Brain mechanisms of unconscious cognitive control

van Gaal, S.

Publication date
2009

4. Dissociable brain mechanisms underlying the conscious and unconscious control of behavior

Abstract

Cognitive control allows humans to overrule and inhibit habitual responses in order 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 distinct event-related brain potential (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 slow-down 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 vs. 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 in patients, e.g. with blindsight or neglect, or in healthy participants, e.g. by the use of masking, attentional blink, binocular rivalry or inattentional 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 (Bar & Biederman, 1999; Dehaene et al.,

2001). 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 (Dehaene et al., 1998; Vorberg et al., 2003).

Although at first controversial (for 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 et al., 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 (Dehaene & Naccache, 2001; Hommel, 2007; Libet, 1999; Mayr, 2004; Umilta, 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 (Eimer & Schlaghecken, 2003) and volition (Libet, 1999; Pisella et al., 2000) can also be initiated unconsciously (van Gaal et al., 2008; van Gaal, Ridderinkhof, van den Wildenberg, & Lamme, 2009). 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 slow-down 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 prefrontal cortex.

In a separate behavioral study, we tested whether stop-signal response inhibition could also be triggered unconsciously (van Gaal et al., 2009). Compared to

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 selective 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) was presented after the choice-stimulus (Logan, 1994), but not when a "go-on" signal was presented after the choice-stimulus. We included visible (unmasked) as well as invisible (masked) stop-signals. 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 prefrontal cortex (Aron & Poldrack, 2006; Chambers et al., 2006). 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 spatiotemporal dynamics of processing masked vs. unmasked stop-signals in the above-outlined selective stop-signal task as a first step towards answering these questions.

In EEG, successful stopping has typically been related to two ERP components: a frontocentral N2 component; a negative peak around 200-300 ms after stop-signal presentation (Dimoska, Johnstone, & Barry, 2003, 2006; Schmajuk, Liotti, Busse, & Woldorff, 2006) and a centroparietal P3 component, a positive peak around 300-500 ms after stop-signal presentation (Bekker, Kenemans, Hoeksma, Talsma, & Verbaten, 2005; Dimoska & Johnstone, 2008; 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 frontoparietal network involved in response inhibition, including middle, inferior and superior frontal cortices, pre-supplementary motor areas and the anterior cingulate cortex (Aron & Poldrack, 2006; Chambers et al., 2006; Ramautar, Slagter, Kok, & Ridderinkhof, 2006; Ray Li, Huang, Constable, & Sinha, 2006; Zheng, Oka, Bokura, & Yamaguchi, 2008). 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 MEG/EEG (Bekker et al., 2005; Boehler et al., 2008; Dimoska & Johnstone, 2008;

Schmajuk et al., 2006) as well as fMRI (Aron & Poldrack, 2006; Ramautar et al., 2006; Ray Li et al., 2006; Zheng et al., 2008) studies revealed a crucial role for sensory processing in response inhibition. For example Boehler et al., (2008) performed magnetoencephalographic recordings while subjects performed a stop task and demonstrated that enhanced (visual) processing of the go-signal facilitated response execution, whereas enhanced attentional allocation to the stop-signal was related to stopping success. This latter effect peaked approximately 160 ms after stop-signal presentation and was localized in the ventral occipitotemporal cortex. Additional evidence for the crucial role of sensory (visual) processing of stop-signals is obtained by Schmajuk et al., (2006). They compared stop trials in which stop-stimuli were task-relevant with a control block in which stop-stimuli were task-irrelevant and observed an enhanced early negative component over occipitoparietal cortex (peaking at ~200-220 ms after stop-signal presentation), which they suggested to reflect early sensory attention to the stop-signal when it is relevant compared to when it is irrelevant. Similarly, using auditory stop-signals, Bekker et al., (2005) observed an enhanced early auditory N1 component (peaking at ~100 ms after stop-signal presentation) for successful compared to failed stop trials. 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 it conscious counterpart.

Materials and Methods

Participants

19 undergraduate Psychology students participated in the experiment for course credits or financial compensation (12 female). 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 4.1a). 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. Bedard et al., 2002; Boehler et al., 2008; Dimoska et al., 2006; van den Wildenberg & van der Molen, 2004). 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, USA) against a black background (2.17 cd/m^2) at the centre of a 17-inch VGA monitor (frequency 70 Hz). Participants viewed the monitor from a distance of approximately 90 cm, so that each cm subtended a visual angle of 0.64 degrees. On masked stop trials, we first presented a white cross (300 ms) followed after 200 ms by a choice-stimulus (29 ms, isoluminant, 9.0 cd/m^2, width 0.64°, height 0.34°, which was either a blue left-pointing arrow or a red right-pointing arrow). This stimulus was followed after a variable SOA by two strings of randomly chosen uppercase consonants (forward masks, presented sequentially, 43 ms per letter string), the stop-signal or go-on signal (see below, 29 ms), and finally two consonant strings (backward masks, both 43 ms). 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 4.1a).

Participants were instructed to respond as quickly and 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, ten of 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 24).

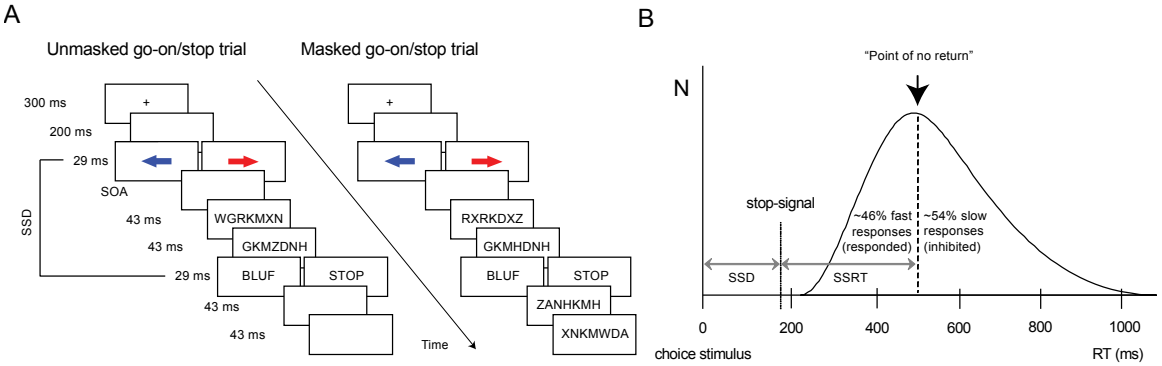

Figure 4.1 Stimulus timing in the masked selective stop-signal paradigm

(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).

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 stop-signal delay (SSD). After an inhibited unmasked stop trial, the SSD in the next trial increased by 14.3 ms, whereas it decreased by 14.3 ms 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 a SSD of 129 ms.

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). 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 inter-trial interval was jittered (2000-3000 ms in steps of 200 ms, drawn randomly from a uniform distribution) to minimize anticipation on the presentation of the choice-stimulus. Participants received performance feedback after every block (mean RT, SD, 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, participants performed one block of a dual task combining choice reaction with a yes-no detection task consisting of 120 trials; 30 for of 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 ms 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. Upon 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). Additionally, 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 ms by a pair of choices presented left (“stop”) and right (control word) of fixation. There was no speed stress on the discrimination response. Upon the second response a new trial started. In both detection tasks SSDs of 129, 157, 186 and 229 ms 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 stop-signal reaction time (SSRT) can be estimated. The SSRT is an estimation of the duration of the stop process; the time it takes to implement inhibitory 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 + the stop-signal delay (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 4.1b). 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% non-inhibited stop trials), SSRT is calculated by subtracting the mean SSD from the 46th percentile of the Go RT distribution (see Figure 4.1b).

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 (Rieger & Gauggel, 1999; Schachar et al., 2004), an adaptive control mechanism referred to here as post-error slowing. Post-error slowing was measured by RTs on correct go-on trials immediately following failed stop trials compared to RTs on correct go-on trials immediately following 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 ms and 1000 ms were incorporated.

Repeated-measures analyses of variance (ANOVA) 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 0.05 (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 was 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. 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 μV or with a voltage step exceeding 50 μV per sampling point. Baseline correction was applied by aligning time series to the average amplitude of the interval from -300 ms to the 0 ms 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, go signal locked ERPs are confounded by variations in SSD. All pre-processing steps were done with Brain Vision Analyzer (Brain Products GmbH, Munich, Germany). Statistical analysis (see below) was conducted using Matlab (The Mathworks, Natick, USA).

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-signal locked ERPs were calculated from the EEG data for all four conditions. Then, difference waveforms were computed by subtracting inhibited unmasked stop trials from responded unmasked go-on 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 stop trials from responded masked go-on 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 regions of interest (ROI's) 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 frontocentral ROI for the N2 (Fz, F1, F2, FCz, FC1, FC2, Cz, C1, C2) and a centroparietal ROI for the P3 (Cz, C1, C2, CPz, CP1, CP2, Pz, P1, P2). All ROI's 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 employed sample-by-sample paired t-tests (two-tailed) on the difference wave obtained from the conscious or 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 4.3b). 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 4.3a) revealed that the topographical distribution of the N2 was slightly more posterior than expected; it peaked at centroparietal, instead of frontocentral electrodes. The unconscious N2 peaked at the expected recording sites; the frontocentral ROI. Therefore, the size of the conscious as well as unconscious N2 is reported for both the centroparietal as well as frontocentral ROI in the results section. Generally, no qualitative differences between these outcomes 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 since these measures were used later to compute correlations between behavioral performance measures. Therefore, SSRT was correlated with the *conscious* N2 calculated for the centroparietal ROI and RT slowing was correlated with the *unconscious* N2 calculated for the frontocentral ROI.

Results

Behavioral performance

Fifteen out of nineteen participants scored at chance-level in a two-alternative forced-choice detection task that we used to gauge the (in)visibility of masked stop signals. Since 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 form behavioral and electrophysiological analyses (see below for further details).

General performance measures are presented in Table 4.1 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 ms). The average SSRT (reflecting the efficiency of response inhibition) in the current paradigm was 315 ms in the first session and 302 ms in the second session. SSRTs were slightly longer than generally reported in non-selective stop-signal tasks (e.g. Aron & Poldrack, 2006; Schmajuk et al., 2006), but comparable to previous studies using the selective stop-signal paradigm (Bedard et al., 2002; de Jong et al., 1995; van den Wildenberg & van der Molen, 2004; van Gaal et al., 2009). That SSRTs in the second session were shorter than in the first session, indicates that participants become slightly more proficient in inhibiting their responses to

unmasked stop-signals as a function of practice ($F(1,14) = 3.25$, $p = 0.046$, one-tailed).

Table 4.1

General performance measures in the stop-signal paradigm

Behavioral measure	*Session 1*	*Session 2*
Inhibition rate masked stop trial	2.83 (2.00)	0.14 (0.07)
Inhibition rate masked go-on trial	1.78 (1.35)	0.06 (0.04)
Inhibition rate unmasked stop trial	54.22 (1.63)	54.47 (1.18)
Inhibition rate unmasked go-on trial	0.31 (0.14)	0.17 (0.10)
Conscious post-error slowing	27.44 (9.1)	15.46 (10.8)
Unconscious post-error slowing	-6.79 (3.35)	0.49 (3.85)
Mean stop-signal delay	184.35 (5.27)	183.04 (5.05)
Stop-signal reaction time	314.92 (6.30)	302.36 (4.81)

Note: S.E.M. reported within brackets.

Although participants did not stop significantly more often on masked stop trials than on masked go-on trials ($F(1,14) = 2.23$, $p = 0.16$), they were significantly slowed down by masked stop-signals compared to masked go-on signals. This was the case across sessions ($F(1,14) = 19.39$, $p = 0.001$), but progressively more in the second session than in the first ($F(1,14) = 9.83$, $p = 0.007$, see Figure 4.2a]. Post-hoc paired t-tests revealed that masked stop-signals slowed down responses in the first ($t(14) = 2.16$; $p = 0.049$), and especially the second session ($t(14) = 6.25$; $p < 0.001$). Thus, masked stop-signals did not trigger complete response termination, but did initiate a general slowing of responses times.

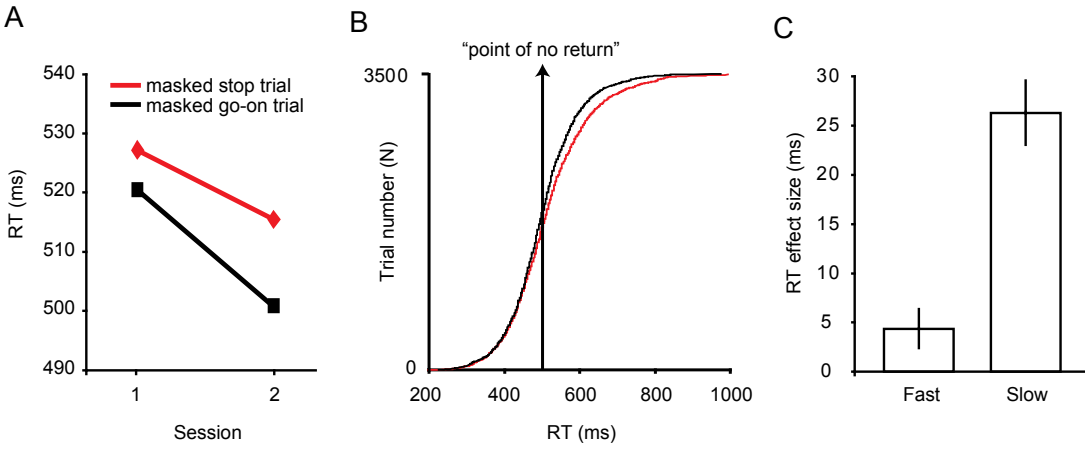

Figure 4.2 Masked stop-signals slow down responses

(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 4.1b. (c) In the second session, the difference between both masked conditions is significantly larger for the 50% slowest responses compared to the 50% fastest responses.

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 4.1b.). 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 ms (SSRT + SSD, see Table 4.1) 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 ms) than for fast responses, this would further support the notion that inhibitory control mechanisms are triggered by masked and unmasked stop-signals alike. Figure 4.2b 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. This figure 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 ($t(14) = 7.08$, $p < 0.001$, see Figure 4.2c). Whereas the 50% fastest responses differed only marginally between both masked conditions ($t(14) = 2.11$, $p = 0.053$), large differences were observed for the 50% slowest responses ($t(14) = 7.74$, $p < 0.001$). 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 ms).

In accordance with our previous behavioral study (van Gaal et al., 2009) conscious commission errors (failures to inhibit the response on an unmasked stop trial) led to considerable post-error slowing ($F(1,14) = 7.00$, $p = 0.019$), whereas unconscious commission errors (failures to inhibit the response on a masked stop trial) did not ($F(1,14) = 1.47$, $p = 0.25$; see Table 4.1).

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; however 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 4.3a shows the differential activity (go-on minus stop) between both conditions (t = 0 is the time of stop/go-on signal presentation). As expected, three electrophysiological events can be observed; the first at occipitoparietal electrodes (~200-300 ms), followed by a second (~300-340 ms) and third event (380-600 ms) peaking at central electrodes (see numbers 1-3 in Figure 4.3a). Figure 4.3b shows the average ERP related to successful inhibition on stop trials compared to responding on go-on trials for the occipitoparietal, the frontocentral as well as the centroparietal ROI.

Conscious response inhibition was associated with an enhanced negative component at occipitoparietal recording sites (Figure 4.3b, left panel; number 1). At the occipitoparietal ROI, the peak difference between both conditions was observed 270 ms after stop-signal presentation [peak difference 5.73 μV], but sample-by-sample paired t-tests revealed significant differences between 70 and 316 ms (see difference waves in blue, significant interval is indicated in black). In line with recent MEG (Boehler et al., 2008) and EEG (Bekker et al., 2005; Schmajuk et al., 2006) studies, this suggests enhanced visual processing of the relevant stop-signal compared to the irrelevant go-on signal.

Somewhat later in time, the ERP to inhibited stop trials showed a sharp negative deflection, peaking at 309 ms after stop-signal presentation at the centroparietal ROI (peak difference 4.19 μV, see Figure 4.3b, 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 ms. Usually, if present, the N2 has a slightly more anterior topographic distribution than observed here (e.g. Pliszka, Liotti, & Woldorff, 2000; Schmajuk et al., 2006). Visual inspection of the difference maps of Figure 4.3a suggests that the

early posterior negativity (70-316 ms) and the N2 (281-336 ms) are slightly overlapping in time, which might have incurred a slightly more posterior scalp maximum for the N2. To be sure, the N2 was also significant at the frontocentral ROI between 313-328 ms, however it was slightly smaller (peak difference 2.72 µV, peak latency 320 ms, middle panel, number 2).

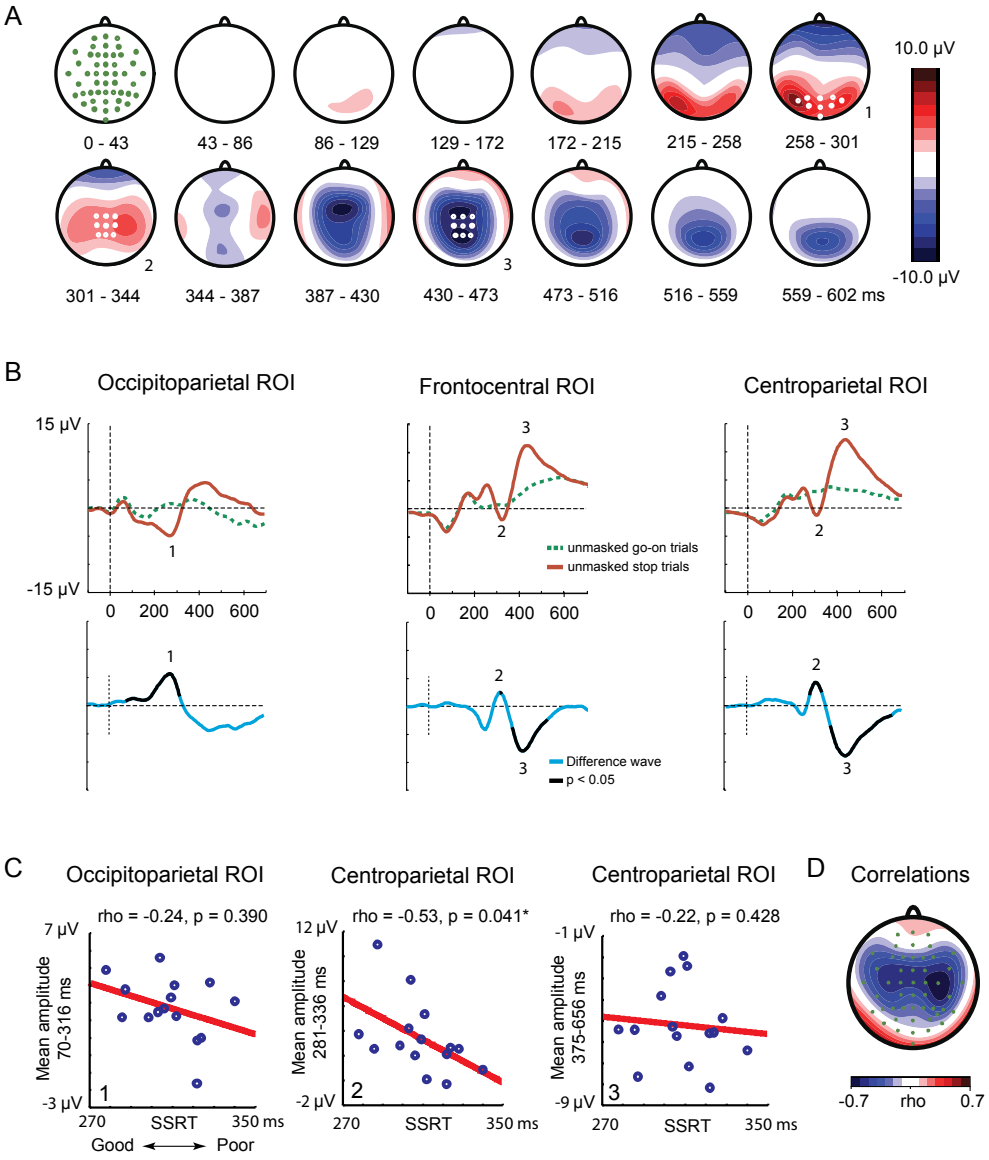

Figure 4.3 Time course of activity associated with consciously initiated response inhibition

a) Voltage scalp maps showing the spatiotemporal differences between the processing of unmasked inhibited stop trials and unmasked responded go-on trials (ERPs in response to the unmasked stop trial have been subtracted from ERPs on unmasked go-on 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 occipitoparietal, the frontocentral and the centroparietal ROI. Difference waves are reported in blue, 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 negative correlation between the centroparietal N2 and SSRT.

The P3 component, arising after the N2, peaked at 445 ms after stop-signal presentation and differed from go-on trials between 375 and 656 ms (peak difference 8.86 μV, Figure 4.3b, right panel, number 3). 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 4.3b) in the significant time-window of each of the three components across subjects. The mean amplitude of the N2 correlated negatively with SSRT (rho = -0.53, p = 0.041, Figure 4.3c). 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 4.3d.) The spatial profile of the observed correlations revealed a central distribution, nicely corresponding to the observed activation maps shown in Figure 4.3a (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 4.4a shows the differential activity between responded masked stop trials and responded masked go-on trials. Again, three electrophysiological events can be observed, peaking at occipitoparietal, centroparietal and frontocentral electrode sites. Figure 4.4b shows the actual ERPs elicited by responded masked stop trials compared to electrophysiological activity on responded masked go-on trials for all three ROI's.

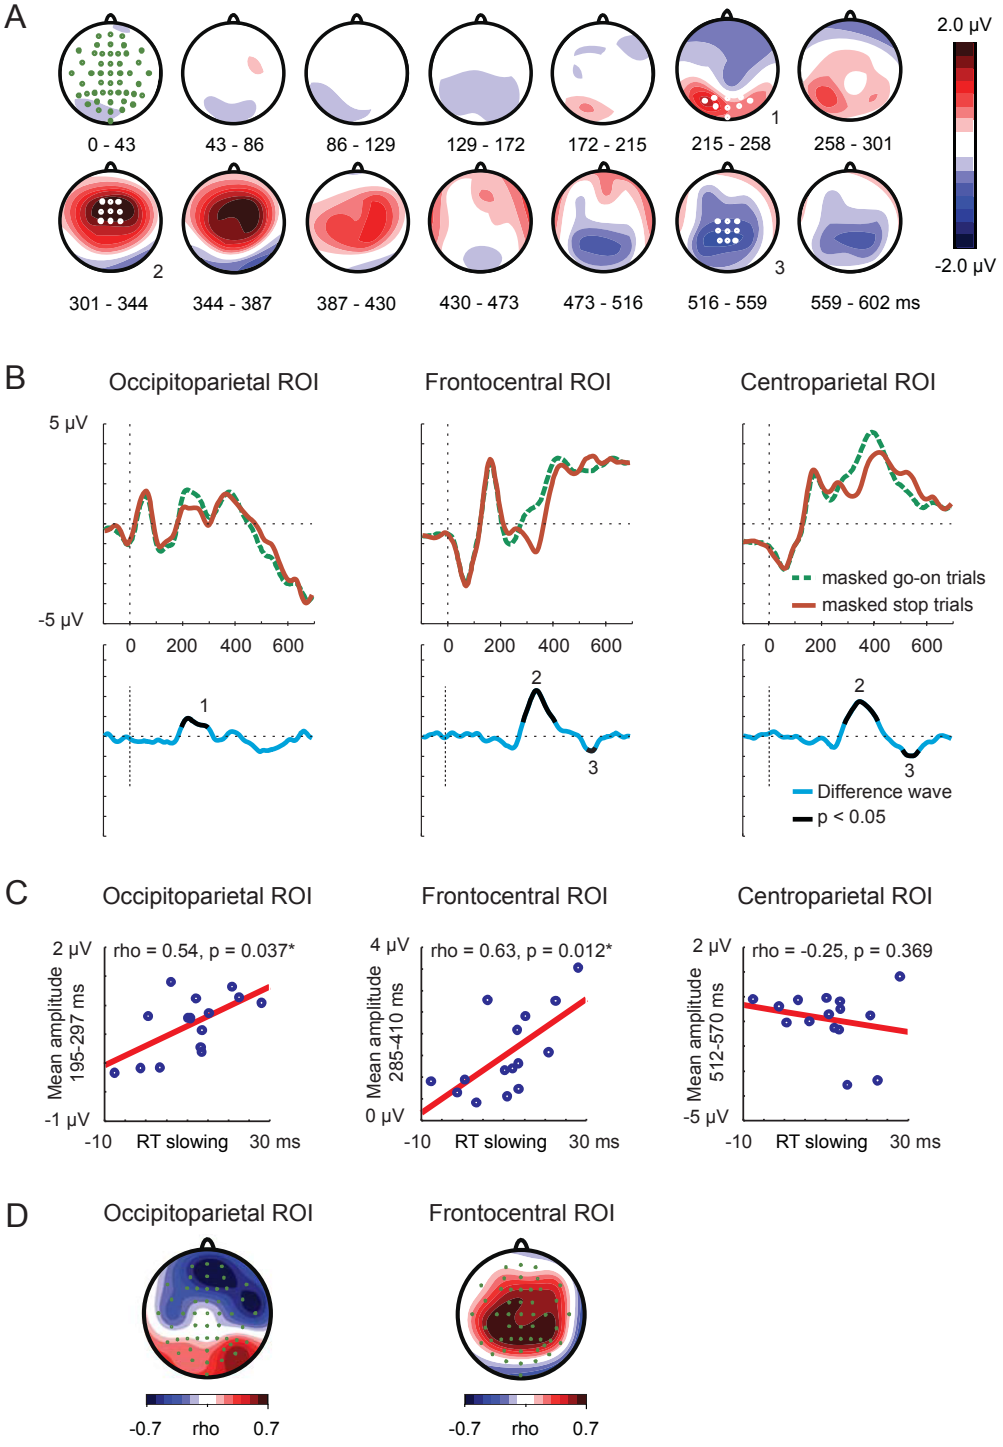

Figure 4.4 Time course of activity associated with unconsciously initiated response inhibition

a) Voltage scalp maps showing the spatiotemporal differences between the processing of masked responded stop trials and masked responded go-on trials (ERPs in response to the masked stop trial have been subtracted from ERPs on masked go-on 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). b) ERPs for masked responded stop trials and masked responded go-on trials for the occipitoparietal ROI, the frontocentral ROI and the centroparietal ROI, at which the expected components were observed to peak (see Figure 4.4a). Difference waves are reported in blue, 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 positive correlation between the occipitoparietal negativity and RT slowing and the frontocentral N2 and RT slowing.

At the occipitoparietal ROI, the neural processing of responded masked stop trials differed significantly from the processing of responded masked go-on trials between 195-297 ms (peak difference 0.90 μV, peak latency 223 ms, Figure 4.4b, left panel, number 1). At the frontocentral ROI, the N2 was significantly larger on masked stop trials than on masked go-on trials between 285 and 410 ms (peak difference 2.30 μV, peak latency 336 ms, Figure 4.4b, middle panel, number 2). In the masked contrast the N2 had a typical frontocentral topographical distribution. Because the N2 was peaking at more centroparietal electrodes in the conscious contrast, we also tested the N2 effect for the centroparietal 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 frontocentral ROI (significant between 285-418 ms, peak difference 1.74 μV, peak latency 348 ms, see Figure 4.4b, right panel). The centroparietal P3 on masked stop-trials was significantly larger than on masked go-on trials between 512 and 570 ms (peak difference 0.99 μV, peak latency 551 ms, Figure 4.4b, 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 4.4b) 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 frontocentral ROI (rho = 0.63, p = 0.012). The correlation was also significant for the early activity observed at the occipitoparietal ROI (rho = 0.54, p = 0.037), but not for the P3 (rho = -0.25, p = 0.369; Figure 4.4c). Again, the spatial profile of the correlations (see Figure 4.4d) nicely corresponded to the observed activity patterns (see Figure 4.4a, number 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 the presence of stop-signals (and go-on signals) in all trials. In the 2-AFC, fifteen out of the nineteen participants scored at chance-level (binominal test). Since 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 fifteen 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 = 0.20, p = 0.49). Additionally, none of the three ERP components elicited by masked stop-signals correlated with stop-signal visibility (smallest p > 0.65). 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 vs. masked stop-signals implies the invisibility of masked stop-signals (Jacoby, 1991; Merikle et al., 2001). 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 vs. 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 compare the spatiotemporal 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. Though present in both sessions, this RT effect was more pronounced in the second session than 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 to when these associations are recently formed, and therefore relatively weak. 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 (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 ms), but relatively large (~26 ms) 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 (Bekker et al., 2005; Boehler et al., 2008; de Jong, Coles, Logan, & Gratton, 1990; Dimoska & Johnstone, 2008; Pliszka et al., 2000; Ramautar et al., 2004; Schmajuk et al., 2006; van Boxtel et al., 2001). 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 occipitoparietal electrodes (compared to 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 occipitoparietal ERP effect, is an important factor in predicting subsequent stopping success (Bekker et al., 2005; Boehler et al., 2008; Schmajuk et al., 2006). This notion is further supported by recent fMRI experiments that showed that successful stopping is associated with increased activity in early visual cortex compared to failed attempts to inhibit the response (Aron & Poldrack, 2006;

Ramautar et al., 2006; Ray Li et al., 2006; Zheng et al., 2008). 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 occipitoparietal 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 prefrontal cortex; 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 ms at the occipitoparietal ROI. At first sight, the first moment of significant deflection seems to arise relatively early compared to previous studies (Bekker et al., 2005; Boehler et al., 2008; Schmajuk et al., 2006). 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 occipitoparietal differences reported in these studies suggests that activity differences also started to deviate from approximately 50-100 ms after stop-signal presentation in these studies. A study that calculated the mean amplitude across time windows of 20 ms observed that the first negative component (the N1) to auditory stop-signals was significantly larger for successful compared to failed inhibitions from 80 ms onwards. In light of these previous findings, the present results suggest that the enhanced visual processing of stop-signals compared to 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 negatively with SSRT. Thus 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 (Falkenstein, Hohnsbein, & Hoormann, 1999; van Boxtel et al., 2001), to our knowledge, this is the first study that reports a negative correlation between the (conscious) N2 and SSRT. The unconscious initiation of inhibitory control was associated with a distinct and relatively large frontocentral N2 together with a centroparietal P3 that was sharply reduced in amplitude and duration compared to its' conscious counterpart. The size of the unconscious N2 correlated positively 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 vs. unconscious control

How can these behavioural 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 recurrent processing in eliciting unconscious vs. conscious vision (e.g. Dehaene et al., 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 sent 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 et al., 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 et al., 2007; Fahrenfort et al., 2007; Lamme et al., 2002).

Unconscious stimuli are capable of triggering many forms of behavior (Lamme, 2006), as evidenced by many masked priming experiments (e.g. Dehaene et al., 1998; Vorberg et al., 2003) and patient studies (Stoerig & Cowey, 1997; Weiskrantz, 1996). A crucial aspect of the unconscious feedforward sweep is that is decays rapidly after travelling 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 not triggering global ignition (Dehaene, 2008; Dehaene & Naccache, 2001; van Gaal et al., 2008). 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 frontoparietal 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 frontocentral 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.

This neural account appears to match with the observed behaviour. Whereas unmasked (visible) stop-signals were capable of triggering response inhibition to the level of complete response termination on many occasions, inhibition rates on masked (invisible) stop trials were low and they only caused a slowing of responses. Thus, in line with current theories, we suggest that although task interruption can be triggered unconsciously by fast feedforward activation, this inhibition process almost never runs to completion due to the absence of global recurrent interactions (see also Dehaene, 2008). The available neural data as well as current theorizing suggests that full-blown, flexible and efficient control requires global recurrent interactions, whereas more automatic cognitive control processes may rely on feedforward activity. 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).