Journal of Personality and Social Psychology
1992, Vol. 63, No. 6, 903–912

0022-3514/92/$3.00

The Hyperaccessibility of Suppressed Thoughts

Daniel M. Wegner
University of Virginia

Ralph Erber
DePaul University

The accessibility of suppressed thoughts was compared with the accessibility of thoughts on which Ss were consciously trying to concentrate. In Experiment 1, Ss made associations to word prompts as they tried to suppress thinking about a target word (e.g., *house*) or tried to concentrate on that word. Under the cognitive load imposed by time pressure, they gave the target word in response to target-related prompts (e.g., *home*) more often during suppression than during concentration. In Experiment 2, reaction times for naming colors of words were found to be greater under conditions of cognitive load when Ss were asked to suppress thinking of the word than under conditions of no cognitive load or when Ss were asked to concentrate on the word. The results support the idea that an automatic search for the suppression target increases the accessibility of the target during suppression.

It is difficult to suppress an unwanted thought. We may try to attend to something else and even succeed for a while; but then we are somehow inevitably reminded, and the thought is with us again. During thought suppression, the thought we try to eliminate from attention remains remarkably near the surface and ready to return to consciousness with minimal prompting. This state of suppression-induced *hyperaccessibility* was the focus of our research.

The Phenomena of Suppression

Everyday life provides many occasions on which people choose to suppress thoughts and on which the sheer difficulty of suppression can be observed. The suppression of thoughts may find a place in the management of social impressions (Gilbert, Krull, & Pelham, 1988) and emotions (Zillman, 1988), for example, even though it is often disappointingly ineffective in these domains. During the attempt to tell a lie, in turn, troubling thoughts about the deception, about the truth, and about whether one will be caught seem to come to mind repeatedly despite efforts to stop them (Pennebaker, 1990; Wegner & Erber, 1993). The intent to overcome prejudicial biases in social judgment can, too, instigate attempts at thought control that are executed only with some difficulty (Devine, 1989; Fiske, 1989). It seems that the very attempt to keep unwanted thoughts out of mind makes them all the more insistent.

Research documenting the paradox of thought suppression has called for people to remove a thought from conscious attention in the laboratory. Wegner, Schneider, Carter, and White (1987) asked subjects to try not to think of a white bear as the subjects reported their thoughts aloud. Although subjects regularly voiced a plan to distract themselves ("Okay, I'll think of something else") and made reports as they thought aloud that revealed intervals of successful absorption in other things, they were incapable of sustaining suppression for very long. The typical participant expressed a replacement thought for a short time, often to the end of a sentence, paragraph, or some other pause in the flow, and then he or she abruptly signaled the occurrence of a thought about a white bear. People show similar vexations in their oral and written reports when they suppress thoughts of depressing or cheerful events (Wenzlaff, Wegner, & Roper, 1988). And, without any reporting requirement at all, the difficulty of suppression persists in that psychophysiological disturbances occur during the suppression of exciting thoughts (Wegner, Shortt, Blake, & Page, 1990).

In light of the difficulty of suppression, it makes sense that successful suppression might require considerable attentional resources (see, e.g., Gilbert, Krull, & Malone, 1990; Shallice, 1972; Wegner & Schneider, 1989). Although no amount of attention may be sufficient to stop a thought that is underway (Logan, 1983; Wegner, 1989), attention devoted to a plan of self-distraction can preempt future occurrences of a thought and so can yield periods of comparatively successful suppression. Attention is needed at this time to search for distracting thoughts, to choose from among them the ones that are likely to be most effective (Wenzlaff, Wegner, & Roper, 1988), and to explore the potentially absorbing avenues for distraction that they provide. Attention to such suppression tasks is not required continuously, of course, as there occur episodes of freedom from the unwanted thought and from the need for suppression that result when a distracter is sufficiently absorbing. These periods of successful distraction alternate with episodes of renewed attention to the suppression task, which occur whenever the thought intrudes into consciousness (Wegner et al., 1990).

The experience of thought intrusions during suppression, on the other hand, seems to require little in the way of conscious

This research was supported by National Science Foundation Grants BNS 88-18611 and BNS 90-96263. We thank Todd Bennington, Stephen Blumberg, Miles Cortez, Rob Dillard, Vickie Fendley, Sonny Marks, and Jenn Taylor for assisting in the conduct of the research and John Bargh, Shelley Chaiken, Daniel Gilbert, Traci Giuliano, Stan Klein, Gordon Logan, Toni Wegner, Richard Wenzlaff, and Sophia Zanakos for their helpful comments on an earlier draft.

Correspondence concerning this article should be addressed to Daniel M. Wegner, Department of Psychology, Gilmer Hall, University of Virginia, Charlottesville, Virginia 22903. Electronic mail may be sent to dmw2m@virginia.edu.

attention or control. Although intrusions might be abetted by relevant attention, they seem to occur against our will and do not seem to require attention for their occurrence. Like our own names called out in a crowd, thoughts we have suppressed command our interest whether we are looking for them or not. And despite all our previous attention devoted to suppression, the suppressed thoughts come back to us inexorably, repeatedly, and sometimes intolerably. Thought suppression is remarkable, in short, for its difficulty, its intermittent strong attentional demands, and its repeated interruption by the ballistic return of the unwanted thought. These characteristics invite an analysis of suppression in terms of automatic and controlled thought processes (Bargh, 1984, 1990; Hasher & Zacks, 1979; Logan, 1988; Posner & Snyder, 1975).

Cognitive Processes in Suppression

We believe that the intention to suppress a thought initiates the operation of two cognitive processes, one controlled and the other automatic. First, there is a controlled process that attempts to avoid the unwanted thought by seeking out distracters. This *controlled distracter search* expends cognitive resources toward the goal of keeping a distracter in consciousness, and it is this process that operates when the person implements the plan to think of something else. Second, there is an automatic process that is initiated in the attempt to suppress as well—a process that operates outside of awareness and without conscious guidance even though it is instigated by the intention to suppress. From the inception of suppression, an *automatic target search* looks for the unwanted thought all the while. If this automatic search finds the target, the controlled distracter search can be launched to carry out the job of eliminating the thought. The ironic property of this automatic target search, however, is that it makes the person continually sensitive to the very thought that is unwanted.

The two processes that operate in suppression can be understood as components of a feedback mechanism aimed at the control of thought (cf. Uleman, 1989). Normally, control systems that operate through feedback are comprised of an operating process and a test process that detects the need for the operating process (Carver & Scheier, 1981; Miller, Galanter, & Pribram, 1960). The controlled search for distracters is the mental operating process that carries out the suppression, and this is the part of suppressing that seems effortful and requires conscious attentional guidance. The automatic target search can be understood as a relatively less effortful mechanism enlisted to test whether the operation of the controlled distracter search is needed at any particular moment. Like the "test" and "operate" components of a test-operate-test-exit or *TOTE* unit (Miller et al., 1960), then, these two processes in combination create a purposive system that functions to suppress a thought.

In normal suppression, the cyclic operation of the two processes is responsible for the phenomena of suppression that have been observed. Beginning with the intention to suppress, both processes are initialized. The automatic target search immediately indicates that the target is in consciousness and thus initiates the controlled distracter search. The controlled distracter search brings a series of distracters to mind until one is selected that absorbs attention. At this point, attention is drawn from the controlled distracter search to the absorbing distracter itself, and in this sense the controlled distracter search is no longer functioning. There occurs here a plateau of indeterminate length that may be called *successful* suppression. In the background, however, the automatic target search is still looking continually for signs of the target in consciousness. When an avenue of association to the thought is encountered, however vague or distant, the automatic target search is tuned to register this discovery, and so to return the unwanted thought to consciousness. This reintroduces the cycle from the start, and the controlled distracter search begins again.

It is important to note that the participation of two subprocesses of this kind is logically inherent in any process of inhibition or suppression. After all, in suppression, two things need to be done. One function of a suppression process is suppression per se and the other is remembering what it is that must be suppressed. Although we normally appreciate only the first function, the second function is critically important because without it, any disturbance of the suppression process in general would also disturb the degree to which it remained focused on its original target. An impeded suppression process would thus produce slight suppression of everything, whereas we propose that impeded suppression undermines only the controlled distracter search and leaves the automatic target search unperturbed. The surviving automatic target search, then, enhances the accessibility of the target.

The automatic target search that is instituted to test for failures of the controlled distracter search, therefore, often produces these very failures itself. The automatic target search should act in much the same way as an externally encountered prime to make the unwanted thought highly accessible. It is usually observed in priming paradigms that the nonconscious priming of a thought promotes assimilation of incoming information to that thought (e.g., Lombardi, Higgins, & Bargh, 1987; Martin, 1986). During successful suppression, the automatic target search is indeed nonconscious, and it therefore could render the suppressed thought highly accessible to consciousness. The thought and its close associates should be activated with minimal prompting, and this should yield frequent remindings that in fact return the thought to mind.

This reasoning leads to the conclusion that during suppression, the unwanted thought is likely to be highly accessible to consciousness. The automatic target search introduced by suppression works to make the person unusually sensitive to any topics that might reintroduce the target thought to mind. If the controlled distracter search could be interrupted, it should be possible to observe the automatic target search operating relatively unimpeded to make the unwanted thought accessible in this way. Indeed, this automatic process could render the thought hyperaccessible, that is, even more accessible than it is during intentional concentration on the thought under similar conditions. This prediction was examined here in two experiments using divergent paradigms for the assessment of cognitive accessibility.

Experiment 1: Thought Suppression Under Time Pressure

This study examined the hyperaccessibility prediction in a setting that has been used often to examine the differential operation of automatic and controlled cognitive processes—the time-pressure paradigm (see, e.g., Bargh & Thein, 1985; Strack, Erber, & Wicklund, 1982). Under the assumption that time pressure disturbs the operation of controlled, resource-dependent processes while interfering only little with the operation of automatic processes, the research examined the nature of word associations made by subjects whose controlled processes were either undermined by time pressure or allowed free reign. The subjects were instructed either to think or not to think about a target word, and over several trials their tendency to respond with the target word to closely associated word prompts was observed. We expected that subjects suppressing under time pressure would show hyperaccessibility by responding frequently with the target word when prompted by an associated word.

Method

Overview and design. Subjects were instructed either to think about or not to think about a target word (e.g., *house*). During this concentration or suppression task, they heard prompts that were either related (e.g., *home*) or unrelated (e.g., *young*) to the concentration or suppression target and gave word associates for each prompt. Subjects responded to half of the prompts under high time pressure and half of the prompts under low time pressure.

Subjects. Undergraduates (46 women and 10 men) in introductory psychology classes at Trinity University in San Antonio participated in the study in return for extra class credit. Each was randomly assigned to a suppression condition or a concentration condition.

Procedure. Subjects participated individually in a room equipped with two tape decks and a microphone. The experimenter said the study was to test "how various mental tasks affect the way people form associations between words" and told subjects they would perform several tasks during the experiment. He or she then gave the instructions for the first task. In the *concentration* condition subjects were asked to think of a target for 5 min: *house, child, mountain,* or *car.* In the *suppression* condition subjects were asked not to think about the target for this period. The experimenter then left the room.

On returning after 5 min, the experimenter asked subjects to continue to think or not to think about the target assigned to them and added that subjects were now to do another task concurrently. This task involved giving associates to words. Subjects were to listen to a list of words on a tape and for each word to generate one associate and verbalize it into the microphone for recording on the second tape deck. In some cases (*low time pressure*), subjects had 10 s to respond with an associate. So, after the prompt word was given, there would be a taped countdown from 10 to 1, before the end of which the response was to be delivered. In other cases (*high time pressure*), the response was to be given before 3 s was up and was signaled by a countdown from 3 to 1. The subject did three practice trials in the presence of the experimenter, and the experimenter then left the room and the subject began the task.

Stimulus materials. Subjects listened to a tape that presented them with 16 different prompts for their word-association task. Four of the prompt words were closely related to the word they were trying to concentrate on or suppress. These *target prompts* had been selected from the 6 most frequent associates of the target word according to the norms of Palermo and Jenkins (1964). Four of the prompt words were unrelated to the target; these *nontarget prompts* were selected as comparable associates of a different word. The remaining prompts were fillers unrelated to any of the targets.

The design included four between-subjects replications varying target and nontarget words. Subjects in the first two replications were given the target word *child* or *house.* The target prompts for *house* (*home, door, brick,* and *roof*) constituted the nontarget prompts for *child* and the target prompts for *child* (*mother, little, adult,* and *young*) were the nontarget prompts for *house.* For subjects in the other two replications, the target word was either *mountain* or *car.* The target prompts for *mountain* (*hill, high, top,* and *climb*) were the nontarget prompts for *car* and the target prompts for *car* (*wheel, bus, truck,* and *drive*) were the nontarget prompts for *mountain.* Within each replication subjects had to respond to half the words within 10 s (low time pressure) and half within 3 s (high time pressure). The manipulation of time pressure was counterbalanced between subjects such that the high time pressure words were switched with the low time pressure words half the time.

Results and Discussion

Prompted responses. The primary dependent measure was the number of prompted responses per trial. For *target-prompted* trials, this was the occurrence of the word the subject was trying to suppress or to concentrate on. For *nontarget-prompted* trials, this was the occurrence of the comparison word for which the subject had received no special instruction. So, for example, the measure for target-prompted trials was the number of times per trial subjects responded with *house* to prompts of *home, door, brick,* or *roof* while suppressing or concentrating on *house.* For nontarget-prompted trials, the measure was how often subjects responded with *child* to prompts of *mother, little, adult,* or *young* while suppressing or concentrating on *house.* This latter set of trials provided a baseline frequency for prompted responses in the absence of any instructions to concentrate on or to suppress a target word related to the prompt.

We submitted the number of prompted responses per trial to a three-way analysis of variance (ANOVA) with instruction (concentration vs. suppression) as a between-subjects variable and prompted word (target vs. nontarget) and time pressure (low vs. high) as within-subject variables. (Initial analyses indicated no significant main or interactive effects of the replication variable or time pressure counterbalancing.) The key result of this analysis was an interaction among all three factors, $F(1, 54) = 6.91$, $p < .02$, and significant lower-order interactions embedded in this one were not analyzed further. Analysis of simple effects indicated that the simple interaction of instruction and time pressure was significant for target-prompted trials, $F(1, 54) = 16.94$, $p < .0001$, whereas this interaction was not significant for nontarget-prompted trials, $F(1, 54) < 1$. Means for the significant simple interaction are shown in Figure 1.

The means for the significant simple interaction were compared as simple simple main effects among the conditions in which subjects were prompted with words related to the target. The hyperaccessibility prediction was clearly supported by these comparisons. Specifically, under high time pressure, subjects who were suppressing the target responded to target prompts with the target word more often ($M = 0.23$) than did subjects who were concentrating on the target ($M = 0.09$), $F(1, 54) = 4.09$, $p < .05$. Under low time pressure, in turn, this difference was significantly reversed. Subjects suppressing the target responded to target prompts with the target word less

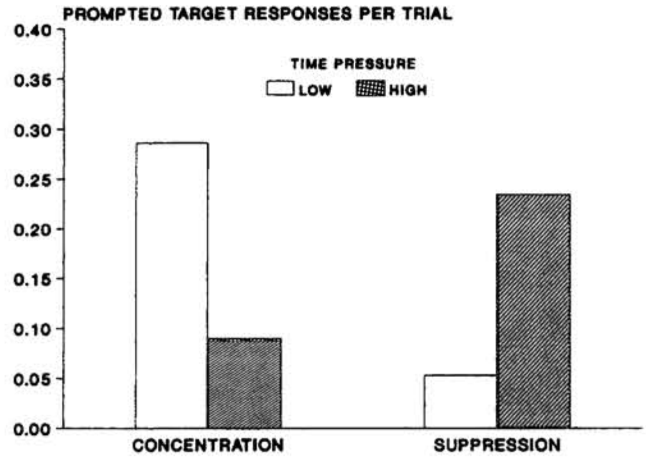

Figure 1. Mean number of prompt-relevant associations per trial: Experiment 1. (Subjects who were suppressing or concentrating on a target thought gave associates to prompts for the target under low or high time pressure.)

often ($M = 0.05$) than did those who were concentrating on the target ($M = 0.29$), $F(1, 54) = 9.29$, $p < .005$.

These four conditions can also be viewed in terms of differences between high and low time pressure within each instruction. Compared in this way, we found that subjects who were suppressing responded to the target prompts with the target more often when time pressure was high ($M = 0.23$) than when time pressure was low ($M = 0.05$), $F(1, 54) = 9.29$, $p < .005$. There was a significant reversal of this for subjects who were concentrating, as they responded to target prompts with the target more often when time pressure was low ($M = 0.29$) than when time pressure was high ($M = 0.09$), $F(1, 54) = 7.68$, $p < .01$.

One other set of comparisons was made to examine whether the means for target prompts departed significantly from the corresponding means for nontarget prompts. Recall that the nonsignificant simple interaction of instruction and time pressure within the nontarget prompt trials revealed a relatively flat pattern of response to nontarget prompts. The mean response to nontarget prompts per trial was .14 for suppressing subjects with high time pressure and .11 for each of the other conditions: suppression with low time pressure, concentration with high time pressure, and concentration with low time pressure. Simple simple main effects computed between target prompts and nontarget prompts indicated that target prompts yielded only marginally more target responses than nontarget responses that were yielded with nontarget prompts under suppression of the target with high time pressure ($M = 0.23$ vs. 0.14), $F(1, 54) = 1.66$, $p < .21$. Concentration with low time pressure produced more target responses to target prompts than nontarget responses to nontarget prompts ($M = 0.29$ vs. 0.11), $F(1, 54) = 7.68$, $p < .01$. No other comparisons in this set reached significance.

In sum, people suppressing a target who were prompted under time pressure to give the target as an associate were more prone to do this than people who were concentrating on the target. They were also more prone to do this than people who were suppressing but giving their response under low time pressure. Although the instruction and time pressure variables had no reliable impact on nontarget responses to nontarget prompts, people suppressing the target under time pressure were only marginally more inclined to give the target as their associate to a target prompt than to give the nontarget as an associate to a nontarget prompt.

Response time check. Response times were clocked by two judges for all trials. Their effective reliability across subjects averaged .98 over the within-subject conditions of the experiment, with a minimum of .96. The mean of the two judges was used as a measure of response time in the ANOVA (conducted on square-root-transformed response times to reduce the heterogeneity of within-cell variance). This analysis revealed a significant main effect for time pressure, $F(1, 54) = 31.30$, $p < .0001$. Subjects responding under high time pressure took significantly less time to respond ($M = 2.37$ s) than did subjects responding under low time pressure ($M = 3.80$ s). This result indicates that the manipulation of time pressure did in fact influence the speed of subjects' responses in the expected direction.

The analysis also indicated a significant interaction of instruction and time pressure, $F(1, 54) = 6.23$, $p < .02$. Subjects concentrating under low time pressure ($M = 4.64$ s) responded more slowly to target and nontarget prompts than did those concentrating under high time pressure ($M = 2.52$ s), those suppressing under low time pressure ($M = 2.95$ s), or those suppressing under high time pressure ($M = 2.20$ s). This effect suggests that concentration without time pressure may have imposed less of a sense of urgency on subjects than other tasks, but it does not qualify any of the primary results.

Finally, it is worth noting that correlations were computed separately between time to respond and response with the target word for each response trial within each condition of the Instruction × Time Pressure × Prompt design. None of these correlations was statistically significant, suggesting that differential response time within condition did not contribute to the observed variation in target responses.

Summary. The instruction to suppress a thought appears to render that thought hyperaccessible to consciousness. The evidence for this consists of word associations that subjects offered during a period of instructed suppression. When subjects were prompted by a word that was a close associate of the target they were suppressing and when they were under pressure to respond quickly, they often found themselves suffering the irony of blurting out exactly the suppressed thought as their word association. This was not as likely to happen when they were concentrating on the thought.

Experiment 2: Thought Suppression and Stroop Interference

Although the time pressure results offer important evidence in favor of the hyperaccessibility hypothesis, they are somewhat incomplete. The data are not entirely conclusive, for example, regarding the greater accessibility of suppressed thoughts in response to target prompts as opposed to the accessibility of nontarget thoughts in response to nontarget prompts. It can be argued more broadly that the time pressure paradigm may not capture response times as brief as those normally understood as indicative of automaticity (Logan, 1988). In Experiment 1, for instance, we observed that subjects with suppression instructions who were responding to target prompts under time pressure were more likely to respond with the target than those who were responding to target prompts without time pressure. This could be explained by the leisurely pace of the low-time-pressure interval, during which subjects who were suppressing thoughts may have made strategic replacements of their initial response. Such an explanation does not involve automatic accessibility at all, but turns instead on the activities individuals perform when they have plenty of response time.

The finding that suppression subjects responded with the target to target prompts under high time pressure more often than did concentration subjects is inconsistent with this weak conclusion, of course, as it contrasts suppression and concentration during the brief intervals for response under high time pressure. With very little time to respond, suppression still yields more frequent target response than does concentration. Nonetheless, it is important to gauge the degree to which suppression can augment accessibility during intervals so brief as to preclude any hint of strategic response replacement on the part of subjects. For this purpose, we chose to examine the influence of thought suppression on responding in a version of the Stroop (1935) color–word interference paradigm.

Method

Overview and design. Subjects were instructed either to think or not to think about one of the four targets used in the first study. During this task, they performed a two-color Stroop reaction time (RT) task. They pressed keys to indicate whether each of a series of words appearing on a computer monitor was printed in red or blue. The words included the target word, nontarget words, and target-related words. Half the subjects performed this task while simultaneously rehearsing a 9-digit number (high-cognitive-load condition), and the remainder did it while rehearsing a 1-digit number (low-cognitive-load condition).

Subjects. Undergraduates (27 men and 19 women) in introductory psychology classes at the University of Virginia participated in the experiment for course credit. Before their participation, subjects were screened for colorblindness and difficulty in reading words on a computer screen. Other subjects who were tested but whose data were not used in the analyses included 4 with high error rates (over 5%) and 5 with high mean RTs (over 750 ms). Subjects were randomly assigned to conditions, and the excluded subjects were evenly distributed among conditions.

Procedure. On arrival at the laboratory, each subject was seated at a computer and told that the experiment was concerned with how various mental tasks affect the way people name colors. The following instructions were shown on the computer screen:

> You will shortly see a series of words on the screen that are printed either in red or blue. If the word is in red, press the key labeled (R) at the righthand pad of the keyboard with your middle finger. If the word is in blue, press the key labeled (B) with your index finger. Your task is to indicate the color of each word as quickly and as accurately as you can. Try not to make mistakes, but try to be fast. Your reactions will be timed. An asterisk (*) will appear on the screen before each word appears to show you where to look.

Subjects then did 12 practice trials and completed another 12 trials if they made any mistakes. On completion of the practice the experimenter asked subjects to sit back from the computer and instructed them either to think or not to think about a house, a car, a mountain, or a child. The experimenter left the room and returned 5 min later to say the following:

> In addition to studying people's ability to name colors, we are also interested in finding out how they do when they perform another task at the same time. Your job will be to respond with the word colors and at the same time remember the following number. I will ask you to recall the number later on. You will have 25 seconds to commit the number to memory. I will ask you for the number at the end of the experiment. It is very important that you give me the number back exactly as it is on this paper. If you fail to remember the number, you will have to be disqualified from this study.

The experimenter then showed the subject a card containing either a 9-digit number or a 1-digit number and gave the subject 25 s to begin rehearsing the number (cf. Gilbert & Osborne, 1989).

Following this manipulation of cognitive load, the experimenter reminded subjects to continue to think or not to think about the object assigned to them, to continue to rehearse the number, and to begin the color–word task on the computer as during the practice trials. The experimenter then started the Stroop task program and left the room. When the task was completed, the experimenter returned and asked the subject to report the rehearsed number. All subjects recalled the number correctly.

Stimuli. The words were presented in the standard 24- × 80-character screen font on a 14" IBM-compatible EGA color monitor. Timing programs by Creeger, Miller, & Paredes (1990) were used to calibrate an IBM-AT-compatible computer for RT measurement in Microsoft Quickbasic 4.5. The program provided for each trial on a white screen to start with a 3-s pause followed by the 2-s presentation of a centered black asterisk. Then, the word appeared in red or blue at the asterisk focus and remained on-screen until a response was recorded.

The set of 64 stimulus words was the same for all subjects, although half the subjects were presented the words in reverse order. The word set contained a random ordering of 8 presentations of each of the 4 possible target words (*house, child, mountain,* and *car*) and 2 presentations of each of the 4 target-related words for each of the 4 possible targets (i.e., the 16 words selected as close associates of the targets in Experiment 1). Color was balanced such that half the presentations for each word were in each color. This set of stimuli allowed for the investigation of color-naming RTs for target words (e.g., *house* when the target is *house*), for target-unrelated words (e.g., *house* when the target is *child*), and for target-related words (e.g., *home* when the target is *house*) within each of the 4 target-word replications.

Results and Discussion

Mean RTs for correct responses to target-unrelated, target-related, and target words are shown in Figure 2. These were examined in an ANOVA with cognitive load (low vs. high), instruction (concentration vs. suppression), specific target word (*house, child, mountain,* and *car*), and stimulus word order (two orders) as between-subjects variables and word type (target, target-related, or target-unrelated word) as a repeated measure. A significant main effect was observed for cognitive load, $F(1, 20) = 6.28$, $p < .03$, indicating that subjects responded with the colors of all words more slowly under conditions of high load (M = 549 ms) than low load (M = 485 ms). A significant main effect was also found for instruction, $F(1, 20) = 5.53$, $p < .03$, in that subjects were slower to indicate colors of all words when

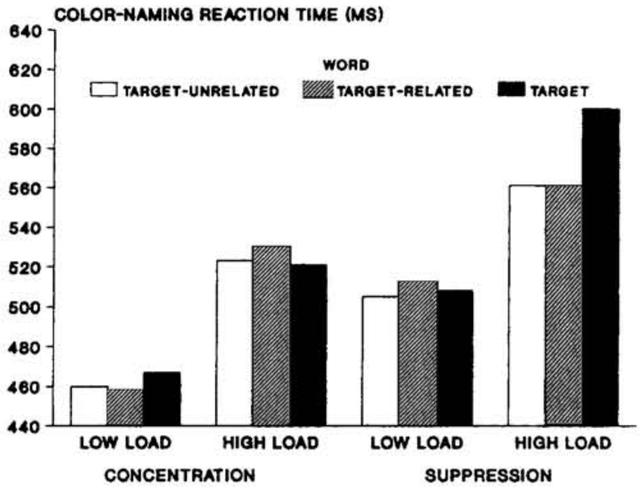

Figure 2. Mean color-naming reaction times: Experiment 2. (Subjects who were suppressing or concentrating on a target thought made keypress responses to name the colors of target-unrelated, target-related, and target words under low or high cognitive load.)

they were suppressing a target (M = 541 ms) than when they were concentrating on it (M = 493 ms).

The interaction of load, instruction, and word type was marginally reliable, $F(2, 40) = 2.84$, $p < .07$, and its form was as predicted. As can be seen in Figure 2, target-unrelated and target-related words showed very similar patterns across the load and instruction conditions, corresponding generally to the main effects mentioned above. Reaction times for the target word itself, however, departed from this pattern with a noteworthy increment in the high-load–suppression condition. Analyses of the simple simple main effects of load on RT for the target word within the suppression condition showed that the high-load, color-naming RT for target words (M = 600 ms) was significantly greater than the low-load mean (M = 508 ms), $F(1, 20) = 5.02$, $p < .04$. The simple simple main effect of instruction on RT for target words within the high-load condition showed a marginally significant effect, $F(1, 20) = 3.36$, $p < .09$, as the color-naming RT for suppression (M = 600 ms) tended to exceed that for concentration (M = 521 ms). No other simple simple main effects were significant.

Another way to view the overall interaction is in terms of comparisons between RTs for target words and target-unrelated words. Contrasts between target words and target-unrelated words within each of the Load × Instruction conditions revealed only one reliable difference. For subjects suppressing the target under high load, RT was significantly greater for target words (M = 600 ms) than for target-unrelated words (M = 561 ms), $F(1, 20) = 9.00$, $p < .01$. The only clear Stroop interference effect observed in this experiment, then, occurred for subjects who were suppressing a target under cognitive load and who were asked to name the color of the target.

The absence of inflated latencies for target words in the other conditions may have occurred for several reasons. Perhaps our implementation of the Stroop paradigm is capable of detecting only strong manipulations of accessibility. The two-color version of the Stroop situation used here may reduce interference in general because there are fewer response options, for example, or the two-finger keypress response (as compared with the more common voice-response method) may introduce sufficient response variance as to obscure interference effects. Given these obstacles, it may have been only under conditions promoting hyperaccessibility—the instruction to suppress under cognitive load—that this paradigm afforded the observation of interference effects. Alternatively, it may be that the absence of interference findings in the concentration conditions echoes research in which target primes that were presented for subjects' conscious perusal reduced rather than increased target accessibility (e.g., Lombardi, Higgins, & Bargh, 1987; Martin, 1986).

Another comment on these results is in order given the absence of clear Stroop effects for target-related words paralleling those for target words. Apparently, the activation of the target that occurs during suppression under high cognitive load does not extend significantly to the target's semantic associates. So, at least as measured in this way, activation due to suppression does not spread to the immediate associates of the suppression target.

This study showed, in sum, that suppression increases the accessibility of an unwanted thought in that color-naming responses for that thought were slowed during an imposed cognitive load. When the target appeared as the colored stimulus, the instruction to suppress the target delayed color-naming responses during high cognitive load relative to low cognitive load. This slowing during high load and suppression was greater for presentations of the target word than for presentations of target-unrelated words. Color-naming of targets also tended to slow during suppression under high load as compared with concentration on the target under high load. This latter finding promotes the inference that suppression creates a state of hyperaccessibility of the target, a state in which the target is even more accessible than it is during the intentional attempt to think about it.

General Discussion

The intention to suppress a thought creates the opposite of what is wanted. This is particularly true when cognitive resources are assigned elsewhere even while the attempted suppression is ongoing. The results of these experiments indicate that a suppressed thought becomes more rather than less accessible to consciousness when time pressures or cognitive loads are imposed during suppression. The effect of this is that people may say the very word they are trying not to think (Experiment 1) or may attend to the very aspect of a stimulus they are trying not to think about (Experiment 2). We believe these hyperaccessibility effects occur as a result of the operation of an automatic search for the suppressed thought that is initiated when a thought is suppressed.

Automatic Target Search

If thought suppression triggers an automatic search for the suppression target, several consequences would be expected. In general, of course, there should be a tendency for the target to regain access to consciousness. Presumably, however, the automatic target search is normally balanced by the effortful activity of a controlled search for distracters, and as a result the unhampered enterprise of suppression can often yield apparent suppression success. This was shown in our studies when suppression was prompted with little or no cognitive load. Under these conditions, subjects in Experiment 1 managed to keep from giving the suppressed word as their associative response and subjects in Experiment 2 showed no greater latency for naming the color of a suppressed word.

The more pointed prediction offered by our formulation is that during suppression, the access of the suppressed thought to consciousness should be enhanced by the imposition of cognitive load. Such a load undermines the continued operation of the controlled distracter search and so unleashes the automatic target search. This prediction was substantiated in both studies, as higher relative levels of accessibility were observed for suppression with load than without load in both cases. Such findings provide substantial support for our conceptualization of the suppression process in terms of controlled distracter search and automatic target search.

The levels of accessibility produced by suppression with load at times even surpassed the accessibility of thoughts on which individuals were consciously trying to concentrate their attention. In Experiment 1, suppression under high time pressure led to more target associations than did concentration under high time pressure; in Experiment 2, suppression under high memory load led to a tendency for greater color-naming latency for the target word than did concentration with high load. It is in this sense that suppression can be most clearly identified as the source of hyperaccessibility. Normally, a thought that is accessible is defined as one that can come readily to mind, and that can thus participate more than other thoughts in judgments, reports, categorizations, or other mental processes (e.g., Bruner, 1957; Higgins & King, 1981). To the degree that intentional concentration marshals as much accessibility as can be engineered by conscious means, the finding that suppression under cognitive load yields yet greater levels of accessibility constitutes evidence for a higher degree of access—hyperaccess.

Eventually, the understanding of intentional concentration might also be approached by means of an analysis of its controlled and automatic components. By extension of the present analysis of suppression, concentration could be viewed as composed of a controlled search for the target and an automatic search for distracters. The controlled target search operates to keep the target in consciousness, whereas the automatic distracter search tests to determine whether the controlled target search is needed. If concentration takes place in this way, specific predictions beyond those tested in the present studies could be made about how cognitive loads influence concentration. For example, load during concentration should undermine a controlled target search and so uncover the automatic search for distracters. This suggests why load during concentration does not produce noteworthy levels of accessibility of the thought. Indeed, the hyperaccessibility observed in these studies for thoughts suppressed during a load may be a function in part of the reduced accessibility that accrues in the comparison conditions of concentration under load. The full exploration of the automatic and controlled processes underlying both concentration and suppression awaits future inquiry in which these instructional sets can be compared with conditions in which subjects become aware of a target thought without any mental control instructions.

Properties of Automaticity

Another focus for future research should be the evaluation of the specific aspects of automaticity that are present in the search process that is initiated in suppression. We have used the term *automatic* quite globally to characterize the search process at this stage of our theorizing, but it has been noted by Bargh (1989) that automaticity as a property of thought processes is comprised of several subproperties, not all of which necessarily covary. Bargh (1989) suggested that a thought process has been viewed as automatic variously because it (a) occurs without awareness, (b) requires relatively little effort, (c) has effects that occur unintentionally, (d) runs autonomously or without conscious monitoring, or (e) has effects that are involuntary or uncontrollable (see also Jonides, Naveh-Benjamin, & Palmer, 1985; Logan, 1988).

To what extent does the "automatic" search process underlying suppression have any of these properties? It should be clear from the outset that this process must be capable of occurring *outside awareness* if suppression is to have any success at all. This is evident because people do reach periods of successful suppression in which they do not volunteer reports of the presence of the suppressed thought even when they have been asked to report it (e.g., Wegner et al., 1987). Thus, even though the search could be quite conscious at its onset, it must leave awareness at some point, by definition, in order to be successful.

The search process that yields suppressed thoughts also appears to require relatively little effort. This process operated in our experiments most clearly under conditions of increased cognitive load. With time pressure and with memory load, we found that the suppressed thought was more inclined to become accessible. These findings suggest that the load operates to undermine the controlled distracter search that is antagonistic to the automatic target search, while diminishing the automatic target search little if at all.

The unintentional effects of the target search during suppression are not quite so clear. Although the effects of this search often appear to be precisely counterintentional, it is nonetheless true that the target search is initiated (along with the distracter search) by an intention, in that it occurs as the result of the intent to engage in suppression. The effect of the automatic target search is consistent in the direction of its effects with the intention only when the controlled distracter search is unimpeded and the two operate in concert. Otherwise, the target search is inconsistent with the intention. Perhaps the best way to understand the intentionality of the target search is to recognize the search as a "goal-dependent" automatic process (cf. Bargh, 1989).

The target search appears to run autonomously as well. When subjects are asked to suppress a thought as they think aloud in an experiment—and report the thought whenever it comes back to mind—they experience phases in which the thought is absent from their think-aloud protocols and is sufficiently out of awareness that they fail to report it as well. Still, then, the unwanted thought returns, apparently without any conscious intention to make it return. By definition, after all, having the conscious intention to make it return would mean that it was not suppressed at the time. Any intrusive return of a thought during suppression could be the result of the autonomous target search operating in the background.

Finally, the involuntary nature of the target search is worth considering. In one sense, the target search must be considered controllable and voluntary. It does occur as the result of a voluntary plan to suppress and might well be controlled by the person's decision to rescind the attempt to suppress. However, the target search as it occurs during an ongoing suppression attempt must be viewed as involuntary. The intrusive return of an unwanted thought is seen by people who suffer from it as entirely out of control. In our model, this ballistic return is due to the background operation of the target search during suppression, and in this sense the target search appears to have involuntary and uncontrollable effects.

Overall, then, it appears that the target search instigated in suppression may have several of the various properties that have been viewed as reasons for calling a process *automatic.* It is perhaps most clear that the search operates autonomously, outside awareness, and with little effort. Its involuntary and unintentional properties are less clear, in that whereas its inception as part of suppression is voluntary and intended, its subsequent effects may be experienced as involuntary and counterintentional.

Suppression and Intrusive Thoughts

There are a variety of clinical and everyday problems in which people exhibit serious difficulties with intrusive thoughts. In depression, obsessive-compulsive disorder, phobia, posttraumatic stress disorder, and during dieting or treatment for addiction, for example, unwanted thoughts appear to recur and dominate the sufferer's mental landscape (cf. Wegner, 1989, 1992). Research on several of these problems has suggested that automaticity of particular classes of thought may be involved in the disorder. Depressive people exhibit automaticity in depressive thinking (Bargh & Tota, 1988; Gotlib & McCann, 1984), for example, and phobics show evidence of this in their phobic thinking (Watts, McKenna, Sharrock, & Trezise, 1986).

To date, automaticity of this kind has been understood by researchers primarily to be a consequence of learning. After all, automaticity is usually seen as a format that cognitive processes assume when they have been practiced repeatedly. Increments in automaticity accrue when a particular kind of social judgment is performed again and again (Smith & Lerner, 1986), for instance, and the automatic accessibility of attitudes follows when people repeat the attitude over and over (Houston & Fazio, 1989). Attributing the development of automaticity to learning suggests that the people who exhibit automaticity in depressive, phobic, or other forms of pathological thinking have reached this pass by virtue of some sort of repetitive or habitual process. They have somehow been pressed to think these things too often, and so now have reached chronic levels of automatic activation of the thoughts (Higgins & King, 1981). This line of reasoning indicates that there must be considerable practice of unwanted thoughts occurring prior to the development of their automatic activation.

The hyperaccessibility of suppressed thoughts suggests that such repetition may not be so necessary. Instead, chronic levels of automatic activation may arise as a result of individuals' attempts to control their thinking through suppression. "Instant" automatic processes might develop without practice at all. The depressed person who is trying hard to suppress negative thoughts, for example, could unwittingly initiate and then suffer from an automatic target search that makes negative thoughts hyperaccessible to consciousness. Perhaps the cognitive stress that accrues from depression itself could exacerbate this process by serving as a continuous load (cf. Sullivan & Conway, 1989). More generally, the stresses that introduce cognitive load at many points in life may have the result of turning our struggle against unwanted, seemingly automatic mental states into an invitation for these states to overwhelm us.

The hyperaccessibility observed here could emerge whenever any psychological equivalent of our load manipulations impinges on the person who is actively suppressing a thought. Stress, distraction, arousal, busyness, and even the time pressures of normal living could introduce intrusions and perhaps even obsession or preoccupation with thoughts that were once suppressed. This reasoning suggests why we often find that the very thing we do not want to say or think comes forward when we are distracted or distressed. The phenomenon of Freudian slips that are precisely the least appropriate thing to say in a given situation might be explained in this way (Baars, 1985). Cognitive busyness or time pressure could interfere with processes of self-presentation, deception, or self-control that depend on thought suppression for their success, and so promote social blunders, unintentional disclosures of deceit, or self-control lapses that are not entirely random. Rather, because the most unwelcome thoughts are typically chosen as targets for suppression, these very thoughts are the ones that are exposed when the controlled processes of suppression are disrupted. Hyperaccessibility of these unwanted thoughts can then come forward to produce chaos.

This suggestion is also consistent with the observation that stress exacerbates problems such as anxiety and phobia (Jacobs & Nadel, 1985). The individual who uses thought suppression as a strategy for the control of anxiety, for example, may effectively put upsetting thoughts out of mind as long as he or she has sufficient cognitive capacity to carry on the controlled search for distracters. This person might become absorbed in a variety of pursuits that aid in the enterprise of suppression and so lead a fairly peaceful life. The occurrence of some major stressor—loss of a job, or of a loved one, for instance—could interrupt this controlled distracter search with some disastrous psychological results. Beyond the negative emotions unleashed by the stressor, the person would now be subject to the hyperaccessibility of the suppressed anxiety-producing thought. A double dividend of negative emotion could be the consequence, because the originally suppressed thought could instigate its own emotional reaction.

Conclusions

The hyperaccessibility of suppressed thoughts has some useful implications, both theoretical and applied. For theory, these results suggest a useful way of understanding the potentially counterintentional automatic processes that may be entailed by the intentional control of thinking. The observation of hyperaccessibility fulfills a key prediction made by our model of automatic and controlled processes in suppression, and it suggests that there may be some value to further research on how automatic processes may undermine the controlled ones that entail them. As for application, the results indicate in a preliminary way how people suffer from the choice to suppress thoughts. Although people may be quite able to distract themselves from unwanted thoughts under normal conditions, they can suffer from hyperaccessibility and the intrusive return of those thoughts when something happens to distract them from their distractions. The thoughts people are most desperate to suppress can come to mind at precisely the point of highest stress and produce the very thought and behavior that the suppression was enlisted to avoid.

References

Baars, B. J. (1985). Can involuntary slips reveal one's state of mind? With an addendum on the problem of conscious control of action. In T. M. Shlechter & M. P. Toglia (Eds.), *New directions in cognitive science* (pp. 242–261). Norwood, NJ: Ablex.

Bargh, J. A. (1984). Automatic and conscious processing of social information. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 3, pp. 1–43). Hillsdale, NJ: Erlbaum.

Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 3–51). New York: Guilford Press.

Bargh, J. A. (1990). Auto-motives: Preconscious determinants of social interaction. In E. T. Higgins & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition* (Vol. 2, pp. 93–130). New York: Guilford Press.

Bargh, J. A., & Thein, R. D. (1985). Individual construct accessibility, person memory, and the recall-judgment link: The case of information overload. *Journal of Personality and Social Psychology, 49,* 1129–1146.

Bargh, J. A., & Tota, M. E. (1988). Context-dependent automatic processing in depression: Accessibility of automatic constructs with regard to self but not others. *Journal of Personality and Social Psychology, 54,* 925–939.

Bruner, J. S. (1957). On perceptual readiness. *Psychological Review, 64,* 123–152.

Carver, C. S., & Scheier, M. F. (1981). *Attention and self-regulation: A control-theory approach to human behavior.* New York: Springer-Verlag.

Creeger, C. P., Miller, K. F., & Paredes, D. R. (1990). Micromanaging time: Measuring and controlling timing errors in computer-controlled experiments. *Behavior Research Methods, Instruments, and Computers, 22,* 34–79.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56,* 5–18.

Fiske, S. T. (1989). Examining the role of intent, toward understanding its role in stereotyping and prejudice. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 253–283). New York: Guilford Press.

Gilbert, D. T., Krull, D. S., & Malone, P. S. (1990). Unbelieving the unbelievable: Some problems in the rejection of false information. *Journal of Personality and Social Psychology, 59,* 601–613.

Gilbert, D. T., Krull, D. S., & Pelham, B. W. (1988). Of thoughts unspoken: Social inference and the self-regulation of behavior. *Journal of Personality and Social Psychology, 55,* 685–694.

Gilbert, D. T., & Osborne, R. E. (1989). Thinking backward: Some curable and incurable consequences of cognitive busyness. *Journal of Personality and Social Psychology, 57,* 940–949.

Gotlib, I. H., & McCann, C. D. (1984). Construct accessibility and depression: An examination of cognitive and affective factors. *Journal of Personality and Social Psychology, 47,* 427–439.

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

Higgins, E. T., & King, G. (1981). Accessibility of social constructs: Information-processing consequences of individual and contextual variability. In N. Cantor & J. F. Kihlstrom (Eds.), *Personality, cognition, and social interaction* (pp. 69–121). Hillsdale, NJ: Erlbaum.

Houston, D. A., & Fazio, R. H. (1989). Biased processing as a function of attitude accessibility: Making objective judgments subjectively. *Social Cognition, 7,* 51–66.

Jacobs, W. J., & Nadel, L. (1985). Stress-induced recovery of fears and phobias. *Psychological Review, 92,* 512–531.

Jonides, J., Naveh-Benjamin, M., & Palmer, J. (1985). Assessing automaticity. *Acta Psychologica, 60,* 157–171.

Logan, G. D. (1983). On the ability to inhibit simple thoughts and actions: I. Stop-signal studies of decision and memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 585–606.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review, 95,* 492–527.

Lombardi, W. J., Higgins, E. T., & Bargh, J. A. (1987). The role of consciousness in priming effects on categorization: Assimilation versus contrast as a function of awareness of the priming task. *Personality and Social Psychology Bulletin, 13,* 411–429.

Martin, L. L. (1986). Set/reset: Use and disuse of concepts in impression formation. *Journal of Personality and Social Psychology, 51,* 493–504.

Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior.* New York: Holt.

Palermo, D. S., & Jenkins, J. J. (1964). *Word association norms.* Minneapolis: University of Minnesota Press.

Pennebaker, J. W. (1990). *Opening up.* New York: Morrow.

Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola Symposium* (pp. 55–88). Hillsdale, NJ: Erlbaum.

Shallice, T. (1972). Dual functions of consciousness. *Psychological Review, 79,* 383–393.

Smith, E. R., & Lerner, M. (1986). Development of automatism of social judgments. *Journal of Personality and Social Psychology, 50,* 246–259.

Strack, F., Erber, R., & Wicklund, R. A. (1982). Effects of salience and time pressure on ratings of social causality. *Journal of Experimental Social Psychology, 18,* 581–594.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18,* 643–662.

Sullivan, M. J. L., & Conway, M. (1989). Negative affect leads to low-effort cognition: Attributional processing for observed social behavior. *Social Cognition, 7,* 315–337.

Uleman, J. S. (1989). A framework for thinking intentionally about unintended thoughts. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 425–449). New York: Guilford Press.

Watts, F. N., McKenna, F. P., Sharrock, R., & Trezise, L. (1986). Colour-naming of phobia-related words. *British Journal of Psychology, 77,* 97–108.

Wegner, D. M. (1989). *White bears and other unwanted thoughts.* New York: Viking/Penguin.

Wegner, D. M. (1992). You can't always think what you want: Problems in the suppression of unwanted thoughts. In M. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 25, pp. 193–225). San Diego, CA: Academic Press.

Wegner, D. M., & Erber, R. (1993). Social foundations of mental control. In D. M. Wegner & J. W. Pennebaker (Eds.), *Handbook of mental control.* (pp. 34–56). Englewood Cliffs, NJ: Prentice-Hall.

Wegner, D. M., & Schneider, D. J. (1989). Mental control: The war of the ghosts in the machine. In J. Uleman & J. Bargh (Eds.), *Unintended thought* (pp. 287–305). New York: Guilford Press.

Wegner, D. M., Schneider, D. J., Carter, S., III, & White, L. (1987). Paradoxical effects of thought suppression. *Journal of Personality and Social Psychology, 53,* 5–13.

Wegner, D. M., Shortt, J. W., Blake, A. W., & Page, M. S. (1990). The suppression of exciting thoughts. *Journal of Personality and Social Psychology, 58,* 409–418.

Wenzlaff, R., Wegner, D. M., & Roper, D. (1988). Depression and mental control: The resurgence of unwanted negative thoughts. *Journal of Personality and Social Psychology, 55,* 882–892.

Zillman, D. (1988). Mood management: Using entertainment to full advantage. In L. Donohew, H. E. Sypher, & E. T. Higgins (Eds.), *Communication, social cognition, and affect* (pp. 147–171). Hillsdale, NJ: Erlbaum.

Received August 6, 1991
Revision received April 16, 1992
Accepted May 12, 1992 ■