

Dissociable Brain Mechanisms Underlying the Conscious and Unconscious Control of Behavior

Simon van Gaal, Victor A. F. Lamme, Johannes J. Fahrenfort,
and K. Richard Ridderinkhof

Abstract

■ Cognitive control allows humans to overrule and inhibit habitual responses to optimize performance in challenging situations. Contradicting traditional views, recent studies suggest that cognitive control processes can be initiated unconsciously. To further capture the relation between consciousness and cognitive control, we studied the dynamics of inhibitory control processes when triggered consciously versus unconsciously in a modified version of the stop task. Attempts to inhibit an imminent response were often successful after unmasked (visible) stop signals. Masked (invisible) stop signals rarely succeeded in instigating overt inhibition but did trigger slowing down of response times. Masked stop signals elicited a sequence of dis-

tinct ERP components that were also observed on unmasked stop signals. The N2 component correlated with the efficiency of inhibitory control when elicited by unmasked stop signals and with the magnitude of slowdown when elicited by masked stop signals. Thus, the N2 likely reflects the initiation of inhibitory control, irrespective of conscious awareness. The P3 component was much reduced in amplitude and duration on masked versus unmasked stop trials. These patterns of differences and similarities between conscious and unconscious cognitive control processes are discussed in a framework that differentiates between feedforward and feedback connections in yielding conscious experience. ■

INTRODUCTION

What are the limits of unconscious cognition? This question can be studied, for example, in patients with blindsight or neglect, or in healthy participants, for example, by the use of masking, attentional blink, binocular rivalry, or inattentive blindness. In a laboratory setting, masking is the most common tool of choice. In typical masking experiments, participants have to respond to or identify a briefly presented stimulus (the prime) that is followed and/or preceded closely in time by a second stimulus (the mask). Under specific conditions, the prime can be difficult or sometimes even impossible to see. However, even if masked stimuli are not perceived, they can still influence perceptual and behavioral processes. An example of unconscious influences on perception is repetition priming; the observation that processing of a conscious stimulus (the target) is facilitated when a masked version of the same stimulus is presented just before the target (Dehaene et al., 2001; Bar & Biederman, 1999). Other examples pertain to unconscious influences on motor responses. Masked primes, briefly presented before a target, that resemble the target (e.g., with respect to location or form) speed up responses and decrease error rates, whereas responses are slowed down and error rates increase when they differ from the

target (Vorberg, Mattler, Heinecke, Schmidt, & Schwarzbach, 2003; Dehaene et al., 1998).

Although at first controversial (for a review, see Kouider & Dehaene, 2007), it is now widely acknowledged that such relatively low-level (e.g., perceptual and motor) processes are affected by unconscious stimuli (but see Hannula, Simons, & Cohen, 2005; Holender & Duscherer, 2004). However, the extent to which higher level cognitive functions (e.g., task preparation, cognitive control) are also influenced by unconscious information remains debated (Hommel, 2007; Mayr, 2004; Eimer & Schlaghecken, 2003; Dehaene & Naccache, 2001; Libet, 1999; Umiltà, 1988). Interestingly, some recent studies have shown that even high-level cognitive processes, such as decision making (Pessiglione et al., 2008), reward prediction (Pessiglione et al., 2007), and task preparation (Lau & Passingham, 2007; Mattler, 2003), can be influenced unconsciously. These recent findings stress the contribution of unconscious processes in shaping everyday, but rather complex, behavior.

Recently, we have shown that inhibitory control processes, which were thought to require conscious experience (for an overview, see Eimer & Schlaghecken, 2003) and volition (Pisella et al., 2000; Libet, 1999), can also be initiated unconsciously (van Gaal, Ridderinkhof, van den Wildenberg, & Lamme, 2009; van Gaal, Ridderinkhof, Fahrenfort, Scholte, & Lamme, 2008). To illustrate, in a modified version of the go/no-go paradigm (van Gaal et al., 2008), participants had to respond as fast as possible to a

go annulus but were instructed to withhold their response when they perceived a no-go circle, preceding the go annulus. By varying the interval between the no-go circle and the metacontrast go signal, no-go signals were either visible (unmasked) or invisible (masked). Under these circumstances, unconscious no-go signals triggered full-blown response inhibition on some occasions and otherwise slowed down those responses that were not withheld. In EEG, unconscious no-go signals elicited two electrophysiological events: (1) an early occipital component and (2) a frontal component somewhat later in time. The amplitude of the frontal ERP component strongly predicted the amount of slowdown across participants. We argued that the first neural event represented the visual encoding of the unconscious no-go stimulus, whereas the second event corresponded to the subsequent initiation of inhibitory control in the pFC.

In a separate behavioral study, we tested whether stop signal response inhibition could also be triggered unconsciously (van Gaal et al., 2009). Compared with the go/no-go task, inhibition in the stop task is considered a more active form of response inhibition because it requires the active inhibition of an already ongoing response at the very last moment (van Boxtel, van der Molen, Jennings, & Brunia, 2001). In that “masked stop signal paradigm,” participants had to respond as fast and accurately as possible to a choice stimulus but cancel their already initiated action when a second stimulus (the stop signal, the word “stop”) was presented after the choice stimulus (Logan, 1994), but not when a “go-on” signal (a control word) was presented after the choice stimulus. We refer to this form of response inhibition as “selective response inhibition” because participants are not instructed to inhibit their response to any stimulus that is presented after the choice stimulus (which is the case for regular global stop tasks). Instead, a stimulus presented after the choice signal sometimes instructs participants to stop (when the word “stop” is presented) and other times to go on (when the control word is presented). We included visible (unmasked) as well as invisible (masked) stop signals. In that task, participants inhibited their response slightly more often on masked stop trials than on masked go-on trials, and they significantly slowed down their responses to masked stop trials that were not inhibited. Again, these results suggest that masked stop signals are also able to influence inhibitory control operations, strongly associated with the pFC (Aron & Poldrack, 2006; Chambers et al., 2006).

Note that the “endogenous” form of inhibitory control that is studied by using the stop signal task and the go/no-go task differs substantially from the more “exogenous” and automatic form of inhibition studied by Eimer and Schlaghecken (1998, 2003) and Eimer (1999) using the masked priming task. They showed that at longer prime-target intervals (>100 msec), initial response facilitation by congruent primes is automatically followed by inhibition leading to longer RTs on congruent trials than on incongruent trials.

If unconscious stimuli are able to influence such high-level cognitive operations, what might then be the additional value of consciousness in this context? And how is this expressed in neural activity? Here, we measured EEG to study the spatio-temporal dynamics of processing masked versus unmasked stop signals in the above-outlined selective stop signal task as a first step toward answering these questions.

In EEG, successful stopping has typically been related to two ERP components: a fronto-central N2 component, a negative peak around 200–300 msec after stop signal presentation (Dimoska, Johnstone, Barry, & Clarke, 2003; Schmajuk, Liotti, Busse, & Woldorff, 2006), and a centro-parietal P3 component, a positive peak around 300–500 msec after stop signal presentation (Dimoska & Johnstone, 2008; Bekker, Kenemans, Hoeksma, Talsma, & Verbaten, 2005; Ramautar, Kok, & Ridderinkhof, 2004). Although the neural generators of the N2 and the P3 have not been localized precisely, numerous neuroimaging experiments have investigated the neural basis of response inhibition in the stop signal task. These studies have revealed a large fronto-parietal network involved in response inhibition, including middle, inferior, and superior frontal cortices, pre-supplementary motor areas, and anterior cingulate cortex (Zheng, Oka, Bokura, & Yamaguchi, 2008; Aron & Poldrack, 2006; Chambers et al., 2006; Li, Huang, Constable, & Sinha, 2006; Ramautar, Slagter, Kok, & Ridderinkhof, 2006). In addition, several basal ganglia structures have also been associated with stop signal inhibition, most prominently the subthalamic nucleus (Aron & Poldrack, 2006; van den Wildenberg et al., 2006).

In addition to these typical inhibition related ERP observations, recent magnetoencephalography (MEG) or EEG studies revealed a crucial role for sensory processing in response inhibition, which is reflected in relatively early effects (~100–200 msec after stop signal onset) observed at occipital/parietal electrode sites (Boehler et al., 2008; Dimoska & Johnstone, 2008; Schmajuk et al., 2006; Bekker et al., 2005). These recent results suggest that the quality of sensory processing or allocation of attentional resources to the stop stimulus is also an important determinant of the likelihood that a response will be inhibited. In the present experiment, we mixed masked and unmasked stop signals in stop signal task to address to what extent unconscious initiated inhibition differs from its conscious counterpart.

METHODS

Participants

Nineteen undergraduate psychology students participated in the experiment for course credits or financial compensation (12 women). All participants had normal or corrected-to-normal vision. All procedures were executed in compliance with relevant laws and institutional guidelines and were approved by the local ethical committee. Subjects gave written informed consent before experimentation.

Stimuli and Task

We masked stop signals with forward masks only or with forward and backward masks, leading to unmasked (visible) and masked (invisible) stop signals, respectively (see Figure 1A). We also included a so-called “go-on” condition, in which a go-on signal instead of a stop signal was presented after the choice stimulus. This stimulus instructed participants to go on and press the button to the direction of the choice stimulus (e.g., Boehler et al., 2008; Dimoska, Johnstone, & Barry, 2006; van den Wildenberg & van der Molen, 2004; Bedard et al., 2002). Inclusion of this additional go-on condition slightly complicates the stop task, as it requires discrimination between two visual stimuli: one requiring the implementation of response inhibition (stop signals) whereas the other does not (go-on signals). An advantage of this experimental design is that we can directly compare behavioral and electrophysiological responses to masked stop signals and masked go-on signals, which occur equally frequently. By this means, any differences between the stop- and the go-on condition can be attributed to inhibition instead of other cognitive processes such as novelty detection, unexpectedness, or attentional selection (Dimoska & Johnstone, 2008).

Stimuli were presented using Presentation (Neurobehavioral Systems, Albany, CA) against a black background (2.17 cd/m^2) at the center of a 17-in. VGA monitor (frequency 70 Hz.). Participants viewed the monitor from a distance of approximately 90 cm, so that each centimeter subtended a visual angle of 0.64° . On masked stop trials, we first presented a white cross (300 msec) followed after 200 msec by a choice stimulus (29 msec, isoluminant, 9.0 cd/m^2), which was either a blue left-pointing arrow or a red right-pointing arrow (width 0.64° , height 0.34°). This stimulus was followed after a variable SOA by two strings of

randomly chosen uppercase consonants (forward masks, presented sequentially, 43 msec per letter string), the stop signal or the go-on signal (see below, 29 msec), and finally two consonant strings (backward masks, both 43 msec). On unmasked stop and go-on trials, the same sequence was used, but the consonant strings at the end (backward masks) were replaced with blank screens (see Figure 1A). We used different colors for the arrows because we observed in pilot studies that participants were sometimes unable to discriminate between right and left pointing arrows when these were presented in black (especially a short stop signal delay [SSD]). On these occasions, the first letter string masked the direction of the arrow. Because the letter strings were unable to mask the color of the arrow, in the present experiment, participants were (almost) always able to figure out whether a left or right pointing arrow was presented when we used different colors.

Participants were instructed to respond as quickly and as accurately as possible to the direction of the choice stimulus, but to inhibit their response when a stop signal was presented after the choice stimulus. Participants were instructed to “keep on going” and press the button as already planned when a go-on signal was presented. The word “STOP” was used as a stop signal, and a control word was used as a go-on signal. For every participant, a different control word was used. The control word set consisted of the following words: BINK, BLUF, DREK, DUNK, FARM, HALM, HARK, KLIM, KNEL, KURK, KWIK, LARF, NERF, NIMF, RANK, VINK, VLEK, ZINK, and ZWAK. The control words were matched to “stop” in terms of frequency of appearance in daily Dutch language (70 vs. 73 per 1 million, respectively, as stated in the Celex database; Baayen, Piepenbrock, & Gulikers, 1995). The stimulus set of consonants used to form the masks consisted of 13 uppercase letters (X, B, K, R, M, H, G, F, D, W, Z, N, and C). For each subject, 10 of




Figure 1. Stimulus timing in the masked selective stop signal paradigm. (A) Participants had to respond to the direction of the arrow but withhold their response when the stop signal (the word “stop”) was presented, but not when the go-on signal (a control word, e.g., the word “bluf”) was presented. In the unmasked conditions, the stop signal (or go-on signal) could be perceived easily, whereas in the masked conditions participants could not (due to the inclusion of backward masks in those conditions). The stop signal could be presented at various delays after the go stimulus (SSD = stop signal delay), which served to vary the difficulty of response inhibition. SOA is the stimulus onset asynchrony between the choice stimulus (the arrow) and the first forward mask. (B) The stop signal task yields an estimate of the duration of the inhibitory process: the stop signal reaction time (SSRT). The “point of no return” reflects the point in time at which the inhibitory process is finished. In theory, in trials at the right side of this point, the stop process wins from the go process and the response will be inhibited. Trials at the left side of the SSRT probably escape inhibition because the go process is finished before the stop process (Logan, 1994).

these were used to form the masks, such that no consonants were used that were also part of the control (go-on) word for that subject. Each mask contained seven randomly chosen letters, which were slightly overlapping to increase the density of the mask. The spacing between the centers of the letters was 12 pixels. Uppercase Courier font was used for all letters and words (white color, font size 24pt).

When the stop signal is presented shortly after the go signal, participants are able to inhibit their responses easily. However, when the interval between go signal and stop signal is increased, participants are less likely to inhibit their response because the go process is closer to completion. Therefore, a staircase-tracking procedure dynamically adjusted the time between the choice stimulus and the stop signal (or go-on signal), the SSD. After an inhibited unmasked stop trial, the SSD in the next trial increased by 14.3 msec, whereas it decreased by 14.3 msec when the participant did not stop. The staircase adjustment of the SSD counteracted strategic slowing of participants (i.e., waiting for the stop signal to appear before executing any choice response) and ascertained that participants would inhibit their response on approximately 50% of the unmasked stop trials, ensuring that we could accurately calculate participants' stop signal reaction time (SSRT; Logan, 1994). The SSRT is an estimate of the duration of the inhibitory process, which can be used to compare the efficiency of inhibitory control processes between conditions or individuals. All blocks started with an SSD of 129 msec.

The experiment consisted of three sessions. In the first two sessions, participants performed the stop signal task; EEG was recorded in the second session only. The third session was dedicated to the assessment of stop signal visibility (see below). We included a behavioral session before the EEG session because we know that the impact of unconscious stop signals on behavior increases with practice (van Gaal et al., 2009, see also Verbruggen & Logan, 2008). By measuring EEG in the second session, we took advantage of this phenomenon. In the first two sessions, participants performed eight experimental blocks of the stop signal task. In the first session, one practice block was included. Each block of the stop task consisted of 30 unmasked stop trials, 30 unmasked go-on trials, 30 masked stop trials, and 30 masked go-on trials. The intertrial interval was jittered (2000–3000 msec in steps of 200 msec, drawn randomly from a uniform distribution) to minimize the effect of anticipation-related processes as well as very slow EEG oscillations (which are not of interest here) on the average ERP. Participants received performance feedback after every block (mean RT, standard deviation, percentage stops on unmasked stop trials) and were not informed about the presence of masked stop signals (or masked go-on signals).

Assessment of Stop Signal Visibility

In the third session, two tests were run to assess the subjective and objective visibility of stop signals. First, par-

ticipants performed one block of a dual task combining choice reaction with a yes–no detection task consisting of 120 trials (30 for each of the four conditions). This block was almost the same as a regular block presented in the two previous sessions, except that each trial was followed after 1000 msec by a pair of choices presented left (“stop”) and right (“no stop”) of fixation. To keep task demands as comparable with the stop task as possible, participants were instructed to respond twice on each trial; they had to respond as quickly as possible to the direction of the arrow, after which they had to determine whether they thought the word “stop” was presented in the preceding trial or not. There was no speed stress on the second (discrimination) response. On the second response, a new trial started.

After this task, participants performed three blocks of a two-alternative forced-choice (2-AFC) task directly aimed at gauging the detectability of the masked control signals. Each block consisted of 64 trials—32 masked stop trials and 32 masked go-on trials. Before running the 2-AFC discrimination task, participants were explained that words were also presented on masked trials in the original stop task (this was not the case in the preceding yes–no detection task). In addition, they were informed about the fact that in the upcoming task exactly half of the trials contained the word “stop” and the other half the control (go-on) word. Again, participants had to respond as fast as possible to the direction of the arrow. Thereafter, participants determined which of the two words was presented in the preceding trial. Each trial was followed after 1000 msec by a pair of choices presented left (“stop”) and right (control word) of fixation. There was no speed stress on the discrimination response. On the second response, a new trial started. In both detection tasks, SSDs of 129, 157, 186, and 229 msec were used. Note that participants were not instructed to inhibit their response on stop signals in both detection tasks.

Calculating SSRT

Performance on the stop signal paradigm can be described in terms of the horse race model (Logan, 1994). According to this model, two cognitive processes run independently while performing this task: a choice process and a stop process. The choice process starts upon presentation of the choice stimulus; the stop process starts slightly later, upon presentation of the stop signal. When the stop process wins the race from the choice process, the response will be inhibited. However, when the choice process is too fast to be caught up by the stop process, the response will be executed. The time it takes to complete the choice process is reflected in the response times to go-on trials. Because response times cannot be calculated on successfully inhibited stop trials, the time it takes to complete the stop process cannot be directly observed. However, when the response-time distribution on go-on trials and the percentage of inhibited stop trials are known, the SSRT can be estimated. The SSRT is an estimation of the duration of the stop process; the time it takes to implement inhibi-

tory control after presentation of the stop signal. It derives logically from the race model that those responses to the choice stimulus that are slower than the SSRT + SSD (the delay between the choice stimulus and the stop signal) will be inhibited, whereas responses faster than this measure will escape inhibition (Logan, 1994, see Figure 1B). SSRT was calculated by rank-ordering RTs on all go-on trials. Then, the n th percentile was selected, where n is the percentage of unmasked stop trials that is not inhibited, which in this experiment was on average 46% (but is determined on a per subject basis). The SSRT can be calculated by subtracting the average SSD from this value (Logan, 1994). For example, given that button-press responses could be withheld in approximately 54% of all unmasked stop trials (46% noninhibited stop trials), SSRT is calculated by subtracting the mean SSD from the 46th percentile of the go RT distribution (see Figure 1B).

Behavioral Data Analysis

Although not always observed (Emeric et al., 2007), participants tend to slow down after they failed to inhibit their response on a stop trial (Schachar et al., 2004; Rieger & Gauggel, 1999), an adaptive control mechanism referred to here as posterror slowing. Posterror slowing was measured by RTs on correct go-on trials immediately after failed stop trials compared with RTs on correct go-on trials immediately after correct go-on trials. Inhibition rates were computed over all trials *without* a response before the start of the next trial. For the RT analyses, RTs between 100 and 1000 msec were incorporated.

Repeated measures ANOVAs were performed on mean RT on correct masked go-on trials, mean RT on responded masked stop trials, SSRT, and square root percentage of responding on masked go-on trials and on masked stop trials with within-subjects' factors of Trial and Session. Detection performance (percentage correct) was tested for significance for each individual participant using a binominal test evaluated at a p value of .05 (two-tailed).

EEG Measurements

EEG was recorded and sampled at 256 Hz using a BioSemi ActiveTwo system (BioSemi, Amsterdam, the Netherlands). Forty-eight scalp electrodes were measured as well as four electrodes for horizontal and vertical eye movements (each referenced to their counterpart) and two reference electrodes on the ear lobes. After acquisition, the EEG data were referenced to the average of both ears and filtered using a high-pass filter of 0.5 Hz, a low-pass filter of 20 Hz, and a notch filter of 50 Hz (to be sure that 50 Hz caused by electrical power lines is entirely removed). Eye movement correction was applied on the basis of the horizontal and vertical EOG, using the algorithm of Gratton, Coles, and Donchin (1983). Thereafter, we applied artifact correction to all channels separately by removing segments outside the range of ± 50 μ V or with a voltage step exceeding

50 μ V per sampling point. Baseline correction was applied by aligning time series to the average amplitude of the interval from -300 to 0 msec preceding the onset of the stop or go-on signal onset. Note that by directly comparing the ERPs from onset of the stop signal with ERPs from onset of the go-on signal, we can isolate activity related to inhibition. On the contrary, choice signal locked ERPs are confounded by variations in SSD. All preprocessing steps were done with Brain Vision Analyzer (Brain Products GmbH, Munich, Germany). Statistical analysis (see below) was conducted using Matlab (MathWorks, Natick, MA).

EEG Analyses

To isolate activity related to the implementation of response inhibition, stop signal locked and go-on signal locked trials were compared directly. First, stop/go-on signal locked ERPs were calculated from the EEG data for all four conditions. Then, difference waveforms were computed by subtracting responded unmasked go-on trials from inhibited unmasked stop trials to isolate activity related to consciously triggered response inhibition. We will refer to this comparison as the *conscious inhibition contrast*. Similarly, to isolate activity related to unconsciously triggered response inhibition, difference waveforms were computed by subtracting responded masked go-on trials from responded masked stop trials, referred to as the *unconscious inhibition contrast*. All subsequent analyses were conducted on difference waves.

A review of the ERP literature indicated three ERP components of interest with different latencies and different topographical distributions (see Introduction). To zoom in on these specific components, three ROIs were defined at which these component generally tend to peak: an occipitoparietal ROI for the early negativity (Iz, Oz, O1, O2, POz, PO3, PO4, PO7, PO8), a fronto-central ROI for the N2 (Fz, F1, F2, FCz, FC1, FC2, Cz, C1, C2), and a centro-parietal ROI for the P3 (Cz, C1, C2, CPz, CP1, CP2, Pz, P1, P2). All ROIs consisted of nine electrode channels, which increases the signal-to-noise ratio. To calculate the precise time frame at which a component differed significantly from zero, we used sample-by-sample paired t tests (two-tailed) on the difference wave obtained from the conscious or the unconscious inhibition contrast. A significant interval was defined by the sequence of all bordering significant samples around the peak of interest. This was done for each component separately.

To test whether any of the components of interest was related to the stop performance, the correlation between ERP activity associated with conscious inhibition and SSRT was calculated. To this end, we calculated the mean amplitude of the difference wave of each of the three ERP components in its significant time interval (see Figure 3B). Then, Spearman's rank correlations (two-tailed) were computed between these measures and the SSRT. Similarly, a correlation between ERP activity associated with unconscious inhibition and RT slowing was calculated. Both behavioral measures were averaged across both sessions to provide the most reliable estimate.

Overall, all expected ERP components were observed in the data and peaked at the anticipated scalp locations. However, with respect to conscious inhibition, visual inspection of the electrophysiological differences between unmasked inhibited stop trials and unmasked go-on trials (see Figure 3A) revealed that the topographical distribution of the N2 was slightly more posterior than expected; it peaked at centro-parietal instead of fronto-central electrodes. The unconscious N2 peaked at the expected recording sites, the fronto-central ROI. Therefore, the size of the conscious as well as the unconscious N2 is reported for both the centro-parietal and the fronto-central ROI in the Results section. Generally, no qualitative differences between these outcomes were obtained. We intended to calculate the mean amplitude in the significant time window of the N2 (as well as the other components) as accurately as possible because these measures were used later to compute correlations between behavioral performance measures. Therefore, SSRT was correlated with the conscious N2 calculated for the centro-parietal ROI, and RT slowing was correlated with the unconscious N2 calculated for the fronto-central ROI.

RESULTS

Behavioral Performance

Fifteen of 19 participants scored at chance level in a 2-AFC detection task that we used to gauge the (in)visibility of masked stop signals. Because we cannot ascertain that the four participants who scored above chance level were truly unable to perceive masked stop signals consciously during the experiment, we excluded them from behavioral and electrophysiological analyses (see below for further details).

General performance measures are presented in Table 1. Participants performed proficiently on the task, as illustrated by typical inhibition rates of ~54%, while still responding fast to the choice stimulus (mean choice RT across both sessions was ~520 msec). The average SSRT (reflecting the efficiency of response inhibition) in the current paradigm was 315 msec in the first session and 302 msec in the second session. SSRTs were slightly longer than generally reported in nonselective stop signal tasks (e.g., Aron & Poldrack, 2006; Schmajuk et al., 2006) but comparable with previous studies using the selective stop signal paradigm (van Gaal et al., 2009; van den Wildenberg & van der Molen, 2004; Bedard et al., 2002; de Jong, Coles, & Logan, 1995). That SSRTs in the second session were shorter than that in the first session indicate that participants become slightly more proficient in inhibiting their responses to unmasked stop signals as a function of practice, $F(1, 14) = 3.25, p = .046$, one-tailed (see also Verbruggen & Logan, 2008).

Although participants did not stop significantly more often on masked stop trials than on masked go-on trials, $F(1, 14) = 2.23, p = .16$, they were significantly slowed down by masked stop signals compared with masked go-on signals. This was the case across sessions, $F(1, 14) = 19.39$,

Table 1. General Performance Measures in the Stop Signal Paradigm

<i>Behavioral Measure</i>	<i>Session 1</i>	<i>Session 2</i>
IR masked stop trial	2.83 (2.00)	0.14 (0.07)
IR masked go-on trial	1.78 (1.35)	0.06 (0.04)
IR unmasked stop trial	54.22 (1.63)	54.47 (1.18)
IR unmasked go-on trial	0.31 (0.14)	0.17 (0.10)
Conscious PES	27.44 (9.1)	15.46 (10.8)
Unconscious PES	-6.79 (3.35)	0.49 (3.85)
Mean SSD unmasked stop trials	184.35 (5.27)	183.04 (5.05)
SSRT	314.92 (6.30)	302.36 (4.81)

IR = inhibition rate (the percentage of inhibited trials); PES = posterror slowing; SSD = mean stop signal delay (msec); SSRT = stop signal reaction time. *SEM* values are reported within parentheses.

$p = .001$, but progressively more in the second session than that in the first, $F(1, 14) = 9.83, p = .007$ (see Figure 2A). Post hoc paired t tests revealed that masked stop signals slowed down responses in the first, $t(14) = 2.16; p = .049$, and especially the second session, $t(14) = 6.25; p < .001$. Thus, masked stop signals did not trigger complete response termination but did initiate a general slowing of responses times.

Because the stop signal is always presented after the choice stimulus, the stop process has to catch up with the choice process. According to the horse race model (for further details, see Methods), the SSRT plus the SSD represents the moment in time that the stop process wins from the choice process (“the point of no return,” see Figure 1B). The horse race model predicts that (conscious) stop signals have their largest impact on the slow end of the RT distribution (Logan, 1994). Thus, in our case, responses on unmasked stop trials slower than ~500 msec (SSRT + SSD, see Table 1) will likely be inhibited, whereas faster responses will probably not. Is this also the case for masked stop signals? If the impact of masked stop signals is also larger for slow responses (>500 msec) than for fast responses, this would further support the notion that inhibitory control mechanisms are triggered by masked and unmasked stop signals alike. Figure 2B shows the RT observations for the second session ranked from fast to slow responses for the masked stop as well as the masked go-on condition. Figure 2B illustrates that the difference between both conditions is relatively small before the “point of no return” but increases substantially after this point in time. This observation was confirmed by post hoc analyses showing that the difference between both masked conditions was significantly larger for the 50% slowest responses than for the 50% fastest responses, $t(14) = 7.08, p < .001$ (see Figure 2C). Whereas the 50% fastest responses differed only marginally between both masked conditions, $t(14) = 2.11, p = .053$, large differences were observed for the 50% slowest

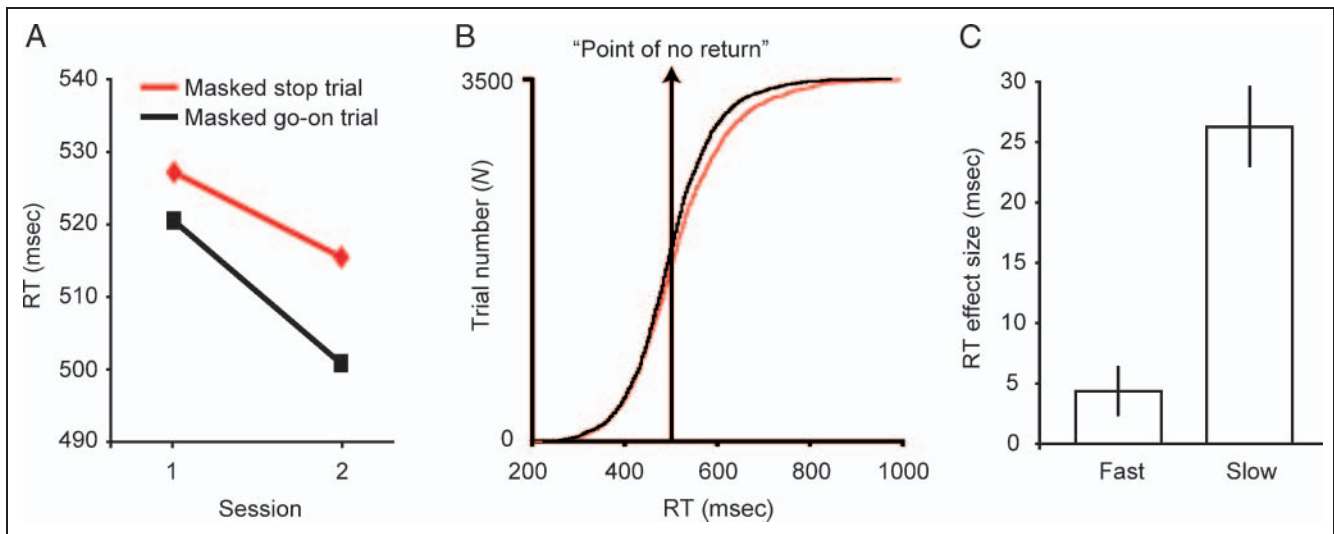


Figure 2. Masked stop signals slowdown responses. (A) Mean RT for masked stop trials and masked go-on trials. Participants responded significantly slower to masked stop trials than to masked go-on trials across sessions and for each session separately. (B) The RT distribution for masked stop trials and masked go-on trials for the second session. The RT difference between masked stop trials and masked go-on trials increases from the moment the stop process wins from the go process (from the vertical line representing the SSRT + SSD). The vertical line in this graph corresponds to the vertical line in Figure 1B. (C) In the second session, the difference between both masked conditions is significantly larger for the 50% slowest responses compared with the 50% fastest responses.

responses, $t(14) = 7.74, p < .001$. These results indicate that masked stop signals become fully operational in the slow part of the RT distribution (as is the case for unmasked ones), and when they do, they have a relatively large effect on the speed of responses (~26 msec).

In accordance with our previous behavioral study (van Gaal et al., 2009), conscious commission errors (failure to inhibit the response on an unmasked stop trial) led to considerable posterror slowing, $F(1, 14) = 7.00, p = .019$, whereas unconscious commission errors (failure to inhibit the response on a masked stop trial) did not, $F(1, 14) = 1.47, p = .25$ (see Table 1).

Taken together, unmasked as well as masked stop signals affected control processes, which led to complete response termination on many occasions when inhibitory control was triggered consciously and led to a considerable increase in response times when it was triggered unconsciously. This indicates that masked stop signals are capable of triggering inhibitory control mechanisms, but not as efficiently as conscious stop signals. These observations raise questions about commonalities and differences between consciously and unconsciously initiated inhibitory control mechanisms and their underlying neural substrates, which are dealt with in the next sections.

Electrophysiological Effects Related to Conscious Inhibition

In conducting ERP analyses, our first aim was to verify whether selective response inhibition in our stop signal task is associated with the same electrophysiological markers as observed in previous studies. To this end, we compared stop signal locked ERPs from successfully inhibited

stop trials with go-on signal locked ERPs from successfully responded go-on trials. Figure 3A shows the differential activity (stop minus go-on) between both conditions ($t = 0$ is the time of stop/go-on signal presentation). Note that the mean SSD on successfully inhibited stop trials (183 msec) is comparable with the mean SSD on responded go-on trials (188 msec). To this end, the degree to which the preceding choice stimulus contributes to the ERPs is similar. As expected, three electrophysiological events can be observed; the first at occipito-parietal electrodes (~200–300 msec), followed by a second (~300–340 msec) and a third event (380–600 msec) peaking at central electrodes (see numbers 1–3 in Figure 3A). Figure 3B shows the average ERP related to successful inhibition on stop trials compared with responding on go-on trials for the occipito-parietal, the fronto-central, and the centro-parietal ROI.

Conscious response inhibition was associated with an enhanced negative component at occipito-parietal recording sites (Figure 3B, left panel; number 1). At the occipito-parietal ROI, the peak difference between both conditions was observed 270 msec after stop signal presentation (peak difference = 5.73 μV), but sample-by-sample paired t tests revealed significant differences between 70 and 316 msec (see difference waves in blue; significant interval is indicated in black). In line with recent MEG (Boehler et al., 2008) and EEG (Schmajuk et al., 2006; Bekker et al., 2005) studies, this suggests enhanced visual processing of the relevant stop signal compared with the irrelevant go-on signal.

Somewhat later in time, the ERP to inhibited stop trials showed a sharp negative deflection, peaking at 309 msec after stop signal presentation at the centro-parietal ROI




Figure 3. Time course of activity associated with consciously initiated response inhibition. (A) Voltage scalp maps showing the spatio-temporal differences between the processing of unmasked inhibited stop trials and unmasked responded go-on trials (ERPs in response to unmasked go-on trials have been subtracted from ERPs on unmasked stop trials). Conscious response inhibition was associated with three neural events at different moments in time after stop signal presentation at different scalp locations (see numbers 1–3). (B) ERPs for unmasked inhibited stop trials and unmasked responded go-on trials for the occipito-parietal, the fronto-central, and the centro-parietal ROI. Difference waves are reported in blue, and the significant time window of each expected component is indicated in black. (C) Correlation between EEG activity and SSRT for each of the three components. (D) Spatial distribution of the significant positive correlation between the N2 and the SSRT.

(peak difference = 4.19 μV ; see Figure 3B, right panel, number 2). Sample-by-sample t tests performed on the difference wave revealed that the N2 component was significantly larger for stop trials than for go-on trials between 281 and 336 msec. Usually, if present, the N2 has a slightly more anterior topographic distribution and deviates

stronger from the 0 μV baseline than observed here (e.g., Schmajuk et al., 2006; Pliszka, Liotti, & Woldorff, 2000). Visual inspection of the difference maps of Figure 3A suggests that the early posterior negativity (70–316 msec) and the N2 (281–336 msec) are slightly overlapping in time, which might have incurred a slightly more posterior scalp

maximum and smaller magnitude for the N2. To be sure, the N2 was also significant at the fronto-central ROI between 313 and 328 msec; however, it was slightly smaller (peak difference = 2.72 μ V, peak latency = 320 msec; Figure 3A, middle panel, number 2). We would like to note that the same pattern of results was obtained using a pre-choice signal baseline instead of a prestop signal baseline.

The P3 component, arising after the N2, peaked at 445 msec after stop signal presentation and differed from go-on trials between 375 and 656 msec (peak difference = 8.86 μ V; Figure 3B, right panel, number 3). For the timing and scalp distribution, the P3 was very similar to stop P3 effects that were reported previously (e.g., Ramautar et al., 2004).

Correlations between EEG and SSRT

These components may reflect processes directly related to response inhibition or ancillary processes less directly related to response inhibition, such as visual processing, attentional selection, response selection, or response evaluation. To further examine the functional significance of the observed ERP components, we examined whether one (or more) of these neural events predicted the individual variability in stopping performance. More specifically, we correlated the average SSRT with the mean amplitude of the difference wave (see Figure 3B) in the significant time window of each of the three components across subjects. The mean amplitude of the N2 correlated positively with SSRT ($\rho = .53, p = .041$; Figure 3C). This indicates that participants with smaller SSRTs, who can be considered “good inhibitors,” display larger N2 components than “poor inhibitors.” To check the spatial specificity of this correlation, it was computed for all 48 measured electrode sites and plotted on a head map (see Figure 3D). The spatial profile of the observed correlations revealed a central distribution, nicely corresponding to the observed activation maps shown in Figure 3A (number 2).

Electrophysiological Effects Related to Unconscious Inhibition

Below we report the electrophysiological correlates of unconsciously initiated inhibitory control. More specifically, we were interested in which of the three components observed on unmasked stop trials are also present on masked stop trials. Figure 4A shows the differential activity between responded masked stop trials and responded masked go-on trials. Again, three electrophysiological events can be observed, peaking at occipito-parietal, centro-parietal, and fronto-central electrode sites. Figure 4B shows the actual ERPs elicited by responded masked stop trials compared with electrophysiological activity on responded masked go-on trials for all three ROIs.

At the occipito-parietal ROI, the neural processing of responded masked stop trials differed significantly from the

processing of responded masked go-on trials between 195 and 297 msec (peak difference = 0.90 μ V, peak latency = 223 msec; Figure 4B, left panel, number 1). At the fronto-central ROI, the N2 was significantly larger on masked stop trials than on masked go-on trials between 285 and 410 msec (peak difference = 2.30 μ V, peak latency = 336 msec; Figure 4B, middle panel, number 2). In the masked contrast, the N2 had a typical fronto-central topographical distribution. Because the N2 was peaking at more centro-parietal electrodes in the conscious contrast, we also tested the N2 effect for the centro-parietal ROI. At this ROI, the N2 was also significantly larger on masked stop trials than masked go-on trials; however, it was slightly smaller than at the fronto-central ROI (significant between 285 and 418 msec, peak difference = 1.74 μ V, peak latency = 348 msec; see Figure 4B, right panel). The centro-parietal P3 on masked stop trials was significantly larger than on masked go-on trials between 512 and 570 msec (peak difference = 0.99 μ V, peak latency = 551 msec; Figure 4B, right panel, number 3).

Correlations between EEG and Unconscious RT Slowing

Next, we analyzed whether the electrophysiological activity on masked stop trials is related to individual differences in the implementation of inhibitory control. To this end, the mean amplitude of the difference wave in each significant time interval (see Figure 4B) was correlated with the amount of slowing observed in response times (mean RT on masked stop trials minus mean RT on masked go-on trials). Based on the conscious inhibition results, one might expect that if any of the observed components would covary with unconscious RT slowing, it would be the N2. Indeed, this analysis revealed significant correlations for the N2 observed at the fronto-central ROI ($\rho = -.63, p = .012$). The correlation was also significant for the early activity observed at the occipito-parietal ROI ($\rho = -.54, p = .037$), but not for the P3 ($\rho = .25, p = .369$; Figure 4C). Again, the spatial profile of the correlations (see Figure 4D) nicely corresponded to the observed activity patterns (see Figure 4A, numbers 1 and 2).

Stop Signal Visibility

In a separate session, we checked whether participants could discriminate masked stop trials from masked go-on trials in a subjective (yes–no detection task) as well as an objective (2-AFC) measurement of stimulus visibility. In the yes–no detection task, participants detected 99.6% of the unmasked stop signals, whereas masked stop signals were never detected. This suggests that participants did not consciously perceive masked stop signals while performing the stop task. Before running the second, more conservative, 2-AFC discrimination task, participants were informed about the precise structure of the trials and were informed about




Figure 4. Time course of activity associated with unconsciously initiated response inhibition. (A) Voltage scalp maps showing the spatio-temporal differences between the processing of masked responded stop trials and masked responded go-on trials (ERPs in response to masked go-on trials have been subtracted from ERPs on masked stop trials). As with conscious inhibition, unconscious response inhibition was also associated with three neural events at different moments in time after stop signal presentation at different scalp locations (see numbers 1-3). (B) ERPs for masked responded stop trials and masked responded go-on trials for the occipito-parietal, the fronto-central, and the centro-parietal ROI, at which the expected components were observed to peak (see A). Difference waves are reported in blue, and the significant time window of each expected component is indicated in black. (C) Correlation between EEG activity and RT slowing for each of the three components. (D) Spatial distribution of the significant negative correlation between the early negativity and RT slowing and the N2 and RT slowing.

the presence of stop signals (and go-on signals) in all trials. In the 2-AFC, 15 of the 19 participants scored at chance level (binominal test). Because we cannot ascertain that the four participants who scored above chance level were truly unable to perceive masked stop signals consciously during the experiment, we excluded them from behavioral and electrophysiological analyses. For the included 15 participants, the mean percentage correct was 52.4% ($SD = 2.6$).

We performed several additional analyses to check whether the unconscious inhibition results could be explained by accidental visibility of masked stop signals. First, a correlational analysis demonstrated that there was no reliable correlation between stop signal visibility (percentage correct in the 2-AFC) and RT slowing ($\rho = .20, p = .49$). In addition, none of the three ERP components elicited by masked stop signals correlated with stop signal visibility (smallest $p > .65$). An additional argument for the invisibility of masked stop signals is that in this experiment as well as in a previous behavioral experiment (van Gaal et al., 2009), participants slowed down their responses after conscious errors, but not after unconscious errors. Such qualitative differences between the processing of unmasked versus masked stop signals implies the invisibility of masked stop signals (Merikle, Smilek, & Eastwood, 2001; Jacoby, 1991). Taken together, although one should be cautious in claiming unconsciousness of stimulus material, it seems that our behavioral as well as electrophysiological effects were not due to accidental visibility of masked stop signals.

DISCUSSION

We mixed unmasked (visible) and masked (invisible) stop signals in a stop task to study the neural activity related to the conscious versus unconscious initiation of inhibitory control. Due to inclusion of stop signals as well as go-on signals, four conditions were created: (1) an unmasked stop condition, (2) an unmasked go-on condition, (3) a masked stop condition, and (4) a masked go-on condition. EEG was measured to track and to compare the spatio-temporal processing of masked and unmasked stop signals in the human brain.

Participants performed the stop task proficiently, as evidenced by typical inhibition rates of ~50% on unmasked stop trials. Responses to masked stop trials were significantly slower than responses to masked go-on trials, as if participants tried to inhibit their response when a masked stop signal was presented but just failed to withhold it completely. Although present in both sessions, this RT effect was more pronounced in the second session than that in the first. This demonstrates that the impact of masked stop signals, like unmasked stop signals (as reflected in a decrease in SSRT across both sessions), increases with task exposure. Apparently, (masked) stop signals trigger inhibitory control more efficiently when stimulus–action associations are strong compared with when these associations

are recently formed and therefore relatively weak (see also Verbruggen & Logan, 2008). This is perfectly in line with previously proposed mechanisms of unconscious information processing, such as the direct parameter specification theory (Neumann, 1990), the action trigger theory (Kunde, 2003), or the evolving automaticity theory (Abrams & Greenwald, 2000). Yet, our results also reveal that extensive learning is not obligatory for unconscious influences on executive processes to unfold (see also van Gaal et al., 2009), as these were present from the first set of trials. In accordance with the predictions of the horse race model (Logan, 1994), the impact of masked stop signals was small on fast responses (~4 msec) but relatively large (~26 msec) on slow responses.

EEG recording revealed that successful inhibition on unmasked stop trials was associated with three ERP components previously associated with response inhibition in the stop signal paradigm (Boehler et al., 2008; Dimoska & Johnstone, 2008; Schmajuk et al., 2006; Bekker et al., 2005; Ramautar et al., 2004; van Boxtel et al., 2001; Pliszka et al., 2000; de Jong, Coles, Logan, & Gratton, 1990). Although all EEG components observed on masked stop trials resembled the corresponding components observed on successfully inhibited unmasked stop trials, several differences were observed. Below, crucial differences as well as commonalities between consciously and unconsciously inhibitory control are discussed.

Visual Processing of the Stop Signal

For one, unmasked inhibited stop signals elicited an early latency negative ERP component at occipito-parietal electrodes (compared with responded unmasked go-on trials). This finding nicely replicates recent EEG and MEG results that demonstrated that the quality of sensory processing of the stop signal, reflected in an early negative occipito-parietal ERP effect, is an important factor in predicting subsequent stopping success (Boehler et al., 2008; Schmajuk et al., 2006; Bekker et al., 2005). This notion is further supported by recent fMRI experiments that showed that successful stopping is associated with increased activity in early visual cortex compared with failed attempts to inhibit the response (Zheng et al., 2008; Aron & Poldrack, 2006; Li et al., 2006; Ramautar et al., 2006). In such a scheme, our data can be easily explained by assuming that stop signals have to be processed more elaborately than go-on signals, which in fact should be ignored and not further processed. Interestingly, a comparable occipito-parietal ERP component was observed on masked stop trials. Although this component was slightly smaller and less prominent, the topographical distribution and timing was highly similar. These results suggest that masked stop signals are (also) processed further and more elaborately than masked go-on trials, which seems to be a prerequisite for the subsequent initiation of control operations in the pFC, a process that might be reflected in the following anterior N2 component.

It should be noted that the conscious inhibition contrast revealed significant differences between 70 and 316 msec at the occipito-parietal ROI. At first sight, the first moment of significant deflection seems to arise relatively early compared with previous studies (Boehler et al., 2008; Schmajuk et al., 2006; Bekker et al., 2005). However, two of these studies (Boehler et al., 2008; Schmajuk et al., 2006) did not run sample-by-sample *t* tests to calculate the first moment of significant deflection but instead tested (a window around) the peak. Therefore, results cannot be compared directly. However, visual inspection of the early occipito-parietal differences reported in these studies suggests that activity differences also started to deviate from approximately 50–100 msec after stop signal presentation in these studies. A study that calculated the mean amplitude across time windows of 20 msec observed that the first negative component (the N1) to auditory stop signals was significantly larger for successful compared with failed inhibitions from 80 msec onward. In light of these previous findings, the present results suggest that the enhanced visual processing of stop signals compared with go-on signals (whether conscious or unconscious) may not only be due to more elaborate processing but also to the stronger processing of stop signals right from the start. This might be explained by subjects setting an attentionally guided sensory template for the stop signal, as if their sensory system is set in advance to selectively process the stop signal. This makes sense as the detection of the stop signal—and not the go-on signal—has behavioral consequences.

The Activation of Inhibitory Control

Response inhibition to unmasked stop trials was associated with two ERP components typically associated with response inhibition; the N2 and P3 component. Whether the N2 or the P3 reflects the “true” inhibition process remains controversial (for reviews, see Band & van Boxtel, 1999; Kok, 1986). In our study, the N2 component correlated with SSRT. Good inhibitors displayed larger N2 components than poor inhibitors, suggesting that it reflects a process related to inhibition. Although it has been shown previously that the N2 is related to inhibition (van Boxtel et al., 2001; Falkenstein, Hoormann, & Hohnsbein, 1999), to our knowledge, this is the first study that reports a correlation between the (conscious) N2 and the SSRT. The unconscious initiation of inhibitory control was associated with a distinct and relatively large fronto-central N2 together with a centro-parietal P3 that was sharply reduced in amplitude and duration compared with its conscious counterpart. The size of the unconscious N2 correlated with the degree to which inhibitory control was triggered by masked stop signals (RT slowing). Thus, the N2 correlated with the efficiency of conscious inhibitory control (SSRT) as well as the strength of the unconscious version of inhibition (RT slowing). Remarkably, in this study, the size of the P3 was not related to conscious as well as unconscious indices of inhibitory control.

Underlying Neural Mechanisms of Conscious versus Unconscious Control

How can these behavioral and electrophysiological effects of conscious and unconscious stop signals be explained? Here we argue that these results can be clarified by theories that differentiate between the role of feedforward and the role of recurrent processing in eliciting unconscious versus conscious vision (e.g., Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006; Lamme, 2006). When a visual stimulus is presented, it travels quickly from the retina through several stages of the cortical hierarchy, which is referred to as the fast feedforward sweep (Lamme & Roelfsema, 2000). Each time information reaches a successive stage in this hierarchy, this higher level area also starts to send information back to lower level areas through feedback connections. Single-cell recordings in monkeys (Super, Spekreijse, & Lamme, 2001) and TMS (Pascual-Leone & Walsh, 2001), fMRI (Haynes, Driver, & Rees, 2005), and EEG (Fahrenfort, Scholte, & Lamme, 2007) experiments in humans have revealed that the feedforward sweep probably remains unconscious, whereas recurrent interactions trigger awareness of a stimulus (for reviews, see Dehaene et al., 2006; Lamme, 2006). Interestingly, masking probably disrupts feedback activations but leaves feedforward activations relatively intact (Del Cul, Baillet, & Dehaene, 2007; Fahrenfort et al., 2007; Lamme, Zipser, & Spekreijse, 2002).

Unconscious stimuli are capable of triggering many forms of behavior (Lamme, 2006), as evidenced by many masked priming experiments (e.g., Vorberg et al., 2003; Dehaene et al., 1998) and patient studies (Stoerig & Cowey, 1997; Weiskrantz, 1996). A crucial aspect of the unconscious feedforward sweep is that it decays rapidly after traveling up the cortical hierarchy. In contrast, a key feature of recurrent interactions is that they promote widespread neural communication between distant brain areas, which initiates a long-lasting, large-scale pattern of neural activation, a phenomenon termed *global ignition* (Dehaene et al., 2006; Dehaene & Naccache, 2001). In EEG, global ignition as well as conscious access has been associated with a highly distributed fronto-parietal-temporal P3-like component (Del Cul et al., 2007).

In light of these ideas, one would have expected that masked stimuli evoke feedforward activation of the same cortical modules as are activated by unmasked stimuli, however, decaying rapidly and therefore weaker (Dehaene, 2008; van Gaal et al., 2008; Dehaene et al., 2001). This is supported by our finding that all three ERP components that are found in response to conscious stop signals are also found when stop signals are masked, albeit smaller and with different relative strength. It seems that both masked and unmasked stop signals trigger (basic) inhibition mechanisms, yet unconscious ones fail to elicit a comparably large, strong, and distributed pattern of activation observed when inhibition is triggered consciously. The spatial resolution of EEG is rather limited, but because it has been repeatedly demonstrated that conscious stop signals trigger a large

fronto-parietal inhibition network (for a review see Aron, 2007), we suggest that masked stop signals can probably also propagate to frontal and parietal cortex. In EEG, this process might be reflected in an enhanced fronto-central N2 component. However, as already suggested by Dehaene (2008), triggering of an information processor, even in frontal cortex, might not lead to global ignition, which could explain the largely absent P3 component (Del Cul et al., 2007), on masked stop trials. Obviously, the exact brain areas involved in unconsciously triggered inhibition should be verified with anatomically more accurate methods, such as fMRI.

Interestingly, others have demonstrated recently that inhibitory control in the stop signal paradigm does not necessarily lead to complete response inhibition but can also produce response slowing (Jahfari, Stinear, Claffey, Verbruggen, & Aron, 2010; Verbruggen & Logan, 2009). Verbruggen and Logan (2009) have demonstrated that when participants expect that a stop signal is presented in the upcoming trial, they proactively increase control and slowdown their go response to increase the likelihood of stopping success. This form of inhibitory control (“responding with restraint”) anyway activates inhibition-related neural networks (Jahfari et al., 2010), however, less strongly as full-blown response inhibition (Aron & Poldrack, 2006), which suggests that the extent to which inhibitory control is triggered can vary across situations.

In the present experiment, unconscious stop signals also seem to trigger inhibition-related neural networks partially (at least less than conscious stop signals), leading to response slowing instead of outright stopping. This seems to be in line with recent theoretical and modelling work concerning the race model (Boucher, Palmeri, Logan, & Schall, 2007). According to the original race model (Logan, 1994), two processes were thought to run independently while performing the stop task: a go process and a stop process. When the stop process wins the race, the response will be inhibited, when the go process wins, the response will be executed. The present data as well as previous work now suggest that the stop process and the go processes do not run entirely independently but interact (at the end) (Boucher et al., 2007), which can lead to response slowing, instead of either complete stopping or going (Jahfari et al., 2010; Verbruggen & Logan, 2009). Thus, the activation of inhibitory control does not necessarily lead to outright stopping but can also produce partial response suppression, either because the signal is not consciously processed (present data) or because the current task set requires it (Jahfari et al., 2010; Verbruggen & Logan, 2009).

In sum, we have shown that unconscious stop signals are able to trigger inhibitory control processes, reflected in a substantial slowdown of response execution. The presented data as well as current theorizing suggest that this form of inhibitory control may rely on fast feedforward activity traveling all the way up to pFC, however, only leading to “partial activation” of the inhibition network. On the contrary, full-blown, flexible, and efficient control (e.g., out-

right stopping) probably requires global recurrent interactions between inhibition-related brain areas (“strong activation” of the entire inhibition network). In that sense, unconscious cognitive control seems to differ substantially from traditional cognitive control processes in that it appears to be less efficient, less flexible, and less durable (Dehaene & Naccache, 2001).

Acknowledgments

We thank Roosmarijn Garben for her help with data acquisition. This work was supported by an advanced investigator grant from the European Research Council to VAF and a VICI grant from the Netherlands Organization of Scientific Research (NWO) to KRR.

Reprint requests should be sent to Simon van Gaal, Department of Psychology, University of Amsterdam, Roetersstraat 15, 1018 WB, Amsterdam, The Netherlands, or via e-mail: s.vangaal@uva.nl.

REFERENCES

- Abrams, R. L., & Greenwald, A. G. (2000). Parts outweigh the whole (word) in unconscious analysis of meaning. *Psychological Science, 11*, 118–124.
- Aron, A. R. (2007). The neural basis of inhibition in cognitive control. *Neuroscientist, 13*, 214–228.
- 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.
- Baayen, H., Piepenbrock, R., & Gulikers, L. (1995). *The CELEX lexical database* [CD-ROM]. Philadelphia: Linguistic Data Consortium, University of Pennsylvania.
- Band, G. P. H., & van Boxtel, G. J. M. (1999). Inhibitory motor control in stop paradigms: Review and reinterpretation of neural mechanisms. *Acta Psychologica, 101*, 179–211.
- Bar, M., & Biederman, I. (1999). Localizing the cortical region mediating visual awareness of object identity. *Proceedings of the National Academy of Sciences, U.S.A., 96*, 1790–1793.
- Bedard, A. C., Nichols, S., Barbosa, J. A., Schachar, R., Logan, G. D., & Tannock, R. (2002). The development of selective inhibitory control across the life span. *Developmental Neuropsychology, 21*, 93–111.
- Bekker, E. M., Kenemans, J. L., Hoeksma, M. R., Talsma, D., & Verbaten, M. N. (2005). The pure electrophysiology of stopping. *International Journal of Psychophysiology, 55*, 191–198.
- Boehler, C. N., Munte, T. F., Krebs, R. M., Heinze, H.-J., Schoenfeld, M. A., & Hopf, J.-M. (2008). Sensory MEG responses predict successful and failed inhibition in a stop-signal task. *Cerebral Cortex, 19*, 134–135.
- 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.
- Chambers, C. D., Bellgrove, M. A., Stokes, M. G., Henderson, T. R., Garavan, H., Robertson, I. H., et al. (2006). Executive “brake failure” following deactivation of human frontal lobe. *Journal of Cognitive Neuroscience, 18*, 444–455.
- de Jong, R., Coles, M. G., & Logan, G. D. (1995). Strategies and mechanisms in nonselective and selective inhibitory motor control. *Journal of Experimental Psychology: Human Perception and Performance, 21*, 498–511.

- de Jong, R., Coles, M. G., Logan, G. D., & Gratton, G. (1990). In search of the point of no return: The control of response processes. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 164–182.
- Dehaene, S. (2008). Conscious and Nonconscious processes: Distinct forms of evidence accumulation? In C. Engel & W. Singer (Eds.), *Decision making, the human mind, and implications for institutions. Strüngmann forum reports* (pp. 21–49). Cambridge, MA: MIT Press.
- Dehaene, S., Changeux, J.-P., Naccache, L., Le Clec'h, G., Koehlin, E., Mueller, M., et al. (1998). Imaging unconscious semantic priming. *Nature*, *395*, 597–600.
- Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences*, *10*, 204–211.
- Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, *79*, 1–37.
- Dehaene, S., Naccache, L., Cohen, L., Le Bihan, D., Mangin, J. F., Poline, J. B., et al. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience*, *4*, 752–758.
- Del Cul, A., Baillet, S., & Dehaene, S. (2007). Brain dynamics underlying the nonlinear threshold for access to consciousness. *PLoS Biology*, *5*, e260. doi:10.1371/journal.pbio.0050260.
- Dimoska, A., & Johnstone, S. J. (2008). Effects of varying stop-signal probability on ERPs in the stop-signal task: Do they reflect variations in inhibitory processing or simply novelty effects? *Biological Psychology*, *77*, 324–336.
- Dimoska, A., Johnstone, S. J., Barry, R. J., & Clarke, A. R. (2003). Inhibitory motor control in children with attention-deficit/hyperactivity disorder: Event-related potentials in the stop-signal paradigm. *Biological Psychiatry*, *54*, 1345–1354.
- Dimoska, A., Johnstone, S. J., & Barry, R. J. (2006). The auditory-evoked N2 and P3 components in the stop-signal task: Indices of inhibition, response-conflict or error-detection? *Brain and Cognition*, *62*, 98–112.
- Eimer, M. (1999). Facilitory and inhibitory effects of masked prime stimuli on motor activation and behavioural performance. *Acta Psychologica*, *101*, 293–313.
- Eimer, M., & Schlaghecken, F. (1998). Effects of masked stimuli on motor activation: Behavioral and electrophysiological evidence. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1737–1747.
- Eimer, M., & Schlaghecken, F. (2003). Response facilitation and inhibition in subliminal priming. *Biological Psychology*, *64*, 7–26.
- Emeric, E. E., Brown, J. W., Boucher, L., Carpenter, R. H. S., Hanes, D. P., Harris, R., et al. (2007). Influence of history on saccade countermanding performance in humans and macaque monkeys. *Vision Research*, *47*, 35–49.
- Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2007). Masking disrupts reentrant processing in human Visual Cortex. *Journal of Cognitive Neuroscience*, *19*, 1488–1497.
- Falkenstein, M., Hoormann, J., & Hohnsbein, J. (1999). ERP components in Go/Nogo tasks and their relation to inhibition. *Acta Psychologica*, *101*, 267–291.
- Gratton, G., Coles, M. G., & Donchin, E. (1983). A new method for off-line removal of ocular artifact. *Electroencephalography and Clinical Neurophysiology*, *55*, 468–484.
- Hannula, D. E., Simons, D. J., & Cohen, N. J. (2005). Imaging implicit perception: Promise and pitfalls. *Nature Reviews Neuroscience*, *6*, 247–255.
- Haynes, J. D., Driver, J., & Rees, G. (2005). Visibility reflects dynamic changes of effective connectivity between V1 and fusiform cortex. *Neuron*, *46*, 811–821.
- Holender, D., & Duscherer, K. (2004). Unconscious perception: The need for a paradigm shift. *Perception and Psychophysics*, *66*, 872–881.
- Hommel, B. (2007). Consciousness and control: Not identical twins. *Journal of Consciousness Studies*, *14*, 155–167.
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*, 513–541.
- Jahfari, S., Stinear, C., Claffey, M., Verbruggen, F., & Aron, A. R. (2010). Responding with restraint: What are the neurocognitive mechanisms? *Journal of Cognitive Neuroscience*, *22*, 1479–1492.
- Kok, A. (1986). Effects of degradation of visual-stimuli on components of the event-related potential (ERP) in go nogo reaction tasks. *Biological Psychology*, *23*, 21–38.
- Kouider, S., & Dehaene, S. (2007). Levels of processing during non-conscious perception: A critical review of visual masking. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *362*, 857–875.
- Kunde, W. (2003). Sequential modulations of stimulus-response correspondence effects depend on awareness of response conflict. *Psychonomic Bulletin & Review*, *10*, 198–205.
- Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, *10*, 494–501.
- Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, *23*, 571–579.
- Lamme, V. A. F., Zipser, K., & Spekreijse, H. (2002). Masking interrupts figure-ground signals in V1. *Journal of Cognitive Neuroscience*, *14*, 1044–1053.
- Lau, H. C., & Passingham, R. E. (2007). Unconscious activation of the cognitive control system in the human prefrontal cortex. *Journal of Neuroscience*, *27*, 5805–5811.
- Li, C.-S., Huang, C., Constable, R. T., & Sinha, R. (2006). Imaging response inhibition in a stop-signal task: Neural correlates independent of signal monitoring and post-response processing. *Journal of Neuroscience*, *26*, 186–192.
- Libet, B. (1999). Do we have free will? *Journal of Consciousness Studies*, *12*, 47–57.
- Logan, G. D. (1994). On the ability to inhibit thought and action: A users' guide to the stop signal paradigm. In D. D. T. H. Carr (Ed.), *Inhibitory processes in attention, memory and language* (pp. 189–239). San Diego, CA: Academic Press.
- Mattler, U. (2003). Priming of mental operations by masked stimuli. *Perception and Psychophysics*, *65*, 167–187.
- Mayr, U. (2004). Conflict, consciousness, and control. *Trends in Cognitive Sciences*, *8*, 145–148.
- Merikle, P. M., Smilek, D., & Eastwood, J. D. (2001). Perception without awareness: Perspectives from cognitive psychology. *Cognition*, *79*, 115–134.
- Neumann, O. (1990). Direct parameter specification and the concept of perception. *Psychological Research*, *52*, 207–215.
- Pascual-Leone, A., & Walsh, V. (2001). Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science*, *292*, 510–512.
- Pessiglione, M., Petrovic, P., Daunizeau, J., Palminteri, S., Dolan, J. D., & Frith, C. D. (2008). Subliminal instrumental conditioning demonstrated in the human brain. *Neuron*, *59*, 561–567.
- Pessiglione, M., Schmidt, L., Draganski, B., Kalisch, R., Lau, H., Dolan, R. J., et al. (2007). How the brain translates money into force: A neuroimaging study of subliminal motivation. *Science*, *316*, 904–906.

- Pisella, L., Grea, H., Tilikete, C., Vighetto, A., Desmurget, M., Rode, G., et al. (2000). An "automatic pilot" for the hand in human posterior parietal cortex: Toward reinterpreting optic ataxia. *Nature Neuroscience*, *3*, 729–736.
- Pliszka, S. R., Liotti, M., & Woldorff, M. G. (2000). Inhibitory control in children with attention-deficit/hyperactivity disorder: Event-related potentials identify the processing component and timing of an impaired right-frontal response-inhibition mechanism. *Biological Psychiatry*, *48*, 238–246.
- Ramautar, J. R., Kok, A., & Ridderinkhof, K. R. (2004). Effects of stop-signal probability in the stop-signal paradigm: The N2/P3 complex further validated. *Brain and Cognition*, *56*, 234–252.
- Ramautar, J. R., Slagter, H. A., Kok, A., & Ridderinkhof, K. R. (2006). Probability effects in the stop-signal paradigm: The insula and the significance of failed inhibition. *Brain Research*, *1105*, 143–154.
- Rieger, M., & Gauggel, S. (1999). Inhibitory after-effects in the stop signal paradigm. *British Journal of Psychology*, *90*, 509–518.
- Schachar, R. J., Chen, S., Logan, G. D., Ornstein, T. J., Crosbie, J., Ickowicz, A., et al. (2004). Evidence for an error monitoring deficit in attention deficit hyperactivity disorder. *Journal of Abnormal Child Psychology*, *32*, 285–293.
- Schmajuk, M., Liotti, M., Busse, L., & Woldorff, M. G. (2006). Electrophysiological activity underlying inhibitory control processes in normal adults. *Neuropsychologia*, *44*, 384–395.
- Stoerig, P., & Cowey, A. (1997). Blindsight in man and monkey. *Brain*, *120*, 535–559.
- Super, H., Spekreijse, H., & Lamme, V. A. F. (2001). Two distinct modes of sensory processing observed in monkey primary visual cortex (V1). *Nature Neuroscience*, *4*, 304–310.
- Umla, C. (1988). The control operations of consciousness. In A. J. Marcel & E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 334–356). Oxford: Oxford University Press.
- van Boxtel, G. J. M., van der Molen, M. W., Jennings, J. R., & Brunia, C. H. M. (2001). A psychophysiological analysis of inhibitory motor control in the stop-signal paradigm. *Biological Psychology*, *58*, 229–262.
- 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.
- van den Wildenberg, W. P. M., & van der Molen, M. W. (2004). Developmental trends in simple and selective inhibition of compatible and incompatible responses. *Journal of Experimental Child Psychology*, *87*, 201–220.
- van Gaal, S., Ridderinkhof, K. R., Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2008). Frontal cortex mediates unconsciously triggered inhibitory control. *Journal of Neuroscience*, *28*, 8053–8062.
- van Gaal, S., Ridderinkhof, K. R., van den Wildenberg, W. P. M., & Lamme, V. A. F. (2009). Dissociating consciousness from inhibitory control: Evidence for unconsciously triggered inhibitory control in the stop-signal paradigm. *Journal of Experimental Psychology: Human Perception and Performance*, *35*, 1129–1139.
- Verbruggen, F., & Logan, G. D. (2008). Automatic and controlled response inhibition: Associative learning in the go/no-go and stop-signal paradigms. *Journal of Experimental Psychology: General*, *137*, 649–672.
- Verbruggen, F., & Logan, G. D. (2009). Proactive adjustments of response strategies in the stop-signal paradigm. *Journal of Experimental Psychology: Human Perception and Performance*, *35*, 835–854.
- Vorberg, D., Mattler, U., Heinecke, A., Schmidt, T., & Schwarzbach, J. (2003). Different time courses for visual perception and action priming. *Proceedings of the National Academy of Sciences, U.S.A.*, *100*, 6275–6280.
- Weiskrantz, L. (1996). Blindsight revisited. *Current Opinion in Neurobiology*, *6*, 215–220.
- Zheng, D., Oka, T., Bokura, H., & Yamaguchi, S. (2008). The key locus of common response inhibition network for no-go and stop signals. *Journal of Cognitive Neuroscience*, *20*, 1434–1442.