Journal of Experimental Psychology:
Learning, Memory, and Cognition
1983, Vol. 9, No. 4, 585-606

On the Ability to Inhibit Simple Thoughts and Actions: I. Stop-Signal Studies of Decision and Memory

Gordon D. Logan
Erindale College, Mississauga, Ontario, Canada

In six experiments, subjects made speeded category and rhyme judgments about pairs of words and were occasionally asked to inhibit their responses. At issue was whether thoughts stop when the corresponding actions are inhibited. The extent to which thoughts were completed was assessed from memory for the words presented in the orienting tasks. Experiments 1 and 2 suggested that thoughts were not inhibited with their corresponding actions: Memory for words whose responses were inhibited was not different from memory for words whose responses were not inhibited, and memory performance did not vary with *stop-signal delay* (i.e., the time between the onset of the word pair and the onset of the stop signal). The same results were obtained in Experiments 3 and 4, though subjects made an overt response to the stop signal (tone) when it occurred. In Experiments 5 and 6, the display changed when the stop signal sounded, and the data suggested that thoughts were inhibited with their actions: Memory for words whose responses were inhibited was much worse than memory for words whose responses were not inhibited, and memory performance improved substantially as stop-signal delay increased. It was concluded that simple thoughts run on to completion whether or not the corresponding action is inhibited, provided that the stimuli that drive the thoughts are not disrupted before the thoughts finish.

This article reports an investigation of the ability to control simple thoughts and actions. Subjects were presented with pairs of words about which they made speeded category or rhyme decisions (e.g., is a frigate a boat? does *sleigh* rhyme with *play*?). Whenever a *stop signal* (a tone) occurred, subjects were to stop their response to the current word pair. Their ability to do so reflects their ability to control action: Responses that can be inhibited are clearly controlled, whereas responses that cannot be inhibited are clearly beyond control, hence ballistic (see Logan, 1981, 1982). The experiments focused primarily on the control of thought—determining whether or not subjects stopped thinking about the word pair when they inhibited the corresponding action and determining whether they inhibited the action by stopping the underlying thought.

Subjects' ability to control thought was assessed from their memory for the words about which they made decisions. I assumed that incomplete thoughts would result in poorer memory than complete thoughts, and I compared memory for words whose thoughts may have been inhibited with memory for words whose thoughts were likely to have gone on to completion. This comparison was made in two ways: First, memory for words that occurred with a stop signal was assessed as a function of *stop-signal delay* (i.e., the time at which the tone occurred relative to the onset of the word pair). If thoughts tend to stop with the action, memory performance should be better, the longer the stop-signal delay because thoughts are more likely to be completed, the longer the delay; if thoughts run on to completion whether or not action stops, memory performance should not depend on stop-signal delay because all thoughts should run on to completion at all delays. Second, memory for words whose responses were inhibited was

This research was supported by Natural Sciences and Engineering Research Council of Canada Grant U0053. I am grateful to Jane Zbrodoff for valuable discussions throughout the development of this project, particularly at its inception. I would like to thank Ken Davis for running the subjects, Fergus Craik, and Jan Rabinowitz for their help throughout the project and Jo Ann Miller, Richard Shiffrin, Donald Bamber, and two anonymous reviewers for comments that substantially improved the article.

Requests for reprints should be sent to Gordon D. Logan, who is now at the Department of Psychology, University of British Columbia, Vancouver, British Columbia, Canada V6T 1W5.

compared with memory for words whose responses were not inhibited (i.e., words whose responses escaped inhibition and words presented without a stop signal). If thoughts tend to stop with the action, memory for words whose responses were inhibited should be worse than memory for words whose responses were not inhibited; if thoughts run on to completion whether or not action stops, memory should be the same whether or not the response was inhibited.

These predictions were tested in six experiments, organized in three sets of two. Experiments 1 and 2 provide basic data on memory for words whose responses were inhibited. The remaining experiments were conducted to determine whether the same results would occur if subjects were given an additional task to perform when the stop signal occurred.

General Method

This section describes aspects of the method that are common to all six experiments. Departures from the general method will be noted when the experiments are introduced individually.

Subjects

Each experiment used a separate group of subjects, recruited primarily from an introductory psychology course subject pool. They received course credit for participation. Occasionally, subjects from other sources were recruited. These subjects were paid for their participation. No subject served in more than one experiment, and all subjects were naive as to the purpose of the experiment. The number of subjects in each experiment is presented in the first column of Table 1.

Apparatus and Stimuli

The stimuli were two sets of 280 words divided into 14 "categories" containing 20 members each. The words for the category task were selected from the Battig and Montague (1969) norms, excluding the 5 most frequent members of the categories, words that were longer than 12 letters, words that could be members of more than 1 category (e.g., *palm*), and category members that were more than 1 word long (e.g., *spanish onion*). The 280 words chosen ranged in length from 3 to 11 letters, averaging 6.29. The category names were *fish, bird, profession, mammal, sport, vegetable,* (musical) *instrument, clothing, body part, weather, tree, boat, fruit,* and (cooking) *utensil.* Each word was paired with its category name for *yes* trials and with the name of a different category for *no* trials.

The rhyme words were chosen from various rhyming dictionaries. Fourteen words were chosen as "rhyme categories," and 20 words that rhymed with each "category name" were selected. The 280 rhyme words ranged in length from 3 to 8 letters, averaging 5.14. The rhyme categories were *sleigh, raid, bear, braille, plane, freight, sign, might, here, plead, scene, tune, too,* and *though.* As in the category task, each word was paired with the category name it rhymed with for *yes* trials and with a category name it did not rhyme with for *no* trials.

The words were displayed on a point-plot CRT (Tektronix Model 604, equipped with P31 phosphor) under the control of a PDP 11/03 laboratory computer. The words were written in uppercase letters, formed by illuminating points in a 5 × 7 matrix. Viewed at a distance of 60 cm, each letter subtended .38° × .57° of visual angle. In the orienting tasks, two words were presented each trial, the category name and the word that was paired with it for that trial. The two words were justified four characters to the left of the fixation point, and the category name appeared one row above the word to be judged.

Each trial began with a fixation dot appearing in the center of the screen. After a 500-msec foreperiod, it was extinguished and replaced by the stimuli for that trial. In orienting blocks, the stimulus was a pair of words, as described above; in the recognition blocks, the stimulus was a single word. The temporal characteristics of the stimuli are presented in column 4 of Table 1 for the orienting task and column 9 of Table 1 for the recognition task. After the stimulus extinguished, the screen remained blank for a 2-sec intertrial interval.

The stop signal, presented in the orienting trials, was a 500-msec, 900 Hz tone presented through a speaker behind the CRT at a comfortable listening level. When it was presented, it occurred 100, 300, 500, or 700 msec after the onset of the word to be judged.

Responses were made by pressing one or more of a panel of eight telegraph keys, mounted on a movable board in front of the subject. The computer recorded which keys were pressed and the time at which each response occurred.

Procedure

Each experiment had two phases, the orienting tasks and the recognition tasks. The orienting tasks generally consisted of two blocks of 100 trials, one block for the category task and one block for the rhyme task (see Table 1, column 6). In each orienting block, half of the trials demanded a *yes* response and half demanded a *no* response. Stop signals occurred on a specified number of orienting trials (see Table 1, column 7), with one quarter of the total number occurring at each of four delays (100, 300, 500, and 700 msec) after the onset of the word pair. Stop signals occurred equally often at each delay with *yes* and *no* responses.

After both the orienting blocks were finished, subjects performed two blocks of recognition trials, one for the category task and one for the rhyme task. The recognition blocks consisted of all the words that had occurred with a stop signal (e.g., 40), an equal number of words that had occurred without a stop signal (e.g., 40), and twice as many words that had not been presented at all (e.g., 80). Because the number of stop signals varied between experiments, so did the number of recognition trials (see Table 1, column 11).

The order of orienting blocks varied between subjects, with half performing the category task before the rhyme task and half doing the opposite. Subjects performed the two recognition tasks in the same order as they performed the orienting tasks. Order of tasks was always orthogonal to other variables that were balanced between subjects

Table 1
Parameter Values for the Six Experiments

		Orienting task					Recognition task			
Experi-ment	*n*	FP (in sec)	Stimulus exposure	ITI (in sec)	No. of trials	No. of stop signals	FP (in sec)	Stimulus exposure (in sec)	ITI (in sec)	No. of trials
1	40	.5	Successive:							
			1st word—1 sec	2	100 (category)	40	.5	1	2	160 (category)
			2nd word—.5 sec		100 (rhyme)	40				160 (rhyme)
2	40	.5	Simultaneous:							
			1.5 sec	2	100 (category)	40	.5	1	2	160 (category)
					100 (rhyme)	40				160 (rhyme)
3	36	.5	Simultaneous:	2						
			1.5 sec		100 (category)	32	.5	1	2	128 (category)
					100 (rhyme)	32				128 (rhyme)
					100 (control)	32				
4	36	.5	Simultaneous:							
			1.5 sec	2	100 (category)	32	.5	1	2	128 (category)
					100 (rhyme)	32				128 (rhyme)
					100 (control)	32				
5	36	.5	Simultaneous:[a]							
			1.5 sec unless stop	2	100 (category)	32	.5	1	2	192 (category)
			signal occurs		100 (rhyme)	32				192 (rhyme)
6	36	.5	Simultaneous:[a]							
			1.5 sec unless stop	2	100 (category)	32	.5	1	2	192 (category)
			signal occurs		100 (rhyme)	32				192 (rhyme)

Note. FP = foreperiod; ITI = intertrial interval.
[a] When a stop signal occurred in Experiments 5 and 6, the first word pair was replaced by a new word pair. The two exposures summed to 1.5 sec (see the Method section for Experiments 5 and 6).

(e.g., stimulus–response mapping). The assignment of words to blocks of trials, the order of words within blocks, the assignment of stop signals to trials, and so on, all varied randomly from subject to subject. The computer constructed a different set of blocks for each subject.

Instructions were given before each block, and subjects were not told about the task for a particular block until they were ready to experience it. Thus, the recognition tests were unexpected; the experiments conformed to the incidental-learning paradigm. In the orienting blocks, the orienting task was described first, and subjects were told to respond as quickly and accurately as they could. Then the stop signal was described, and subjects were told to inhibit their responses to the orienting task when they heard it. They were told that stop-signal delay varied and that delays had been chosen so that sometimes they would be able to inhibit their responses and sometimes they would not. They were told not to let the stopping task interfere with the orienting task.

After both orienting tasks were completed, the first recognition task was described. Subjects were told to respond "yes" if they had seen the word in the orienting trials and "no" if they had not. They were told that half of the items in the recognition task had been presented in the orienting task and half had not. They were also told that the old items came from only one orienting task (they were told which one) and that the category names from the orienting task would not be presented.

Data analysis. The data analyses for the orienting tasks and the recognition tasks include orienting-task responses that were incorrect as well as responses that were correct. This departs from the common procedure of excluding erroneous responses from analysis because there was no way to determine whether a response that was inhibited would have been correct. Thus, it seemed best to analyze all the data.

In each experiment, data from the recognition tasks were analyzed in two ways. One analysis compared memory for words whose responses were inhibited (*signal-inhibit* trials) with memory for words whose responses were not inhibited. Words whose responses were not inhibited came from two kinds of trials, *signal-respond* trials, in which the signal occurred but subjects failed to inhibit their responses, and *no-signal* trials, in which no stop signal was presented. The signal-respond data and the signal-inhibit data were collapsed across stop-signal delay because there were too few observations in some conditions for analysis to be meaningful (i.e., some subjects inhibited all the time in some conditions and none of the time in other conditions). The other analysis assessed the effects of stop-signal delay on memory for words that occurred with a stop signal. There, the data were collapsed across signal-respond and signal-inhibit trials to provide enough observations for meaningful analysis. Both analyses were done on proportions of correct responses (i.e., hit rates), even though the cateogry and rhyme tasks generally had different false-alarm rates. The effects we are primarily concerned with (i.e., stop-signal delay and signal-inhibit vs. signal-respond and no-signal conditions) varied within tasks and so had the same false-alarm rates. The analyses were repeated using d' scores to correct for differences in false-alarm rate (calculating d' for each subject in each condition) and generally showed the same pattern of significance. Thus, the differences in false-alarm rate are not important.

For clarity of exposition, various data will be presented in appendixes rather than in the main text. Appendix A contains the orienting reaction times and accuracies for each experiment; Appendix B shows that there were no interactions between delay and conditions in the recognition data, suggesting that it was appropriate to collapse the data as described above; and Appendix C deals with a possible confound with processing time in the comparison between signal-respond and signal-inhibit recognition data, showing that the confound is not likely to affect conclusions drawn from the data.

Experiments 1 and 2

In the first two experiments, subjects were asked to stop their response in the orienting task when the stop signal sounded but were not asked to do anything else. In Experiment 1, the category name or the rhyme word was presented 1 sec before the word about which a decision was made. In Experiment 2 and in the experiments that followed it, the two words were presented simultaneously. Intuitively, it seemed that subjects would be more likely to complete thoughts when there was a delay between the category name or rhyme word and the word to be judged than when they were simultaneous. In the former case, thought can begin earlier and should be closer to finishing when the stop signal occurs.

Method

Apparatus and stimuli. These were the same as described in the General Method section, except that in Experiment 1, in which the category name or rhyme word preceded the word to be judged, stop-signal delay was defined relative to the onset of the word to be judged rather than the onset of the category name or rhyme word. In Experiment 2, in which the two words were exposed simultaneously, stop-signal delay was defined relative to the onset of the word pair.

Procedure. The procedure was as described in the General Method section. In the orienting tasks, half of the subjects pressed the rightmost key to indicate a *yes* response and the leftmost key to indicate a *no* response, and half did the opposite. The same mapping was used in the recognition tasks, where *yes* meant the word had appeared in the orienting tasks and *no* meant it had not.

Results and Discussion

Orienting tasks. The data on response inhibition are plotted in the left-hand panel of Figure 1 for Experiment 1 and the left-hand panel of Figure 2 for Experiment 2. The figures display the probability of failing to inhibit a response when a stop-signal occurs, *P* (re-

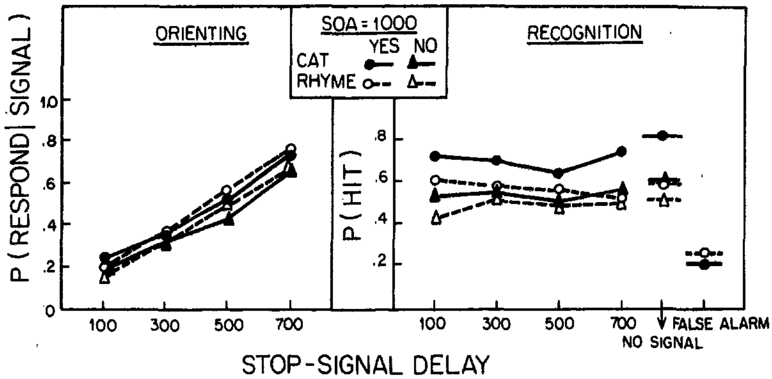

Figure 1. *Left:* Probability of responding given a stop signal in the orienting task of Experiment 1 as a function of stop-signal delay. (Type of orienting task [category vs. rhyme judgment] and response type [yes vs. no] are the parameters.) *Right:* Hit rate in the recognition task of Experiment 1 for words presented with a stop signal in the orienting task as a function of stop-signal delay. (Hit rates for no-signal trials are also plotted. Type of orienting task [category vs. rhyme judgment] and response type [yes vs. no] are the parameters. False alarm rates for each orienting task are also plotted. SOA = stimulus-onset asynchrony; cat = category.)

sponse/signal), as a function of stop-signal delay. In both experiments, the probability of responding given a stop signal increased with stop-signal delay, $F(3, 234) = 195.61, p < .01$, $MS_e = .0666$, corroborating previous findings (Lappin & Eriksen, 1966; Lisberger, Fuchs, King, & Evinger, 1975; Logan, 1981, 1982; Ollman, 1973). The Delay × Response Type interaction was significant, $F(3, 234) = 3.54$, $p < .05$, $MS_e = .0284$, reflecting floor effects with *no* responses. The probability of responding given a signal was lower in Experiment 2 than in Experiment 1, $F(1, 78) = 4.79$, $p < .05$, $MS_e = .7902$, and lower when subjects were to respond "no" than when they were to respond "yes," $F(1, 78) = 45.97$, $p < .01$, $MS_e = .0306$. The Response Type × Orienting Task interaction was significant, $F(1, 78) = 5.75$, $p < .05$, $MS_e = .0331$.

Recognition tasks. The effects of stop-signal delay are relevant to the issue of whether thought continues when the corresponding action stops. If thought stops with the action, memory performance should improve with

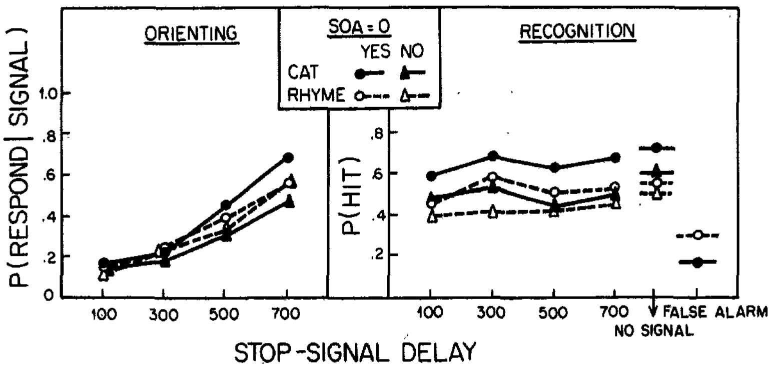

Figure 2. *Left:* Probability of responding given a stop signal in the orienting task of Experiment 2 as a function of stop-signal delay. (Type of orienting task [category vs. rhyme judgment] and response type [yes vs. no] are the parameters.) *Right:* Hit rate in the recognition task of Experiment 2 for words presented with a stop signal in the orienting task as a function of stop-signal delay. (Hit rates for no-signal trials are also plotted. Type of orienting task [category vs. rhyme judgment] and response type [yes vs. no] are the parameters. False alarm rates for each orienting task are also plotted. SOA = stimulus-onset asynchrony; cat = category.)

Table 2
Proportion-Correct Scores From the Recognition Tasks in Experiments 1 and 2

Orienting task	Response type: Yes NS	Yes SR	Yes SI	No NS	No SR	No SI	False alarm
Experiment 1							
Category	.82	.74	.68	.61	.60	.51	.21
Rhyme	.60	.61	.57	.52	.46	.49	.26
Experiment 2							
Category	.73	.64	.64	.60	.45	.50	.17
Rhyme	.56	.52	.51	.50	.40	.43	.30

Note. NS = no-signal trials; SR = signal-respond trials; SI = signal-inhibit trials.

stop-signal delay; if thought goes on to completion, memory performance should be independent of stop-signal delay. Hit rate in each combination of task and response type is plotted as a function of delay in the right-hand panel of Figure 1 for Experiment 1 and the right-hand panel of Figure 2 for Experiment 2. In neither experiment was hit rate dependent on stop-signal delay. In particular, there was no trace of an increase in hit rate with delay. In an analysis of variance (ANOVA) that included data from both experiments, the main effect of delay was significant, $F(3, 117) = 3.45$, $p < .01$, $MS_e = .0461$, but the linear trend reflecting the tendency for hit rate to increase was not significant, $F(1, 117) < 1$, $MS_e = .0461$, and no significant differences ($p < .05$) between pairs of means could be found using Fisher's Least Significant Difference (LSD) test (see Kirk, 1968). These null effects are clearly contrary to the notion that thoughts stop when action does, suggesting instead that thought tends to go on to completion independent of action.

The means across subjects for the analysis comparing signal-inhibit with signal-respond and no-signal conditions, collapsed across stop-signal delay, are presented in Table 2. The primary results, for the present purposes, were comparisons between signal-inhibit, signal-respond, and no-signal conditions: If thought stops when action does, memory should be poorer when actions were inhibited than when actions were completed; if thought goes on to completion, memory should be the same whether or not the actions were inhibited. Overall, subjects remembered words from no-signal trials, $P(\text{hit}) = .62$, better than words from signal-respond or signal-inhibit trials, average $P(\text{hit}) = .55$, but they remembered signal-inhibit words, $P(\text{hit}) = .54$, about as well as they remembered signal-respond words, $P(\text{hit}) = .55$. On the one hand, the superior performance in the no-signal condition suggests that thoughts were inhibited with the actions. Indeed, the main effect of conditions was significant, $F(2, 156) = 16.33$, $p < .01$, $MS_e = .0319$, as was a contrast comparing the no-signal condition with the average of the signal-inhibit and signal-respond conditions, $F(1, 156) = 8.34$, $p < .01$, $MS_e = .0319$. On the other hand, the equivalence of performance in the signal-inhibit and signal-respond conditions suggests that thoughts were not inhibited with the actions but, rather, went on to completion. A contrast comparing signal-inhibit and signal-respond performance was not significant, $F(1, 156) < 1$, $MS_e = .0319$.

The question is, which of the two contrasts provides the better basis for evaluating signal-inhibit performance? No-signal trials involve only one task, deciding about the word pair, whereas stop-signal trials involve switching from one task to another, namely, responding to the word pair and trying to inhibit the response. The second task may compete for "resources" with the first, resulting in a somewhat weaker trace (e.g., Johnston, Greenberg, Fisher, & Martin, 1970) or the second task may be associated with the first, resulting in a more distinctive trace that is less accessible to the recognition probe (cf. Tulving & Thompson, 1973). These possibilities will be discussed in more detail in the General Discussion section. From either perspective, however, it may be more appropriate to compare signal-inhibit and signal-respond trials, because both involve two tasks and differ only in whether or not an overt response occurred. In both experiments, this comparison showed that memory was the same whether or not the response was inhibited.

Memory performance was better in Experiment 1 than in Experiment 2, $F(1, 78) = 6.31$, $p < .05$, $MS_e = .1349$, but there were no significant interactions involving experiments. Memory performance was better with the category task than with the rhyme task, $F(1, 78) = 53.89, p < .01, MS_e = .0581$, better with *yes* responses in the orienting task than with *no* responses, $F(1, 78) = 107.75, p < .01$, $MS_e = .0372$, and the difference between *yes* and *no* responses was stronger in the category task than in the rhyme task, $F(1, 78), = 8.99$, $p < .01, MS_e = .0327$. These effects reflect typical findings in the literature on levels of processing (e.g., Craik & Tulving, 1975) and need not be discussed in depth. None of the interactions involving orienting task or response type were significant in this analysis or in the other, which assessed the effects of delay.

Experiments 3 and 4

Experiments 1 and 2 showed that subjects tend to complete simple thoughts even when they inhibit the corresponding actions. The incidental-learning procedure makes it unlikely that subjects completed the thoughts because they anticipated a memory test. However, they may have completed the thoughts deliberately for other reasons, for example, to alleviate boredom. Indeed, there was nothing in the procedure to prevent them from deliberately continuing the thought after they had inhibited the response.

Experiments 3 and 4 were designed to determine whether subjects would continue to complete thoughts if their attention was diverted to another task when the stop signal occurred. Thus, when the tone sounded in Experiment 3, subjects were to inhibit their response to the word pair and begin a different response to the tone. Because there was only one tone and one response, this was a *simple reaction time* task. In Experiment 4, the stop signal was one of two tones that differed in frequency. When a tone sounded, subjects were to inhibit their response to the word pair and respond to the tone, indicating which one had occurred. Because there were two different tones and two different responses, this was a *choice reaction time* task. There is evidence that it is difficult to perform an attention-demanding task while performing simple or choice reaction time tasks concurrently (e.g., Logan, 1980), so the new procedure should divert attention from processing the word pair. If it takes deliberate attention to complete thoughts, the pattern of results observed in Experiments 1 and 2 should not replicate. Instead, we should see evidence that thought stops when the corresponding responses are inhibited.

Method

Apparatus and stimuli. The apparatus and stimuli were the same as in Experiments 1 and 2 with the following exceptions: First, the word to be judged appeared simultaneously with the category name in both Experiments 3 and 4. Second, Experiment 3 used the same 900 Hz tone that was used in Experiments 1 and 2, but Experiment 4 used a 400 Hz tone as well as the 900 Hz tone. Third, in both Experiments 3 and 4, there were control trials to assess reaction time to the tone under single-task conditions. The control trials were the same as the regular orienting-task trials, except that two rows of six Xs appeared instead of the word pairs.

Procedure. Each subject completed the orienting tasks and the control task before beginning the recognition tasks. The order of these three tasks (category, rhyme, and control) was balanced between subjects, with six subjects in each experiment receiving each of the six possible orders of tasks.

Subjects reported their judgments about the word pairs by pressing the two rightmost keys in the panel of eight. In both experiments, every subject pressed the next-to-rightmost key to indicate a *yes* response and the rightmost key to indicate a *no* response. In Experiment 3, subjects responded to the tone by pressing the leftmost key. In Experiment 4, every subject pressed the next-to-leftmost key to indicate a high tone and the leftmost key to indicate a low tone. In the recognition tasks in both experiments, every subject pressed the rightmost key to indicate a *yes* response and the leftmost key to indicate a *no* response.

In Experiment 3, the same stop signal (tone) occurred in all conditions. In Experiment 4, the high tone occurred half of the time in each condition, and the low tone occurred the other half of the time.

The instructions were the same as those in Experiments 1 and 2, except that subjects were told to respond overtly to the tone when it sounded as well as inhibiting their response to the word pair.

Results and Discussion

Orienting tasks. The probability of responding to the word pair when the tone sounded (i.e., failing to inhibit the response) is plotted as a function of tone delay in the left-hand panel of Figure 3 for Experiment 3 and the left-hand panel of Figure 4 for Experiment 4. The probability of responding given a stop signal increased with stop-signal delay, $F(3, 210) = 103.84$, $p < .01$, $MS_e = 1.2159$, and was lower for *no* responses than for *yes* responses, $F(1, 70) = 78.02$, $p < .01$, $MS_e = .5784$. The experiments effect was marginally significant, $F(1, 70) = 3.83$, $p <$

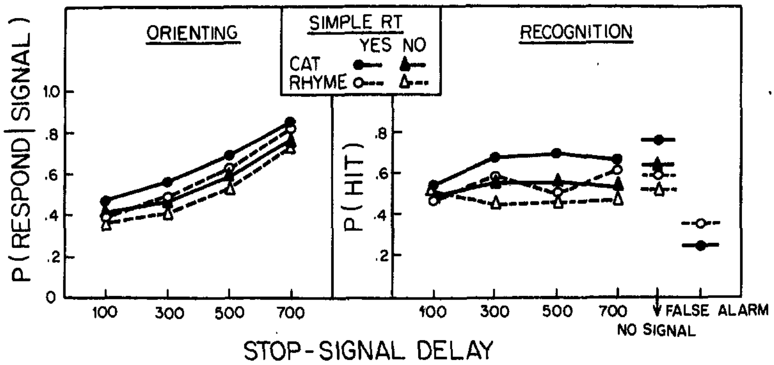

Figure 3. *Left:* Probability of responding given a stop signal in the orienting task of Experiment 3 as a function of stop-signal delay. (Type of orienting task [category vs. rhyme judgment] and response type [yes vs. no] are the parameters.) *Right:* Hit rate in the recognition task of Experiment 3 for words presented with a stop signal in the orienting task as a function of stop-signal delay. (Hit rates for no-signal trials are also plotted. Type of orienting task [category vs. rhyme judgment] and response type [yes vs. no] are the parameters. False alarm rates for each orienting task are also plotted. RT = reaction time; cat = category.)

.06, $MS_e = 21.4889$, and the category-rhyme effect was significant, $F(1, 70) = 4.63, p < .05$, $MS_e = 3.0722$.

Recognition tasks. Hit rates in the various conditions are plotted as a function of stop-signal delay in the right-hand panel of Figure 3 for Experiment 3 and the right-hand panel of Figure 4 for Experiment 4. Memory performance was relatively stable over delay, but it did tend to increase somewhat as delay increased. The increase was substantially smaller than the increase in the probability of an overt response in the orienting task as a function of delay, and it did not receive the same statistical support: In an analysis that included both experiments, the main effect of delay was significant, $F(3, 210) = 6.28, p < .01, MS_e = .0552$, but the linear trend, which tested the tendency for improvement with delay, was only marginal, $F(1, 210) = 3.25, p < .10$. The only pairwise difference to reach significance ($p < .05$) by Fisher's LSD test was between the shortest delay and the longest delay, and it barely exceeded the critical value. These effects were not significant when the experiments were analyzed separately. Further, there were no significant interactions involving delay in any analysis. Thus, on the balance, the delay effects

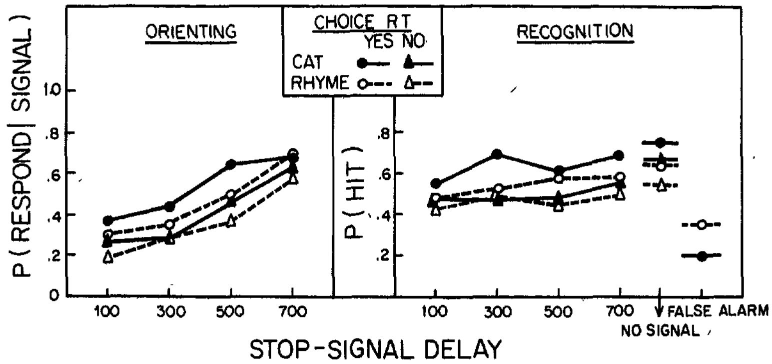

Figure 4. *Left:* Probability of responding given a stop signal in the orienting task of Experiment 4 as a function of stop-signal delay. (Type of orienting task [category vs. rhyme judgment] and response type [yes vs. no] are the parameters.) *Right:* Hit rate in the recognition task of Experiment 4 for words presented with a stop signal in the orienting task as a function of stop-signal delay. (Hit rates for no-signal trials are also plotted. Type of orienting task [category vs. rhyme judgment] and response type [yes vs. no] are the parameters. False alarm rates for each orienting task are also plotted. RT = reaction time; cat = category.)

suggest that thoughts went on to completion independent of the response.

The mean hit rates in no-signal, signal-respond, and signal-inhibit conditions appear in Table 3. The conditions effect seemed stronger in these experiments than in the previous ones, largely because performance in the signal-respond condition was much worse than performance in the no-signal and signal-inhibit conditions, *P*(hit)s were .45, .64, and .57, respectively. The difference between signal-inhibit and no-signal hit rates was about the same as in the previous experiments, however (.07 in the present experiments vs. .08 in Experiments 1 and 2). The important point is that signal-inhibit performance is in between no-signal and signal-respond performance, indicating that memory was no better or no worse for words whose responses were inhibited than for words whose responses were executed.

An ANOVA that included data from both experiments revealed a significant main effect of conditions, $F(2, 140) = 40.63$, $p < .01$, $MS_e = .0671$, and contrasts revealed a significant difference between the no-signal condition and the mean of the signal-respond and signal-inhibit conditions, $F(1, 140) = 12.64$, $p < .01$, $MS_e = .0671$, and a significant difference between signal-inhibit and signal-respond conditions, $F(1, 140) = 7.68$, $p < .01$, $MS_e = .0671$. The latter contrast was not significant in the analysis of d' scores, $F(1, 140) = 2.94$, $p < .10$, $MS_e = .7867$. The interaction between tasks and conditions was significant in the hit-rate data, $F(2, 140) = 3.18$, $p < .05$, $MS_e = .0510$, but not in the d's, $F(2, 140) = 1.22$, $MS_e = .7177$. The interaction in the hit rates was primarily due to poor performance in the signal-respond condition; it was not significant when only the signal-inhibit and no-signal conditions were tested, $F(1, 140) = 2.16$, $MS_e = .0510$.

Subjects in Experiment 3 who performed a simple reaction time task when the tone sounded had the same average hit rate as subjects in Experiment 4 who performed a choice reaction time task when the tone sounded, $F(1, 70) < 1$, $MS_e = .1554$, but the false-alarm rate was higher in Experiment 3 than in Experiment 4, resulting in significantly lower d's (.77 vs. .95), $F(1, 70) = 5.13$, $p < .05$, $MS_e = 1.4346$. However, the pattern of results was the same in the two experiments (i.e., there were no significant interactions involving experiments). The levels-of-processing effects replicated once again. Memory was better following category decisions than following rhyme decisions, $F(1, 70) = 20.07$, $p < .01$, $MS_e = .0559$, better for *yes* responses than for *no* responses, $F(1, 70) = 23.53$, $p < .01$, $MS_e = .0587$, and the task effect was stronger with *yes* responses than with *no* responses, $F(1, 70) = 8.36$, $p < .01$, $MS_e = .0308$.

Table 3
Proportion-Correct Scores From the Recognition Tasks in Experiments 3 and 4

	Response type						
	Yes			No			
Orienting task	NS	SR	SI	NS	SR	SI	False alarm
	Experiment 3						
Category	.76	.50	.66	.63	.41	.56	.24
Rhyme	.59	.38	.57	.51	.42	.49	.35
	Experiment 4						
Category	.76	.55	.64	.67	.40	.50	.20
Rhyme	.65	.46	.57	.54	.45	.54	.35

Note. NS = no-signal trials; SR = signal-respond trials; SI = signal-inhibit trials.

The conclusions drawn here corroborate the conclusions drawn from Experiments 1 and 2. The fact that memory was no poorer in the signal-inhibit condition than in the signal-respond condition and that memory did not improve as stop-signal delay increased suggests that thoughts went on to completion even when the accompanying responses were inhibited. This is important because the requirement to respond overtly to the tone makes it less likely that subjects deliberately continued thinking about the word pair after the stop signal. This suggests that simple thoughts run on to completion ballistically without attentional control once they begin.

Experiments 5 and 6

Experiments 1 and 2 suggested that thoughts are completed ballistically once begun, but there was nothing in the procedure to prevent the subjects from completing their thoughts deliberately. Experiments 3 and 4 cast doubt on this possibility by showing evidence for ballistic thinking when attention was distracted to different tasks that were dissimilar to the orienting tasks. Experiments 5 and 6 were designed to extend the findings in Experiments 3 and 4 by requiring subjects to switch to tasks that were very similar to the orienting tasks.

When a stop signal occurred in Experiments 5 and 6, the word pair for that trial was extinguished and replaced by a new word pair. In Experiment 5, subjects were required to make a decision about the new word pair and report it with an overt response that was different from the response to the first word pair. Experiment 6 was conducted as a control for Experiment 5. The word pair changed when the stop signal occurred, but subjects were only required to inhibit their response to the first word pair. The second word pair did not require a response. Thus, the contrast between Experiments 5 and 6 should reveal the extent to which having to begin a new train of similar thought interferes with the completion of the current one.

Method

Apparatus and stimuli. These were the same as before, except that on stop-signal trials, the first word pair was exposed only until the tone sounded (100, 300, 500, or 700 msec after the onset of the pair). It was then extinguished and replaced by the new word pair, which was exposed for 1,500 msec minus the stop-signal delay (i.e., 1,400, 1,200, 1,000, or 800 msec).

The stop signal was the word change and the 900 Hz, 500-msec tone used previously. Two redundant signals were used because previous experiments had shown that response latency to a word change as a stop signal was substantially slower than response latency to a tone as a stop signal (see Logan, 1982). In order to make the response to the signal comparable to that from the other experiments in this series, a tone was presented with the word change.

Procedure. In both experiments, every subject pressed the next-to-rightmost key to indicate a *yes* response and the rightmost key to indicate a *no* response. Experiment 5 required a response to the new word pair on stop-signal trials. Half of the subjects pressed the leftmost key to indicate a *yes* response to the second word pair and the next-to-leftmost to indicate a *no* response, and half did the opposite.

In both experiments, the word pair was replaced by a new word pair when the stop signal sounded. The task for the new word pair was the same as for the old (i.e., in category-task blocks the new word pairs required category judgments; in rhyme-task blocks the new word pairs required rhyme judgments). Half of the new word pairs required a *yes* response and half required a *no* response, and response type for the new word pair was orthogonal to the response type for the first word pair.

In both experiments, the recognition tasks involved 192 trials. This included the 32 words that occurred with a stop signal, the 32 "new" words that had replaced the first word pair on stop-signal trials, 32 of the 68 words that were presented without stop signals, and 96 words that had not appeared in the orienting task.

Instructions were largely the same as in the previous experiments, except that in Experiment 5, subjects were told to respond to the new word pair on stop-signal trials.

Results and Discussion

Orienting tasks. The probability of responding given a stop signal is plotted as a function of delay in the left-hand panel of Figure 5 for Experiment 5 and the left-hand panel of Figure 6 for Experiment 6. Again, the probability of responding given a signal increased with stop-signal delay, $F(3, 210) = 82.68$, $p < .01$, $MS_e = 1.3874$. The probability of responding given a signal was lower in Experiment 5 than in Experiment 6, $F(1, 70) = 4.94$, $p < .05$, $MS_e = 16.5538$, and lower with *no* responses than with *yes* responses, $F(1, 70) = 17.38$, $p < .01$, $MS_e = .7312$. The main effect of orienting task was not significant, $F(1, 70) = 2.77$, $MS_e = 2.9493$. There were significant interactions between experiments and response type, $F(1, 70) = 5.02$, $p < .05$, $MS_e = .7312$, and delay and response type, $F(3, 210) = 7.61$, $p < .01$, $MS_e = .4970$, indicating floor effects.

Recognition tasks. The delay effects are plotted in the right-hand panel of Figure 5 for Experiment 5 and the right-hand panel of Figure 6 for Experiment 6. Hit rates increased with delay in Experiment 5, suggesting that the change in the display or the requirement to stop one task and begin another similar one inhibited or otherwise impaired the underlying thought. The increase in hit rate with delay was just as strong in Experiment 6, in which subjects only stopped, suggesting that it was the change in the display rather than the switch to a similar task that inhibited thought.

An ANOVA showed that the main effect of delay was significant, $F(3, 210) = 42.25$, $p < .01$, $MS_e = .0567$, and contained a highly significant linear component, $F(1, 210) = 30.30$, $p < .01$, $MS_e = .0567$. Fisher's LSD test revealed significant differences ($p < .05$) between the shortest delay and the longest two and between the second shortest delay and the longest delay. In addition, there were interactions be-

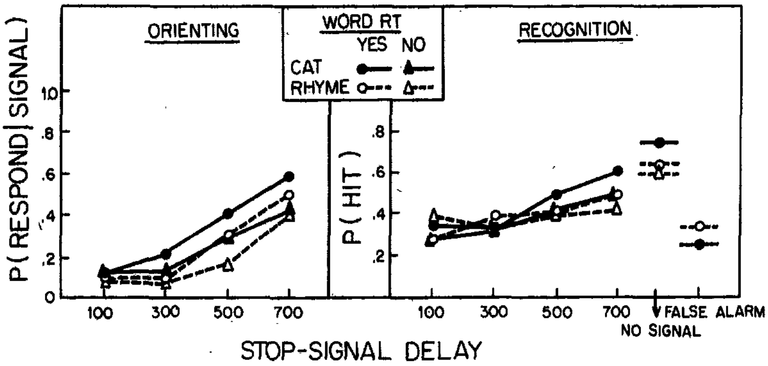

Figure 5. Left: Probability of responding given a stop signal in the orienting task of Experiment 5 as a function of stop-signal delay. (Type of orienting task [category vs. rhyme judgment] and response type [yes vs. no] are the parameters.) *Right:* Hit rate in the recognition task of Experiment 5 for words presented with a stop signal in the orienting task as a function of stop-signal delay. (Hit rates for no-signal trials are also plotted. Type of orienting task [category vs. rhyme judgment] and response type [yes vs. no] are the parameters. False alarm rates for each orienting task are also plotted. RT = reaction time; cat = category.)

tween tasks and delay, $F(3, 210) = 5.56$, $p < .01$, $MS_e = .0479$, and between response type and delay, $F(2, 210) = 5.36$, $p < .01$, $MS_e = .0608$.

Another possible interpretation is that memory for words that occurred with a signal was impaired because the second word masked the first. It is well established that the perceptibility of a visual stimulus can be impaired by another visual stimulus in close temporal and spatial contiguity, and there is evidence that the effect may extend over intervals of several hundred milliseconds between stimulus onsets if the task is a difficult one (Merikle, 1977). According to the masking hypothesis, memory performance would increase with delay because masking would be weaker, the longer the delay. The masking hypothesis can also account for the finding that in the orienting task, accuracy was substantially lower on signal-respond trials than on no-signal trials; the former were subject to masking by a subsequent word pair, and the latter were not (see Appendix A).

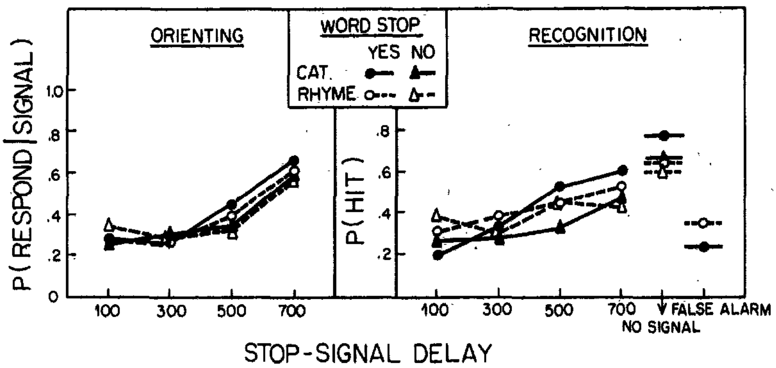

Figure 6. Left: Probability of responding given a stop signal in the orienting task of Experiment 6 as a function of stop-signal delay. Type of orienting task [category vs. rhyme judgment] and response type [yes vs. no] are the parameters.) *Right:* Hit rate in the recognition task of Experiment 6 for words presented with a stop signal in the orienting task as a function of stop-signal delay. (Hit rates for no-signal trials are also plotted. Type of orienting task [category vs. rhyme judgment] and response type [yes vs. no] are the parameters. False alarm rates for each orienting task are also plotted. cat = category.)

Table 4
Proportion-Correct Scores From the Recognition Tasks in Experiments 5 and 6

Orienting task	Yes NS	Yes SR	Yes SI	No NS	No SR	No SI	False alarm
Experiment 5							
Category	.76	.37	.53	.64	.32	.45	.27
Rhyme	.61	.35	.43	.53	.35	.29	.36
Experiment 6							
Category	.78	.33	.49	.67	.31	.40	.25
Rhyme	.60	.38	.52	.55	.32	.35	.36

Note. NS = no-signal trials; SR = signal-respond trials; SI = signal-inhibit trials.

The hit rates from the analysis of conditions are presented in Table 4. The conditions effects were much stronger in these experiments than in the previous ones. In Experiment 5, where subjects switched to a new task when the stop signal occurred, performance was much lower on signal trials than on no-signal trials, P(hit)s were .39 and .64, respectively. This could be due to switching to a new task or to changing to a new word. In Experiment 6, in which subjects simply stopped when the word changed, performance was much lower on signal trials than on no-signal trials, P(hit)s were .38 and .65, respectively, suggesting that depressed memory performance on signal trials was due to changing the display.

An ANOVA revealed no significant differences between experiments, neither main effect, $F(1, 70) < 1$, $MS_e = .1390$, nor interactions. The conditions effect was highly significant, $F(2, 140) = 122.10$, $p < .01$, $MS_e = .0563$. It was analyzed by planned contrasts, which revealed significant differences between the no-signal condition and the average of signal-respond and signal-inhibit conditions, $F(1, 140) = 55.71$, $p < .01$, $MS_e = .0563$, and significant differences between signal-inhibit and signal-respond conditions, $F(1, 140) = 5.56$, $p < .05$, $MS_e = .0563$. Conditions interacted significantly with tasks, $F(2, 140) = 8.82$, $p < .01$, $MS_e = .0520$, and with response type, $F(2, 140) = 3.18$ $p < .05$, $MS_e = .0498$. The interactions between conditions and encoding effects were primarily due to floor effects in the signal-respond conditions.

Note that the conditions effects are more consistent with the masking hypothesis than with the hypothesis that thoughts were interrupted. If thoughts were interrupted, signal-inhibit performance should be worse than signal-respond performance because the former represent incompleted thoughts and the latter represent completed thoughts. However, if the first word pair was masked by the second, signal-inhibit and signal-respond performance should not differ. The finding that performance was worse in the signal-respond condition than in the signal-inhibit condition clearly contradicts the prediction of the hypothesis that thought was interrupted, though it need not be inconsistent with the masking hypothesis.

The level-of-processing effects replicated once again: The category-rhyme effect was significant, $F(1, 70) = 14.52$, $p < .01$, $MS_e = .0608$, as was the response type effect (*yes* vs. *no*), $F(1, 70) = 28.94$, $p < .01$, $MS_e = .0490$. The Orienting Task × Response Type interaction was not significant, $F(1, 70) < 1$, $MS_e = .0470$.

Hit rates for the new words that replaced the old ones on stop-signal trials are presented in Table 5. In Experiment 5, where the new words received a full-blown response, memory for the new words was worse than memory for no-signal words and was not much better than memory for words on signal-inhibit trials. This suggests that the inhibitory influence between the word pairs on stop-signal trials was mutual; the first word impaired the second, and the second impaired the first. Similar effects were found in Experiment 6, although subjects were not required to make an overt response to the new word. Note that the response type effect is reversed in the new words

Table 5
Proportion-Correct Scores for the Second Word in Experiments 5 and 6 as a Function of Orienting Task and Response Type

Experiment	Category task Yes	Category task No	Rhyme task Yes	Rhyme task No
5	.67	.46	.53	.51
6	.43	.56	.45	.59

in Experiment 6. I cannot offer a reasonable interpretation for the reversal.

General Discussion

Major Findings

There were two major effects in the memory data from each experiment that bore on the control of thought: the effects of stop-signal delay and the effects of conditions (no signal vs. signal-inhibit vs. signal-respond). These effects will be discussed and qualified by other aspects of the data. Response inhibition was the third major effect in each experiment. It will be discussed and possible mechanisms underlying it will be described.

Delay effects. In Experiments 1–4, memory for words presented with a stop signal was relatively independent of stop-signal delay. This suggests that simple thoughts tend to go on to completion even when their responses are inhibited. The fact that the same pattern was found whether subjects simply stopped when the signal occurred (Experiments 1 and 2) or whether subjects stopped one task to begin another (Experiments 3 and 4) is important because it suggests that the tendency to complete thought does not depend on subjects' sustained attention to the word pair.

The pattern of performance was different in Experiments 5 and 6, in which the word pair changed when the stop signal sounded. Memory for words that occurred with a stop signal improved substantially as delay increased. In contrast with the first four experiments, the improvement suggests that thoughts were inhibited with the responses or impaired by a masking effect from the second word pair. The pattern was the same whether or not subjects had to respond to the new word pair (Experiment 5 vs. Experiment 6). This is important because it suggests that it was the change in the display rather than the switch to a new task that inhibited thought. The findings from Experiments 5 and 6 place an important qualification on the conclusions drawn from Experiments 1–4: Thoughts tend to go on to completion independent of action provided that the stimuli that drive them are not disrupted before the thoughts reach completion.

Conditions effects. In Experiments 1–4, memory for signal-inhibit words was no worse and often better than memory for signal-respond words, suggesting that thoughts tended to go on to completion whether or not the responses were inhibited, corroborating the analysis of delay effects. Signal-inhibit words were remembered rather poorly in Experiments 5 and 6, relative to no-signal words, but still, they were remembered better than signal-respond words. The former is consistent with the idea that thoughts were inhibited with their actions, whereas the former and the latter are consistent with the idea that the new word pair that accompanied the stop signal masked or otherwise interfered with the processing of the first word pair. In either case, the results corroborate the delay analysis, suggesting that thoughts tend to go on to completion independent of action unless the stimuli that drive them are disrupted.

In each experiment, there was a tendency for words that occurred with a stop signal (signal-inhibit and signal-respond words) to be remembered less accurately than words that occurred without a stop signal (no-signal words). The effect was stronger in Experiments 3 and 4 than in Experiments 1 and 2, and stronger in Experiments 5 and 6 than in Experiments 3 and 4. Clearly, the signal did affect thought in some way. Possibly, the requirement to process the signal drew "resources" away from the decision at hand and resulted in a poorer trace (cf. Johnston et al., 1970). This would account for the effect becoming stronger as the stopping task became more complex (e.g., Experiments 1 and 2 vs. Experiments 3 and 4 vs. Experiment 5), but it could not account for the difference between Experiments 1 and 2 on the one hand and Experiment 6 on the other. An account in terms of resources would predict the same deficit in Experiment 6 as in Experiments 1 and 2 because subjects simply stopped on signal in all three experiments.

Another possibility is that thoughts about the stop signal are incorporated into the thoughts about the word pair. Consequently, a memory trace from a stop-signal trial is less similar to the single-word recognition probe than is a trace from a no-signal trial. If memory performance depends on the similarity between the encoded trace and the recognition probe (cf. Tulving & Thompson, 1973), the progressive strengthening of the effect over experiments can be accounted for: The trace becomes less similar to the recognition probe as the stopping task becomes more complex. Signal trials in Experiments 5 and 6 would yield equivalent performance because in both cases the second word would be incorporated into the trace of the first.

The present data cannot provide a definitive test of these alternative accounts. However, the alternatives do make the point that the difference between signal and no-signal trials can be accounted for without assuming that the signal inhibits thought. Moreover, they suggest that it is more appropriate to compare signal-inhibit trials with signal-respond trials than with no-signal trials because signal-inhibit trials are more like signal-respond trials than no-signal trials.

A potentially serious problem in comparing signal-inhibit and signal-respond trials is that signal-inhibit responses were slower than signal-respond responses, and slower responses may lead to better memory. This problem was ruled out by comparing memory for the fastest and slowest halves of the no-signal orienting trials in each experiment (see Appendix C). In general, speed in the orienting task made little difference to memory performance (also see Craik & Tulving, 1975), suggesting that there should be no bias in comparing (fast) signal-respond and (slow) signal-inhibit trials. However, there may be the other sources of bias (see the discussion of the Zeigarnik effect below).

How action stops. The response-inhibition data were related closely to the orienting reaction times reported in Appendix A. In each experiment, factors that affected reaction time had corresponding effects on the probability of responding given a stop signal. In each experiment, reaction time was faster for *yes* responses than for *no* responses, and the probability of responding given a signal was higher for *yes* responses than for *no* responses. Reaction time was faster in Experiment 1 than in Experiment 2, and the probability of responding given a signal was higher in Experiment 1 than in Experiment 2. The same was true of the contrast between Experiments 5 and 6. These findings are consistent with other data that suggest that responses can be stopped up to the point of execution (e.g., Lisberger et al., 1975; Logan, 1981, 1982). In general, response-inhibition data can be accounted for by a model in which execution or inhibition of a response depends on a race between the processes generating the response and the processes responding to the stop signal. If the response-generation process finishes first, the response is executed; if the stopping process finishes first, the response will be inhibited. The model predicts that response inhibition will be locked in time to the expected occurrence of the response rather than the occurrence of the stimulus. Empirically, this means that the interval between the onset of the stop signal and the expected occurrence of the response (i.e., reaction time minus stop-signal delay) will be a better predictor of the probability of responding given a stop signal than will be the interval between the onset of the stimulus and the onset of the stop signal (i.e., stop-signal delay). This prediction has been confirmd in studies of simple eye movements (Lisberger et al., 1975), choice reaction time (Logan, 1981), and typewriting (Logan, 1982), and it also appears to be true in the present studies. The probability of responding given a signal is plotted as a function of the orienting reaction time minus stop-signal delay in Figure 7. In each experiment, the points from different orienting tasks and response types (*yes* vs. *no*) seem to line up, as if they were generated from the same function.

The memory data have some implications about the mechanisms by which actions are inhibited. The finding that thoughts tend to run on to completion whether or not the action is executed suggests that actions are not stopped by inhibiting the thoughts that drive them. If actions were stopped by inhibiting thought, thoughts should stop when actions do and run on to completion when actions do; memory performance should correlate with response inhibition. Contrary to the prediction, the data from Experiments 1–4 showed no evidence of a correlation. This rules out the possibility that subjects stopped their actions by inhibiting the underlying thoughts, leaving two broad alternatives: Actions could be stopped directly, by inhibiting or reprogramming the motor system, or indirectly, by disrupting communication between the processes responsible for thought and action (Logan, 1982).

Both of these alternatives predict that thoughts would run on to completion whether or not action stopped. In the former case, sub-

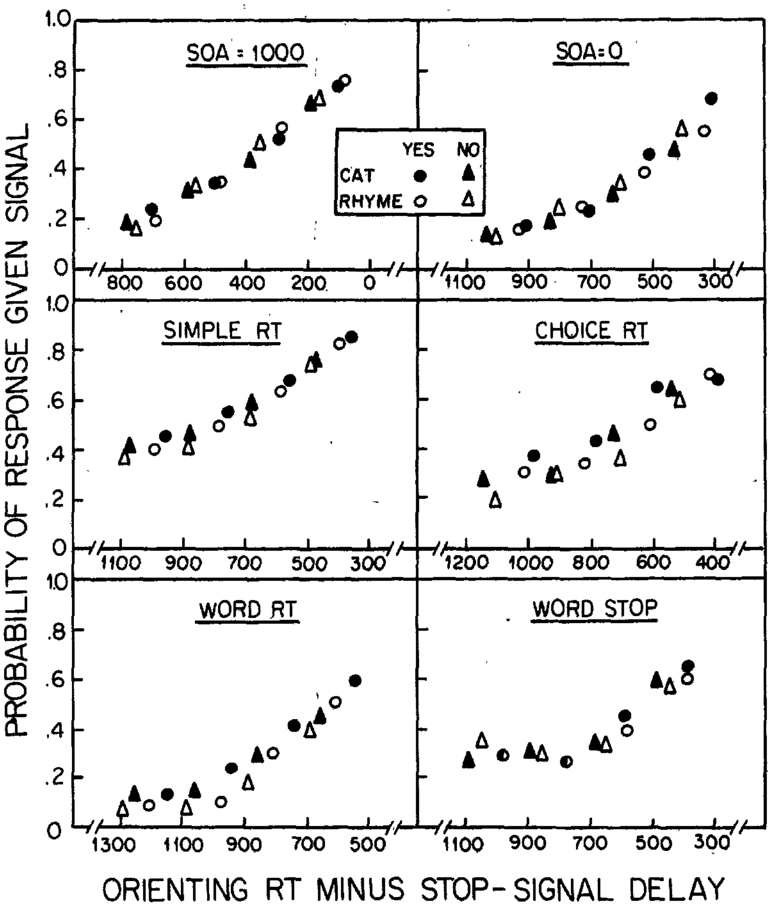

Figure 7. Probability of responding given a stop signal in each experiment as a function of the time elapsing between the onset of the stop signal and the expected occurrence of the response (orienting reaction time [RT] minus stop signal delay; stimulus-onset asynchrony [SOA] = 1,000 is Experiment 1; SOA = 0 is Experiment 2; simple RT is Experiment 3; choice RT is Experiment 4; word RT is Experiment 5; and word stop is Experiment 6; cat = category).

jects could complete the thought and inhibit the response just before execution. In the latter case, subjects could complete the thought without communicating the result to the motor system. The memory data do not distinguish between these alternatives. Nevertheless, the experiments rule out the possibility that subjects stop action by inhibiting the underlying thought.

Implications for Future Research

Automaticity and ballistic processing. Many theorists distinguish between two modes of processing, *attentionally controlled* processing, which is relatively slow, serial, and subject to interference from concurrent tasks, and *automatic* processing, which is relatively fast, parallel, and resistant to interference from concurrent tasks (Hasher & Zacks, 1979; James, 1890; LaBerge & Samuels, 1974; Logan, 1978, 1979; Posner & Snyder, 1975; Shiffrin & Schneider, 1977). In addition to these well-documented properties, theorists often claim that automatic processes are harder to inhibit than attentionally controlled processes, citing as evidence the ubiquitous Stroop effect (e.g., Posner & Snyder, 1975; Shiffrin & Schneider, 1977).

The present experiments provide direct evidence on the difficulty of inhibiting various processes. The general finding that thoughts ran on to completion even when responses were inhibited suggests that thoughts are harder to inhibit than actions. This, in turn, suggests that thoughts are automatic and actions are attentionally controlled. Contrasts between the experiments in which subjects stopped and switched to another task (Experiments 3, 4, and 5) and the experiments in which subjects simply stopped (Experiments 1, 2, and 6) provide converging evidence. On the one hand, thoughts were completed to the same extent whether subjects switched to a new task or simply stopped responding when the signal occurred, suggesting that thoughts are not subject to dual-task interference, hence automatic. On the other hand, orienting reaction times were substantially slower when subjects switched to a new task than when they simply stopped, suggesting that the processes generating a response are subject to dual-task interference, hence attentionally controlled.

These conclusions are consistent with the idea, developed from a dual-task analysis of reaction time data, that the components of a task are automatic but their organization as a set or program to perform the task at hand is not (Logan, 1978, 1979, 1980). From this perspective, the components are driven by the available data and run on to completion once begun, just as the simple thoughts in the present experiments ran on to completion. However, the organization of components is driven by an intention to perform, and the components will work together to produce task-relevant action only as long as the intention persists. Thus, the stop signal inhibits action by changing the subject's intention. (Note the similarity between this view and the idea that action is inhibited by disrupting communication between processes.)

Complexity and ballistic thought. Is it the case that all thoughts are completed ballistically? Probably not. The present experiments focused on simple thoughts that likely represent single links in the more complex chains that we normally characterize as thinking. I suspect that chains of simple thoughts should be easy to stop before they reach their natural conclusion. Thus, studies of the inhibition of complex thoughts should show the effects of stop-signal delay and conditions (no signal vs. signal-respond vs. signal-inhibit) that were absent in the present studies.

In the present studies, there was no relation between processing time and memory. Stop-signal delay had little effect on memory performance, and in the no-signal trials, memory did not depend on how fast the orienting task was performed (see Appedix C; also see Craik & Tulving, 1975, Figure 5). By contrast, there is a large literature showing a strong relation between processing time and memory in more complex tasks. There is evidence that slower presentation rates lead to better memory (Murdock, 1962, 1965), that more time spent in rehearsal leads to better memory (Rundas, 1971; Woodword, Bjork, & Jongeward, 1973), and that repeated presentations lead to better memory (see Murdock, 1974).

The difference between the present studies and the ones showing positive effects of processing time suggests that memory is independent of the time taken to complete a single thought but improves as more thoughts are associated with the item or event being encoded. Thus, the present studies reflect memory for single thoughts, whereas the studies of presentation rate, rehearsal, and repeated presentations reflect memory that results from several thoughts. This suggests that the stop-signal method, applied to memory as it was in this article, can be used to determine empirically what constitutes a single thought. This possibility is potentially important and warrants further research.

The conditions effects may be less informative in complex tasks than they were in the present ones because they may be more susceptible to motivational biases in complex tasks than in simple ones. Zeigarnik (1927) investigated memory for complex tasks that had been interrupted and found that subjects recalled interrupted tasks much better than they recalled completed tasks (also see Marrow, 1938a, 1938b; Pachauri, 1935a, 1935b, 1936). Superior recall of interrupted tasks depends critically on the subject's attitude, however. Marrow (1938b) told subjects that they would be interrupted if they showed they had mastered the task, with the implication that they would have to complete tasks they did not do well on. He found that completed tasks were recalled much better than interrupted tasks, reversing the usual Zeigarnik effect. Zeigarnik herself found no difference in the recall of interrupted and completed tasks when the tasks were performed in the context of a tour of the laboratory instead of the usual context of a formal experiment.

It may be possible to counter the motivational effects by testing recognition for the materials used in the task rather than recall of the tasks themselves. Pachauri (1935a) found that subjects tended to recall interrupted tasks first, and it is possible that their subsequent recall of completed tasks was impaired by output interference. Such recall strategies can be eliminated by using a recognition test in which interrupted and completed tasks are mixed randomly. Testing recognition of the materials may also help; Caron and Wallach (1957) found no difference between interrupted and completed tasks when subjects' only task was to unscramble scrambled sentences, and recall of the sentences was tested. Nevertheless, the Zeigarnik effect suggests caution in interpreting conditions effects in complex tasks.

Fortunately, the delay effects are not susceptible to the same criticisms as the conditions effects. Indeed, Zeigarnik (1927) found that tasks that were interrupted close to being completed were recalled better than tasks that were interrupted earlier, suggesting that more complete thoughts tend to be remembered better.

Conclusions

The experiments suggest that the simple thoughts underlying category and rhyme judgments tend to go on to completion whether or not the corresponding response is inhibited, unless the stimuli driving them are disrupted. This was true whether or not subjects switched to a new task when they stopped their response to the category or rhyme task. The experiments also suggest that the inhibition of simple actions (key presses) can be described by a model in which the processes generating a response race against the processes responding to the stop signal. The memory data rule out the possibility that subjects stop their actions by inhibiting the underlying thought.

References

Battig, W. F., & Montague, W. E. Category norms for verbal items in 56 categories: A replication and extension of the Connecticut category norms. *Journal of Experimental Psychology Monograph,* 1969, *80*(3, Pt. 2).

Caron, A. J., & Wallach, M. A. Recall of interrupted tasks under stress: A phenomenon of memory or of learning? *Journal of Abnormal and Social Psychology,* 1957, *55,* 372-381.

Craik, F. I. M., & Tulving, E. Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General,* 1975, *104,* 268–294.

Hasher, L. T., & Zacks, R. T. Automatic and effortful processes in memory. *Journal of Experimental Psychology: General,* 1979, *108,* 356–388.

James, W. *Principles of psychology.* New York: Holt, 1890.

Johnston, W. A., Greenberg, S. N., Fisher, R. P., & Martin, D. W. Divided attention: A vehicle for monitoring memory processes. *Journal of Experimental Psychology,* 1970, *83,* 164–171.

Kirk, R. E. *Experimental design: Procedures for the behavioral sciences.* Monterey, Calif.: Brooks/Cole, 1968.

LaBerge, D., & Samuels, S. J. Toward a theory of automatic information processing in reading. *Cognitive Psychology,* 1974, *6,* 293–323.

Lappin, J. S., & Eriksen, C. W. Use of a delayed signal to stop a visual reaction-time response. *Journal of Experimental Psychology,* 1966, *72,* 805–811.

Lisberger, S. G., Fuchs, A. F., King, W. M., & Evinger, L. C. Effect of mean reaction time on saccadic responses to two-step stimuli with horizontal and vertical components. *Vision Research,* 1975, *15,* 1021–1025.

Logan, G. D. Attention in character classification tasks: Evidence for the automaticity of component stages. *Journal of Experimental Psychology: General,* 1978, *107,* 32–63.

Logan, G. D. On the use of a concurrent memory load to measure attention and automaticity. *Journal of Experimental Psychology: Human Perception and Performance,* 1979, *5,* 189–207.

Logan, G. D. Short-term memory demands of reaction time tasks that differ in complexity. *Journal of Experimental Psychology: Human Perception and Performance,* 1980, *6,* 375–389.

Logan, G. D. Attention, automaticity, and the ability to stop a speeded choice response. In J. Long & A. D. Baddeley (Eds.), *Attention and performance IX.* Hillsdale, N.J.: Erlbaum, 1981.

Logan, G. D. On the ability to inhibit complex movements: A stop-signal study of typewriting. *Journal of Experimental Psychology: Human Perception and Performance,* 1982, *8,* 778–792.

Marrow, A. J. Goal tensions and recall: I. *Journal of General Psychology,* 1938, *19,* 3–35. (a)

Marrow, A. J. Goal tensions and recall: II. *Journal of General Psychology,* 1938, *19,* 37–64. (b)

Merikle, P. M. On the nature of metacontrast with complex targets and masks. *Journal of Experimental Psychology: Human Perception and Performance,* 1977, *3,* 607–621.

Murdock, B. B., Jr. The serial position effect of free recall. *Journal of Experimental Psychology,* 1962, *64,* 482–488.

Murdock, B. B., Jr. A test of the "limited capacity" hypothesis. *Journal of Experimental Psychology,* 1965, *69,* 237–240.

Murdock, B. B., Jr. *Human memory: Theory and data.* Potomac, Md.: Erlbaum, 1974.

Ollman, R. T. Simple reactions with random countermanding of the "go" signal. In S. Kornblum (Ed.) *Attention and performance IV*. New York: Academic Press, 1973.

Pachauri, A. R. A study of Gestalt problems in completed and interrupted tasks: Part I. *British Journal of Psychology*, 1935, *25*, 365–381. (a)

Pachauri, A. R. A study of Gestalt problems in completed and interrupted tasks: Part II. *British Journal of Psychology*, 1935, *25*, 447–457. (b)

Pachauri, A. R. A study of Gestalt problems in completed and interrupted tasks: Part III. *British Journal of Psychology*, 1936, *27*, 170–180.

Posner, M. I., & Snyder, C. R. R. Attention and cognitive control. In R. L. Solso (Ed.), *Information Processing and Cognition: The Loyola Symposium*. New York: Halstead Press, 1975.

Rundas, D. Analysis of rehearsal processes in free recall. *Journal of Experimental Psychology*, 1971, *89*, 63–77.

Shiffrin, R. M., & Schneider, W. Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 1977, *84*, 127–190.

Tulving, E., & Thompson, D. M. Encoding specificity and retrieval processes in episodic memory. *Psychological Review*, 1973, *80*, 352–373.

Woodward, A. E., Bjork, R. A., & Jongeward, R. H. Recall and recognition as a function of primary rehearsal. *Journal of Verbal Learning and Verbal Behavior*, 1973, *12*, 608–617.

Zeigarnik, B. Ueber das Behalten von erledigten und unerledigten Handlungen. *Psycholgische Forschung*, 1927, *9*, 1–85.

Appendix A

Orienting-Task Data

Experiments 1 and 2

The mean reaction times from the no-signal trials appear in Table A1 with the corresponding error rates. Overall, reaction times were much faster in Experiment 1 than in Experiment 2, $F(1, 78) = 33.76$, $p < .01$, $MS_e = 130{,}703.03$. However, there were no significant interactions involving experiments. In both experiments, subjects responded "yes" faster than they responded "no," $F(1, 78) = 137.93$, $p < .01$, $MS_e = 4{,}892.37$, but they were no faster on the rhyme task than on the category

Table A1
Reaction Time (RT; in msec) and Percentage Correct in the Orienting Task

	No signal				Signal respond			
	Category		Rhyme		Category		Rhyme	
Experiment	Yes	No	Yes	No	Yes	No	Yes	No
1								
RT	800	892	795	858	632	791	632	743
% correct	90	91	89	93	88	87	87	83
2								
RT	1,005	1,140	1,030	1,108	823	990	846	886
% correct	91	90	90	91	86	85	83	80
3								
RT	1,058	1,181	1,086	1,190	970	1,120	994	1,065
% correct	90	92	94	92	88	77	90	72
4								
RT	1,095	1,227	1,119	1,205	1,064	1,198	1,063	1,190
% correct	93	95	94	94	87	87	90	78
5								
RT	1,241	1,358	1,311	1,393	1,123	1,245	1,216	1,261
% correct	91	92	91	90	80	71	72	63
6								
RT	1,083	1,197	1,082	1,143	1,309	1,445	1,225	1,335
% correct	89	91	92	93	65	62	65	62

task, $F(1, 78) < 1$, $MS_e = 21{,}474.94$. The effect of response type was larger in the category task than in the rhyme task, $F(1, 78) = 8.67$, $p < .01$, $MS_e = 4{,}220.50$.

Reaction times from signal-respond trials are also presented in Table A1. The pattern is similar to the one observed in the no-signal trials, except that signal-respond trials tended to be faster and less accurate than no-signal trials. This is consistent with previous findings (Logan, 1981), suggesting that signal-respond trials came from the fast, inaccurate tail of the reaction time distribution.

Experiments 3 and 4

Mean reaction time and accuracy from the no-signal trials are presented in Table A1. Overall, there was no difference between experiments, $F(1, 70) < 1$, $MS_e = 95{,}710.76$. *Yes* responses were faster than *no* responses, $F(1, 70) = 73.60$, $p < .01$, MS_e 5,134.54, but category judgments were no faster than rhyme judgments, $F(1, 70) < 1$, $MS_e = 22{,}074.98$. Response type did have a stronger effect in the category task than in the rhyme task, $F(1, 70) = 4.73$, $p < .05$, $MS_e = 4{,}151.82$.

Reaction time and accuracy for signal-respond trials also appears in Table A1. As in the previous experiments, signal-respond trials tended to be faster and less accurate than no-signal trials, but reflected the same trends.

Reaction times to the tones presented during the orienting tasks and the control trials are presented in Table A2. There was a large difference between experiments, reflecting the usual difference between simple- and choice-reaction time, $F(1, 70) = 91.15$, $p < .01$, $MS_e = 472{,}460.90$. Also, there were large differences between tones presented during the orienting tasks and tones presented during the control trials, $F(2, 140) = 213.37$, $p < .01$, $MS_e = 138{,}202.09$, reflecting interference from performing the orienting tasks concurrently. There was a tendency for tone reaction times to be longer in the context of the category task than in the context of the rhyme task, but the difference was relatively small, $F(1, 140) < 1$, $MS_e = 138{,}202.09$. Finally, reaction time declined gradually with tone delay in the orienting tasks but not in the control trials, producing a significant main effect of delay, $F(3, 210) = 7.26$, $p < .01$, $MS_e = 30{,}629.30$, and a significant Delay × Tasks interaction, $F(6, 420) = 3.63$, $p < .01$, $MS_e = 181{,}10.71$.

Experiments 5 and 6

Reaction times and accuracy from no-signal trials appear in Table A1. Reaction times were much longer in Experiment 5 than in Experiment 6, $F(1, 70) = 23.82$, $p < .01$, $MS_e = 120{,}607.76$, possibly because of the requirement to respond to the second word on stop-signal trials. Apart from this, the pattern of results was much the same in the two experiments. *Yes* responses were faster than *no* responses, $F(1, 70) = 58.73$, $p < .01$, $MS_e = 10{,}804.62$, but category judgments were no faster than rhyme judgments, $F(1, 70) < 1$, $MS_e = 27{,}226.62$. The Response Type × Task interaction was marginally significant, $F(1, 70) = 3.79$, $p < .06$, $MS_e = 9{,}004.62$. The category-rhyme effect interacted with experiments, $F(1, 70) = 4.24$, $p < .05$, $MS_e = 27{,}226.62$, indicating that category judgments tended to be faster than rhyme judgments in Experiment 6 and slower than rhyme judgments in Experiment 5.

Table A2
Reaction Times to the Tones Presented During the Orienting Tasks in Experiments 3 and 4 as a Function of Tone Delay, Type of Response in the Orienting Task (Yes/No), and Orienting Task

	Tone delay (in msec)							
	Yes				No			
Orienting task	100	300	500	700	100	300	500	700
				Experiment 3				
Category	653	632	609	642	640	603	609	632
Rhyme	633	588	569	595	621	569	577	627
Control	244	229	235	253	238	233	243	250
				Experiment 4				
Category	1,054	976	982	930	1,042	961	978	932
Rhyme	1,026	883	871	906	923	867	817	872
Control	525	537	524	549	554	514	548	530

(*appendix continued*)

Signal-respond reaction times and accuracy also appear in Table A1. Again, they tend to reflect the same trends as the no-signal data.

In Experiment 5, subjects had to make an overt reponse to the new word pair on stop-signal trials. The mean reaction times (accuracy in parentheses) were 1,879 (.87) and 1,947 (.89) for *yes* and *no* responses in the category task, and 1,905 (.83) and 1,899 (.86) for *yes* and *no* responses in the rhyme task, respectively. These reaction times are somewhat longer than the no-signal reaction times, possibly because of interference from processing the first word pair, but they reflected the same general trends.

Appendix B

No Interactions Between Delay and Conditions

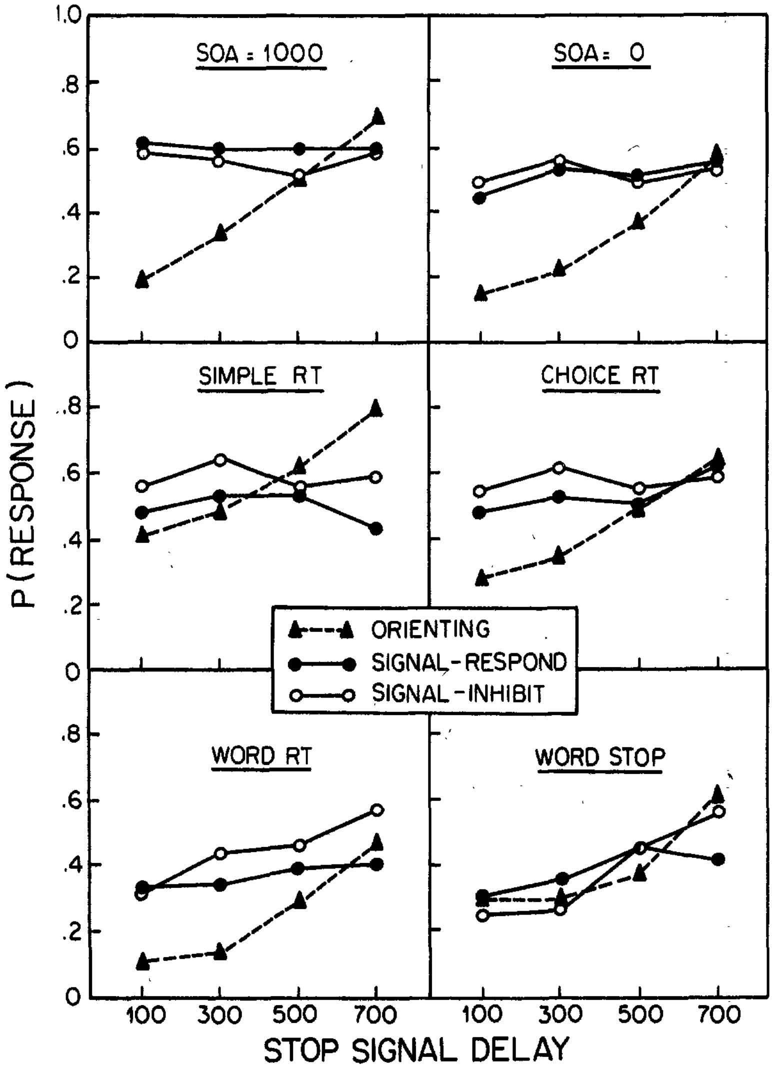

Figure B1. Probability of responding given a stop signal in the orienting tasks (broken lines) and hit rate in the recognition tasks (solid lines) as a function of stop-signal delay in each experiment. (The data are averaged across encoding conditions and are from stop-signal trials only; stimulus-onset asynchrony [SOA] = 1,000 is Experiment 1; SOA = 0 is Experiment 2; simple reaction time [RT] is Experiment 3; choice RT is Experiment 4; word RT is Experiment 5; and word stop is Experiment 6.)

In each experiment, the memory data at each delay were collapsed across signal-inhibit and signal-respond conditions, and the memory data in signal-respond and signal-inhibit conditions were collapsed across delay to increase their stability. Collapsing is appropriate only if there are no interactions between delay and signal-inhibit versus signal-respond conditions. It was not feasible to test the interactions with ANOVA because the extreme delays had too many missing observations (i.e., from subjects who inhibited all the time or responded all the time). Nevertheless, it is possible to get an impression of the interactions by plotting the data and inspecting them visually. Figure B1 displays the mean hit rates for signal-inhibit and signal-respond conditions in each experiment as a function of stop-signal delay. The data are collapsed across tasks (category vs. rhyme) and response type (*yes* vs. *no*). The data in the figure agree well with the previous analyses. Stop-signal delay had a negligible effect in Experiments 1–4 and a relatively strong effect in Experiments 5 and 6. Moreover, the trends seemed the same for signal-inhibit and signal-respond data, except perhaps in Experiments 5 and 6. In those experiments, signal-inhibit performance seemed to improve with delay, whereas signal-respond performance remained relatively constant. On the balance, however, the data in the figure suggest that it may be appropriate to collapse across signal-inhibit and signal-respond conditions to analyze the delay effects and to collapse across delay to compare signal-inhibit and signal-respond conditions.

Appendix C

Processing Time and Memory

There is a possible confound in comparing memory performance from signal-inhibit and signal-respond trials, which must be addressed: Signal-respond data represent trials on which the response to the orienting task was faster than the internal response to the stop signal, whereas signal-inhibit data represent trials on which the response to the orienting task was slower than the internal response to the stop signal (see Logan, 1981, 1982). Assuming that the response to the stop signal has the same latency in the two conditions, signal-respond latencies are faster than signal-inhibit latencies. If mem-

Table C1
Hit Rates (in %) for No-Signal Words Whose Orienting Reaction Times (RT; in msec) Were Faster and Slower Than the Median Orienting Reaction Times

	Category				Rhyme			
	Yes		No		Yes		No	
Experiment	Fast	Slow	Fast	Slow	Fast	Slow	Fast	Slow
1								
Hits	82	81	61	60	58	61	54	50
RT	635	960	723	1,059	637	973	711	1,010
2								
Hits	73	72	61	60	58	55	47	54
RT	814	1,198	925	1,328	846	1,198	910	1,284
3								
Hits	76	75	61	65	61	59	49	52
RT	867	1,277	987	1,413	899	1,276	988	1,401
4								
Hits	72	80	64	70	66	65	53	56
RT	897	1,315	1,023	1,429	926	1,324	994	1,447
5								
Hits	70	79	64	64	58	60	50	55
RT	972	1,489	1,098	1,629	1,036	1,563	1,144	1,667
6								
Hits	75	78	62	68	58	61	50	58
RT	887	1,302	986	1,433	884	1,298	961	1,359

(*appendix continued*)

ory performance depends on the time spent in the orienting task, signal-respond words should be remembered less well than signal-inhibit words.

To assess the seriousness of this confound, the no-signal words were partitioned according to orienting-task reaction time, and hit rates were calculated for words whose orienting reaction times were below the median (fast) and for words whose orienting reaction times were above the median (slow). The mean reaction times and hit rates for fast and slow responses in each condition of each experiment are presented in Table C1. In general, there was very little difference in hit rates for fast and slow responses (also see Craik & Tulving, 1975, Figure 5). The difference was not significant in any experiment.

Received July 12, 1982
Revision received December 8, 1982 ■