of responding would not be 0.0, as with a normal inhibition function. Subjects would respond whenever the inhibition process failed to be triggered, so the lowest probability of responding would be the probability that the inhibition process failed to be triggered, P_f . The inhibition function itself would be compressed. It would range from P_f to 1.0, compressed by a factor of $1 - P_f$. Formally,

$$P_o(t_d) = (1 - P_f) \cdot P_t(t_d) + P_f \quad (9)$$

where $P_o(t_d)$ is the observed inhibition function and $P_t(t_d)$ is the "true"

inhibition function for those trials on which the stopping process is triggered by the stop signal.

A difference in slope after ZRFT correction does not discriminate between these alternatives. It may be possible to discriminate them by using stop signal delays short enough to estimate the lower asymptote of the inhibition function, which can be interpreted as a measure of P_f . Equation 9 can then be used to separate the effects of failing to trigger from the effects of increased variability; residual differences in slope not accounted for by Eq. 9 would reflect differences in variability. It may also be possible to use some other method to estimate the variability of the stop signal reaction times. The analyses of signal respond reaction times suggested by Colonius (1990) can be used to estimate the distribution of stop signal times, from which the variance can be calculated. Equation (3) can then be used to separate the effects of differences in variability from differences in triggering; slope differences that remained after subtracting out differences due to variability would reflect differences in triggering probability.

XI. STOPPING IN THE FUTURE

What lies ahead for the stop signal paradigm? I think the next important step will be to understand the nature of the stopping process. This can be done in several ways, but four deserve special mention. First, more research must be done to discover the relations between stop signal inhibition and other forms of inhibition. Relations between stop signal inhibition and behavioral measures of inhibition, such as negative priming and Stroop interference, must be worked out empirically and theoretically. Second, more research must be done on the brain systems underlying stop signal inhibition. The hypothesized distinction between central and peripheral, cortical and midbrain inhibition systems needs further investigation. Drug studies may reveal the psychopharmacology of the processes involved. Studies of patients with inhibitory problems would be interesting in this regard. Third, more research must be done on deficiencies in the stopping process. Encouraged by our work with hyperactive children, Schachar, Tannock, and I are looking for inhibitory deficits in other clinical groups. The stop signal paradigm might be used more generally to study inhibitory psychopathology, such as that evidenced by people with psychopathic personalities. It would be quite interesting to look at the stopping performance of patients with frontal lobe damage, who have trouble with inhibitory control. Fourth, more research must be done on life span developmental changes in stopping performance. It is important to look at children much younger than those Schachar and Logan (1990a)

examined to learn when the ability to inhibit is first acquired and to study its early development. The work of Kramer et al. (1992) suggests that investigations of stop signal inhibition in aging may be very fruitful.

ACKNOWLEDGMENTS

I wrote this chapter while I was a visiting professor at the University of Amsterdam. I am grateful to the Developmental Psychology group for their hospitality and stimulation. I was lucky enough to get reactions from many of the people whose research has shaped the current state of stopping research. I am grateful to Ritske De Jong, Darryl Humphrey, Art Kramer, Allen Osman, Russell Schachar, Rosemary Tannock, and Maurits van der Molen for valuable comments and suggestions. They may not endorse everything I say here, but their comments made it clearer, broader, and more representative. Tom Carr and Dale Dagenbach provided excellent editorial comments, for which I am grateful.

REFERENCES

- Berlyne, D. E. (1957). Uncertainty and conflict: A point of contact between information-theory and behavior-theory concepts. *Psychological Review*, 64, 329-339.
- Bjorklund, D. F., & Harnishfeger, K. K. (1990). The resources construct in cognitive development: Diverse sources of evidence and a theory of inefficient inhibition. *Developmental Review*, 10, 48-71.
- Brooks, V. D. (1986). *The neural basis of motor control*. New York: Oxford University Press.
- Bullock, D., & Grossberg, S. (1988). Neural dynamics of planned arm movements: Emerging invariants and speed-accuracy properties during trajectory formation. *Psychological Review*, 95, 49-90.
- Carr, T. H., & Dagenbach, D. (1990). Semantic priming and repetition priming from masked words: Evidence for a center-surround mechanism in perceptual recognition. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 16, 341-350.
- Colomius, H. (1990). A note on the stop-signal paradigm, or how to observe the unobservable. *Psychological Review*, 97, 309-312.
- Dagenbach, D., Carr, T. H., & Barnhardt, T. (1990). Inhibitory semantic priming of lexical decisions due to failure to retrieve weakly activated primes. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 16, 328-339.
- De Jong, R. (in press). Multiple bottlenecks in overlapping-task performance. *Journal of Experimental Psychology: Human Perception and Performance*.
- De Jong, R., Coles, M. G. H., Logan, G. D., & Gratton, G. (1990). Searching for the point of no return: The control of response processes in speeded choice reaction performance. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 164-182.
- Dell, G. S. (1986). A spreading-activation theory of sentence production. *Psychological Review*, 93, 283-321.
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception and Psychophysics*, 16, 143-149.
- Fletcher, B., & Rabbitt, P. M. A. (1978). The changing pattern of perceptual analytic strategies and response selection with practice in a two-choice reaction time task. *Quarterly Journal of Experimental Psychology*, 30, 417-427.
- Hasher, L. T., & Zacks, R. T. (1989). Working memory, comprehension, and aging: A review and a new view. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22). Orlando, FL: Academic Press.
- Henry, F. M., & Harrison, J. S. (1961). Refractoriness of a fast movement. *Perceptual and Motor Skills*, 13, 351-354.
- Jennings, J. R., van der Molen, M. W., Brock, K., & Somsen, R. J. M. (1992). On the synchrony of stopping motor responses and delaying heartbeats. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 422-436.
- Jennings, J. R., & Wood, C. C. (1977). Cardiac cycle time effects on performance, phasic cardiac responses and their intercorrelation in choice reaction time. *Psychophysiology*, 13, 297-307.
- Kornblum, S. (1973). Sequential effects in choice reaction time: A tutorial review. In S. Kornblum (Ed.), *Attention and performance IV* (pp. 259-288). New York: Academic Press.
- Kornblum, S., Hasbroucq, T., & Osman, A. (1990). Dimensional overlap: Cognitive basis for stimulus-response compatibility—A model and taxonomy. *Psychological Review*, 97, 253-270.
- Kramer, A. F., Humphrey, D. G., Larish, J., Logan, G. D., & Strayer, D. L. (1992, March). *Aging and inhibition*. Paper presented at the Conference on Cognition and Aging, Atlanta, GA.
- Ladefoged, P., Silverstein, R., & Papcun, G. (1973). Interruptibility 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.
- Lashley, K. S. (1951). The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior* (pp. 112-136). New York: Wiley.
- Levelt, W. J. M. (1983). Monitoring and self-repair in speech. *Cognition*, 14, 41-104.
- 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*. Hillsdale, NJ: Lawrence Erlbaum.
- Logan, G. D. (1982). On the ability to inhibit complex actions: 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: 1. Stop signal studies of decision and memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 9, 585-606.
- Logan, G. D. (1985). On the ability to inhibit simple thoughts and actions: 2. Stop-signal studies of repetition priming. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 11, 675-691.
- Logan, G. D. (1990). Repetition priming and automaticity: Common underlying mechanisms? *Cognitive Psychology*, 22, 1-35.
- Logan, G. D., & Burkell, J. (1986). Dependence and independence in responding to double stimulation: A comparison of stop, change, and dual-task paradigms. *Journal of Experimental Psychology: Human Perception and Performance*, 12, 549-563.
- Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, 91, 295-327.
- Logan, G. D., Cowan, W. B., & Davis, K. A. (1984). On the ability to inhibit responses in simple and choice reaction time tasks: A model and a method. *Journal of Experimental Psychology: Human Perception and Performance*, 10, 276-291.