



UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Brain mechanisms of unconscious cognitive control

van Gaal, S.

Publication date
2009

[Link to publication](#)

Citation for published version (APA):

van Gaal, S. (2009). *Brain mechanisms of unconscious cognitive control*.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

2. Dissociating consciousness from inhibitory control: Evidence for unconsciously triggered response inhibition in the stop-signal task

Abstract

Theories about the functional relevance of consciousness commonly posit that higher-order cognitive control functions, such as response inhibition, require consciousness. To test this assertion, we designed a masked stop-signal paradigm to examine whether response inhibition could be triggered and initiated by masked stop-signals, a signal informing a participant to stop an action they have begun. In two experiments, masked stop-signals were observed to occasionally result in full response inhibition as well as yield a slow-down in the speed of responses that were not inhibited. The magnitude of this subliminally triggered RT slowing effect correlated with the efficiency measure (stop-signal reaction time) of response inhibition across subjects. Thus we show that response inhibition can be triggered unconsciously, the more so in individuals and under conditions that are associated with efficient response inhibition. These results indicate that, in contradiction to common theorizing, inhibitory control processes can operate outside awareness.

Introduction

Many perceptual and motor processes can occur in the absence of consciousness, as evidenced by recent subliminal priming reports (Dehaene et al., 1998; Eimer & Schlaghecken, 1998; Kunde, 2003; Vorberg et al., 2003) and patient studies (Driver & Mattingley, 1998; Stoerig & Cowey, 1997; Weiskrantz, 1996). Apparently, a substantial amount of processing can occur unconsciously and yet affect behavior. However, the question which specific cognitive processes can be triggered by unconscious information, and which cannot, is subject to a lively debate (Dehaene & Naccache, 2001; Eimer & Schlaghecken, 2003; Jacoby, 1991; Kunde, 2003; Mayr, 2004; Nieuwenhuis et al., 2001; Rossetti, Pisella, & Vighetto, 2003). Many authors have suggested that cognitive control functions, associated with prefrontal cortices, such as conflict detection and response inhibition, require consciousness (Baars, 2002; Dehaene & Naccache, 2001; Eimer & Schlaghecken, 2003; Libet, 1999; Rossetti et al., 2003; Tsushima, Sasaki, & Watanabe, 2006; Umiltà, 1988). Cognitive control processes can be defined as those processes that regulate and monitor ongoing actions to optimize goal-directed behavior. The logic behind the consciousness-control relationship is the idea that we usually become aware of stimuli that interfere or interrupt routine action, which are the same stimuli that call for adaptive control operations. Therefore, it has been proposed that higher-level control

operations, such as response inhibition, depend on the conscious detection of response-relevant warning signals (Dehaene & Naccache, 2001; Eimer & Schlaghecken, 2003). Following this line of reasoning, it should not be possible to trigger inhibitory control processes when the instruction stimulus itself is presented subliminally. In two studies we put this claim to a direct test.

The ability to cancel an already initiated response is an important cognitive control function that allows people to meet complicated task demands and quickly adapt to environmental changes. One experimental measure of response inhibition is provided by the stop-signal paradigm (Logan, 1994). The stop-signal task provides a direct behavioral test of the ability to stop a planned or ongoing action in a voluntary fashion. In our adapted version of the task, participants perform speeded right- or left-hand responses to go-signals. On a small proportion of those trials, the onset of the go-signal is followed by a stop-signal, instructing the subjects to refrain from responding. 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. By varying the stop-signal delay, the stop-signal task yields an estimate of the duration of the inhibitory process, the stop-signal reaction time (SSRT). The SSRT can be used to compare the efficiency of inhibitory control processes between conditions or individuals.

To examine the role of consciousness in response inhibition, we conducted two experiments in which we masked stop-signals optimally, which caused stop-signals to be invisible, and sub-optimally, which caused stop-signals to be visible. For simplicity, we call optimally masked stop-signals *masked* stop-signals and sub-optimally masked stop-signals *non masked* stop-signals. We reasoned that if response inhibition could be triggered unconsciously, this could result in small differences in inhibition rates on stop trials containing a masked stop-signal compared to inhibition rates on go trials, trials without a stop-signal. Additionally, we expected response times (RTs) on stop trials containing a masked stop-signal (escaping inhibition) to be slower than RTs on go trials due to the triggering of inhibitory control processes (however, not to the level of complete response termination). Furthermore, participants with relatively short SSRTs (reflecting efficient response inhibition on stop trials containing a *non masked* stop-signal) could be more susceptible to the influence of *masked* stop-signals, and hence show greater RT slowing effects to masked stop-signals than those with slower SSRTs. By using this experimental set-up we were able to show that inhibitory control can be

triggered unconsciously, contradicting the proposed direct association between inhibitory control functions and consciousness.

Experiment 1

Materials and methods

Participants

Twenty-four students from the University of Amsterdam participated in this study (16 female; age ranging from 18 to 25). The local ethical committee approved the experiment, and subjects gave written informed consent before experimentation. All participants had normal or corrected-to-normal vision.

Design and materials

Stimuli were presented on a grey frame (59.10 cd/m², visual angle of 3.78°) against a black background (2.17 cd/m²) at the center of a 15-inch BenQ TFT monitor. Trials started with a 1000 ms blank screen, after which a go-signal appeared. The primary Go-task involved the discrimination between two isoluminant (16.84 cd/m²) colored circles (blue and green, visual angle of 0.60°), mapped onto responses with the left and right hand respectively (thumb responses). Go-signals appeared pseudo randomly with equal frequency. During stop trials, a grey circle, the stop-signal (41.85 cd/m², visual angle of 0.60°), appeared after the presentation of the go-signal. The stop-signal delay (SSD) between the go-signal and the stop-signal was one of five equiprobable intervals (100, 200, 300, 400 or 500 ms). After presentation of a stop-signal a black circular nonoverlapping metacontrast mask was presented (2.17 cd/m², visual angle of 1.30°). Go-signal and stop-signal duration was 16.7 ms, mask duration was 150 ms (see Figure 2.1). The above-mentioned parameters (target duration, mask duration, stimulus contrasts) typically result in a linear masking curve, which means that masking is most effective with short SOAs between targets and metacontrast masks (Di Lollo, von Muhlenen, Enns, & Bridgeman, 2004; Francis, 1997). In this experiment, the SOA between the stop-signal and the mask could either be short (16.7 ms) or long (100, 150 and 200 ms, equal frequency). When the SOA between the stop-signal and the mask was short, the stop-signal was masked optimally and could not be perceived consciously (as revealed by a post-experimental forced-choice detection task). However, when the SOA between the stop-signal and the mask was long, the stop-signal was masked sub-optimally, which caused stop-signals to be visible. We will refer to stop trials with a *short* stop-signal-mask SOA as *short* stop trials, and to stop trials with a *long* stop-signal-mask SOA as

long stop trials. Short stop trials were matched to short go trials to reliably compute performances measures. In short go trials, stimulus timing between the go-signal and the mask was exactly the same as in short stop trials, since a blank screen was presented instead of a stop-signal on these trials. By this means, the only difference between these trials is that short stop trials contain an additional stop-signal, whereas short go trials do not. Similarly, long stop trials were matched to long go trials (see Figure 2.1). Trials were presented in a random order with equal frequency (all 25%).

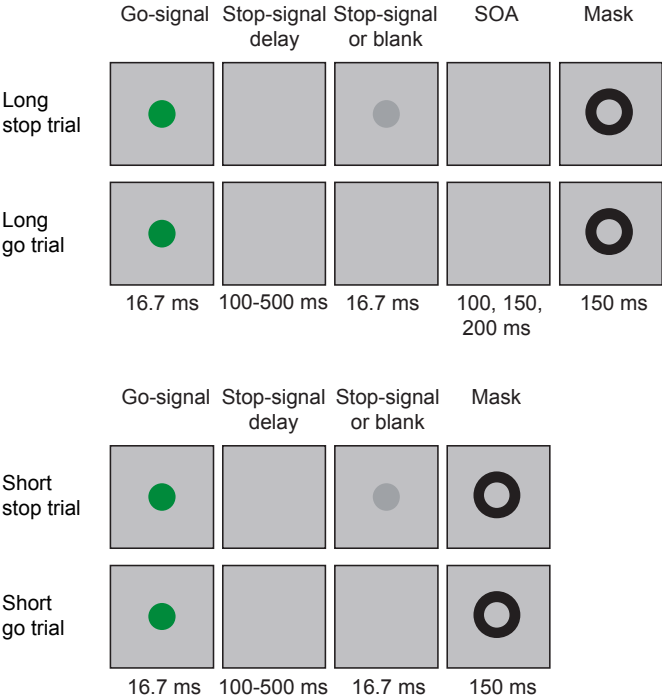


Figure 2.1 Stimulus timing of the masked stop-signal paradigm

Participants did three experimental sessions held on separate days on three consecutive weeks in a row. Total test duration was 5 hours for each participant (first session two hours, second session one hour, third session two hours). Experimental blocks consisted of 120 trials. The inter-trial interval was jittered (2500-3000 ms in steps of 100 ms, pseudo randomly) to minimize anticipation on the presentation of the go-signal. Before the first experimental blocks, participants performed one block of choice RT only (80 go trials) and one practice block of the stop-signal task. Participants performed 840 to 960 trials in each experimental session. Performance feedback was given after each block (Mean RT, standard deviation, go error rate, inhibition rate on go trials, and inhibition rate on long stop trials). Participants were not informed about the presence of trials containing a

masked stop-signal and did not receive any feedback about performance on these trials during testing. To equalize the inhibition rate on long stop trials between participants, the experimenter monitored participants' performance by stressing either the primacy of responding or inhibiting, aiming at an overall inhibition rate on long stop trials of approximately one third.

Forced-choice detection task

To obtain an objective measure of stop-signal visibility participants performed five blocks of a forced-choice detection task with single-trial error feedback after the final experimental block in the third session. None of the participants reported to be aware of the masked stop-signals before being informed about them at the beginning of forced-choice detection. Trial timing was exactly the same as in the experimental sessions, except that visual feedback (the words "goed" or "fout", Dutch for right or wrong) was centrally presented after every trial (for 1000 ms) in uppercase, Arial font in black color (visual angle of 0.45°). Each block consisted of 40 short stop trials, 40 short go trials (400 trials in total) and to check whether subjects were fully attending during the detection task 10 long stop trials were included (stop-signal-mask SOAs of 100 and 150 ms only). Participants were instructed to ignore go-signals and press left if there was a stop-signal presented and right if there was no stop-signal presented. Detection responses were to be given after mask presentation. Before running the detection task participants were fully informed about the precise structure of the trials. Additionally, participants were informed about the fact that this task contained 10 "easy" trials (long stop trials) and 80 "difficult" trials (40 short go trials and 40 short stop trials) and that exactly half of the difficult trials contained a stop-signal. Summary feedback (percentage correct) was given after every block for easy and difficult trials separately. Between blocks participants were encouraged to perform the best they could, even though they were unable to detect the stop-signals.

Data analysis

This is a quite difficult task, especially because stop-signals in long stop trials are pretty difficult to perceive at the beginning. Therefore participants will make false alarms on go trials, which means that they sometimes stop on trials on which no stop-signal is presented (but participants thought that there was one presented). To reliably measure unconsciously triggered response inhibition, *unconscious* response inhibition is quantified in terms of a relative inhibition rate on *short* stop trials compared to *short* go trials. These trials are perceptually similar (see Results) and

the timing between go-signal and mask presentation is exactly the same. So, if subjects inhibit more often on short stop trials compared to short go trials this means that masked stop-signals triggered response inhibition on these trials. This yields a more conservative measure than just comparing the number of inhibited trials on short stop trials with zero as baseline. The same logic accounts for RT slowing. If participants are slower on responded short stop trials than on responded short go trials this means masked stop-signals triggered inhibitory control, however not to the level of complete response termination. Effect sizes (reported in figures) were calculated by subtracting the inhibition rates or mean RTs on short go trials from the inhibition rates or mean RTs on short stop trials. All responses before the appearance of the next trial are incorporated in our RT analysis. Inhibition rates are computed by taking all trials *without* a response. When no response was given before the appearance of the next trial, this trial was considered to be inhibited.

Repeated measures analyses of variance (ANOVA) were performed on mean RT on correct short go trials, responded short stop trials, SSRT and square root inhibition rates on short go trials and on short stop trials with within-subjects' factors of Trial and Session. The behavioral data of the second session of one subject was lost due to technical problems. Group averages were used for this participants' second session in the analyses. Detection performance was analyzed by signal detection methods, a bias-free measure of stimulus visibility, and tested for significance using a one-sample t-test.

For each behavioral session, an estimate of the SSRT of each single subject was calculated. The SSRT can be estimated by collapsing RTs on go trials (on which no stop-signal occurred) into a single distribution. RTs are rank ordered and the RT on the n th percentile is selected, where n is the proportion of failed inhibitions on long stop trials at a given stop-signal delay (SSD). SSRT can be calculated by subtracting the SSD from this value (Logan, 1994). This process can be repeated several times for each delay or once across delays (using the average SSD). When the experiment contains sufficient trials, there is no reason to weight one method more than another (Band, van der Molen, & Logan, 2003). Because some subjects did not stop on long stop trials on which the stop-signal was presented 500 ms after the go-signal, the SSRT cannot be calculated for these SSDs. Therefore, we used the average SSD to compute SSRTs. For example, given that button-press responses could be withheld in approximately 34% of all long stop trials (66% noncancelled long stop trials), SSRT is calculated by subtracting the mean SSD (300 ms) from the 66th percentile of the Go RT distribution (~640 ms). This participant would have a SSRT of 340 ms.

Results

SSRTs ($F(2,46) = 12.31, p < 0.001$) decreased across sessions, reflecting an increased efficiency of response inhibition, clearly indicating that participants benefit from repeated task experience. Because we degraded go-signal and stop-signal visibility, we strongly increased the difficulty of our task compared to regular stop-signal studies. Therefore, Go RTs (± 620 ms) and SSRTs (± 330 ms) in our task were slightly larger than usually found in stop-signal tasks (Logan, 1994; but see Rubia, Smith, Brammer, & Taylor, 2003).

Unconsciously triggered response inhibition

Table 2.1 summarizes the inhibition rates on short go trials and short stop trials separately. What can be seen in this table is that participants sometimes inhibit their response on short go trials as well as on short stop trials. Inhibition rates decreased across sessions for both trials ($F(2,46) = 19.26, p < 0.001$), which indicates that participants become better at detecting *non masked* stop-signals (on long stop trials) by doing this task and therefore stop less on trials on which no stop-signal is perceived. However, across sessions participants stopped *relatively more* on short stop trials than on short go trials ($F(1,23) = 7.65, p = 0.011$, see Figure 3.2a for effect sizes). Although unconscious stopping behavior increased slightly across sessions, the interaction did not reach significance ($F(2,46) = 1.74, p = 0.19$). To specifically test whether training plays a role in unconscious stopping behavior, separate follow-up analyses per session were performed, revealing that participants significantly stopped more on short stop trials than on short go trials in the final session only ($t(23) = 3.04, p = 0.006$). Thus, perhaps surprisingly, masked stop-signals occasionally trigger response inhibitory processes to the level of complete response termination.

Unconsciously triggered RT slowing

On the whole, participants did not slow down their responses on short stop trials compared to short go trials ($F < 1$; see the white bars in Figure 2.2b). To test whether good inhibitors are more susceptible to the influence of masked stop-signals than poor inhibitors, we used a median split on participants' SSRT in the final session. It has been shown previously that good inhibitors show larger stopping related ERP components (Falkenstein, Hoormann, & Hohnsbein, 1999) and recruit stopping related areas, such as the right inferior frontal gyrus and the subthalamic nucleus more than poor inhibitors (Aron & Poldrack, 2006). Additionally, activity in these stopping related areas correlated with SSRT in the latter experiment. Good inhibitors

($N = 12$) slowed down their responses on short stop trials more than poor inhibitors ($N = 12$) ($F(1,22) = 10.34, p = 0.004$). For good inhibitors, RTs on short stop trials were significantly longer than RTs on short go trials, not only across all sessions ($F(1,11) = 14.21, p = 0.003$), but also progressively so for successive sessions ($F(2,22) = 3.98, p = 0.034$; see grey bars in Figure 2.2b). Separate t-test analyses confirmed this interaction, revealing significant RT slowing in the second ($t(11) = 3.35, p = 0.006$) and third session ($t(11) = 2.62, p = 0.024$). In contrast to good inhibitors, poor inhibitors did not slow down their responses on short stop trials compared to short go trials ($F(1,11) = 1.92, p = 0.19$; see black bars in Figure 2.2b). To rule out the possibility that the stop-signal was perceptually more similar to one of the go-signals and therefore might produce response congruency effects, we analyzed the RT slowing effects for both go trials separately (blue or green). RT slowing effects did not differ between go stimuli (overall, $p = 0.93$; good inhibitors, $p = 0.19$; poor inhibitors, $p = 0.53$).

Table 2.1

General performance measures for each separate session averaged across all participants

	Short Go Trial		Short Stop Trial		SSRT
	Inhibition rate	Reaction time	Inhibition rate	Reaction time	
Session 1	4.5 (0.9)	630 (6.2)	4.6 (0.8)	629 (6.2)	353 (7.2)
Session 2	2.5 (0.6)	614 (5.5)	2.7 (0.6)	616 (5.2)	337 (7.0)
Session 3	1.0 (0.2)	599 (6.5)	1.4 (0.2)	601 (5.5)	328 (7.3)

Note: SSRT is the stop-signal reaction time (ms) (\pm SEM).

To verify the role of evolving Stimulus-Response (S-R) associations due to training in unconsciously triggered inhibitory control, the RT slowing effect in the final session, which reflects the eventual impact of masked stop-signals on RT slowing behavior, was correlated with the decrease/increase in SSRT across sessions (SSRT in the first session – SSRT in the final session) which reflects the increase in performance over sessions for each individual participant ($r = 0.41, p = 0.049$, see Figure 2.2c). This significant correlation indicates that the more participants increase their performance on long stop trials across sessions, the more they slow down their responses on short stop trials (compared to short go trials) in the final session.

To test the specific hypothesis that the influence of masked stop-signals increases with individual levels in the efficiency of inhibitory control, we correlated the unconscious RT slowing effect with SSRT for all sessions separately (see Figure 2.2d). In our experiment unconscious RT slowing correlated significantly with SSRT

in the final session ($r = -0.53, p = 0.007$). The highly significant negative correlation in the final session suggests that the extent to which masked stop-signals trigger inhibitory control processes depends on the participants' efficiency at inhibiting ongoing responses. So, participants with relatively short SSRTs (reflecting efficient response inhibition on long stop trials containing a *non masked* stop-signal) are more susceptible to the influence of *masked* stop-signals, and hence show greater RT slowing effects to masked stop-signals than those with slower SSRTs. That this relation strengthens over sessions is an indication for the role of S-R associations due to repeated task experience in unconsciously triggered response inhibition. Because the impact of masked stop-signals increases across sessions due to learning (even being absent in the first session) and because the individual ability of participants to inhibit responses on non masked stop-signals determines the impact of masked stop-signals on behavior we rule out the possibility that the effect we found is due to a general disruption of ongoing processes by an additional stimulus in the short stop condition compared to the short go condition.

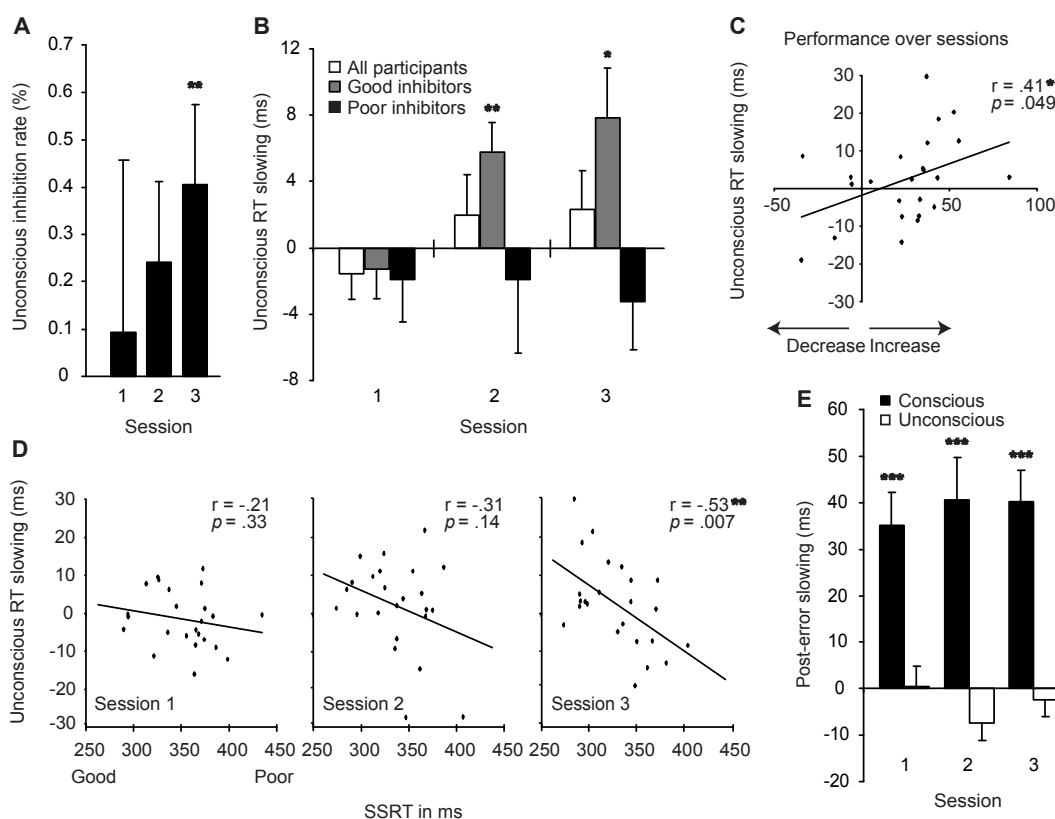


Figure 2.2 Behavioral measures of unconsciously triggered inhibitory control processes

(a) Unconscious inhibition rate per session averaged across all participants (\pm SEM). Participants stopped significantly more often on masked stop trials than on masked go trials. (b) Unconscious RT slowing per session for all participants and for good and poor inhibitors separately (\pm SEM). Good inhibitors were significantly slower on masked stop trials than on masked go trials. (c) Correlation between performance increase/decrease (SSRT in the first session – SSRT in the final session) and unconscious RT slowing in the

*final session. (d) Correlations between unconscious RT slowing and SSRT for all three sessions. (e) Conscious- and unconscious post-error slowing averaged across all participants (\pm SEM). Participants slowed down after conscious errors only and not after unconscious errors. *, $p < 0.05$; **, $p < 0.01$; ***, $p < 0.001$.*

An unexpected finding of this study was that poor inhibitors actually seemed to be faster on short stop trials than on short go trials. How can this be explained? We hypothesize that the slowing down of RTs found in our study may have been counteracted by another effect known to occur in the conditions used here. It has been shown previously that responses to masks speed up when preceded by undetected primes (Fehrer & Raab, 1962) and that these masks appear subjectively earlier than single masks (Bachman, 1989). In our experiment all trials contain masks, but on short stop trials, masks are preceded by task-relevant stop-signals, whereas short go trials contain single masks only. Although participants were not instructed to respond to the masks, but to the preceding go-signal, masks do indicate “the end of the trial” and thus are highly relevant in this task. Consequently, baseline RTs to short stop trials, disregarding any unconscious inhibition effect, could be faster than to short go trials. Although this pattern was not observed as a significant effect for the 50% poor inhibitors group, it was observed for the more extreme group of the one-third worst inhibitors who speeded up their responses to short stop trials (compared to short go trials) across sessions ($F(1,7) = 4.76, p = 0.033$, one-tailed). The one-third best inhibitors in fact overcome this speeding bias and slow down their responses ($F(1,7) = 8.92, p = 0.020$). This indicates that masked stop-signals initially speed up responses, to be counteracted and reversed only by response inhibition on masked stop-signals. Good inhibitors seem able to do so by training, whereas poor inhibitors do not, or less so. Because of these opposite effects of masked stop-signals on RTs, the ‘true’ inhibition effect triggered by masked stop-signals might be even larger than we were able to demonstrate here.

Post-error slowing

A third behavioral measure of control in the stop-signal task is post-error slowing, the slowing that occurs after a failed attempt to inhibit the response on a stop trial (Rieger & Gauggel, 1999; Schachar et al., 2004). Post-error slowing is measured by RTs on correct go trials immediately following failed stop trials compared to RTs on correct go trials immediately following go trials. Although post-error slowing effects are not always observed in the stop-signal task (Emeric et al., 2007) we looked into whether such trial-by-trial adaptive control mechanisms were triggered consciously and/or unconsciously in our task. Recent studies did not find regulatory trial-to-trial

behavioral adjustments after unconscious errors (Nieuwenhuis et al., 2001) or unconscious conflict (Kunde, 2003; but see Praamstra & Seiss, 2005).

In our task, failed long stop trials were associated with substantial post-error slowing across sessions ($F(1,23) = 40.74, p < 0.001$; all $ps < 0.001$ for each separate session, see Figure 2.2e). This effect did not increase upon training ($F < 1$). On the contrary, participants did not slow down their responses after failed short stop trials ($F(1,23) = 1.64, p = 0.21$). There was no sign of a training related increase in post-error slowing across sessions ($F(2,46) = 1.13, p = 0.33$; all $ps > 0.05$ for separate sessions). Good and poor inhibitors did not differ in conscious ($F(1,22) = 1.95, p = 0.18$) as well as unconscious post-error slowing ($F(1,22) = 0.11, p = 0.74$). So, trial-by-trial adaptive control processes (post-error slowing) were clearly triggered by non masked stop-signals, but not by masked stop-signals. Contrary to the online effects of masked stop-signals, post-error slowing was not dependent on task training. Apparently, automation of S-R associations was not sufficient for unconscious regulatory trial-by-trial control processes to emerge in our masked stop-signal task.

Stop-signal visibility

The forced-choice detection task with single-trial error feedback after the final session yielded a hit rate of 45.1% and a false alarm rate of 45.0%. The resulting d-score of -0.002 did not deviate significantly from zero ($t(23) = 0.07, p = 0.95$). Thus, short stop trials could not be discriminated from short go trials, and participants were exactly at chance level (50.0%) in detecting the presence or absence of stop-signals when masked with a stop-signal-mask SOA of 16.7 ms. An analysis for good ($d' = 0.003, t(11) = 0.06, p = 0.95$) and poor inhibitors ($d' = -0.007, t(11) = 0.17, p = 0.86$) separately revealed that detection performance was at chance level for both groups. Good and poor inhibitors did not differ in detecting masked stop-signals ($t(22) = 0.16, p = 0.88$).

Furthermore, the correlation between each participants' individual d-score and unconscious RT slowing effect ($r = 0.27, p = 0.21$) or unconscious inhibition rate ($r = -0.40, p = 0.053$, note that this is a trend towards a negative relation) in the third session was not significant. These results clearly demonstrate that the RT slowing effects and the inhibition effects could not be due to accidental perception of stop-signals. Any effect of perceptual learning would have revealed itself after the final session. After the final session participants easily detected stop-signals when masked with a long SOA, as expressed in an average detection rate of 93.4% (SD = ± 5.6).

Experiment 2

In experiment 1 we showed that, after training, a masked stop-signal resulted in a small increase of actually inhibited responses and a slowdown of non-inhibited responses compared to a condition without a stop-signal. This latter effect was present only for the good inhibitors and not for poor inhibitors. As we have previously noted, the comparison of short stop trials and short go trials is slightly complicated by the fact that the former contain an additional stimulus (the stop-signal), whereas the latter do not. This additional stimulus might have rather unspecific effects on responding, such as facilitating response production (Fehrer & Raab, 1962) or interfering with it. Interference as an explanation of our previous results is unlikely on the grounds that the impact of masked stop-signals correlated with participants' ability to inhibit responses on long stop trials (containing a non masked stop-signal). However, initial response facilitation by masked stop-signals was more likely, because the worst eight inhibitors actually speeded up their responses on short stop trials compared to short go trials.

To test whether unconsciously triggered inhibitory control processes do necessarily require training to emerge, and to test whether it can be triggered in all participants, we designed a somewhat different version of the previous experiment, namely a *selective* stop-signal task (de Jong, Coles, & Logan, 1995; van den Wildenberg & van der Molen, 2004). In this task, one signal (the stop-signal) instructs the participant to inhibit their response, whereas another signal (the “go-on” signal) instructs the participant to continue and press the button. In this experiment, on stop trials the word “STOP” was presented, and on go-on trials the word “BLUF” was presented. The stop-signal as well as the go-on signal could be masked by forward masks only or by forward *and* backward masks. By this means four conditions were created: (1) go-on trials *without* backward masks, (2) stop trials *without* backward masks, (3) go-on trials *with* backward masks, and (4) stop trials *with* backward masks (see Figure 2.3). The stop-signal (or go-on signal) was masked optimally by introducing forward *and* backward masks, which caused these stimuli to be invisible (as revealed by a post-experimental forced-choice detection task). The stop-signal (or go-on signal) was masked sub-optimally when the backward masks were absent, which caused them to be visible in these conditions. We were particularly interested in the comparison between stop trials with backward masks and go-on trials with backward masks, since this would reveal whether inhibitory control processes can be triggered unconsciously in the selective stop-signal paradigm. The advantage of this paradigm is that a subliminal stimulus is presented on stop trials as well as go-on trials. This should serve to equalize the

Fehrer-Raab effect across conditions. Furthermore, because we used a highly automatized stimulus (the word STOP) as a stop-signal in this experiment, RT slowing and response inhibitory effects could potentially appear in the first session already, because the stimulus-response association is more easily formed during testing.

Materials and methods

Participants

Twelve volunteers participated in the experiment for course credits (10 female, ages ranging from 18 to 25). The local ethical committee approved the experiment, and subjects gave written informed consent before experimentation. All participants had normal or corrected-to-normal vision.

Stimuli and procedure

Stimuli were presented against a black background (2.17 cd/m^2) at the centre of a 17-inch VGA monitor. Participants viewed the monitor from a distance of approximately 90 cm, so that each cm subtended a visual angle of 0.64 degrees. On trials with backward masks, we first presented a go-signal (29 ms, blue left pointing arrow or red right pointing arrow [isoluminant at 9.0 cd/m^2], width 0.64° , height 0.34°), followed after a variable delay (the SOA) by two randomly chosen uppercase consonant strings (both 43 ms), the stop-signal or go-on signal (29 ms), and finally two consonant strings (both 43 ms). On trials without backward masks, the same sequence was used, but the consonant strings at the end (backward masks) were replaced with blank screens (see Figure 2.3). Before each trial, a warning signal (a white cross presented for 300 ms) appeared 500 ms before the go-signal. The inter-trial interval was jittered (2000-3000 ms in steps of 200 ms, pseudo randomly) to minimize anticipation on the presentation of the go-signal.

The word "STOP" was used as a stop-signal and the word "BLUF" as a go-on signal (uppercase Courier font in white color, font size 24). The control word "BLUF" was matched to the "STOP" word on frequency of appearance in daily Dutch language (91 vs. 78 respectively per 1 million, as stated in the Celex database). The stimulus set of the consonants used as masks consisted of 10 uppercase letters (X, K, R, M, H, G, D, W, Z and N). No consonants were used that were also part of the experimental words. The mask contained seven letters, which were slightly overlapping to increase the density of the mask. The spacing between the centres of the letters was 12 pixels.

In this experiment, a staircase-tracking procedure dynamically adjusted the delay between the onset of the go-signal and the onset of the stop-signal to control the probability of inhibition. After a successfully inhibited stop trial without backward masks, the stop-signal delay (SSD) in the next trial increased by 14.3 ms, whereas the SSD decreased by 14.3 ms in the next trial when the participant was unable to stop. This algorithm ensured that responses were successfully inhibited in about 50% of the stop trials without backward masks. Every block started with a SSD of 129 ms.

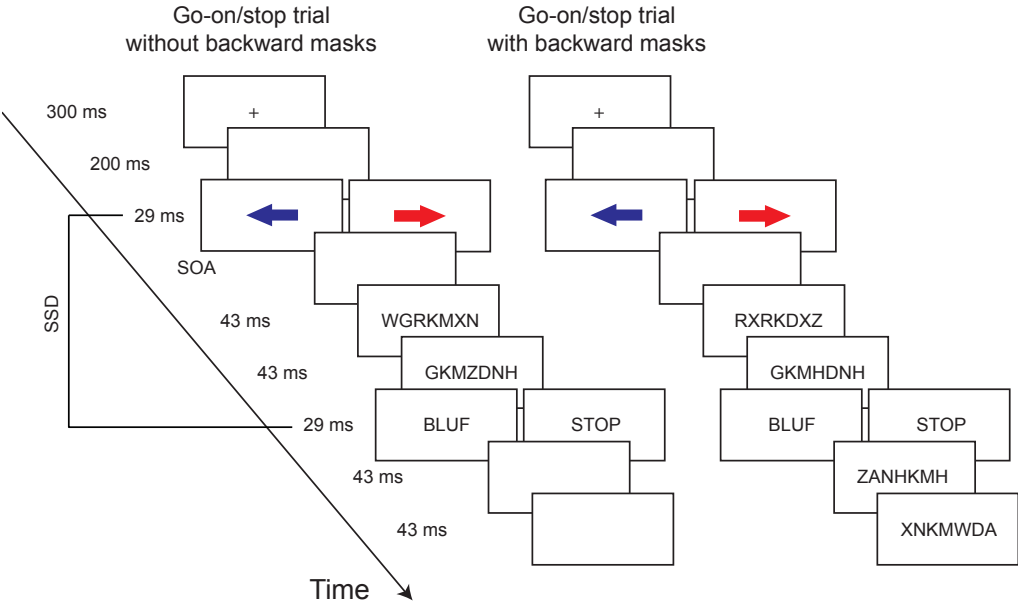


Figure 2.3 Stimulus timing of the masked selective stop-signal paradigm
SOA = stimulus onset asynchrony, SSD = stop-signal delay.

Participants were tested in a 1½-hour session in which they performed ten experimental blocks of 120 trials (25% of each of the four conditions). The first two blocks were considered practice. Participants received performance feedback after every block (mean RT, SD, and percentage inhibited trials on stop trials without backward masks). Participants were instructed to press as fast as possible to the presentation of the go-signal, but inhibit their response when the word “STOP” was presented. They were instructed to press the button to the go-signal when the word “BLUF” was presented or when random letters only were presented after the go-signal. Participants were not informed about the presence of masked stop-signals (or masked go-on signals). To equalize the inhibition rate on stop trials without backward masks between participants, the experimenter monitored participants’ performance by stressing either the primacy of responding or inhibiting, aiming at

an overall inhibition rate on stop trials without backward masks of approximately 50%.

Forced-choice detection task

After the stop-signal task participants performed 64 trials (32 stop trials with backward masks and 32 go-on trials with backward masks) of forced-choice detection. In this forced-choice detection task, trials were randomly drawn from a uniform distribution, rendering the presentation of the trials unpredictable. On each trial a whole sequence of events was presented and the subjects were instructed to press the left button when they thought the word STOP was presented and press the right button when they thought the word BLUF was presented. For simplicity, a trial was followed after 1000 ms by a pair of choices presented left (STOP) and right (BLUF) of fixation. Participants determined which of the two words was presented in the preceding trial. There was no stress to speed up responses in this task. Trial timing was exactly the same as in the stop-signal task and SSDs of 129, 157, 186 and 229 ms were used. Before running the detection task participants were fully informed about the precise structure of the trials and the fact that exactly half of the trials contained the word STOP and the other half the word BLUF. None of the participants reported to be aware of the masked stop-signals (or go-on signals) before being informed about them at the beginning of forced-choice detection.

Data analysis

The statistical analyses were similar to those in experiment 1.

Results

The statistical analyses revealed a similar pattern of results as in experiment 1. Participants stopped more often on stop trials with backward masks than on go-on trials with backward masks ($t(11) = 2.21$; $p = 0.049$, see Figure 2.4a). Additionally, RTs on stop trials with backward masks were significantly longer than RTs on go-on trials with backward masks ($t(11) = 3.56$; $p = 0.004$, see Figure 2.4b). In this experiment, there was no reliable correlation between RT slowing and SSRT ($p = 0.65$). Failed stop trials without backward masks were associated with substantial post-error slowing ($t(11) = 6.17$, $p < 0.001$), whereas participants did not slow down their responses after failed stop trials with backward masks ($t(11) = 1.01$, $p = 0.33$, see Figure 2.4c). The average SSRT in this experiment was 411 ms (SD = 19.3).

Stop-signal visibility

The forced-choice detection task after the final session yielded a hit rate of 52.7% (SD = 6.6). The resulting d-prime score of 0.14 did not deviate significantly from zero, ($t(11) = 1.47, p = 0.17$). Furthermore, the correlations between each participant's individual d-prime score and unconscious RT slowing effect ($r = -0.27, p = 0.39$) or unconscious inhibition rate ($r = 0.11, p = 0.74$) were not significant. In sum, the forced-choice detection task confirmed that our masking procedure rendered stop-signals largely invisible. Although our d-prime measure was slightly above zero (however, not significantly), the absence of reliable positive correlations between d-prime scores and inhibition measures suggests that it is unlikely that our results are due to accidental visibility of masked stop-signals.

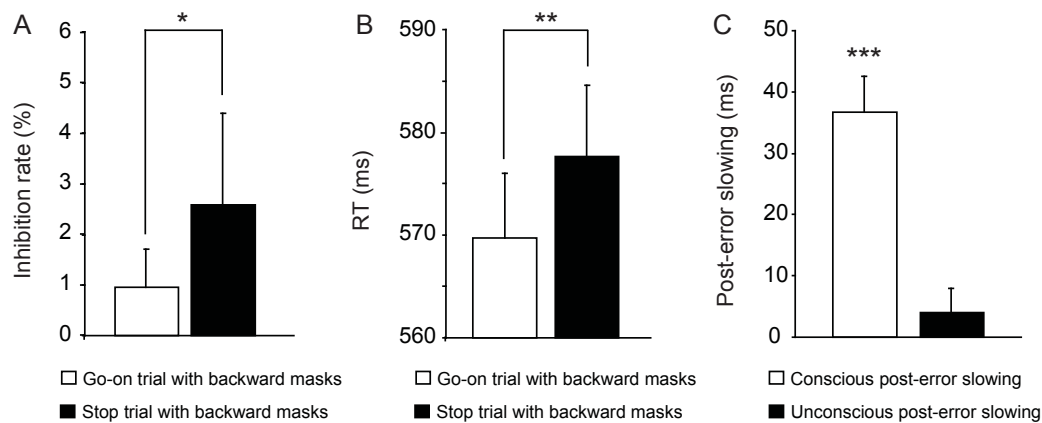


Figure 2.4 Behavioral measures of unconsciously triggered inhibitory control processes

(a) Unconscious inhibition rate on go-on trials with backward masks and stop trials with backward masks averaged across all participants (\pm SEM). Participants stopped significantly more often on stop trials with backward masks than on go-on trials with backward masks. (b) Mean RTs on go-on/stop trials with backward masks (\pm SEM). Participants were significantly slower on stop trials with backward masks than on go-on trials with backward masks. (c) Conscious- and unconscious post-error slowing averaged across all participants (\pm SEM). Participants slowed down after conscious errors only and not after unconscious errors. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

General discussion

Our study provides the first empirical evidence that unconscious task-relevant signals can actively trigger and initiate response inhibition, thereby breaking the alleged intimate relationship between consciousness and inhibitory control (Dehaene & Naccache, 2001; Eimer & Schlaghecken, 2003; Libet, 1999). Additionally, our results show that cognitive control functions are differentially affected by awareness. In our masked version of the stop-signal task, on-line inhibitory control operations could be triggered unconsciously, whereas regulatory trial-by-trial

control operations (post-error slowing) could not. Although the primary focus of our study was testing whether high-level cognitive control processes, such as response inhibition (governed by prefrontal brain areas), can be triggered unconsciously, our data also shows that unconsciously triggered inhibitory control is not as efficient as consciously triggered inhibitory control. While non masked stop-signals lead to complete response inhibition on the majority of trials, this is the exception rather than the rule on trials containing a masked stop-signal. Additionally, trial-by-trial regulations (post-error slowing) are clearly present after conscious errors, but absent after unconscious errors in two different versions of the masked stop-signal paradigm. So, unconsciously triggered inhibitory control seems to differ from some other cognitive control processes in that it appears to be less flexible and less efficient and probably takes more time to develop. These results converge with studies showing that awareness seems crucial for some types of (trial-by-trial) cognitive control regulations (Dehaene et al., 2003; Kunde, 2003; Nieuwenhuis et al., 2001; but see Praamstra & Seiss, 2005), but also demonstrates the possibility of unconsciously triggered inhibitory control.

Our results contradict a previously proposed function of consciousness; “conscious vetoing” (Haggard & Libet, 2001; Libet, 1999). In the famous Libet experiments (Libet, Gleason, Wright, & Pearl, 1983), it is nicely demonstrated that humans become aware of their intention to act *after* the emergence of electrophysiological markers in the brain that reflect action preparation (see also Lau, Rogers, Haggard, & Passingham, 2004; Soon, Brass, Heinze, & Haynes, 2008). Therefore, it is likely that motor acts are initiated in the brain unconsciously, which undermines the concept of conscious free will. In such a view intention is just a consequence of neural activity, and not the cause (Haggard & Libet, 2001; Libet, 1999). Although our actions are probably initiated in the brain unconsciously, Libet proposed that stopping or vetoing an action requires voluntary inhibition (see also Pisella et al., 2000; Rossetti et al., 2003). Thus, consciousness might be needed to control the final execution of actions by vetoing or stopping inappropriate action plans (Libet, 1999). This idea is supported by Rossetti and colleagues (Pisella et al., 2000; Rossetti et al., 2003) who showed that ongoing motor actions (e.g. pointing movements) can be redirected on-line automatically and involuntarily. However, they argued that stopping (inhibiting) such an action requires a conscious and voluntary process, which seems to be much slower. The present results extend this idea by showing that inhibitory control processes can be triggered by masked stop-signals. Thus, not only the initiation of an act might unfold unconsciously, it seems also possible to initiate the veto process unconsciously (see also Brass & Haggard,

2007). However, as stated previously, non masked stop-signals are able to trigger response inhibition much more efficiently than masked stop-signals, which suggest that consciously and unconsciously triggered inhibitory control (or veto) mechanisms differ qualitatively from each other. Future studies are needed to unravel the exact role of conscious and unconscious processes in inhibiting or vetoing our actions.

There is a range of methods in the literature to demonstrate unconsciousness of stimulus material, ranging from simple subjective assessments ('were you aware') to rigorous forced-choice detection tasks. Demonstrating chance performance in a forced-choice detection task is usually considered sufficient (and the strongest) evidence for non-conscious perception (e.g. Dehaene et al., 1998; Vorberg et al., 2003). However, whether detection performance exceeds chance-level performance (reaches significance or not) depends on several task aspects, such as the number of trials used for detection performance and the number of participants tested (for a more elaborate discussion on this issue see Dehaene, 2008; Hannula, Simons, & Cohen, 2005). Although detection performance did not differ significantly from chance-level in both experiments, it cannot be ruled out that detection performance might have exceeded chance-level eventually when more trials and/or participants were measured. However, our rationale for believing that participants were probably unable to perceive masked stop-signals consciously in both experiments is not based on detection performance only, but on many more observations. First, masked stop-signals do not yield response time slowing in all participants in experiment 1. RT slowing was correlated negatively with SSRT and in fact, RT slowing effects were found only for good inhibitors and not for poor inhibitors. If RT slowing effects had been caused by slight visibility of masked stop-signals, our observed RT slowing effect would not differentiate between participants. Note that the very worst performers in experiment 1 even speeded up their responses on stop trials containing a masked stop-signal, probably because they did not overcome the Feher-Raab effect. Second, we have demonstrated qualitative differences in the effects of masked vs. non masked stop-signals on trial-by-trial post-error slowing. Such qualitative differences in processing masked vs. non masked stimuli are considered as convincing evidence for unconscious perception (Jacoby, 1991; Merikle, Smilek, & Eastwood, 2001). Last but not least, there were no reliable correlations between participants' d-scores and unconscious inhibition rates or RT slowing effects in both experiments. This indicates that unconscious RT slowing and unconscious response inhibition are probably independent of stop-signal visibility.

It is important to stress that this study is essentially different from previous work on inhibition effects in the masked priming task (Eimer, 1999; Eimer & Schlaghecken, 1998). Eimer and colleagues have shown that at longer prime-target intervals (> 100 ms.), initial response facilitation by congruent *task-irrelevant* primes is automatically followed by inhibition (longer RTs on congruent than on incongruent trials). They called this form of inhibition “exogenous”, because in situations like the masked priming task, suppression of incorrect or premature response tendencies seems to be largely under automatic, stimulus-driven control (Eimer & Schlaghecken, 2003), mediated mainly by subcortical structures (Aron, Schlaghecken et al., 2003; but see Sumner et al., 2007). Although these findings challenged the view that all inhibitory processes are necessarily endogenous, suppression of motor responses is still thought to be largely under voluntary, cognitive control, mediated by prefrontal structures (Aron, Fletcher, Bullmore, Sahakian, & Robbins, 2003) in situations like the stop-signal task (Dehaene & Naccache, 2001; Eimer & Schlaghecken, 2003).

Additionally, Ansorge (Ansorge, 2004) showed that a masked *task-irrelevant* prime presented at the same location as a subsequent target speeds up responses, whereas a masked prime presented at the noncorresponding location slows the target responses down. In one condition of one experiment (Experiment 5), a prime was sometimes presented *without* the following target, in which case the response was to be withheld. In this condition, a prime presented at the same location as a subsequent target had a less beneficial effect, compared to a control condition.

In the above-mentioned experiments (Ansorge, 2004; Eimer & Schlaghecken, 1998), awareness of control-signals was manipulated in a task in which an action that should be *executed* was primed (validly or invalidly) by a *task-irrelevant* masked signal that *preceded* the go-signal. By contrast, in our experiments, awareness of control-signals was manipulated in a task in which an ongoing action should be *aborted* upon a *task-relevant* stop-signal that *followed* the go-signal. Our present experiments extend these previous findings by providing a direct test of the notion that masked control signals can elicit (inhibitory) control over an already ongoing action.

In the stop-signal task a particular arbitrary stimulus is assigned to a “withhold” response without being mapped onto a motor response. In line with previously proposed mechanisms of unconscious processing, such as the direct parameter specification theory (Neumann, 1990), the action trigger theory (Kunde, 2003) or the evolving automaticity theory (Abrams & Greenwald, 2000; Damian, 2001), the association between the unconscious stimulus (the stop-signal) and the (withhold) response evolves during learning. In experiment 1, masked stop-signals

did not have sufficient inhibiting power immediately, at least not enough to overcome the Fehrer-Raab effect (no unconscious RT slowing or inhibition in the first session). During training the stop-signal becomes associated with inhibition and therefore masked stop-signals can increasingly activate response inhibition without a mediating conscious percept. However, unconscious inhibitory control does not necessitate extensive training to evolve. We have shown in a selective stop-signal paradigm (Experiment 2) that a stop-signal with inherent stopping properties is able to trigger inhibitory control unconsciously in the first session already.

These data fit in nicely with recent experimental reports (using various subliminal priming paradigms) showing that unconscious processes are more elaborate than previously supposed. These studies have revealed various top-down effects on the processing of unconscious stimuli such as temporal attention (Naccache et al., 2002), spatial attention (Sumner et al., 2006), cognitive intentions (Ansorge & Neumann, 2005; Kunde, Kiesel, & Hoffmann, 2003) and strategy (Greenwald, Abrams, Naccache, & Dehaene, 2003), thereby showing that people can exert *conscious* control over processes that are triggered *unconsciously*. In line with our results, recently one study showed that even very high-level cognitive control functions could be influenced without awareness. Lau and Passingham (2007) used functional MRI to test whether the cognitive control network in the prefrontal cortex can be activated by unconscious primes. In their task, participants had to prepare to perform either a phonological judgment task or a semantic judgment on an upcoming word. On a single-trial basis, participants were instructed which task to perform. However, on all trials a prime was presented just before the instruction stimulus, half of which triggering the alternative (“wrong”) task, and half of them triggering the same (“right”) task. When participants were unconsciously primed to perform the wrong task, there was increased activity in the network associated with the wrong task, and decreased activity in the consciously instructed task set. So, participants were actually engaged in the wrong task when they were primed to do so. Additionally, activity in the mid-dorsolateral prefrontal cortex was associated with this effect.

This is very strong evidence that cognitive control processes in the prefrontal cortex can indeed be triggered unconsciously. However, one limitation of this study is that participants were not asked to respond to the primes. So, the unconscious stimuli were not task-relevant for participants. Our results further extend these results by showing that unconscious task-relevant stimuli can trigger cognitive control processes and influence behavior. Apparently, we do not have to become aware of a stimulus for it to initiate (probably prefrontal) “endogenous” inhibitory

control mechanisms. An important next step to understand the function of consciousness would be systematically searching for precise boundaries between cognitive (control) functions that require consciousness, and cognitive functions that do not.