Journal of Experimental Psychology: General
2008, Vol. 137, No. 4, 649–672

Copyright 2008 by the American Psychological Association
0096-3445/08/$12.00 DOI: 10.1037/a0013170

Automatic and Controlled Response Inhibition: Associative Learning in the Go/No-Go and Stop-Signal Paradigms

Frederick Verbruggen
Vanderbilt University and Ghent University

Gordon D. Logan
Vanderbilt University

In 5 experiments, the authors examined the development of automatic response inhibition in the go/no-go paradigm and a modified version of the stop-signal paradigm. They hypothesized that automatic response inhibition may develop over practice when stimuli are consistently associated with stopping. All 5 experiments consisted of a training phase and a test phase in which the stimulus mapping was reversed for a subset of the stimuli. Consistent with the automatic-inhibition hypothesis, the authors found that responding in the test phase was slowed when the stimulus had been consistently associated with stopping in the training phase. In addition, they found that response inhibition benefited from consistent stimulus–stop associations. These findings suggest that response inhibition may rely on the retrieval of stimulus–stop associations after practice with consistent stimulus–stop mappings. Stimulus–stop mapping is typically consistent in the go/no-go paradigm, so automatic inhibition is likely to occur. However, stimulus–stop mapping is typically inconsistent in the stop-signal paradigm, so automatic inhibition is unlikely to occur. Thus, the results suggest that the two paradigms are not equivalent because they allow different kinds of response inhibition.

Keywords: cognitive control, automaticity, response inhibition, stop-signal paradigm, go/no-go paradigm

Inhibition plays a central role in theorizing about human cognition. Inhibition refers to the suppression of thoughts, actions, and emotions and is often regarded as a key component of executive control (e.g., Aron, Robbins, & Poldrack, 2004; Baddeley, 1996; Logan, 1985; Miyake et al., 2000; Nigg, 2000; Ridderinkhof, van den Wildenberg, Segalowitz, & Carter, 2004; Sylvester et al., 2003). Researchers have used the concept of inhibition to explain a broad range of phenomena in clinical psychology, cognitive psychology, neuropsychology, lifespan development, and individual differences. The role of inhibitory processes in several paradigms is still debated (see e.g., MacLeod, Dodd, Sheard, Wilson, & Bibi, 2003), but most researchers have agreed that some kind of inhibition is involved in deliberately stopping a prepared motor response (e.g., Andres, 2003; Aron et al., 2004; Boucher, Palmeri, Logan, & Schall, 2007; Logan, 1994; Logan & Cowan, 1984; MacLeod et al., 2003; Miyake et al., 2000; Nigg, 2000; Ridderinkhof et al., 2004; Rubia et al., 2001; Stuphorn & Schall, 2006; Verbruggen, Liefooghe, & Vandierendonck, 2004). Response inhibition supports flexible behavior in a constantly changing environment: When actions are no longer relevant, they can be stopped and possibly replaced by other actions when needed.

Two paradigms that are frequently used to study response inhibition are the go/no-go paradigm (Donders, 1868/1969) and the stop-signal paradigm (Lappin & Eriksen, 1966; Logan & Cowan, 1984; Vince, 1948). In both paradigms, subjects are instructed to withhold a response when a no-go or stop signal is presented. In the literature, the two paradigms have often been treated as equivalent (e.g., Aron et al., 2004; Nigg, 2000), and researchers have presumed that the same inhibitory processes are responsible for cancelling a response. However, in the present study, we show that this presumption is false and that response inhibition can be achieved through either automatic or controlled processing, depending on the consistency of associations between stimuli and stopping.

The Go/No-Go Paradigm and the Stop-Signal Paradigm

In the go/no-go paradigm, subjects are presented with a series of stimuli and are told to respond when a go stimulus is presented and to withhold their response when a no-go stimulus is presented (e.g., press the response key for the letter *K* but do not press the response key for the letter *L*). The mapping of stimuli onto go and no-go responses is explained at the beginning of the experiment and typically remains the same throughout the experiment. In this paradigm, the index of inhibitory control is the probability of executing a response on a no-go trial [$p(respond|no\text{-}go)$]. In the stop-signal paradigm, subjects usually perform a choice reaction task on *no-stop-signal trials* (e.g., press the left response key for the letter *K* and press the right response key for the letter *L*). On a random selection of the trials (*stop-signal trials*), a stop signal is presented after a variable delay (*stop-signal delay;* SSD), which instructs subjects to withhold the response to the go stimulus on those trials. The first index of

Frederick Verbruggen, Department of Psychology, Vanderbilt University, and Department of Experimental Psychology, Ghent University; Gordon D. Logan, Department of Psychology, Vanderbilt University.

Frederick Verbruggen is a Postdoctoral Fellow of the Research Foundation—Flanders (FWO-Vlaanderen). This research was also supported by Grant BCS 0446806 from the National Science Foundation, Grant FA9550-07-1-0192 from the Air Force Office of Scientific Research, and Grant R01-MH073879-01 from the National Institute of Mental Health to Gordon D. Logan. We thank Julie Delheimer for help with Experiments 1–3 and Darryl W. Schneider, Stephen Monsell, David Huber, and two anonymous reviewers for their helpful comments on a draft of this article.

Correspondence concerning this article should be addressed to Frederick Verbruggen, Department of Psychology, Vanderbilt University, Nashville, TN 37203. E-mail: frederick.verbruggen@ugent.be

inhibitory control is the probability of responding on stop-signal trials [$p(respond|signal)$], which is often evaluated as a function of SSD. The second index of inhibitory control is an estimate of the covert latency of the stop process, *stop-signal reaction time* (SSRT).

The go/no-go and stop-signal paradigms are popular tools for investigating response inhibition in basic research in cognitive science and cognitive neuroscience. Both paradigms have been used to study response-inhibition deficits in clinical populations such as children and adults with attention-deficit/hyperactivity disorder (Bekker et al., 2005; Iaboni, Douglas, & Baker, 1995; McLean et al., 2004; Schachar & Logan, 1990), patients with schizophrenia (Badcock, Michie, Johnson, & Combrinck, 2002; Kiehl, Smith, Hare, & Liddle, 2000), Parkinson's disease (Gauggel, Rieger, & Feghoff, 2004; Hershey et al., 2004), and autism (Ozonoff & Strayer, 1997; Schmitz et al., 2006). In addition, researchers have used the paradigms to study effects of aging and development (Kramer, Humphrey, Larish, Logan, & Strayer, 1994; Levin et al., 1991; Nielson, Langenecker, & Garavan, 2002; Williams, Ponesse, Schachar, Logan, & Tannock, 1999) and to test individual differences in constructs such as impulsivity (Logan, Schachar, & Tannock, 1997; Reynolds, Ortengren, Richards, & de Wit, 2006).

In the response-inhibition literature, it is common to generalize the results obtained in the go/no-go paradigm to the stop-signal paradigm, and vice versa. The rationale is that response inhibition is achieved in the same way in the two paradigms. Performance in both paradigms is typically described in terms of a race between two competing processes: a go process and a stop process (Logan, 1981; Logan & Cowan, 1984). In the go/no-go paradigm, the go process is triggered by stimulus presentation because of a prepotent response tendency, and the stop process is triggered by the identification of the no-go stimulus. In the stop-signal paradigm, the go process is triggered by the presentation of the go stimulus, and the stop process is triggered by the presentation of the stop signal. The probability of responding on a no-go trial or a stop-signal trial depends on the relative finishing time of the go process and the stop process. When the stop process finishes before the go process, response inhibition is successful and no response is emitted; when the go processes finishes before the stop process, response inhibition is unsuccessful and the response is incorrectly emitted.

In their executive act of control model, Logan and Cowan (1984) suggested that an executive system inhibits a response by replacing the go goal, which is associated with responding to the go stimulus, with a stop goal, which is associated with withholding a response. When the go goal is no longer active, processing in the subordinate systems that execute the response stops relatively quickly, often within 200 to 300 ms of stop-signal presentation. In the present study, we investigated whether activating the stop goal in the go/no-go paradigm is indeed an executively controlled process. More specifically, we investigated whether the stop goal can be automatically activated through memory retrieval of consistent associations between the no-go stimulus and stopping.

Automatic Response Inhibition

Theories of automaticity assume that automatic processes develop over practice. Schneider and Shiffrin (1977) distinguished between consistent and varied mappings of stimuli onto responses. In *consistent mapping*, the stimulus is (consistently) mapped onto the same response throughout practice, whereas in *varied mapping*, the stimulus is (inconsistently) mapped onto different responses throughout practice. In consistent mapping, associations between the stimulus and response are formed, and automatic processing develops across practice. In varied mapping, inconsistent stimulus–response associations are formed, thereby preventing automatic processing (Schneider & Shiffrin, 1977). In a similar vein, Logan (1988) proposed that attending to a stimulus leads to the storage of a new processing episode, which consists of a specific combination of the stimulus, the interpretation given to the stimulus, the response, and the task goal. When the stimulus is repeated, previous processing episodes are retrieved, facilitating performance if the mapping is consistent and impairing performance if the mapping is inconsistent.

The distinction between consistent and varied mapping suggests that automatic response inhibition could develop in the go/no-go paradigm but is unlikely to develop in the stop-signal paradigm, suggesting an important difference between the paradigms. In the go/no-go paradigm, the go stimuli are consistently associated with going and the no-go stimuli are consistently associated with stopping. This consistent mapping of stimuli onto goals may allow automatic inhibition to develop over the course of the experiment: Go stimuli may automatically activate the go goal, while no-go stimuli may automatically activate the stop goal (Logan, 1988; Schneider & Shiffrin, 1977). When the stop goal is activated through memory retrieval, go processing may be suppressed without the need for additional control processes. By contrast, in the stop-signal paradigm, each choice stimulus is usually associated with going (on no-stop-signal trials) and stopping (on stop-signal trials), so automatic inhibition could not support successful stop-signal performance (Logan, 1988; Schneider & Shiffrin, 1977), and executive control processes would be needed to stop the response (Logan & Cowan, 1984).

The consistent mapping in the go/no-go paradigm and the inconsistent mapping in the stop-signal paradigm suggests that automatic, bottom-up inhibition and controlled, top-down inhibition may be differentially involved in the two paradigms. This suggests that the two paradigms may not be equivalent because they put different demands on cognitive control. Given the widespread assumption that the two paradigms are equivalent, this issue has important implications for many literatures, including research on response inhibition in cognitive science, cognitive neuroscience, psychopathology, aging, and individual differences. The purpose of the present study is to explore the *automatic-inhibition hypothesis*—the hypothesis that consistent stimulus–stop mappings can lead to automatic retrieval of the stop goal—in the go/no-go paradigm (Experiments 1–4) and in a modified version of the stop-signal paradigm (Experiment 5).

The Present Study

The automatic-inhibition hypothesis assumes that automatic inhibition will develop across practice with consistent mappings of stimuli onto the stop goal (Logan, 1988; Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977). The stop goal is an abstract cognitive representation of not responding to a stimulus, and we assume that this representation will be associated with no-go stimuli. When a no-go stimulus is repeated, the stop goal will be activated through memory retrieval, which may facilitate response inhibition and, ultimately, may be sufficient to stop the response

without intervening control processes. We tested this hypothesis in five experiments in which we focused on the role of stimulus–stop associations. This focus does not imply that we assume that no other learning effects can occur, and as our data will show, stimulus–category (Experiment 2) and category–goal associations (Experiment 3) most likely also play a role.

In Experiments 1–4, subjects made semantic judgments (e.g., living or nonliving) about the referents of words in a go/no-go paradigm. Whether the stimulus required a response was determined by the stimulus category (e.g., living = go, and nonliving = no-go). Each experiment consisted of a training phase, in which we repeated each stimulus several times, and a test phase, in which we reversed the go and the no-go categories (e.g., living = go and nonliving = no-go in the training phase became nonliving = go and living = no-go in the test phase). By reversing the go/no-go categories, stimuli that were associated with stopping in the training phase required a go response in the test phase and vice versa. The automatic-inhibition hypothesis predicts slower responses to go stimuli in the test phase that were associated with stopping in the training phase. Consistent mapping of these stimuli onto stopping during training should cause the stop goal to be retrieved automatically when the go/no-go categories are reversed in the test phase, and the retrieved stop goal should suppress responding. In Experiments 1–4, we cannot distinguish between *stimulus–go–response associations* (i.e., associations between the stimulus and cognitive representation of the to-be-executed response, namely pressing the space bar) and *stimulus–go–goal associations* (i.e., associations between the stimulus and a more abstract cognitive representation of the action, namely the goal of responding to a word on the screen, without specifying the exact response). In Experiment 5, we can distinguish between go–response and go–goal associations, and we will show that at least part of the effects is due to the retrieval of stimulus–go–goal associations. For consistency and simplicity, we refer to both types of association as stimulus–go associations (similarly, we will refer to stimulus–stop–goal associations as stimulus–stop associations).

In Experiment 5, we tested whether automatic inhibition is a specific characteristic of the go/no-go paradigm or a more general phenomenon that depends on the consistency of the stimulus–stop associations. To this end, we asked if automatic response inhibition could also occur in a modified version of the stop-signal paradigm in which some of the stimuli were associated consistently with going and stopping. The automatic-inhibition hypothesis predicts that under these circumstances, automatic response inhibition would also occur in the stop-signal paradigm.

Experiment 1

There were two conditions: the *inconsistent-test-phase condition*, which assessed the effect of automatic inhibition, and the *consistent-test-phase condition*, which served as a within-subjects control condition. The inconsistent-test-phase condition consisted of two phases: a *training phase*, in which we presented go and no-go stimuli for several repetitions, and a *test phase*, in which we reversed the go/no-go mapping (e.g., living = go and nonliving = no-go in the training phase whereas nonliving = go and living = no-go in the test phase) and added a set of new stimuli to assess learning. The automatic-inhibition hypothesis predicts that go stimuli will be associated with going and no-go stimuli will be associated with stopping in the training phase. Consequently, when go/no-go mapping is reversed in the test phase, go reaction times (RTs) will be longer for old stimuli that were associated with stopping than for new stimuli that were not presented in the training phase. An alternative, *no-stop-learning hypothesis*, assumes that subjects learn associations between go stimuli and going but not between no-go stimuli and stopping. Consequently, this hypothesis predicts no difference in go RT between old and new stimuli in the test phase.

In addition to go RTs, we looked at $p(respond|no\text{-}go)$. Both the automatic-inhibition hypothesis and the no-stop-learning hypothesis predict that $p(respond|no\text{-}go)$ will be higher for old stimuli than for new stimuli in the test phase. Old stimuli were previously associated with going, so go RTs should be shorter for old stimuli than for new stimuli. Because response inhibition on no-go trials depends on the relative finishing time of the go process and stop process (Logan, 1981; Logan & Cowan, 1984), $p(respond|no\text{-}go)$ will increase when the finishing time of the go process decreases, resulting in higher $p(respond|no\text{-}go)$ for old stimuli than for new stimuli.

In addition to the inconsistent-test-phase condition, we included a within-subjects control condition—the consistent-test-phase condition—which also consisted of a training phase and a test phase, in which new items were introduced along with the old ones. The go/no-go mapping was the same in the training phase and the test phase. The consistent-test-phase condition allowed us to test stimulus-specific learning effects under consistent mapping. The automatic-inhibition hypothesis and the no-stop-learning hypothesis both predict that go RTs will be shorter for old stimuli than for new stimuli in the test phase: Old stimuli will be associated with going in the training phase, whereas new stimuli will not. Thus, go RTs should be shorter for old stimuli than for new stimuli in the test phase. The two hypotheses make different predictions for $p(respond|no\text{-}go)$. The automatic-inhibition hypothesis assumes that old stimuli are associated with stopping but new stimuli are not. Consequently, the stop process will finish sooner for old stimuli than for new stimuli, resulting in lower $p(respond|no\text{-}go)$. By contrast, the no-stop-learning hypothesis assumes that no stimulus–stop associations are learned, so $p(respond|no\text{-}go)$ should be the same for old and new stimuli.

It is possible that in addition to stimulus–go associations, subjects may learn other stimulus-specific associations, such as stimulus–category associations. In that case, the no-stop-learning hypothesis also predicts that $p(respond|no\text{-}go)$ would be lower for old items than for new items in the consistent-test-phase condition because old items would be categorized faster than would new items. More generally, this would imply that some of the effects (e.g., speeding of the go RT for consistent go items) could be caused partly by the retrieval of stimulus–category associations. We will discuss this in more detail when it is relevant for the interpretation of our data.

Note that the inconsistent and consistent test phases had the same procedure except for the reversal of mapping of categories onto go and no-go responses in the test phase. The two conditions were compared within subjects, so we used different judgment tasks and different stimulus sets for the two conditions (living/nonliving judgments for one; large/small judgments for the other) to avoid transfer effects between conditions. The judgment tasks assigned to each condition were counterbalanced across subjects.

Method

Subjects. Sixteen subjects from Vanderbilt University participated for monetary compensation ($12). All subjects reported normal or corrected-to-normal vision and were native speakers of English.

Apparatus and stimuli. The experiment was run on a PC running Tscope (Stevens, Lammertyn, Verbruggen, & Vandierendonck, 2006), and the stimuli were presented on a 21-in. monitor. A list of 128 words was drawn from a list of 640 words used by Arrington and Logan (2004). Word length and word frequency data appear in Table 1. For every subject, 4 different subsets of 32 words were selected. For each set, there were 8 large living stimuli, 8 small living stimuli, 8 large nonliving stimuli, and 8 small nonliving stimuli. The first subset was presented in the training and test phases of the inconsistent-test-phase condition, and the second subset was presented only in the test phase of the inconsistent-test-phase condition. The third subset was presented in the training and test phases of the consistent-test-phase condition, and the fourth subset was presented only in the test phase of the consistent-test-phase condition. All stimuli were presented in a white lower case Courier font on a black background and ranged from 12 to 52 mm in width (approximately 1.1° to 5.0°) and 4 to 7 mm (approximately 0.4° to 0.7°) in height.

Procedure. Subjects participated in both the inconsistent- and consistent-test-phase conditions. The order of the conditions, the order of the judgment tasks, and the go/no-go mapping were completely counterbalanced. Subjects were seated individually in private testing rooms after providing informed consent. The experimenter left the room after giving instructions and watching the first few practice trials. Subjects received instructions for the second condition after they had finished the first condition. In the inconsistent test phase, subjects received new instructions after the training phase, explaining the new go/no-go mapping rules. Instructions for the second condition and instructions for the test phase of the inconsistent-test-phase condition were presented on the screen. Subjects could decide when to continue by pressing the space bar.

The training phase of the inconsistent-test-phase condition consisted of eight blocks of 64 trials. In each training block, the words from the first subset of 32 words were presented two times. The training phase was followed by a test phase in which the go/no-go mapping was reversed (e.g., living = go and nonliving = no-go in the training phase and nonliving = go and living = no-go in the test phase). The test phase consisted of three blocks of 64 trials. In each test block, the words from the first subset of 32 words (i.e., the old items) and the second subset of 32 words (i.e., the new items) were presented once in random order.

The training phase of the consistent-test-phase condition consisted of eight blocks of 64 trials. In each training block, the words from the third subset of 32 words were presented two times. The test phase consisted of 3 blocks of 64 trials. In each test block, the words from the third subset of 32 words (i.e., old stimuli) and the fourth subset of 32 words (i.e., new stimuli) were presented once in random order. The go/no-go mapping was the same (i.e., consistent) for the training phase and the test phase. There were no other differences between the two conditions.

In the inconsistent-test-phase condition, the go/no-go mapping was defined by the stimulus category. Half of the subjects made living/nonliving judgments about the referents of words. The other half of the subjects decided whether the referent was smaller or larger than a basketball. This resulted in four possible go/no-go mappings, which are summarized in Table 2. In the consistent-test-phase condition, the go/no-go mapping was also defined by the stimulus category (see Table 2). If subjects made living/nonliving judgments in the inconsistent-test-phase condition, then they made small/large judgments in the consistent-test-phase condition. If subjects made small/large judgments in the inconsistent-test-phase condition, then they made living/nonliving judgments in the consistent-test-phase condition.

In both phases of both conditions, all trials started with the presentation of the word in the center of the screen. Subjects were instructed to press the space bar of a QWERTY keyboard with the index finger of the dominant hand as quickly as possible when a go stimulus was presented and to refrain from pressing it when a no-go stimulus was presented. The word remained on the screen for 1,000 ms, regardless of go RT, in order to equate study time for go and no-go stimuli. A response could be given only while the stimulus was on the screen. Regardless of whether a response was given, the trial ended at stimulus offset. The intertrial interval was 750 ms. At the end of each block, the mean RT on go trials, the number of missed responses on go trials, and the number of incorrect responses on no-go trials were

Table 1
Word Length and Word Frequency (per Million; see Kučera & Francis, 1967) for the Different Stimulus Categories in Experiments 1 and 2

	Word length				Word frequency			
	Experiment 1		Experiment 2		Experiment 1		Experiment 2	
Stimuli	L	NL	L	NL	L	NL	L	NL
Small								
M	5.3	5.9	4.9	5.2	10.6	11.4	15.1	15.7
Range	3–9	3–11	3–8	4–8	2–37	1–34	6–37	6–34
Large								
M	5.0	5.6	4.6	4.7	10.9	11.3	15.5	15.4
Range	3–8	3–9	3–8	3–7	1–32	2–32	7–32	6–32

Note. L = living; NL = non-living.

Table 2
The Four Possible Mappings for the Inconsistent-Test-Phase Condition and the Four Possible Mappings for the Consistent-Test-Phase Condition

	Training phase		Test phase	
Mapping	Go category	No-go category	Go category	No-go category
Inconsistent-test-phase condition				
Mapping 1	living	non-living	non-living	living
Mapping 2	non-living	living	living	non-living
Mapping 3	small	large	large	small
Mapping 4	large	small	small	large
Consistent-test-phase condition				
Mapping 1	living	non-living	living	non-living
Mapping 2	non-living	living	non-living	living
Mapping 3	small	large	small	large
Mapping 4	large	small	large	small

Note. Mapping was counterbalanced with the restriction that for the same subject, different judgment tasks were used for the inconsistent-test-phase condition and the consistent-test-phase condition.

displayed and subjects had to pause for 10 s, after which they could continue by pressing the space bar.

Results and Discussion

Mean go RT and *p*(*respond*|*no-go*) are depicted in Figure 1. The percentage of correct go trials (i.e., go trials on which a response was executed) was 96% and was not further analyzed. Mean RTs for correct go trials were calculated after removal of RTs longer than 2.5 *SD*s above the mean for each trial type (2.5%). In this experiment (and all following experiments), *p*(*respond*|*no-go*) was calculated without removal of outlying RTs.

Global analyses of the training phase data and test phase data and an overview of the relevant planned comparisons appear in Appendix A. The training phase data show that go RTs and *p*(*respond*|*no-go*) decreased over practice for both conditions, suggesting that learning occurred (see Figure 1 and Appendix A). The main question is whether there was evidence of stimulus-specific learning in the test phases of the inconsistent-test-phase and consistent-test-phase conditions.

Inconsistent-test-phase condition. The automatic-inhibition hypothesis predicts that go RTs should be longer for old stimuli than for new stimuli, whereas the no-stop-learning hypothesis predicts that go RTs should be similar for old and new stimuli. Overall, go RTs were longer for old stimuli (559 ms) than for new stimuli (547 ms), although this main effect just failed to reach significance (see Appendix A). The difference between stimulus types was influenced by practice in the test phase. As can be seen in the top left panel of Figure 1, go RTs were shorter for old stimuli (566 ms) than for new stimuli in Block 9 (590 ms), suggesting a benefit that carried over from the training phase. However, go RTs for new items sped up substantially over practice in the test phase, whereas go RTs for old items remained relatively stable. As a result, go RTs were significantly longer for old stimuli than for new stimuli in Blocks 10–11 (Block 10: old = 565 ms, new = 535 ms; Block 11: old = 546 ms, new = 517 ms). This result confirms the automatic-inhibition hypothesis, which predicts that learned associations between old stimuli and stopping impairs go performance in the test phase and disconfirms the no-stop-learning hypothesis, which predicts similar go RTs for old and new stimuli in all blocks of the test phase.

For *p*(*respond*|*no-go*), the difference between the stimulus types was also influenced by practice in the test phase. As can be seen in the bottom left panel of Figure 1, *p*(*respond*|*no-go*) for old stimuli remained relatively stable over the test blocks whereas *p*(*respond*|*no-go*) for new stimuli decreased substantially over Blocks 9–11. As a result, *p*(*respond*|*no-go*) was significantly lower for old stimuli (8.9%) than for new stimuli (14%) in the beginning of the test phase (i.e., Block 9), whereas the opposite was observed at the end of the test phase (i.e., Block 11, old = 9.8%, new = 4.3%). Mirroring the go RT data, this finding suggests that subjects learned associations between the new stimuli and stopping relatively quickly in the test phase, whereas learning these new stimulus–stop associations was much slower for old items.

Consistent-test-phase condition. To test the learning effects under consistent mapping, we compared performance with old and new items in the test phase of the consistent-test-phase condition. As can be in seen in Figure 1, we found that go RTs decreased over blocks and were consistently shorter for old stimuli (506 ms) than for new stimuli (547 ms). This result confirms the idea that associations between old stimuli and going were learned during the training phase and facilitated performance in the test phase. However, the difference between old and new stimuli became smaller over Blocks 9–11: Whereas go RTs for old stimuli remained relatively stable, go RTs for new stimuli sped up substantially (see Figure 1), suggesting that subjects quickly learn to associate the new stimuli with going. We found similar effects for *p*(*respond*|*no-go*). Overall, *p*(*respond*|*no-go*) was lower for old stimuli (3.4%) than for new stimuli (9.2%). This finding suggests that response inhibition benefited from a consistent stimulus–stop mapping and confirms the predictions of the automatic-inhibition hypothesis. In line with the go RT data, we found that the difference between old and new items became smaller over Blocks 9–11.

Discussion. In Experiment 1, we tested the automatic-inhibition hypothesis and no-stop-learning hypothesis by comparing performance for old and new items in the test phase of the inconsistent- and the consistent-test-phase conditions. In the inconsistent-test-phase condition, we found that go RTs were longer for stimuli that were associated with stopping (i.e., the old stimuli) than for stimuli that were not associated with going or stopping (i.e., the new stimuli) in Blocks 10–11. This finding confirms the automatic-inhibition hypothesis, which predicts that responding is automatically suppressed when the stop goal is activated through the retrieval of stimulus–stop associations, and it disconfirms the no-stop-learning hypothesis, which predicts no difference between old and new stimuli. We also found that subjects were less likely to respond to old no-go stimuli than to new no-go stimuli in the consistent-test-phase condition in Blocks 10–11, suggesting that response inhibition benefited from consistent associations between the stimulus and stopping.

In the training phase of both conditions, we found that performance improved substantially over practice. As can be seen in Figure 1, the greatest gain was observed in the beginning of the training phase, which is a characteristic of learning curves in skill acquisition (Newell & Rosenbloom, 1981). Similar learning

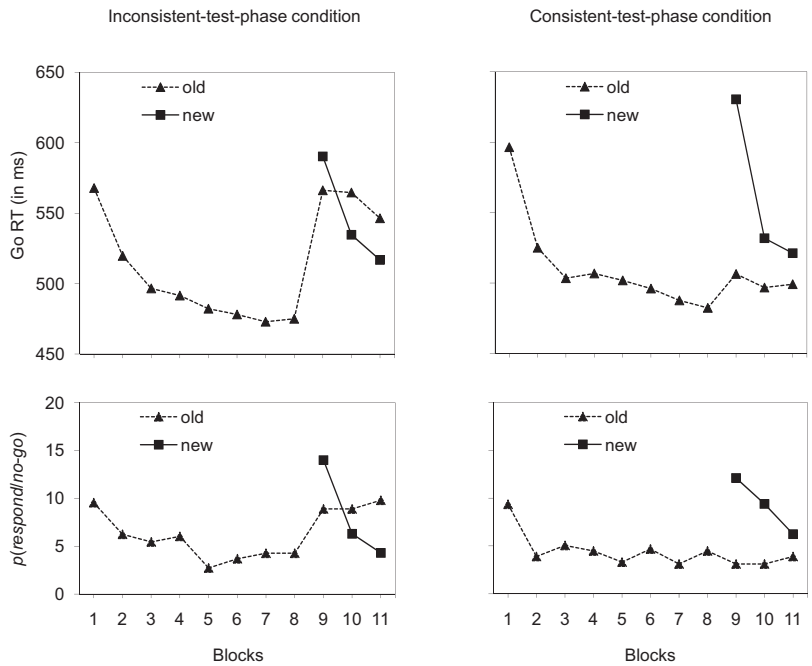

Figure 1. Go reaction times (RTs; in ms; upper panels) and *p*(*respond*|*no-go*) data (in %; lower panels) for the training phase (Blocks 1–8) and test phase (Blocks 9–11) of the inconsistent-test-phase condition (left panels) and the consistent-test-phase condition (right panels) in Experiment 1.

effects were observed for the new stimuli in the test phase: In both conditions, go RT and *p*(*respond*|*no-go*) decreased substantially over Blocks 9–11. Combined, these findings suggest that subjects quickly learned the associations between new no-go stimuli and stopping as well as between new go stimuli and going. The results of the inconsistent-test-phase condition suggest that additional learning effects could also have played a role. In Block 9, both go and stop performance was slightly better for old stimuli than for new stimuli. This suggests an initial learning benefit for old items. This benefit could be due to the retrieval of stimulus–category associations (see e.g., Fisk & Schneider, 1984; Logan, 1990; Shiffrin & Schneider, 1977), which works against the cost from retrieving the stimulus–stop associations. For new stimuli, there are initially no stimulus–category associations, so go RT is longer for new stimuli than for old stimuli at the beginning of the test phase. However, subjects quickly learn the stimulus–category associations for new stimuli, which speed up go RT substantially. This stimulus–category repetition benefit is not counteracted by a stimulus–stop association, so go RT actually becomes shorter for new stimuli than for old stimuli over practice in the test phase.

Experiment 2

In Experiment 1, we examined learning effects under inconsistent and consistent stimulus–stop and stimulus–go mapping. The results of the inconsistent-test-phase condition confirmed the predictions of the automatic-inhibition hypothesis, which assumes that subjects learn associations between no-go stimuli and stopping. In Experiment 2, we further investigated the role of specific stimulus–stop associations in the go/no-go paradigm. In contrast to Experiment 1, in which the judgment task was always the same for the training and test phases of the same condition (see Table 2), we used different judgment tasks for the training and test phase: large/small judgments and living/nonliving judgments (e.g., living = go and nonliving = no-go in the training phase, but small = go and large = no-go in the test phase). This manipulation produced consistent stimuli and inconsistent stimuli. For the *consistent stimuli*, the stimulus–stop and stimulus–go mappings were the same in the training phase and in the test phase (e.g., small living stimuli), even though the go category had changed (i.e., living = go in the training phase changed to small = go in the test phase). For the *inconsistent stimuli*, the stimulus–stop and stimulus–go mappings were different for the training phase and the test phase (e.g., large living stimuli). The automatic-inhibition hypothesis predicts (and the results of Experiment 1 suggest) that the stimuli are directly associated with going and stopping. In the test phase, these stimulus–stop and stimulus–go associations would be retrieved, even though the judgment tasks and relevant stimulus categories were different from the training phase. Consequently, the automatic-inhibition hypothesis predicts longer go RTs for inconsistent stimuli that are associated with stopping in the training phase than for consistent stimuli that are associated with going in the training phase. Similarly, the automatic-inhibition hypothesis predicts higher *p*(*respond*|*no-go*) for inconsistent no-go stimuli that are associated with going in the training phase than for consistent no-go stimuli that are associated with stopping in the training phase.

Method

Subjects. Twenty-four subjects from Vanderbilt University participated for monetary compensation ($12). All subjects reported normal or corrected-to-normal vision and were native speakers of English. None of them participated in Experiment 1.

Apparatus and stimuli. These were similar to Experiment 1, except for the following: A list of 60 words was drawn from a list of 640 words used by Arrington and Logan (2004). Word length and word frequency data appear in Table 1.

Procedure. The procedure was similar to Experiment 1 except for the following: All subjects started with a training phase, which consisted of 20 blocks, followed by a test phase, which consisted of 2 blocks. Each block consisted of 60 trials, and each stimulus was presented one time per block. Half of the subjects decided whether the referents of the words were living or nonliving in the training phase and decided whether the referents were smaller or larger than a computer screen in the test phase (we told the subjects to use the computer screen on which the stimuli were presented as a reference point). For the other half, the order was reversed. Subjects received instructions for the test phase after they finished the training phase. The go/no-go mappings and the order of the conditions were counterbalanced.

In the test phase, one quarter of the stimuli were consistent go stimuli (i.e., go stimulus in the training and test phases), one quarter were consistent no-go stimuli (i.e., no-go stimulus in the training and test phases), one quarter were inconsistent go stimuli (no-go stimulus in the training phase and go stimulus in the test phase), and one quarter were inconsistent no-go stimuli (i.e., go stimulus in the training phase and no-go stimulus in the test phase).

Results and Discussion

Mean go RT and *p*(*respond*|*no-go*) data are depicted in Figure 2. The percentage of correct go trials was 98% and was not further analyzed. Mean RTs for correct go trials were calculated after removal of RTs longer than 2.5 *SD*s above the mean for each trial type (2.3%). Global analyses for the data of the test and training phases appear in Appendix B. As in Experiment 1, there was substantial learning during the training phase. Mean RT decreased as a negatively accelerated function of practice, which is characteristic of changes in RT with practice. A similar learning effect was observed for *p*(*respond*|*no-go*). Again, we will focus on the test phase only.

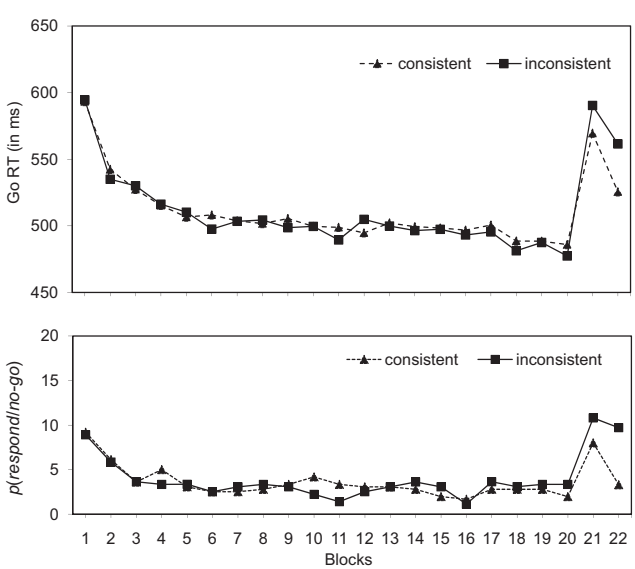

Figure 2. Go reaction times (RTs; in ms; upper panel) and *p*(*respond*|*no-go*) data (in %; lower panel) for the training phase (Blocks 1–20) and test phase (Blocks 21–22) as a function of consistency in Experiment 2.

As can be seen in Figure 2, go RTs were longer for inconsistent stimuli (576 ms) than for consistent stimuli (547 ms), suggesting that subjects learned stimulus–stop associations in the training phase, which are then retrieved in the test phase. Similarly, we found that *p*(*respond*|*no-go*) was higher for inconsistent stimuli (10.2%) than for consistent stimuli (5.2%), suggesting that response inhibition benefited from a consistent stimulus–stop mapping.

To further test whether the longer go RTs for inconsistent stimuli were due to the retrieval of stimulus–stop associations (i.e., the automatic-inhibition hypothesis) or to the absence of stimulus–go–response association (i.e., the no-stop-learning hypothesis), we compared go RT in the first two blocks of the training phase (Blocks 1–2) with go RT in the two test blocks (Blocks 21–22). The automatic-inhibition hypothesis predicts that go RTs for inconsistent stimuli should speed up more in Blocks 1–2 than in Blocks 21–22. In Blocks 21–22, prior stimulus–stop associations would compete with learning new stimulus–go associations, whereas in Blocks 1–2 there are no prior stimulus–stop associations. Consequently, learning new stimulus–go associations would be faster in Blocks 1–2 than in Blocks 21–22. By contrast, the no-stop-learning hypothesis predicts similar learning for inconsistent go stimuli in Blocks 1–2 and Blocks 21–22 because there are no prior stimulus–stop associations to compete with learning new stimulus–go associations. As can be seen in Figure 2, the speed-up in the test phase was slower than the speed-up in the training phase for inconsistent stimuli, which is in accord with the automatic-inhibition hypothesis. This observation is supported by a significant interaction between phase (i.e., training phase vs. test phase) and block (i.e., the first block of the phase vs. the second block of the phase), $F(1, 23) = 5.0$, $p < .05$, $\eta_p^2 = .18$. No such interaction was found for consistent stimuli ($F < 1$; see Figure 2). Combined, these findings suggest that the longer RTs for inconsistent stimuli in the test phase are due to the retrieval of stimulus–stop associations, which interfere with learning new stimulus–go associations, and are not due to the absence of learning of stimulus–go associations in the training phase.

Discussion. In the test phase of Experiment 2, we found that go RTs were longer for inconsistent go stimuli that were associated with stopping in the training phase than for consistent go stimuli that were associated with going in the training phase. Moreover, a comparison of the learning effects in the training phase and the test phase suggests that the longer go RTs for inconsistent stimuli are not due to the absence of an association between the stimulus and going. Combined, these findings support the automatic-inhibition hypothesis and suggest that responding to inconsistent stimuli was slowed down because of the retrieval of stimulus–stop associations. Similarly, the *p*(*respond*|*no-go*) data suggest that stopping performance was impaired for inconsistent no-go stimuli that were associated with going in the training phase. Together, the go RT and *p*(*respond*|*no-go*) results suggest that stimulus–stop associations influence task performance even when the relevant stimulus categories of the training and test phases are different.

As in Experiment 1, the results suggest that retrieval of stimulus–category associations and category–go associations may also have played a role. For consistent stimuli, go RT was longer and *p*(*respond*|*no-go*) was higher in the test phase than in the last blocks of the training phase (see Figure 2), suggesting that category-learning effects also played a role. When new stimulus categories were introduced, performance could no longer benefit from the retrieval of the stimulus–category associations learned in the training phase (Fisk & Schneider, 1984; Logan, 1990; Shiffrin & Schneider, 1977). In addition, associations between categories and going could have sped up go RTs and improved stop performance throughout the training phase. When new categories were introduced in the test phase, performance could no longer benefit from the retrieval of these associations. In sum, category learning may also play a role in the go/no-go paradigm. Because the training phase of Experiment 2 was similar to the training phases of the other experiments, we assume that similar learning effects may have occurred in these experiments as well. Consequently, learned stimulus–category associations may contribute to the positive effect of retrieving consistent stimulus–stop and stimulus–go associations and may reduce the negative effects of retrieving inconsistent stimulus–stop and stimulus–go associations (see *Discussion,* Experiment 1).

Experiment 3

In Experiments 1 and 2, the stimulus–stop mappings (e.g., *hawk* = stop) and the category–stop mappings (e.g., living = stop) were consistent throughout practice in the training phase, which allowed automatic inhibition to develop. Experiment 3 was designed to determine whether automatic inhibition could develop when the stimulus–stop mappings were consistent while the category–stop mappings were inconsistent. To do this, we compared a consistent-stimulus-mapping condition with an inconsistent-stimulus-mapping condition.

In the both conditions, the go/no-go mapping changed after every block (e.g., living = go in Block 1, nonliving = go in Block 2, living = go in Block 3, nonliving = go in Block 4, and so on), making the category–go and category–stop mappings inconsistent. In the consistent-stimulus-mapping condition, the stimulus set also changed after every block (i.e., whenever the go/no-go mapping changed). Half of the stimuli were presented in Blocks 1, 3, 5, and so on, and the other half of the stimuli were presented in Blocks 2, 4, 6, and so on. Thus, particular stimuli were associated consistently with stopping and going even though the categories were not consistently associated with stopping and going (e.g., *goat* was a go stimulus in Blocks 1, 3, 5, etc., but did not occur in Blocks 2, 4, 6, etc.; *horse* was a no-go stimulus in Blocks 2, 4, 6, etc., but did not occur in Blocks 1, 3, 5, etc.). In the inconsistent-stimulus-mapping condition, each stimulus was presented in every block. Consequently, the stimulus mapping and the category mapping were both inconsistent (e.g., *hawk* was a go stimulus in Blocks 1, 3, 5, etc., and a no-go stimulus in Blocks 2, 4, 6, etc.; similarly, living was the go category in Blocks 1, 3, 5, etc., and the no-go category in Blocks 2, 4, 6, etc.).

The automatic-inhibition hypothesis predicts that stimulus–stop associations would develop in the consistent-stimulus-mapping condition (Logan, 1988; Shiffrin & Schneider, 1977). To test this idea, we reversed the stimulus mapping in a test phase at the end of the consistent-stimulus-mapping condition, and we performed two comparisons. First, we compared performance in the test phase of the consistent-mapping condition with performance in the last two blocks of the training phase of the consistent-mapping condition. The stimulus mapping is reversed in the test phase, so the automatic-inhibition hypothesis predicts longer go RT and higher *p*(*respond*|*no-go*) in the test phase than in the last blocks of the training phase. Second, we compared performance in the test phase of the consistent-stimulus-mapping condition with performance in Block 2 of the inconsistent-stimulus-mapping condition. For go RTs, this comparison examines the first go response to a repeated stimulus in both conditions, pitting 12 no-go trials (and 0 go trials) in the consistent-stimulus-mapping condition against one no-go trial (and 0 go trials) in the inconsistent-stimulus-mapping condition. The automatic-inhibition hypothesis predicts that the probability that stimulus–stop associations are retrieved should be higher after 12 repetitions than after 1, so go RT should be longer in the test phase of the consistent-stimulus-mapping condition than in Block 2 of the inconsistent-stimulus-mapping condition (Logan, 1988; Shiffrin & Schneider, 1977). By contrast, the no-stop-learning hypothesis predicts no difference in go RT because subjects respond to the stimulus for the first time. (Note that the no-stop-learning hypothesis predicts shorter go RTs in the test phase of the consistent-stimulus-mapping condition when stimulus–category associations are learned). For *p*(*respond*|*no-go*), this second comparison examines the first no-go response to a repeated stimulus in both conditions, pitting 12 go trials (and 0 no-go trials) in the consistent-stimulus-mapping condition against 1 go trial (and 0 no-go trials) in the inconsistent-stimulus-mapping condition. The automatic-inhibition hypothesis and the no-stop-learning hypothesis predict that *p*(*respond*|*no-go*) should be higher in the test phase of the consistent-stimulus-mapping condition than in Block 2 of the inconsistent-stimulus-mapping condition.

Method

Subjects. Twenty-eight subjects (14 subjects in the consistent-stimulus-mapping condition and 14 subjects in the inconsistent-stimulus-mapping condition) from Vanderbilt University participated for monetary compensation ($6). All subjects reported normal or corrected-to-normal vision and were native speakers of English. None of them participated in Experiments 1 or 2.

Apparatus and stimuli. These were similar to Experiment 1, except for the following: A list of 60 words (this list was different from the list used in Experiment 2) was drawn from a list of 640 words used by Arrington and Logan (2004). Word length ranged from 3 to 8 letters and averaged 4.43 and 4.40 for living and nonliving stimuli, respectively. The word frequency ranged from 6 to 127 per million and averaged 24.7 and 25.1 for living and nonliving stimuli, respectively (Kučera & Francis, 1967).

Procedure. The procedure was similar to Experiment 1 except for the following:

In both the consistent-stimulus-mapping and inconsistent-stimulus-mapping conditions, the mapping of categories (living or nonliving) onto go and no-go was reversed after each block. The consistent-stimulus-mapping condition consisted of 26 blocks of 30 trials, divided into a training phase (Blocks 1–24) and a test phase (Blocks 25–26). In the training phase, half of the stimuli were presented in the even blocks, and the other half were pre-

sented in the odd blocks. Consequently, half of the stimuli always required a go response and half always required a no-go response (i.e., stimulus mapping was consistent). In the test phase, we changed the stimulus mapping: Stimuli that were presented in the even blocks (Blocks 2, 4, 6, etc.) were presented in Block 25, and stimuli that were presented in the odd blocks (Blocks 1, 3, 5, etc.) were presented in Block 26. Consequently, stimuli that were associated with stopping in the training phase were go stimuli in the test phase and vice versa. The inconsistent-stimulus-mapping condition consisted of 12 blocks of 60 trials. Each stimulus was presented once in each block, and the go/no-go mapping changed after every block. Consequently, each stimulus required a go response as often as it required a no-go response (i.e., stimulus mapping was inconsistent).

For half of the subjects in each mapping condition, living stimuli were go stimuli in the odd blocks, and nonliving stimuli were go stimuli in the even blocks. This order was reversed for the other half of the subjects. In both conditions, subjects received instructions at the beginning of each block, and the go category was presented on the top and bottom of the screen during the whole block. Unlike Experiments 1 and 2, subjects did not receive feedback after each block. This was done to avoid having subjects receive more feedback in the consistent-stimulus-mapping condition than in the inconsistent-stimulus-mapping condition.

Results and Discussion

Go RTs and *p*(*respond*|*no-go*) are depicted in Figure 3. The percentage of correct go trials was 98% and was not further analyzed. Mean RTs for correct go trials were calculated after removal of RTs longer than 2.5 *SD*s above the mean for each trial type (2.6%). Global analyses for the training and test phases appear in Appendix C. As can be seen in Figure 3, performance did not improve much over practice in the training phase (Shiffrin & Schneider, 1977). There was no difference in go RT between consistent- and inconsistent-stimulus-mapping conditions, suggesting that consistent associations between the stimulus and going did not improve performance when the category–go mapping changed after each block. For the *p*(*respond*|*no-go*) data, we found that no-go performance was numerically better in the consistent-stimulus-mapping condition than in the inconsistent-stimulus-mapping condition (see Figure 3). However, this difference failed to reach significance (see Appendix C).

To test the idea that stimulus–stop associations developed in the consistent-stimulus-mapping condition, we compared performance in the test phase (Blocks 25 and 26) with performance in the last two blocks (Blocks 23 and 24) of the training phase of the consistent-stimulus-mapping condition. For this analysis, we collapsed across Blocks 23 and 24 (training) and across Blocks 25 and 26 (test). Go RTs were longer in the test phase (571 ms) than in the last blocks of the training phase (547 ms), $F(1, 13) = 7.5$, $MSE = 549$, $p < .05$, $\eta_p^2 = .36$, confirming the predictions of the automatic-inhibition hypothesis. As can be seen in Figure 3, *p*(*respond*|*no-go*) numerically increased in the test phase, but this difference was not significant, $F(1, 13) = 1.6$. To test whether the longer go RTs in the test phase were due to the retrieval of the stimulus–stop associations (i.e., the automatic-inhibition hypothesis) or to the absence of stimulus–go associations (i.e., the no-stop-learning hypothesis), we compared performance in the test phase of the consistent-stimulus-mapping condition with performance in Block 2 of the inconsistent-stimulus-mapping condition, pitting 12 no-go trials (and 0 go trials) against 1 no-go trial (and 0 go trials). As can be seen in Figure 3, go RTs were significantly longer in Blocks 25 and 26 of the consistent-stimulus-mapping condition (571 ms) than in Block 2 (530 ms) of the inconsistent stimulus-mapping condition, $F(1, 26) = 8.8$, $MSE = 1{,}382$, $p < .01$, $\eta_p^2 = .25$. This finding disconfirms the predictions of the no-stop-learning hypothesis and suggests that the longer go RTs in the test phase are due to the retrieval of stimulus–stop associations and not to the absence of stimulus–go associations. We did not observe a significant difference for *p*(*respond*|*no-go*) data ($F < 1$).

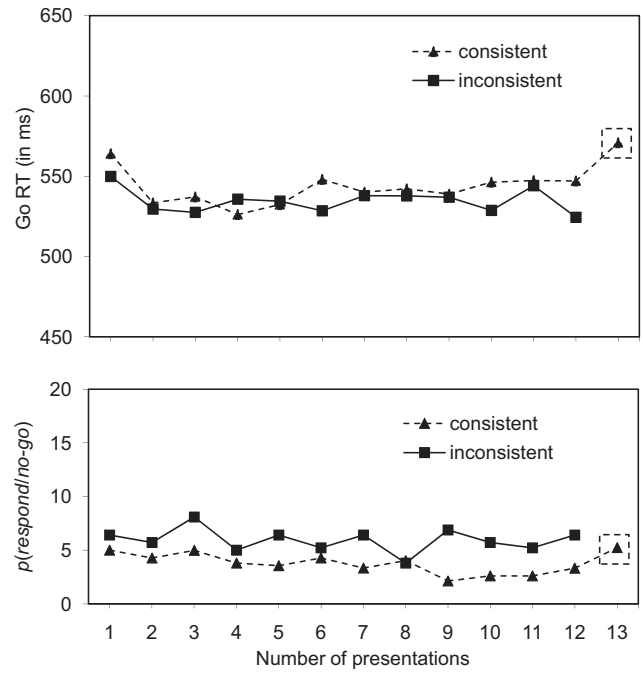

Figure 3. Go reaction times (RTs; in ms; upper panel) and *p*(*respond*|*no-go*) data (in %; lower panel) for the consistent-stimulus-goal-mapping condition and the inconsistent-stimulus-goal-mapping condition as a function of the number of presentations in Experiment 3. Data of the test phase in the consistent-stimulus-goal-mapping condition appear in the small dashed box.

We also tested whether automatic inhibition could occur when the stimulus–stop mappings were inconsistent by comparing performance in the test phase of the consistent-stimulus-mapping condition with performance in the last block (i.e., Block 12) of the inconsistent-stimulus-mapping condition. As can be seen in Figure 3, go RTs were significantly longer in the test phase of the consistent-stimulus-mapping condition (571 ms) than in Block 12 (524 ms) of the inconsistent-stimulus-mapping condition, $F(1, 26) = 7.6$, $MSE = 2{,}037$, $p < .05$, $\eta_p^2 = .22$. This finding suggests that stimulus–stop associations were acquired in the consistent-stimulus-mapping condition but not in the inconsistent-stimulus-mapping condition (Logan, 1988; Shiffrin & Schneider, 1977).

Discussion. In Experiment 3, we further tested the role of consistent stimulus–stop associations in automatic inhibition. In the consistent-stimulus-mapping condition, stimuli were consistently associated with stopping and going whereas the categories

were inconsistently associated with stopping and going. In the inconsistent-stimulus-mapping condition, both the stimuli and categories were inconsistently associated with stopping and going. Go RT was longer in the test phase of the consistent-stimulus-mapping condition than in Block 2 of the inconsistent-stimulus-mapping condition. Again, this finding confirms the automatic-inhibition hypothesis and disconfirms the no-stop-learning hypothesis because subjects responded to the stimulus for the first time in both the test block of the consistent-stimulus-mapping condition and Block 2 of the inconsistent-stimulus-mapping condition.

We did not find differences in go performance in the training phase of the consistent-stimulus-mapping condition and the inconsistent-stimulus-mapping condition (see Appendix C). This suggests that go performance is influenced by the consistency of associations between the category and going or stopping. Indeed, there was greater improvement with practice in Experiments 1 and 2, which involved consistent category mappings, than in the present experiment, which did not (compare Figures 2 and 3 with Figure 3). This suggests that consistent category mappings may play an important role in addition to consistent stimulus mappings: Categories may also be associated with stopping and going, and the retrieval of category–go and category–stop associations may influence task performance. Possibly, the stimulus–go associations were weaker than the stimulus–stop associations in this experiment, explaining why changing the category–go mapping undid the benefit of consistent stimulus–go associations in the training phase, whereas changing the category–stop mapping did not undo the cost of inconsistent stimulus–stop associations in the test phase.

Experiment 4

In Experiments 1 and 2, we demonstrated that subjects learned associations between go stimuli and going and between no-go stimuli and stopping in the training phase. When the stimuli are presented in the test phase, stimulus–go and stimulus–stop associations are retrieved, which impairs performance if the associations are inconsistent but facilitates performance if the associations are consistent. Experiment 3 demonstrated that automatic inhibition could develop through practice when the associations between stimuli and stopping were consistent even when the associations between categories and stopping were inconsistent.

In Experiment 4, we examined the amount of practice that is necessary for automatic-inhibition effects to emerge by varying the number of repetitions in the training phase. Studies of skill acquisition show that automaticity develops as a function of the number of repetitions in the training phase (Logan, 1988; Shiffrin & Schneider, 1977). Thus, the probability that a stimulus–stop association is retrieved should increase with the number of repetitions in the training phase. To test this idea, we manipulated the number of repetitions of go and no-go stimuli. There were four repetition conditions: 1-training-block condition, the 4-training-blocks condition, the 8-training-blocks condition, and the 16-training-blocks condition. The training phase of each condition consisted of a variable number of training blocks (i.e., 1, 4, 8, or 16). All words were presented one time per training block. The training phase of each condition was followed by a test block in which the go/no-go mapping was reversed. In the test block, all stimuli were inconsistent. If the probability that a stimulus–stop association is retrieved increases with practice, then automatic-inhibit effects are more likely to emerge when the number of repetitions in the training phase increases.

Method

Subjects. Twenty-four subjects from Vanderbilt University participated for monetary compensation ($12). All subjects reported normal or corrected-to-normal vision and were native speakers of English. None of them had participated in Experiments 1–3.

Apparatus and stimuli. These were similar to Experiment 1, except for the following: A list of 160 words was drawn from a list of 640 words used by Arrington and Logan (2004). For every subject, 4 different subsets of 40 words were selected. The first subset was presented in the 1-training-block condition, the second subset was presented in the 4-training-blocks condition, the third subset was presented in the 8-training-blocks condition, and the fourth subset was presented in the 16-training-blocks condition. The word length ranged from 3 to 11 letters and averaged 5.1 and 5.3 for living and nonliving stimuli, respectively. The word frequency ranged from 4 to 117 per million and averaged 15.4 and 15.7 for living and nonliving stimuli, respectively (Kučera & Francis, 1967).

Procedure. The procedure was similar to Experiment 1, except for the following: There were four conditions in the experiment. Each condition consisted of a variable number of training blocks (i.e., 1, 4, 8, or 16 training blocks) and one test block. The order of the conditions was counterbalanced according to a balanced Latin square. During the training and test phase, subjects made living/nonliving judgments about the referents of words. For half of the subjects, living stimuli were go stimuli and nonliving stimuli were no-go stimuli in the training phase. For the other half, this mapping was reversed. The go/no-go mapping was the same for the four training phases and was always reversed in the test phase (e.g., living = go and nonliving = no-go in the training phases, and nonliving = go and living = no-go in the test phases).

For each condition, a nonoverlapping list of 40 words was selected from the set of 160. We used a different stimulus list for each condition to avoid transfer effects between conditions. Each training block and the test block consisted of 40 trials and each word was presented once per block in random order. We used the same words for the training and test phases of each condition. In each condition, subjects received instructions explaining the go/no-go mapping rules before the first block of the training phase and before the test block.

Results and Discussion

Go RTs and $p(respond|no\text{-}go)$ are depicted in Figure 4. The percentage of correct go trials was 98% and was not further analyzed. Mean RTs for correct go trials were calculated after removal of RTs longer than 2.5 *SD*s above the mean for each trial type (2.5%). Global analyses for the training and test phases and planned comparisons of interest appear in Appendix D. Replicating Experiments 1 and 2, there was significant learning in each training phase that did not differ significantly between conditions.

To test whether automatic inhibition depends on the amount of practice in the training phase, we compared performance in each test block with performance in the last block of the training phase that preceded it. For the 1-training-block condition, we found that go RTs were shorter in the test block (536 ms) than in the training

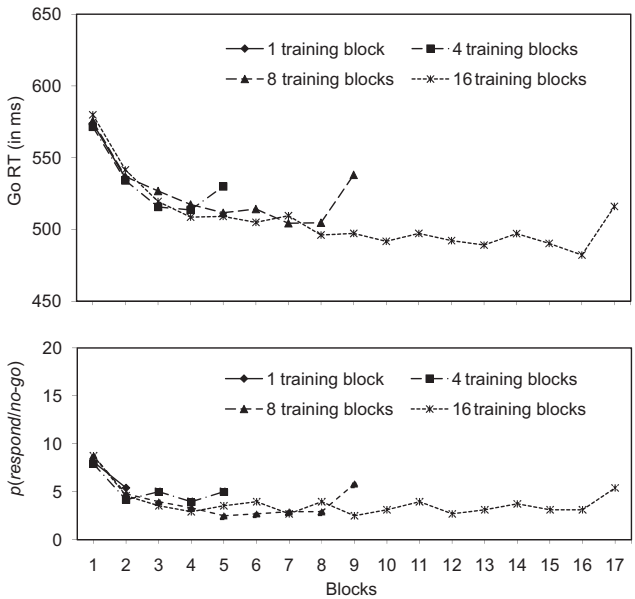

Figure 4. Go reaction times (RTs; in ms; upper panel) and *p*(*respond*|*no-go*) data (in %; lower panel) for the training phase and test phase of every condition (1 training block, 4 training blocks, 8 training blocks, 16 training blocks) in Experiment 4. In every condition, the test phase is the last block of that condition.

block (573 ms), suggesting that automatic-inhibition effects do not emerge after one presentation. This could be because subjects had not learned the association between stimulus and stopping or that the effect of the stimulus–stop association was counteracted completely by the retrieval of stimulus–category associations (see *Discussion* in Experiments 1 and 2). We found longer go RTs in the test block for the other conditions (see Figure 4), suggesting that associations between the stimulus and stopping developed after 4, 8, or 16 repetitions. For the 4-training-blocks condition, go RTs were longer in the test block (530 ms) than in the last block of the training phase (514 ms). The slowing was more pronounced for the 8-training-blocks condition (training: 504 ms; test: 538 ms) and the 16-training-blocks condition (training: 482 ms; test: 516 ms). This observation was further supported by a linear contrast (see Appendix D). Note that there was still a linear trend when we included only the 4-, 8-, and 16-training-blocks conditions, although the contrast no longer reached significance, $p = .10$ (see Appendix D). Combined, these findings suggest that the strength of associations between stimuli and stopping increases with the number of repetitions (Logan, 1988; Shiffrin & Schneider, 1977). However, this conclusion should be interpreted with caution because strength of association is assessed by comparing the last training block with the test block. Go performance in the last training block differs between the training conditions because go performance improves over practice, changing the baseline against which transfer is assessed. Consequently, the difference between the last block of the training phase and the test block may not allow a direct test of the strength of stimulus–stop associations.

For the *p*(*respond*|*no-go*) data, we found a similar effect of practice. For the 1-training-block condition, we found that *p*(*respond*|*no-go*) was higher in the training block (8.1%) than in the test block (5.4%), suggesting that stopping performance was not impaired after only one presentation. For the 4-training-blocks condition, we did not find a difference between the last block of the training phase and the test block. In line with Experiments 1 and 2, we found that for the 8-training-blocks condition, *p*(*respond*|*no-go*) was lower in the last block of the training phase (2.9%) than in the test block (5.8%). We found a similar difference for the 16-training-blocks condition (training: 3.1%; test = 5.4%), although this difference failed to reach significance (see Appendix D). Combined, the findings for the 8-training-blocks and 16-training-blocks conditions suggest that stopping performance was impaired when the stimulus was associated frequently with going in the training phase. In line with the go RT results, the linear contrast including all training-block conditions was significant (see Appendix D).

Discussion. In Experiment 4, we tested whether automatic inhibition was influenced by the number of repetitions in the training phase. The automatic-inhibition hypothesis predicted that the probability that stimulus–stop associations are learned and subsequently retrieved should increase with the number of repetitions in the training phase (Logan, 1988; Shiffrin & Schneider, 1977). In accord with this prediction, we found longer go RTs in the test block than in the last block of the training phase for the 4-, 8-, and 16-training-blocks conditions but not for the 1-training-block condition. Similarly, *p*(*respond*|*no-go*) was higher in the test phase than in the last block of the training phase for the 8-training-blocks condition and the 16-training-blocks condition, suggesting that the probability that stimulus–go associations are learned and subsequently retrieved increased with the number of practice trials.

We found that the go RT slowing and *p*(*respond*|*no-go*) increase were similar for the test phase of the 8-training-block condition and the test phase of the 16-training-block condition (see Figure 4). This finding is consistent with the general finding that learning curves are negatively accelerated (Newell & Rosenbloom, 1981), showing substantial gains in early practice that diminish with further experience (i.e., the difference between Block $n - 1$ and Block n becomes smaller when n increases). The idea of diminishing gains in Experiment 4 is further supported by the data from the training phases of each part, showing that the largest learning effects were observed in the first few training blocks (see Figure 4).

Experiment 5

In Experiments 1–4, we demonstrated that responses can be automatically inhibited in the go/no-go paradigm through the retrieval of consistent stimulus–stop associations. In Experiment 5, we tested whether automatic inhibition could also occur in a modified version of the stop-signal paradigm in which some of the choice stimuli were consistently associated with stopping and going.

In the stop-signal paradigm, choice stimuli are usually associated with both stopping and going: On no-stop-signal trials (i.e., trials on which no stop signal is presented), the choice stimulus is associated with going, whereas on stop-signal trials (i.e., trials on which a stop signal is presented), the choice stimulus is associated with stopping (although this may depend on the outcome of the race between stopping and going; Verbruggen & Logan, 2008, in press-a; Verbruggen, Logan, Liefooghe, & Vandierendonck,

2008). Because of the inconsistent associations between choice stimuli and going and stopping, automatic inhibition is unlikely to develop and subjects would have to rely on controlled inhibition to stop their responses throughout the whole experiment (cf. Shiffrin & Schneider, 1977). However, stop-signal studies that examined aftereffects of response inhibition showed that stimuli can be associated with stopping (Verbruggen & Logan, 2008, in press-a; Verbruggen, Logan, Liefooghe, & Vandierendonck, 2008). Go RT on no-stop-signal trials was longer after successful stopping when the stimulus from the stop trial was repeated but not when the stimulus alternated (for similar findings, see Rieger & Gauggel, 1999). Verbruggen and Logan (in press-a) found that slowing after successful inhibition can be observed up to 20 trials after the first presentation of the stimulus. Aftereffects of unsuccessful inhibition were smaller or even absent (Verbruggen & Logan, 2008, in press-a; Verbruggen, Logan, Liefooghe, & Vandierendonck, 2008). Together, these results suggest that stimulus–stop associations are less likely to be learned when inhibition is unsuccessful.

The long-term aftereffects of successful stopping suggest that automatic inhibition may be possible in the stop-signal paradigm if stimulus–stop associations are consistent over practice. To test this hypothesis, in Experiment 5, we used a modified version of the stop-signal paradigm that consisted of a training phase and a test phase. In the training phase, a subset of the choice stimuli was consistently associated with stopping or going, and another subset was inconsistently associated with stopping and going, as is typical in stop-signal experiments. In the test phase, we reversed the stimulus–stop and stimulus–go mappings for consistent stimuli, similar to the procedure of Experiments 1–4.

We distinguished between four stimulus types:

1. *Stop–go stimuli* always occurred with a stop signal during training and so were consistently associated with stopping in the training phase. They were always presented without a stop-signal trial during testing and so were consistently associated with going in the test phase. These stimuli correspond to the inconsistent-go stimuli in the test phase of Experiments 1–4.

2. *Go–stop stimuli* never occurred with a stop signal during training and so were consistently associated with going in the training phase. They were always presented with a stop-signal during testing and so were consistently associated with stopping in the test phase. These stimuli correspond to the inconsistent no-go stimuli in the test phase of Experiments 1–4.

3. *Stop + go stimuli* occurred with and without a stop signal in training and testing and so were inconsistently associated with both stopping and going in both phases. These stimuli correspond to the stimuli typically used in the stop-signal procedure.

4. *Go–go stimuli* never occurred with a stop signal and so were consistently associated with going in both the training and test phases. We will use these stimuli as a baseline to assess the automatic-inhibition effects on go trials.

The automatic-inhibition hypothesis predicts longer go RTs for stop–go stimuli than for stop + go and go–go stimuli in the test phase, suggesting that automatic inhibition can also occur in the stop-signal paradigm when the stimuli are consistently associated with stopping. It also predicts that the probability of responding on stop-signal trials [$p(respond|signal)$] during the training phase should be lower for stop–go stimuli than for stop + go stimuli because stop–go stimuli were consistently associated with stopping. Finally, we compared go RTs for stop + go stimuli and go–go stimuli to test whether go performance improves over practice when stimuli are consistently associated with going.

Method

Subjects. Twenty subjects from Vanderbilt University participated for monetary compensation ($12). All subjects reported normal or corrected-to-normal vision and were native speakers of English. None of them had participated in Experiments 1–4.

Apparatus and stimuli. These were similar to Experiment 1, except for the following: We used a subset of the 100 words used in Experiment 1. The word length ranged from 3 to 11 letters and averaged 5.2 and 5.7 for living and nonliving stimuli, respectively. The word frequency ranged from 1 to 32 per million and averaged 11.6 and 12.2 for living and nonliving stimuli, respectively (Kučera & Francis, 1967). The assignment of words to the different stimulus types was different for each subject.

Subjects made living/nonliving judgments about the referents of words and responded by pressing the *Z* (living) or the *slash* key (nonliving) of a QWERTY keyboard with the left and right index fingers, respectively. On stop-signal trials, a loud and clear auditory stop signal (750 Hz, 80dB, 75 ms) was presented through closed headphones (Sennheiser eH 150), using stop-signal presentation functions of the program STOP-IT (Verbruggen, Logan, & Stevens, 2008).

Procedure. Instructions emphasized both accuracy and speed. Subjects were not informed about the different stimulus types and the different phases of the experiment. All trials started with the presentation of the choice stimulus. The stimulus was removed after 1,000 ms and required a response on no-stop-signal trials within 2,500 ms, after which time the new stimulus was presented. On stop-signal trials, an auditory tone was presented, instructing subjects to stop their responses. The stop signal was presented after a variable SSD: 100, 200, 300, 400, or 500 ms after the onset of the choice stimulus. Subjects were told not to wait for the stop signal and that the stop signal would be presented at several delays, so it would be easy to stop on some trials and difficult or impossible to stop on others.

The experiment consisted of 10 blocks and each word was presented once per block. Stimulus presentation was pseudorandomized. First, stop–go stimuli (15 per block) were presented on stop-signal trials in the training phase (Blocks 1–8) and on no-stop-signal trials in the test phase (Blocks 9–10). Second, go–stop stimuli (15 per block) were presented on no-stop-signal trials in the training phase and on stop-signal trials in the test phase. Third, stop + go stimuli (30 per block) were presented on both stop-signal trials and on no-stop-signal trials in both phases. In every block, half of the stop + go stimuli were presented on stop-signal trials (15 per block), and the other half were presented on no-stop-signal trials (15 per block). Every stop + go stimulus was presented four times on a stop-signal trial in the training phase. Fourth, go–go stimuli (40 per block) were presented on no-stop-signal trials in both the training and test phases. Note that we used more go–go stimuli than other stimulus types to ensure that the overall proportion of stop-signal trials (30%) was comparable with the proportion of stop-signal trials used in other studies (see Logan, 1994). At the end of each block, the mean RT on go trials, the number of choice errors on no-stop-signal trials, the number of missed responses on no-stop-signal trials, and the

probability of stopping on a stop-signal trial were displayed, and subjects had to pause for at least 10 s.

Results and Discussion

Global analyses of go RT, choice error rates and *p*(*respond*|*signal*) appear in Appendix E. In the analysis of the test phase, we included only stop–go stimuli for which *p*(*respond*|*signal*) was less than 50% after the training phase to ensure that stop–go stimuli were actually associated with stopping. Previous research has shown that subjects are less likely to learn stimulus–stop associations when stopping is unsuccessful (Verbruggen & Logan, 2008, in press-a; Verbruggen, Logan, Liefooghe, & Vandierendonck, 2008). This was important to allow a comparison with the go/no-go experiments, in which inhibition was successful on the majority of the trials. We show in Appendix F that the effects observed in the test phase were not caused by item selection artifacts.

Mean no-signal RTs, choice error rates for no-stop-signal trials and *p*(*respond*|*signal*) are depicted in Figure 5. Mean RTs for correct no-stop-signal trials were calculated after removal of RTs longer than 2.5 *SD*s above the mean for each trial type (2.8%). For each subject, there were on average at least 12 trials for each cell of the design matrix. We also report stop-signal reaction times (SSRTs) in Appendix E. For choice error rates, no significant differences were found, so we will focus on go RT and *p*(*respond*|*signal*) only.

Go RT for all stimulus types decreased as a negatively accelerated function of practice in the training phase. As can be seen in Figure 5 and Appendix E, comparisons of the difference between Block 1 and Block 8 for the three stimulus types showed that this speed-up was more pronounced for go–go stimuli and go–stop stimuli than for stop + go stimuli, suggesting that go performance benefited from the retrieval of consistent stimulus–go associations. Similarly, *p*(*respond*|*signal*) was significantly lower for stop–go stimuli than for stop + go stimuli, suggesting that stopping performance in the training phase benefited from a consistent stimulus–stop mapping. We also analyzed *p*(*respond*|*signal*) as a function of SSD (see Appendix E). In accord with the "horse race" model (Logan & Cowan, 1984), *p*(*respond*|*signal*) increased with increasing SSD in the training phase. For stop + go stimuli, *p*(*respond*|*signal*) = 0.05, 0.11, 0.27, 0.58, and 0.80 for SSD of 100–500 ms, respectively. For stop–go stimuli, *p*(*respond*|*signal*) = 0.06, 0.09, 0.25, 0.56, and 0.73 for SSD of 100–500 ms, respectively. Thus, it appears that the overall *p*(*respond*|*signal*) difference between stop + go and stop–go stimuli was mainly due to a difference between the stimulus types at the longest SSD.

To test whether automatic inhibition could also occur in the stop-signal paradigm, we compared go RTs for the different stimulus types in the test phase. In accord with the automatic-inhibition hypothesis, go RTs were longer for stop–go trials (668 ms) than for go–go trials (634 ms) and stop + go trials (651 ms), $F(1, 38) = 18.3$, $p < .001$, $\eta_p^2 = .32$; and $F(1, 38) = 4.5$, $p < .05$, $\eta_p^2 = .11$,

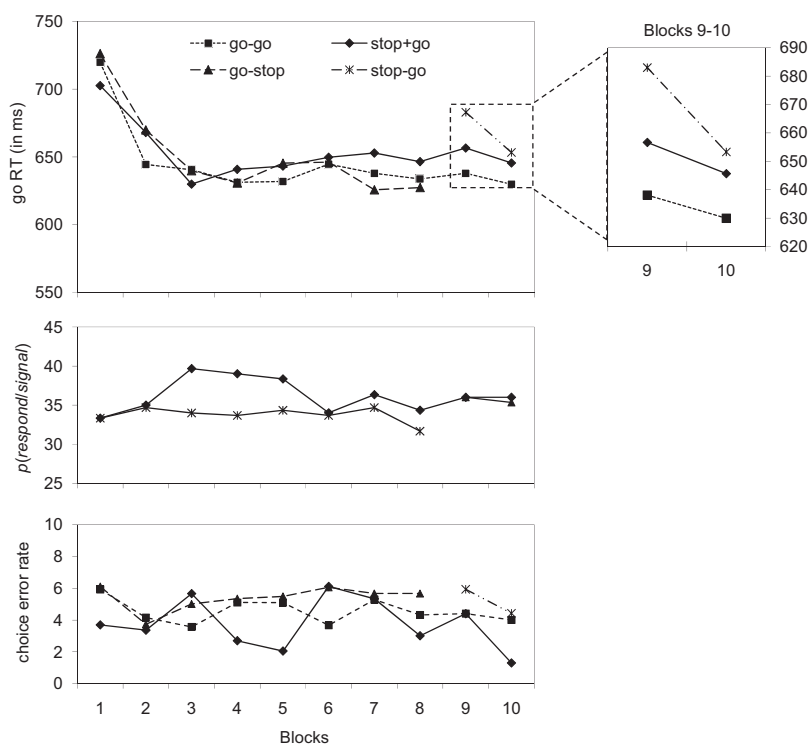

Figure 5. Go reaction times (RTs; in ms; upper panel) and *p*(*respond*|*signal*) data (in %; lower panel) for the four stimulus types as a function of the block (Blocks 1–8 = training phase; Blocks 9–10 = test phase) in Experiment 5.

respectively (see Figure 5). These differences suggest that previously learned stimulus–stop associations are retrieved, which then activates the stop goal even when no stop signal is presented. Go RTs were shorter for go–go stimuli than for stop + go stimuli, $F(1, 23) = 4.6$, $p < .05$, $\eta_p^2 = .11$, suggesting that consistent stimulus–go associations improved performance with go–go stimuli or that inconsistent stimulus–stop and stimulus–go associations impaired stop + go performance through automatic inhibition. The other effects were nonsignificant (see Appendix E). *p*(*respond*|*signal*) in the test phase was 0.36 on average and was not influenced by block or trial type (see Figure 5). Possibly, we did not find the same difference we observed in training because there were only two test blocks, compared with eight training blocks.

Discussion. Experiment 5 replicated the main findings of Experiments 1–4: We found that responding in the test phase slowed when stimuli were consistently associated with stopping in the training phase and that response inhibition in the training phase benefited from consistent associations between the stimulus and stopping (see Figure 5 and Appendix E). These findings support the automatic-inhibition hypothesis, suggesting that under the right circumstances, retrieval of associations between the stimulus and stopping can support performance in the stop-signal paradigm. Also, the findings of Experiment 5 suggest that automatic inhibition can develop when the probability of no-go trials is low. Stop signals occurred on only 30% of the trials, yet automatic inhibition developed with consistent stimulus–stop mapping. This suggests that the results of our go/no-go experiments can be generalized to versions of the go/no-go task in which the proportion of no-go trials is lower than 50% (see e.g., Hershey et al., 2004; Rubia et al., 2001; Schulz et al., 2004; Wager et al., 2005).

In the training phase, we found that the speed-up in go RT was more pronounced for go–go and go–stop stimuli than for stop + go stimuli, suggesting that the retrieval of stimulus–go–goal associations, stimulus–go–response associations, or both facilitated performance. In the previous experiments, we could not distinguish between these two possibilities. However, in this experiment, we can use choice accuracy data to distinguish them. When subjects retrieve associations between the stimulus and the correct response (e.g., *hawk* = left), both go RT and error rates should decrease over practice. By contrast, when subjects retrieve associations between the stimulus and the go goal (i.e., the goal of responding to a word on the screen without specifying the exact response), RT should decrease but error rate need not. As can be seen in Figure 5 and Appendix E, go error rates did not improve over practice, and there were no significant differences between the different stimulus types. This finding suggests that at least part of the speed-up is due to associations between the stimulus and the go goal, whereas associations between stimuli and go responses seem to play a lesser role.

In the training phase, we also found a learning benefit for stop performance. Even though the difference was relatively small, *p*(*respond*|*signal*) was significantly lower for stop–go stimuli (which were consistently associated with stopping in the training phase) than for stop + go stimuli (which were inconsistently associated with stopping in the training phase). This finding suggests that response inhibition in the stop-signal paradigm also benefited from consistent stimulus–stop mappings. The SSRT results reported in Appendix E further supported this idea.

A final question we will address before the General Discussion is whether subjects also learned associations between the stop signal and stopping over practice. Only few studies looked at practice effects in the standard stop-signal paradigm, and results are mixed. Logan and Burkell (1986) found that SSRT decreased over practice (see also e.g., Fillmore, Rush, Kelly, & Hays, 2001), and this suggests that response inhibition could become more automatic. However, Cohen and Poldrack (2008) did not find shorter SSRTs over practice. These inconsistent practice effects for SSRT suggest that response inhibition may benefit less from learning associations between the stop signal and stopping than from learning associations between choice-task stimuli and stopping (as shown in the present study). Associations between choice stimuli and stopping can be retrieved and activated before the stop signal is presented. Therefore, consistent associations turn the choice stimulus itself into a stop signal, speeding inhibition by effectively shortening SSD.

General Discussion

In the present study, we investigated whether a response can be inhibited automatically through the retrieval of consistent stimulus–stop associations. The automatic-inhibition hypothesis assumes that subjects learn associations between the stimulus and stopping on no-go or stop trials. When the stimulus is repeated, the stop goal is activated through the retrieval of stimulus–stop associations, and it suppresses the go response (Logan & Cowan, 1984). If the stimulus is consistently mapped onto stopping, this inhibition can become automatic (Logan, 1988; Shiffrin & Schneider, 1977), thereby reducing the need for top-down executive control processes.

Experiments 1–4 tested the automatic-inhibition hypothesis in the go/no-go paradigm, and Experiment 5 tested the hypothesis in a modified version of the stop-signal paradigm. In all five experiments, we reversed the stimulus–stop mapping after a training phase. In Blocks 10–11 of the test phase of Experiment 1, we found that go RTs were longer for stimuli that were previously associated with stopping than for stimuli that were not associated with stopping or going. We replicated these findings in the test phases of Experiments 2–5. In Experiments 1, 2, 4, and 5, we also observed that the probability of responding on a no-go or stop-signal trial decreased when the stimulus was previously associated with stopping. Combined, these findings demonstrate that automatic inhibition can develop through practice when stimuli are consistently mapped onto stopping (Logan, 1988; Shiffrin & Schneider, 1977).

Automaticity of Go and Stop Performance

Based on the results of the present study, we argue that both go and stop performance can become automatic over practice. Initially, go and stop performance depend on controlled processing only. Over practice, different types of associations, such as stimulus–goal, stimulus–category, and category–goal associations, are formed, and this will lead to a mixture of controlled and automatic processing. Task performance depends on which process (i.e., controlled or automatic) finishes first (Logan, 1988). When the number of practice trials increases, so will the proba-

bility that automatic processing finishes first. Ultimately, subjects can rely on automatic processing exclusively.

The idea of automatic go performance and automatic stop performance can be understood in an interactive race model of response inhibition (Boucher et al., 2007). Elaborating on the independent race model of Logan and Cowan (1984), Boucher et al. assumed a race between a go process and a stop process. The go process is initiated by the presentation of the go stimulus and activates a go unit (viz., the go goal), whereas the stop process is initiated by the presentation of the stop signal and activates the stop unit (viz., the stop goal). The go process and the stop process are independent during the initiation stage, which constitutes the greater part of their durations[1] (see Figure 6). However, once the stop unit is activated, it inhibits the go unit strongly and rapidly. The interactive race is depicted in Figure 6C. On no-stop-signal trials, activation of the go response begins to accumulate in the go unit after the initiation stage is finished, and the response is executed when a threshold is reached. On stop-signal trials, activation of the go unit begins in the same way it begins on no-stop-signal trials; the initiation phase of the stop process begins when the stop signal is presented; and activation of the stop unit begins after the initiation stage finishes (also see Gomez, Perea, & Ratcliff, 2007). On signal-inhibit trials, the stop unit is activated before the threshold for the go response is reached, and the stop unit suppresses activation in the go unit, thereby preventing activation of the go response from reaching the threshold (see Figure 6C). On signal-respond trials, activation of the go response reaches the threshold before the stop unit is activated; consequently, inhibition fails.

Based on the results of Experiments 1–5, we propose that the go unit can be activated through associations between the go stimulus and going. Similarly, the stop unit can also be activated through associations between the go stimulus and stopping (see Figure 6B). When this happens, activation of the stop unit slows the rate at which activation of the go response accumulates, which slows go RT relative to stimuli that were not associated with the stop goal (see Figure 6D). The data of Experiments 1–5 suggest that subjects did not rely on automatic processing entirely because the go unit often reached threshold when the go stimulus was associated with the stop goal. The data also suggest that stop unit activation was weaker when the stop unit is incorrectly activated on go trials than when it is correctly activated on no-go or stop-signal trials, which explains why responding was slowed down but not completely inhibited. Combined, these findings suggest that go and stop performance depended on a mixture of automatic and controlled processing in Experiments 1–5.

Alternative Hypotheses

We argued that the results of the present study confirmed the automatic-inhibition hypothesis and disconfirmed the alternative no-stop-learning hypothesis. In this section, we briefly discuss three other alternative hypotheses: an alternative-response-association hypothesis, a response-threshold-adjustment hypothesis, and an error-tag hypothesis.

The alternative-response-association hypothesis predicts that no-go and stop stimuli are associated with an alternative response. This hypothesis assumes that subjects stop go responses by preparing an opposing response (e.g., by activating opposing muscles). When the stimulus is repeated in the test phase, the association with the opposing response will be retrieved, causing interference with the correct response alternative. However, recent behavioral evidence suggested that subjects do not stop a response through the activation of an alternative response (Camalier et al., 2007; Verbruggen, Schneider, & Logan, in press). These studies required subjects to replace one response with another, so in principle, the first response could be stopped by activating the second response. However, modeling and the effects of experimental manipulations showed that a stop process was required to inhibit the first response before the second one could be activated. Moreover, neuroimaging studies (Aron, Behrens, Smith, Frank, & Poldrack, 2007; Aron & Poldrack, 2006; Rubia et al., 2001) and a recent TMS study (Coxon, Stinear, & Byblow, 2006) suggested the involvement of a global inhibitory mechanism on stop-signal and no-go trials. Combined, these findings undermine the alternative-response-association hypothesis.

The second hypothesis is the response-threshold-adjustment hypothesis, which assumes that the threshold in the go unit is increased when associations between stimuli and the stop goal are retrieved. Subjects may expect that they need to stop when the stop goal is retrieved, so they may increase the response threshold to reduce the probability of responding. This would be consistent with the finding that subjects proactively change the response threshold when they expect a stop signal to occur on the next few trials (Verbruggen & Logan, in press-b). However, other studies showed that adjusting decision criteria is a time-consuming process that subjects find difficult to do proactively on a trial-by-trial basis (see e.g., Brown & Steyvers, 2005; Rotello & Macmillan, 2007; Strayer & Kramer, 1994). In this present study, subjects would have to make adjustments within a trial after the retrieval of stimulus–stop associations. Therefore, we think that it is unlikely that subjects adjust response thresholds when they retrieve stimulus–stop associations.

A third hypothesis is the error-tag hypothesis, which assumes that no-go stimuli are associated with an error tag when an erroneous response is executed, and retrieval of the error tag interferes with responding when the stimulus is repeated. In the stop-signal paradigm, we found long-term aftereffects of successful stopping but no long-term aftereffects of unsuccessful stopping (Verbruggen & Logan, in press-a). This finding undermines the error-tag hypothesis. Therefore, we think that it is unlikely that the automatic-inhibition effect in the go/no-go paradigm and the stop-signal paradigm can be attributed to the retrieval of an error tag.

In sum, we believe that none of the alternative hypotheses can explain the stimulus-specific learning effects of Experiments 1–5 as adequately as the automatic-inhibition hypothesis. It is possible that several factors, including retrieval of different types of associations and strategic adjustments, played a role to some degree, but we believe that the stimulus-specific learning effects are more likely to be due to the retrieval of stimulus–stop associations.

[1] SSRT primarily reflects the period before the stop unit is activated, during which stop and go processing are independent (see Figure 6), so its predictions approximate those of the independent race model (Logan & Cowan, 1984).

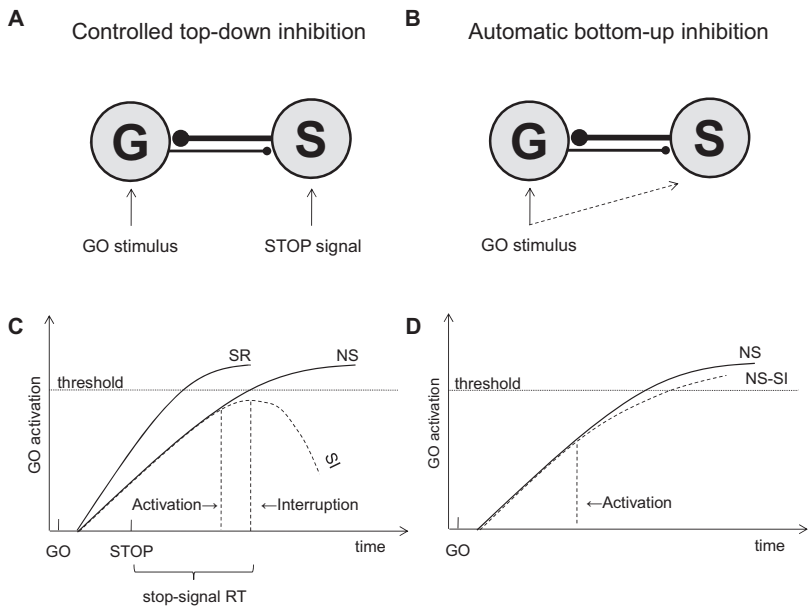

Figure 6. The interactive race model architecture. The go (G) unit is activated via the presentation of the go stimulus. The go response is executed once the threshold is reached. Interruption of the go process starts once the stop (S) unit is activated. Panel A: On stop-signal trials, the S unit is activated via the presentation of the stop signal (left panel). Panel B: On no-stop-signal trials for which the stimulus was previously associated with the stop goal, the S unit is activated via the presentation of the primary-task stimulus. Panel C: Go unit activation for no-stop-signal trials (NS), signal-respond trials (SR), and signal-inhibit trials (SI) as a function of the primary-task stimulus presentation (GO) and stop-signal presentation (STOP). The moment the stop unit is activated is indicated by the leftmost vertical dashed line; the moment go processing is successfully interrupted is indicated by the rightmost vertical dashed line. Panel D: Go unit activation for consistent no-stop-signal trials (NS) and no-stop-signal trials on which the stop goal is activated through memory retrieval (NS-SI) as a function of the primary-task stimulus presentation (GO). The moment the stop unit is activated is indicated by the vertical dashed line. RT = reaction time.

Differences in Automatic Inhibition in the Go/No-Go and Stop-Signal Paradigms

Many researchers assume that response inhibition is a top-down executive control process (e.g., Andres, 2003; Logan, 1994; Ridderinkhof et al., 2004; Stuphorn & Schall, 2006). The go/no-go paradigm and the stop-signal paradigm are considered equivalent, and researchers generalize the results obtained in one paradigm to the other (e.g., Aron et al., 2004; Nigg, 2000; Ridderinkhof et al., 2004; but see Rubia et al., 2001). The critical assumption underlying this practice is that response inhibition is achieved the same way in the two paradigms. However, the present study demonstrated that responses can be inhibited two ways: a controlled top-down way and an automatic bottom-up way, depending on the consistency of the stimulus–stop associations. Stimuli are typically mapped consistently onto stopping and going in the go/no-go paradigm, so automatic inhibition is likely to occur in the go/no-go paradigm. However, stimuli are typically mapped inconsistently onto stopping and going in the stop-signal paradigm, so automatic inhibition is unlikely to occur in the stop-signal paradigm. Thus, our results suggest that the two paradigms are not equivalent because they allow different kinds of response inhibition.

The distinction between automatic and controlled inhibition has important implications for many literatures. The go/no-go paradigm and the stop-signal paradigm are used to study basic cognitive control processes, inhibitory deficits in psychiatric and neurological disorders, lifespan development, and individual differences. The present study suggests that researchers who use the go/no-go paradigm should consider the possibility that response inhibition can be achieved automatically through practice with consistent stimulus–goal associations. Consequently, performance differences between clinical and control groups could be due to differences in controlled inhibition, automatic inhibition, or both. If the clinical group has difficulties with learning, automatic inhibition may not develop, so they will need to rely on controlled inhibition throughout the experiment. Control subjects who do not have learning difficulties can rely on automatic and controlled inhibition. For example, patients with schizophrenia, Alzheimer's disease, and Huntington's disease show impaired performance in the go/no-go task compared with that of control subjects (Crawford et al., 2005; Kiehl et al., 2000; Sprengelmeyer, Lange, & Homberg, 1995). However, these patient groups also show impaired performance in associative learning tasks (e.g., Faust, Balota, & Spieler, 2001; Rich, Campodonico, Rothlind, Bylsma, & Brandt, 1997; Rushe, Woodruff, Murray, & Morris, 1999). Consequently, the apparent deficit in inhibition may really reflect a deficit in learning: Control subjects can rely on automatic inhibition as practice progresses, whereas patients with schizophrenia, Alzheimer's disease, and Huntington's disease may have to rely on controlled inhibition throughout the experiment. More generally, brain regions, such as the basal ganglia, are involved in response

inhibition and associative learning, so lesions to these regions could produce response-inhibition deficits by impairing the learning of associations between stimuli and the stop goal. When consistent stimulus–goal associations are used—as in the typical go/no-go paradigm—researchers cannot distinguish between deficiencies in automatic inhibition, controlled inhibition, or both.

Studies that show impaired performance in the go/no-go paradigm should consider the possibility that the observed deficits really reflect learning deficits. However, the opposite pattern of results is also possible, namely, that existing inhibitory deficits are not observed because the need for executive control processes is reduced throughout practice with consistent stimulus–goal associations. Indeed, several studies reported group differences in the stop-signal paradigm but not in the go/no-go paradigm (see e.g., Rush, Barch, & Braver, 2006; van den Wildenberg et al., 2006). This discrepancy in results could be due to the different involvement of controlled inhibition in the two paradigms. In the go/no-go paradigm, automatic inhibition develops and the need for top-down inhibition is reduced as practice progresses. By contrast, in the stop-signal paradigm, the inconsistent mapping prevents automatic inhibition from developing and the need for controlled inhibition remains high throughout the experiment (although some small practice effects can be observed in the stop-signal paradigm as well; Cohen & Poldrack, 2008; Logan & Burkell, 1986). Consequently, researchers are more likely to observe deficits in controlled inhibition in the stop-signal paradigm than in the go/no-go paradigm.

In sum, our results suggest that researchers who want to investigate controlled inhibition should use inconsistent stimulus–goal mappings to preclude automatic inhibition. Stimulus–stop associations are usually inconsistent in the stop-signal paradigm, so it may be a better procedure for studying controlled inhibition than is the go/no-go paradigm. However, it is possible to prevent the development of automatic inhibition in the go/no-go paradigm by increasing the number of go and no-go stimuli so there are no repetitions. This would weaken the associations between the stimuli and stopping (Shiffrin & Schneider, 1977) and reduce the probability that the stop goal is retrieved automatically (Logan, 1988).

Automatic Control and Cognitive Control of Goal-Directed Behavior

In the present study, we propose that automatic inhibition occurs because stimuli that are associated with the stop goal retrieve the stop goal when they are repeated, and that interferes with go processing. This idea resembles accounts of inhibitory aftereffects, such as the negative priming effect (Neill & Valdes, 1992; Neill, Valdes, Terry, & Gorfein, 1992), long-term inhibition of return (Tipper, Grison, & Kessler, 2003), and certain accounts of task-switching effects (e.g., Koch & Allport, 2006; Mayr & Bryck, 2005; Waszak, Hommel, & Allport, 2003). The common idea in these accounts is that the stimulus is associated with higher-level information such as an inhibitory tag, an inhibitory state, a task rule, an action plan, or a task goal. When the stimulus is repeated, this higher-level information is retrieved and influences ongoing processing.

More generally, the automatic retrieval of associations between stimuli and task goals allows bottom-up control of goal-directed behavior (e.g., Ach, 1935; Bargh & Ferguson, 2000; G. A. Miller, Galanter, & Pribram, 1960). People may activate and manipulate task goals without conscious choice or controlled processing. This way, habits develop and people no longer rely on executive processes to perform certain actions. For example, if you played soccer in your childhood, seeing a soccer ball during a walk in the park may provoke you to kick it or do your favorite soccer trick. However, the possibility that stimulus-driven behavior occurs does not imply that executive control processes are never required in everyday life. For novel stimuli and novel situations, the stimulus, context, or situation may not provide sufficient information to guide goal-directed behavior automatically. In this case, people may rely on an executive system that manipulates goal representations (e.g., Logan & Cowan, 1984; Meyer & Kieras, 1997; E. K. Miller & Cohen, 2001) or activates action schemas (Norman & Shallice, 1986). If you have never seen a soccer ball in your life, it is unlikely that you will automatically kick the ball away. However, after inspecting the ball, you may touch it gently with your foot to find out if it is heavy. Executive control processes may also be needed when stimulus-driven or familiar actions are inappropriate in a certain situation or context. In this case, people may rely on executive control processes to suppress habitual responses because the current situation requires them to do so. Going back to the soccer example, if someone tells you not to touch the ball, you may suppress your habit of kicking the ball and give it to the owner instead. In sum, we argue that the relative contribution of bottom-up and top-down control depends on previous experience and the context in which the stimulus is repeated (for similar ideas, see Koechlin & Summerfield, 2007).

Conclusion

The major contribution of the present study is to show that goal-directed performance in the go/no-go and stop-signal paradigms can rely on both bottom-up control and top-down control. We propose that automatic and controlled inhibition can work together to guide goal-directed behavior; therefore, they should not be regarded as opposites (see Logan, 1988). When retrieval of consistent associations is sufficient to activate goal representations, performance may rely completely on bottom-up control. However, when consistent associations are not available or when higher-level goals require subjects to oppose familiar actions, an executive system may be needed to control actions (also see Bargh & Chartrand, 1999; Koechlin & Summerfield, 2007).

References

Ach, N. (1935). Analyse des Willens [Analysis of the will]. In E. Abderhalden (Ed.), *Handbuch der biologischen Arbeitsmethoden* (Vol. VI, pp. 1–459). Berlin, Germany: Urban & Schwarzenberg.

Andres, P. (2003). Frontal cortex as the central executive of working memory: Time to revise our view. *Cortex, 39,* 871–895.

Aron, A. R., Behrens, T. E., Smith, S., Frank, M. J., & Poldrack, R. A. (2007). Triangulating a cognitive control network using diffusion-weighted magnetic resonance imaging (MRI) and functional MRI. *Journal of Neuroscience, 27,* 3743–3752.

Aron, A. R., & Poldrack, R. A. (2006). Cortical and subcortical contributions to stop signal response inhibition: Role of the subthalamic nucleus. *Journal of Neuroscience, 26,* 2424–2433.

Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2004). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences, 8,* 170–177.

Arrington, C. M., & Logan, G. D. (2004). Episodic and semantic components of the compound-stimulus strategy in the explicit task-cuing procedure. *Memory & Cognition, 32,* 965–978.

Badcock, J. C., Michie, P. T., Johnson, L., & Combrinck, J. (2002). Acts of control in schizophrenia: Dissociating the components of inhibition. *Psychological Medicine, 32,* 287–297.

Baddeley, A. (1996). Exploring the central executive. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 49*(A), 5–28.

Bargh, J. A., & Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist, 54,* 462–479.

Bargh, J. A., & Ferguson, M. J. (2000). Beyond behaviorism: On the automaticity of higher mental processes. *Psychological Bulletin, 126,* 925–945.

Bekker, E. M., Overtoom, C. C., Kenemans, J. L., Kooij, J. J., De Noord, I., Buitelaar, J. K., & Verbaten, M. N. (2005). Stopping and changing in adults with ADHD. *Psychological Medicine, 35,* 807–816.

Boucher, L., Palmeri, T. J., Logan, G. D., & Schall, J. D. (2007). Inhibitory control in mind and brain: An interactive race model of countermanding saccades. *Psychological Review, 114,* 376–397.

Brown, S., & Steyvers, M. (2005). The dynamics of experimentally induced criterion shifts. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 587–599.

Camalier, C. R., Gotler, A., Murthy, A., Thompson, K. G., Logan, G. D., Palmeri, T. J., & Schall, J. D. (2007). Dynamics of saccade target selection: Race model analysis of double step and search step saccade production in human and macaque. *Vision Research, 47,* 2187–2211.

Cohen, J. R., & Poldrack, R. A. (2008). Automaticity in motor sequence learning does not impair response inhibition. *Psychonomic Bulletin & Review, 15,* 108–115.

Coxon, J. P., Stinear, C. M., & Byblow, W. D. (2006). Intracortical inhibition during volitional inhibition of prepared action. *Journal of Neurophysiology, 95,* 3371–3383.

Crawford, T. J., Higham, S., Renvoize, T., Patel, J., Dale, M., Suriya, A., & Tetley, S. (2005). Inhibitory control of saccadic eye movements and cognitive impairment in Alzheimer's disease. *Biological Psychiatry, 57,* 1052–1060.

Donders, F. C. (1969). On the speed of mental processes. In W. G. Koster, *Attention and performance II* (pp. 412–431). Amsterdam: North-Holland. (Original work published in 1868).

Faust, M. E., Balota, D. A., & Spieler, D. H. (2001). Building episodic connections: Changes in episodic priming with age and dementia. *Neuropsychology, 15,* 626–637.

Fillmore, M. T., Rush, C. R., Kelly, T. H., & Hays, L. (2001). Triazolam impairs inhibitory control of behavior in humans. *Experimental and Clinical Psychopharmacology, 9,* 363–371.

Fisk, A. D., & Schneider, W. (1984). Consistent attending versus consistent responding in visual search: Task versus component consistency in automatic processing development. *Bulletin of the Psychonomic Society, 22,* 330–332.

Gauggel, S., Rieger, M., & Feghoff, T. A. (2004). Inhibition of ongoing responses in patients with Parkinson's disease. *Journal of Neurology Neurosurgery and Psychiatry, 75,* 539–544.

Gomez, P., Perea, M., & Ratcliff, R. (2007). A model of the go/no-go task. *Journal of Experimental Psychology: General, 136,* 389–413.

Hershey, T., Revilla, F. J., Wernle, A., Gibson, P. S., Dowling, J. L., & Perlmutter, J. S. (2004). Stimulation of STN impairs aspects of cognitive control in PD. *Neurology, 62,* 1110–1114.

Iaboni, F., Douglas, V. I., & Baker, A. G. (1995). Effects of reward and response costs on inhibition in ADHD children. *Journal of Abnormal Psychology, 104,* 232–240.

Kiehl, K. A., Smith, A. M., Hare, R. D., & Liddle, P. F. (2000). An event-related potential investigation of response inhibition in schizophrenia and psychopathy. *Biological Psychiatry, 48,* 210–221.

Koch, I., & Allport, A. (2006). Cue-based preparation and stimulus-based priming of tasks in task switching. *Memory & Cognition, 34,* 433–444.

Koechlin, E., & Summerfield, C. (2007). An information theoretical approach to prefrontal executive function. *Trends in Cognitive Sciences, 11,* 229–235.

Kramer, A. F., Humphrey, D. G., Larish, J. F., Logan, G. D., & Strayer, D. L. (1994). Aging and inhibition: Beyond a unitary view of inhibitory processing in attention. *Psychology and Aging, 9,* 491–512.

Kučera, H., & Francis, W. N. (1967). *Computational analysis of present day American English.* Providence, RI: Brown University Press.

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

Levin, H. S., Culhane, K. A., Hartmann, J., Evankovich, K., Mattson, A. J., Harward, H., et al. (1991). Developmental changes in performance on tests of purported frontal-lobe functioning. *Developmental Neuropsychology, 7,* 377–395.

Logan, G. D. (1981). Attention, automaticity, and the ability to stop a speeded choice response. In J. Long & A. D. Baddeley (Eds.), *Attention and performance IX* (pp. 205–222). Hillsdale, NY: Erlbaum.

Logan, G. D. (1985). Executive control of thought and action. *Acta Psychologica, 60,* 193–210.

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

Logan, G. D. (1990). Repetition priming and automaticity: Common underlying mechanisms. *Cognitive Psychology, 22,* 1–35.

Logan, G. D. (1994). On the ability to inhibit thought and action: A user's guide to the stop signal paradigm. In D. Dagenbach & T. H. Carr (Ed.), *Inhibitory processes in attention, memory and language* (pp. 189–239). San Diego, CA: Academic Press.

Logan, G. D., & Burkell, J. (1986). Dependence and independence in responding to double stimulation: A comparison of stop, change, and dual-task paradigms. *Journal of Experimental Psychology: Human Perception and Performance, 12,* 549–563.

Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review, 91,* 295–327.

Logan, G. D., Schachar, R. J., & Tannock, R. (1997). Impulsivity and inhibitory control. *Psychological Science, 8,* 60–64.

MacLeod, C. M., Dodd, M. D., Sheard, E. D., Wilson, D. E., & Bibi, U. (2003). In opposition to inhibition. In B. H. Ross (Ed.), *Psychology of learning and motivation: Advances in research and theory* (Vol. 43, pp. 163–214).

Mayr, U., & Bryck, R. L. (2005). Sticky rules: Integration between abstract rules and specific actions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 337–350.

McLean, A., Dowson, J., Toone, B., Young, S., Bazanis, E., Robbins, T. W., & Sahakian, B. J. (2004). Characteristic neurocognitive profile associated with adult attention-deficit/hyperactivity disorder. *Psychological Medicine, 34,* 681–692.

Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance. I. Basic mechanisms. *Psychological Review, 104,* 3–65.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 24,* 167–202.

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

Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wager, T. D. (2000). The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: A latent variable analysis. *Cognitive Psychology, 41,* 49–100.

Neill, W. T., & Valdes, L. A. (1992). Persistence of negative priming: Steady-state or decay. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 565–576.

Neill, W. T., Valdes, L. A., Terry, K. M., & Gorfein, D. S. (1992). Persistence of negative priming: II. Evidence for episodic trace retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 993–1000.

Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 1–55). Hillsdale, NJ: Erlbaum.

Nielson, K. A., Langenecker, S. A., & Garavan, H. (2002). Differences in the functional neuroanatomy of inhibitory control across the adult life span. *Psychology and Aging, 17,* 56–71.

Nigg, J. T. (2000). On inhibition/disinhibition in developmental psychopathology: Views from cognitive and personality psychology and a working inhibition taxonomy. *Psychological Bulletin, 126,* 220–246.

Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation: Advances in research and theory* (pp. 1–18). New York: Plenum.

Ozonoff, S., & Strayer, D. L. (1997). Inhibitory function in nonretarded children with autism. *Journal of Autism and Developmental Disorders, 27,* 59–77.

Reynolds, B., Ortengren, A., Richards, J. B., & de Wit, H. (2006). Dimensions of impulsive behavior: Personality and behavioral measures. *Personality and Individual Differences, 40,* 305–315.

Rich, J. B., Campodonico, J. R., Rothlind, J., Bylsma, F. W., & Brandt, J. (1997). Perseverations during paired-associate learning in Huntington's disease. *Journal of Clinical and Experimental Neuropsychology, 19,* 191–203.

Ridderinkhof, K. R., van den Wildenberg, W. P. M., Segalowitz, S. J., & Carter, C. S. (2004). Neurocognitive mechanisms of cognitive control: The role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain and Cognition, 56,* 129–140.

Rieger, M., & Gauggel, S. (1999). Inhibitory after-effects in the stop signal paradigm. *British Journal of Psychology, 90,* 509–518.

Rotello, C. M., & Macmillan, N. A. (2007). Response bias in recognition memory. In A. S. Benjamin & B. H. Ross (Eds.), *The psychology of learning and motivation: Skill and strategy in memory use* (pp. 61–94). London: Academic Press.

Rubia, K., Russell, T., Overmeyer, S., Brammer, M. J., Bullmore, E. T., Sharma, T., et al. (2001). Mapping motor inhibition: Conjunctive brain activations across different versions of go/no-go and stop tasks. *NeuroImage, 13,* 250–261.

Rush, B. K., Barch, D. M., & Braver, T. S. (2006). Accounting for cognitive aging: Context processing, inhibition or processing speed? *Aging Neuropsychology and Cognition, 13,* 588–610.

Rushe, T. M., Woodruff, P. W. R., Murray, R. M., & Morris, R. G. (1999). Episodic memory and learning in patients with chronic schizophrenia. *Schizophrenia Research, 35,* 85–96.

Schachar, R., & Logan, G. (1990). Are hyperactive children deficient in attentional capacity? *Journal of Abnormal Child Psychology, 18,* 493–513.

Schmitz, N., Rubia, K., Daly, E., Smith, A., Williams, S., & Murphy, D. G. M. (2006). Neural correlates of executive function in autistic spectrum disorders. *Biological Psychiatry, 59,* 7–16.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing. I. Detection, search and attention. *Psychological Review, 84,* 1–66.

Schulz, K. P., Fan, J., Tang, C. Y., Newcorn, J. H., Buchsbaum, M. S., Cheung, A. M., & Halperin, J. M. (2004). Response inhibition in adolescents diagnosed with attention deficit hyperactivity disorder during childhood: An event-related fMRI study. *American Journal of Psychiatry, 161,* 1650–1657.

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

Sprengelmeyer, R., Lange, H., & Homberg, V. (1995). The pattern of attentional deficits in Huntingtons disease. *Brain, 118,* 145–152.

Stevens, M., Lammertyn, J., Verbruggen, F., & Vandierendonck, A. (2006). Tscope: A C library for programming cognitive experiments on the MS Windows platform. *Behavior Research Methods, 38,* 280–286.

Strayer, D. L., & Kramer, A. F. (1994). Strategies and automaticity: 2. Dynamic aspects of strategy adjustment. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 342–365.

Stuphorn, V., & Schall, J. D. (2006). Executive control of countermanding saccades by the supplementary eye field. *Nature Neuroscience, 9,* 925–931.

Sylvester, C. Y. C., Wager, T. D., Lacey, S. C., Hernandez, L., Nichols, T. E., Smith, E. E., & Jonides, J. (2003). Switching attention and resolving interference: fMRI measures of executive functions. *Neuropsychologia, 41,* 357–370.

Tipper, S. P., Grison, S., & Kessler, K. (2003). Long-term inhibition of return of attention. *Psychological Science, 14,* 19–25.

van den Wildenberg, W. P. M., van Boxtel, G. J. M., van der Molen, M. W., Bosch, D. A., Speelman, J. D., & Brunia, C. H. M. (2006). Stimulation of the subthalamic region facilitates the selection and inhibition of motor responses in Parkinson's disease. *Journal of Cognitive Neuroscience, 18,* 626–636.

Verbruggen, F., Liefooghe, B., & Vandierendonck, A. (2004). The interaction between stop signal inhibition and distractor interference in the flanker and Stroop task. *Acta Psychologica, 116,* 21–37.

Verbruggen, F., & Logan, G. D. (2008). Aftereffects of goal shifting and response inhibition: A comparison of the stop-change and dual-task paradigms. *Quarterly Journal of Experimental Psychology, 61A*, 1151–1159.

Verbruggen, F., & Logan, G. D. (in press-a). Long-term aftereffects of response inhibition: Memory retrieval, task goals, and cognitive control. *Journal of Experimental Psychology: Human Perception and Performance*.

Verbruggen, F., & Logan, G. D. (in press-b). Proactive adjustments of response strategies in the stop-signal paradigm. *Journal of Experimental Psychology: Human Perception and Performance*.

Verbruggen, F., Logan, G. D., Liefooghe, B., & Vandierendonck, A. (2008). Aftereffects of response inhibition: Repetition priming or between-trial control adjustments. *Journal of Experimental Psychology: Human Perception and Performance, 34,* 413–426.

Verbruggen, F., Logan, G. D., & Stevens, M. A. (2008). STOP-IT: Windows executable software for the stop-signal paradigm. *Behavior Research Methods, 40,* 479–483.

Verbruggen, F., Schneider, D. W., & Logan, G. D. (in press). How to stop and change a response: The role of goal activation in multitasking. *Journal of Experimental Psychology: Human Perception and Performance*.

Vince, M. A. (1948). The intermittency of control movements and the psychological refractory period. *British Journal of Psychology, 38,* 149–157.

Wager, T. D., Sylvester, C. Y. C., Lacey, S. C., Nee, D. E., Franklin, M., & Jonides, J. (2005). Common and unique components of response inhibition revealed by fMRI. *NeuroImage, 27,* 323–340.

Waszak, F., Hommel, B., & Allport, A. (2003). Task-switching and long-term priming: Role of episodic stimulus-task bindings in task-shift costs. *Cognitive Psychology, 46,* 361–413.

Williams, B. R., Ponesse, J. S., Schachar, R. J., Logan, G. D., & Tannock, R. (1999). Development of inhibitory control across the life span. *Developmental Psychology, 35,* 205–213.

(*Appendixes follow*)

Appendix A

Analyses for Experiment 1

The data of the inconsistent-test-phase and the consistent-test-phase conditions were analyzed by means of separate 2 (stimulus type: old vs. new) × 3 (block: 9–11) repeated measures analysis of variance (ANOVA; see Table A1). Planned comparisons appear in Table A2. For all planned comparisons in Experiments 1–5, we used the error term of the interaction. The data of the training phase of the inconsistent-test-phase and the consistent-test-phase conditions were analyzed by means of separate repeated measures ANOVAs with training block (Blocks 1–8) as within-subjects factor (see Table A3).

Table A1
Overview of the Global Analyses of the Test Phase in Experiment 1

	Go reaction time				*p*(*respond*|*no-go*)			
Variable	*df*	*F*	*MSE*	η_p^2	*df*	*F*	*MSE*	η_p^2
			Inconsistent-test-phase condition					
Block	2, 30	33.5***	528	.69	2, 30	2.9†	64	.16
Stimulus type	1, 15	3.6†	922	.19	1, 15	0.3	73	.02
Block × Stimulus Type	2, 30	10.7***	718	.42	2, 30	6.2***	38	.29
			Consistent-test-phase condition					
Block	2, 30	34.3***	983	.70	2, 30	3.1†	17	.17
Stimulus type	1, 15	60.5***	1,443	.80	1, 15	14.5***	57	.49
Block × Stimulus Type	2, 30	57.9***	426	.79	2, 30	3.1†	29	.17

** $p < .05$. *** $p < .01$. † $p < .08$.

Table A2
Overview of the Planned Comparisons for the Inconsistent-Test-Phase Condition and the Consistent-Test-Phase Condition in Experiment 1

	Go reaction time			*p*(*respond*|*no-go*)		
Block	*df*	*F*	η_p^2	*df*	*F*	η_p^2
		Old vs. new items: Inconsistent-test-phase condition				
Block 9	1, 30	6.4**	.18	1, 30	5.5**	.15
Block 10	1, 30	10.0***	.25	1, 30	1.1	.04
Block 11	1, 30	9.7***	.24	1, 30	6.3**	.17
		Old vs. new items: Consistent-test-phase condition				
Block 9	1, 30	288.8***	.91	1, 30	22.4***	.42
Block 10	1, 30	22.9***	.43	1, 30	10.9***	.27
Block 11	1, 30	9.1***	.23	1, 30	1.5	.05

** $p < .05$. *** $p < .01$.

Table A3
Overview of the Global Analyses of the Training Phase in Experiment 1

	Go reaction time				*p*(*respond*|*no-go*)			
Phase	*df*	*F*	*MSE*	η_p^2	*df*	*F*	*MSE*	η_p^2
Inconsistent test phase	7, 105	38.3***	428	.72	7, 105	4.4***	16	.23
Consistent test phase	7, 105	37.9***	555	.72	7, 105	3.6***	17	.20

*** $p < .01$.

Appendix B

Analyses for Experiment 2

The data of the test phase were analyzed by means of 2 (consistency: consistent or inconsistent) × 2 (block: Test Block 1 or Test Block 2) repeated measures ANOVA (see Table B1). The data of the training phase were analyzed by means of a 2 (consistency: consistent or inconsistent) × 20 (block: 1–20) repeated measures ANOVA (see Table B1).

Table B1
Overview of the Global Analyses in Experiment 2

	Go reaction time				*p*(*respond*|*no-go*)			
Variable	*df*	*F*	*MSE*	η_p^2	*df*	*F*	*MSE*	η_p^2
			Training phase					
Consistency	1, 23	0.8	1,445	.04	1, 23	0.0	63	.00
Block	19, 437	22.7***	1,263	.50	19, 437	5.3**	23	.19
Consistency × Block	19, 437	0.7	462	.03	19, 437	0.5	21	.02
			Test phase					
Consistency	1, 23	12.5***	1,538	.35	1, 23	3.7†††	41	.14
Block	1, 23	14.5***	2,189	.39	1, 23	6.3**	96	.21
Consistency × Block	1, 23	2.7	491	.11	1, 23	2.7	25	.11

** $p < .05$. *** $p < .01$. ††† $p = .06$.

Appendix C

Analyses for Experiment 3

The data of the training phase were analyzed by means of a 2 (condition: consistent or inconsistent) × 12 (number of presentations: 1–12) mixed ANOVA (see Table C1). For the consistent-stimulus-goal-mapping condition, we collapsed the data of Blocks 1 and 2 (in both blocks, each stimulus was presented for the first time), the data of Blocks 3 and 4 (in both blocks, each stimulus was presented for the second time), and so on.

Table C1
Analysis of Block and Condition in Experiment 3

	Go reaction time				*p*(*respond*|*no-go*)			
Grouping	*df*	*F*	*MSE*	η_p^2	*df*	*F*	*MSE*	η_p^2
Condition	1, 26	0.3	17,574	.01	1, 26	2.5	173.5	.09
Block	11, 286	1.7	874	.06	11, 286	0.8	18.9	.03
Condition × Block	11, 286	0.7	874	.03	11, 286	0.7	18.9	.03

(*Appendixes continue*)

Appendix D

Analyses for Experiment 4

The data of the test phase were analyzed by means of 2 (phase: training or test) × 4 (training blocks: 1, 4, 8, 16) repeated measures ANOVA (see Table D1). The data of the training phases of the 4-, 8-, and 16-training-blocks conditions were analyzed by means of separate repeated measures ANOVAs with block as within-subjects factor (see Table D2).

Table D1
Overview of Global Analyses and Planned Comparisons of Interest in Experiment 4 by Performances in Test Block and Final Training Block

	Go reaction time				*p*(*respond*|*no-go*)			
Performance	*df*	*F*	*MSE*	η_p^2	*df*	*F*	*MSE*	η_p^2
Global analysis								
Test block condition	3, 69	17.0***	1,495	.43	3, 69	2.7*	25.9	.11
Phase	1, 24	3.0	2,205	.11	1, 24	1.8	21.3	.07
Test Block Condition × Phase	3, 69	20.5***	665	.47	3, 69	3.7	20.7	.14
Planned comparison								
1-training-block condition	1, 69	25.2***		.27	1, 69	4.2**		.06
4-training-block condition	1, 69	5.0**		.07	1, 69	0.6		.01
8-training-block condition	1, 69	20.3***		.23	1, 69	4.9**		.07
16-training-block condition	1, 69	20.5***		.23	1, 69	3.0		.04
Linear contrast								
All conditions	1, 69	47.5***		.41	1, 69	16.6***		.20
1 training block excluded	1, 69	2.6		.04	1, 69	0.9		.01

* $p = .05$. ** $p < .05$. *** $p < .01$.

Table D2
Global Analysis of the Training Phase for the 4-Block

	Go reaction time				*p*(*respond*|*no-go*)			
Condition	*df*	*F*	*MSE*	η_p^2	*df*	*F*	*MSE*	η_p^2
4-training-block condition	3, 69	15.5***	1,115	.40	3, 69	3.6**	22.3	.13
8-training-block condition	7, 161	15.3***	871	.40	7, 161	5.2***	19.7	.18
16-training-block condition	15, 345	18.9***	743	.45	15, 345	3.3***	15.4	.13

** $p < .05$. *** $p < .01$.

Appendix E

Analyses and SSRTs for Experiment 5

The global analyses appear in Table E1. Go performance in the test phase was analyzed by means of a 2 (block: 9–10) × 3 (trial type: stop + go, stop–go, and go–go) repeated measures ANOVA. Stop performance was analyzed by means of a 2 (block: 9–10) × 2 (trial type: stop + go and go–stop) repeated measures ANOVA. Go performance in the training phase was analyzed by means of an 8 (block: 1–8) × 3 (trial type: stop + go, go–stop, and go–go) repeated measures ANOVA. Stop performance was analyzed by means of an 8 (block: 1–8) × 2 (trial type: stop + go and stop–go) repeated measures ANOVA, and by means of a 5 (SSD: 100, 200, 300, 400, 500 ms) × 2 (trial type: stop + go and stop–go) repeated measures ANOVA. For the latter analysis, we collapsed *p*(*respond*|*signal*) across Blocks 1–8 (note that there were not enough stop-signal trials in the test phase to do such an analyses). Planned comparisons for the training phase appear in Table E2.

We also estimated SSRT for the training phase (again, there were not sufficient stop signals to estimate SSRTs reliably in the test phase). SSRT was estimated via the integration method (Logan & Cowan, 1984). Because all stop–go stimuli were present on stop-signal trials, we used two different RT distributions for SSRT estimation: (a) the RT distribution of the go–go stimuli in the training phase and (b) the RT distribution

of stop + go stimuli in the training phase. When we used the RT distribution of the go–go stimuli for the SSRT estimations, SSRT was 308 ms for stop + go stimuli and 279 ms for stop–go stimuli, $F(1, 19) = 11.9$, $p < .01$. When we used the RT distribution of the stop + go stimuli for SSRT estimations, SSRT was 307 ms for stop + go stimuli and 283 ms for stop–go stimuli, $F(1, 19) = 8.6$, $p < .01$. The SSRT differences between stop + go and stop–go stimuli are consistent with the significant difference in *p*(*respond*|*signal*) and suggest that response inhibition benefited from a consistent stimulus–stop–goal mapping.

Table E1
Overview of the Global Analyses in Experiment 5

Grouping	Training phase, Blocks 1–8				Test phase, Blocks 9–10			
	df	*F*	*MSE*	η_p^2	*df*	*F*	*MSE*	η_p^2
Go reaction time								
Block	7, 133	15.9***	2,874	.46	1, 19	3.5	2,246	.16
Stimulus type	2, 38	1.4	1,094	.07	2, 38	7.0***	1,664	.27
Block × Stimulus Type	14, 266	2.0**	965	.10	2, 38	1.1	1,276	.05
p(*respond*|*signal*) as a function of block								
Block	7, 133	0.6	0.01	.03	1, 19	0.02	0.01	.001
Stimulus type	1, 19	5.2**	0.01	.21	1, 19	0.03	0.01	.002
Block × Stimulus Type	7, 133	0.9	0.01	.05	1, 19	0.02	0.01	.001
p(*respond*|*signal*) as a function of stop-signal delay								
Stimulus type	1, 19	5.1	0.01	.21				
Stop-signal delay	4, 76	137.8	0.02	.88				
Stimulus Type × Stop-Signal Delay	4, 76	1.7	0.01	.08				
Go error rates								
Block	7, 133	1.3	17	.06	1, 19	3.5††	18	.16
Stimulus type	2, 38	1.7	45	.08	2, 38	2.4	28	.11
Block × Stimulus Type	14, 266	1.5	16	.07	2, 38	1.1	19	.05

** $p < .05$. *** $p < .01$. †† $p = .07$.

Tabl E2
Planned Comparisons for the Training Phase of Experiment 5 by Difference Between Block 1 and Block 8 for Different Stimulus Types

Stimulus	Difference	*df*	*F*	η_p^2
Stop + go vs. go–go	56 ms vs. 86 ms	1, 266	4.6**	.02
Stop + go vs. go–stop	56 ms vs. 99 ms	1, 266	0.8	.00
Go–go vs. go–stop	86 ms vs. 99 ms	1, 266	9.5***	.04

** $p < .05$. *** $p < .01$.

Appendix F

Effects of Excluding Items in Experiment 5

Consistent with the idea that learning of stimulus–stop associations is influenced by the outcome of the stop process, we found a slightly different data pattern when we included all stop–go items. Similar to the main analysis (i.e., the analysis after exclusion of a subset of the stop–go data for which stopping failed most of the time), we found that the main effect of stimulus type was significant when we included all stop–go items (i.e., even the items for which response inhibition failed on most of the trials), $F(2, 38) = 4.3$, $MSE = 1{,}469$, $p < .05$. Most importantly, the stop–go items (658 ms) were significantly slower than the go–go items (633 ms) when we included all items, $F(1, 38) = 8.5$, $p < .01$. This finding replicates the differences observed in the go/no-go experiments. There was still a small numerical RT difference between stop–go items (658 ms) and stop + go items (651 ms), but this was no longer significant ($F < 1$; unlike in the main analysis). Thus,

(*Appendixes continue*)

it appears that exclusion of items does play a role; the crucial question is whether this is due to less learning of stimulus–stop associations (as we argued in the main text) or because we included only the more difficult stop–go items. Indeed, the probability of stopping increases when go RT increases; consequently, $p(respond|signal)$ may be lower for more difficult items than for easier items. Of course, a selection of stop–go items on the basis of $p(respond|signal)$ could then imply that the included trials were more difficult than the excluded trials. In other words, the difference between stop–go stimuli and the other stimulus types (for which we did not exclude any items) could be due to item difficulty. We tested this item-difficulty hypothesis in several ways.

First, we performed an item analysis (i.e., an ANOVA with items as the random effect). We used the same exclusion criteria as for the subject analysis, but for this analysis, trial type is a within-item factor, so item difficulty is not an issue here. All items that occurred on go–go, stop + go, and stop–go trials across all subjects (after exclusion) were analyzed; this resulted in a set of 87 items. Mean RTs were analyzed by means of an ANOVA with trial type (go–go, stop + go, stop–go) as within-item factor. We found a main effect of trial type, $F(2, 172) = 5.0$, $MSE = 7{,}558$, $p < .01$, which suggests that the differences between go–go, stop–go, and stop + go generalize across items as well as subjects. More importantly, this finding also suggests that the differences observed in the subject analyses are not due to item difficulty. Planned comparisons showed that stop–go RTs (688 ms) were longer than go–go RTs (647 ms) and stop + go RTs (663 ms): $F(1, 172) = 9.7$, $p < .01$; and $F(1, 172) = 3.7$, $p = .057$, respectively. The difference between go–go RTs and stop + go RTs did not reach significance in the item analyses, $F(1, 172) = 1.5$, $p > .23$. Note that we found similar effects when we did not exclude items on the basis of $p(respond|signal)$. For this analysis, we had a set of 94 items. We found a main effect of stimulus type, $F(2, 186) = 3.4$, $MSE = 7{,}488$, $p < .05$. Planned comparisons showed that stop–go RTs (679 ms) were longer than go–go RTs (647 ms) and stop + go RTs (658 ms): $F(1, 186) = 6.5$, $p < .01$; and $F(1, 186) = 2.8$, $p = .098$, respectively (although the latter difference was nonsignificant; $p < .05$, one-tailed).

Second, we compared mean go RT for stop–go items that were either included or excluded on the basis of $p(respond|signal)$. We performed both a subject analysis and an item analysis. For the subject analysis, we compared signal-respond RTs (i.e., RTs for trials on which subjects erroneously executed the response when a stop signal was presented) in Blocks 1–8 for stop–go items that were included in the test phase analyses with stop–go items that were excluded in the test phase analyses. We found that signal-respond RT for excluded items (611 ms) was similar to signal-respond RT for included items (601 ms). This numerical difference did not reach significance ($F < 1$), although it is opposite to what one would expect if the included items are more difficult than the excluded items. For the item analysis, we analyzed the difference between included and excluded items with item as the random effect. For this item analysis, we calculated mean RT for every item on the basis of subjects for which this item occurred on go–go trials in the test phase. Items were categorized as *included items* ($N = 56$) when they were always included as a stop–go item across all subjects; items were categorized as *excluded items* ($N = 40$) when they were excluded as a stop–go item for at least 1 subject. We found that mean go RT was similar for excluded and included items; 649 ms versus 656 ms respectively ($F < 1$). Combined, these findings suggest that the slower stop–go RTs observed in the main subject analyses (i.e., after exclusion of the stop–go items for which response inhibition was unsuccessful most of the time) were not due to the selection of more difficult items.

Received November 14, 2007
Revision received May 22, 2008
Accepted May 31, 2008 ■

Retraction of Hard, Lozano, and Tversky (2006)

The following article from the November 2006 issue has been retracted:

> Hard, B. M., Lozano, S. C., and Tversky, B. (2006). Hierarchical Encoding of behavior: Translating perception into action. *Journal of Experimental Psychology: General, 135,* 588-608.

All authors retract this article. Co-author Tversky and co-author Hard believe that the research results cannot be relied upon; Sandra C. Lozano takes full responsibility for the need to retract this article.

DOI: 10.1037/a0013753