



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*. [Thesis, fully internal, Universiteit van Amsterdam].

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.

9. Summary and discussion

Traditionally, consciousness and cognitive control are supposed to be intimately related, in such a way that people often refer to “conscious cognitive control”, as if there would be no alternative (Hommel, 2007). In this thesis I put this long-held assumption to a direct test. To this end, I combined several traditional inhibition paradigms (the Go/No-Go task, the stop task and conflict tasks) with the masking paradigm to investigate the existence and potential neural mechanisms of unconscious cognitive control. In the following sections I will summarize and interpret the obtained results and argue that unconscious cognitive control is “real” and associated with prefrontal computations in the human brain; as is also the case for conscious cognitive control. Yet, the present research also revealed substantial differences between consciously and unconsciously triggered cognitive control. Commonalities and differences between both versions of cognitive control will be discussed in the light of influential models concerning the neural correlates and function of consciousness.

Summary of the results: the existence of unconscious cognitive control

The research described in this thesis can be divided roughly in two parts. In the first part of this thesis, I focused mainly (however not exclusively) on control operations that I called *online cognitive control* processes. This was done to highlight that behavioral adaptations based on unconscious control signals are implemented directly; in the current trial. These processes are sometimes also referred to as *reactive control* processes (Braver et al., 2007). In the second part of this thesis, I shifted my focus towards *regulatory control* processes (sometimes also referred to as *proactive control*). The implementation of such control processes are reflected in behavioral adaptations after the experience of conflict/errors; thus in the next trial. Below I will first summarize the obtained results.

Unconsciously triggered online cognitive control (Chapter 2-6)

The results presented in Chapter 2 demonstrate that unconscious (masked) stop-signals are able to trigger response inhibition in a stop-signal task. In two different experiments, unconscious stop-signals were observed to result occasionally in full-blown response inhibition and to slow down in the speed of responses that were not inhibited; as if participants attempted to inhibit their response but just failed to withhold it completely. Interestingly, this effect increased with practice and correlated with an efficiency measure of response inhibition (stop-signal reaction

time; SSRT). Thus we showed that response inhibition can be triggered unconsciously, the more so in individuals and under conditions that are associated with efficient response inhibition. Interestingly, trial-by-trial changes in cognitive control (post-error slowing) were clearly triggered by response-errors on conscious stop trials, but not by response-errors on unconscious stop-trials. These results showed that cognitive control functions are differentially affected by awareness. Whereas *online* control operations could be triggered unconsciously, *regulatory* control operations could not; at least not in the two versions of the masked stop-signal paradigm used in this chapter.

In two follow-up experiments, we recorded event-related potentials (ERPs) while participants performed a masked Go/No-Go task (chapter 3) and a stop-signal task (chapter 4) to track the fate and time course of unconscious control signals in the human brain. Both experiments revealed that unconscious control stimuli elicit a sequence of distinct ERP components that are also observed on conscious control signals, albeit smaller and with different relative strength. Importantly, in both studies, unconscious control stimuli triggered a prefrontal ERP component, arising around 300 ms post-stimulus, that correlated with the extent to which participants slowed down their responses to these stimuli (online response inhibition effects). Source imaging suggested that unconsciously triggered inhibitory control was associated with activity in the lateral PFC, close to the cortical areas thought to be key in (conscious) response inhibition, such as the inferior frontal cortex (IFC) and the pre-SMA (Blasi et al., 2006; Chikazoe et al., 2007; Garavan et al., 2002; Konishi et al., 1999; Rubia et al., 2003; Simmonds et al., 2008; Wager et al., 2005). These results suggest that unconscious control signals are able to trigger prefrontal control networks, traditionally assumed to be exclusively associated with conscious information processing.

Furthermore, the results presented in chapter 3 provided evidence for the notion that the processing of unconscious stimuli is boosted by top-down interactions; in this case task-set. We showed that the same unconscious stimuli were processed more elaborately (probably activating brain areas further downstream) when they were task-relevant (associated with No-Go inhibition), than when they are irrelevant for the task at hand. This is in line with the idea that the depth and scope of neural processing of masked stimuli is modulated by top-down settings of the cognitive system (Dehaene et al., 2006; Nakamura et al., 2006; Nakamura et al., 2007).

In a subsequent study (chapter 5), we tried to further illuminate the possible depth of processing of unconscious information using methods that allow more

spatial (rather than temporal) precision in the neuroanatomical inferences. Therefore we used functional MRI to investigate to what extent unconscious No-Go stimuli are capable of activating brain areas involved in inhibitory control. In a new version of the masked Go/No-Go task, consciously initiated response inhibition was associated with a typical (mostly right-lateralized) frontoparietal network associated with response inhibition. Importantly, unconscious No-Go signals also triggered a prefrontal network of brain regions, most prominently the IFC and the pre-SMA. Again, unconsciously triggered neural activation correlated with the extent to which participants slowed down their responses to unconscious No-Go signals, which further supports the notion that the observed activations are “functional” in the sense that they predict behavior. However, neural differences between conscious and unconscious control were also revealed. Whereas conscious response inhibition was associated with a widespread frontoparietal inhibition network, unconscious response inhibition was more localized and specific to crucial No-Go processors such as the IFC and pre-SMA.

In chapter 6, we aimed to investigate the link between brain function and brain structure and to test whether individual differences in the ability to select the appropriate action (and inhibit the inappropriate action) during response conflict are related to individual differences in brain structure, especially in the pre-SMA. Furthermore, we were interested in the relation between overcoming response conflict and the awareness-level of such conflict. To this end, we used a masked priming paradigm in which conflicting response tendencies were elicited either consciously (through primes that were weakly masked) or unconsciously (strongly masked primes). Consistent with our predictions, whole-brain voxel-based morphometry revealed that individual differences in pre-SMA grey-matter density predicted subjects’ ability to select the correct action in the face of response conflict, irrespective of the awareness-level of conflict-eliciting stimuli. These results link structural anatomy to cognitive control ability, and provide support for the role of the pre-SMA in the flexible adjustment of ongoing behavior by inhibiting irrelevant action plans (see also chapter 5). Furthermore, these results suggest that flexible and voluntary behavior requires efficiently dealing with competing motor plans, including those that are activated automatically and unconsciously.

Unconsciously triggered regulatory changes in cognitive control (chapter 7-8)

In the last two chapters of this thesis, I shifted my focus away from the immediate online effects of unconscious control stimuli on behavior to the potential, more long-term, regulatory effects that these stimuli might have for subsequent occasions. To

this end, in chapter 7, I looked into conflict adaptation effects in the masked priming task: the phenomenon that the presence of conflict in the previous trial reduces the impact of conflicting information in the current trial. We have shown that conflict adaptation can be triggered by unconscious conflict-inducing stimuli. This suggests that participants engender a more cautious response-strategy and increase cognitive control to reduce the impact of conflict-inducing stimuli after the experience of conscious as well as unconscious conflict.

In the final experimental chapter (chapter 8), we further zoomed in on the long-term effects of unconscious information using a masked Go/No-Go task (see also chapter 3). In EEG, response-errors (responses to trials containing a No-Go signal) were observed to elicit increased oscillatory synchrony (2-12 Hz) between the prefrontal and occipital cortex, which began before the next trial started (during the inter-trial-interval) and occurred even in absence of conscious awareness of the No-Go signal. Consistent with our predictions, Granger causality analysis revealed that the dominant direction of these long-distance frontal-occipital interactions was from frontal to occipital electrodes, suggesting that after a response error, prefrontal cortex engages control by increasing the tuning or efficiency of stimulus processing. The strength of these long-range interactions correlated with an overall measure of performance efficiency. In combination, the results obtained in these last two chapters suggest that an unconscious stimulus can have relatively long-lasting effects on behavior and brain activity and highlight that unconscious stimuli can trigger regulatory cognitive control processes.

The fact that we observed trial-by-trial changes in cognitive control in these latter two chapters somewhat contradicts the absence of similar effects in the masked stop-signal paradigm (chapter 2, 4). Below, I will first highlight commonalities and differences between conscious and cognitive control; thereafter, I will return to this point.

Conscious vs. unconscious cognitive control

Although we have demonstrated that prefrontal cognitive control functions can be initiated in the absence of awareness (which was the main goal of this thesis), the presented data also revealed substantial differences between conscious and unconscious versions of cognitive control. Let's first focus on the behavioral differences between both types of control. First of all, the effect of unconscious control signals strongly depends on task practice, inhibitory performance (SSRT) or familiarity with the (unconscious) stimulus material (chapter 2 and 4). Second, unconsciously triggered inhibitory control seems not as efficient as consciously

triggered inhibitory control since conscious control signals lead to complete response inhibition on the majority of trials, whereas this is the exception rather than the rule on trials containing a masked control signal (chapter 2-5). Third, regulatory trial-by-trial control regulations (post-error slowing) were clearly present after conscious errors in all experiments, whereas post-error slowing was less consistently observed after unconscious conflict/errors. So behaviorally, unconsciously triggered cognitive control seems to differ from consciously triggered cognitive control in that it appears to be less efficient, less flexible, less durable and probably takes more training to develop.

Also from a neural perspective, the experiments presented in this thesis revealed substantial differences between conscious and unconscious versions of cognitive control. In several experiments, unconscious control stimuli were observed to trigger ERP components or activate brain regions that were also present during consciously triggered cognitive control. However, in all experiments, the neural signature of unconsciously triggered control differed from its conscious counterpart in that conscious control-initiating stimuli initiated stronger, more wide-spread and longer-lasting neural activity than unconscious-initiating control stimuli.

In part, these differential results could be explained by the fact that the stimulus-onset asynchrony between the control signal and the mask (chapter 2-5, 8) or the stimulus duration (chapter 6-7) was always shorter for the unconscious, than for the conscious conditions (see also Cleeremans, 2005; Overgaard, Rote, Mouridsen, & Ramsøy, 2006). It has been shown previously that when the time between the stimulus and the mask is increased (while the prime remains unconscious) the impact of this stimulus on behavior also increases (Mattler, 2003; Vorberg et al., 2003). So generally, unconscious effects might be weaker than conscious effects, because the stimulus of interest has usually less bottom-up stimulus strength (Dehaene et al, 2006) and less time to exert its influence (e.g. Mattler, 2003; Schmidt, Niehaus, & Nagel, 2006; Vorberg et al., 2003). Although this holds for many masked priming studies, some authors cleverly used a masking design in which the unconscious stimulus could be presented longer than the conscious stimulus (U-shaped masking design). In that case, unconscious stimuli can have larger effects on behavior and brain activity than conscious stimuli, highlighting the importance bottom-up stimulus strength (Lau & Passingham, 2007).

Although acknowledging this issue, which deserves more attention in future work, I have tried to explain the differential results between consciously and unconsciously triggered cognitive control from influential theoretical frameworks. Especially, online cognitive control processes might be well explained by theories

that differentiate between the role of feedforward and recurrent processing in conscious/unconscious vision (e.g. Dehaene et al., 2006; Lamme, 2003, 2006).

Feedforward vs. recurrent processing of visual information

It has been shown that when a visual stimulus is presented, it travels quickly from the retina through several stages of the cortical hierarchy, referred to as the fast feedforward sweep (Lamme & Roelfsema, 2000). Each time information reaches a successive stage in this hierarchy, a higher-level area starts to send information back to lower-level areas through feedback connections. Single-cell recordings in monkeys (Super et al., 2001) and TMS (Pascual-Leone & Walsh, 2001), fMRI (Haynes et al., 2005) and EEG (Fahrenfort et al., 2007) experiments in humans have revealed that the feedforward sweep probably remains unconscious, whereas recurrent interactions (multiple feedback loops) trigger the conscious awareness of a stimulus (for reviews see Dehaene et al., 2006; Lamme, 2006; Lamme & Roelfsema, 2000; Tononi & Koch, 2008).

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). Therefore, unconscious stimuli are capable of triggering many forms of behavior (Lamme, 2006), as evidenced by many masked priming experiments (e.g. Binsted et al., 2007; Leuthold & Kopp, 1998; Vorberg et al., 2003). A crucial aspect of the (unconscious) feedforward sweep is that it probably decays rapidly after travelling up the cortical hierarchy (Dehaene et al., 2001). In contrast, a key feature of recurrent interactions is that they promote widespread neural communication between distant cortical 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 have been associated with a highly distributed frontoparietal P3-like component (Del Cul et al., 2007) as well as increased long-range oscillatory synchrony (Gaillard et al., 2009). In fMRI, it is thought to be reflected in stronger, more widespread neural activations and increased functional connectivity between brain areas (Dehaene et al., 2001; Haynes et al., 2005).

Two different versions of cognitive control?

With this in mind, let's now take a closer look at the data presented in this thesis. Although fMRI and EEG cannot easily distinguish between feedforward and feedback influences, our EEG findings demonstrated that ERP components as found in response to conscious control signals are also found when control signals are

masked, albeit smaller and with different relative strength (chapters 3, 4). In chapter 4, we have seen that especially the late ERP components (e.g. the P3) are radically reduced in strength (see also Dehaene et al., 2001). (f)MRI measurements demonstrated that unconscious control signals can propagate to (medio)frontal cortex (chapter 5, 6). However, whereas conscious No-Go signals triggered a large-scale frontoparietal inhibition network, unconscious control signals only activated a specific subset of local neural processors therein. Activity in the unconscious inhibition network was reduced and clusters were sufficiently smaller than conscious activations. Also, BOLD time-course analysis revealed that conscious activations last longer than unconscious activations.

In sum, it seems reasonable to assume that unconscious control stimuli evoke feedforward activation of the same cortical modules as conscious stimuli, even in far distant areas (such as the IFC and pre-SMA). However, this activation might die out rather quickly because it is not supported by recurrent interactions. Therefore, it seems that both conscious and unconscious control signals trigger basic/online inhibition mechanisms, yet unconscious ones fail to elicit a comparably large, strong and distributed pattern of activation as observed when inhibition is triggered consciously. Although speculative, this sustained form of recurrent activity might cause information to become integrated across a number of (high-level) brain areas which might be necessary to initiate more flexible and durable cognitive control operations (Dehaene, 2008; Dupoux et al., 2008). In our case, this might be reflected in rapidly deployed full-blown response inhibition, available without extensive training.

As mentioned by Dehaene (2008), in light of recent experimental evidence published in this thesis (van Gaal et al., 2008; van Gaal et al., 2009) as well as elsewhere (Lau & Passingham, 2007; Pessiglione et al., 2007) it seems reasonable to conclude that cognitive control processes can (also) be triggered unconsciously. However, he also points out that these findings are mainly observed when task performance is automatized, because participants usually perform hundreds of trials before unconscious stimuli start to affect behavior. In contrast, conscious stimuli can rapidly trigger a novel response-strategy that seems impossible for unconscious stimuli. Recently, we performed a new experiment (not included in this thesis) to test this assumption (Wokke, van Gaal et al., in preparation). To this end, we designed a new version of the experiment presented in chapter 5 (the fMRI experiment) in which we introduced a cue (square/diamond), which preceded each trial. This cue instructed the participants which stimulus (square/diamond) was the No-Go signal in the upcoming trial. By this means the nature of the square/diamond

varied on a trial-by-trial basis; a square could be a No-Go signal in one trial, but not in another trial. Interestingly, in behavior, we still observed RT slowing to unconscious No-Go signals (summarized across squares and diamonds), which suggests that S-R mappings can be established on a trial-by-trial basis and that unconscious stimuli can also relatively flexibly influence our behavior. Although intriguing, future studies are required to test this issue more extensively in order to draw firm conclusions about the potential flexibility of unconscious information processing.

What about unconsciously triggered trial-by-trial changes in cognitive control?

Until chapter 6 the presented results can be explained fairly well by recent theoretical models concerning the underlying neural correlates of conscious and unconscious vision (with some extensions). In short, it seems that the feedforward sweep is able to penetrate all the way up to the prefrontal cortex to trigger *online* cognitive control processes. However since the bottom-up stimulus strength of unconscious stimuli is usually weaker compared to conscious stimuli (in experimental tasks), unconscious stimuli trigger online control processes to a lesser extent than conscious stimuli. In addition, *regulatory* control processes require that information is carried across a temporal gap of several seconds (the ITI in experimental tasks), which requires information to reverberate in the system; therefore it is absent after unconscious conflict/error signals. The absence of post-error slowing in chapter 2 and 4 nicely fits within this scheme. However, the data presented in chapter 7 and 8 challenge this latter assumption since, in those chapters, we observed that unconscious stimuli can trigger trial-by-trial control regulations and (plausibly) also evoke reentrant activation from prefrontal to occipital cortex. At the moment, an outstanding and intriguing question is why we observed such effects on some occasions, but not on others. Although I do not have a definite answer yet (after considerable thought), it seems that the differential results might be explained by task-specific differences between the tasks that we used.

First, as argued by others (Verbruggen & Logan, 2008b; Verbruggen, Logan, Liefoghe, & Vandierendonck, 2008), in the Go/No-Go task (as well as the masked priming task) there is a more consistent S-R mapping than in the stop-task. In the Go/No-Go task the same stimulus is always associated with going and another with stopping (Logan, 1988), whereas in the stop-signal task, the go-signal is sometimes associated with going (when it is not followed by a stop-signal) and on others with stopping (trials containing a stop-signal). These differences in S-R mappings suggest that the demands on cognitive control processes vary between both tasks, which

might (partially) explain the differential results that we observed. Several authors have argued that No-Go inhibition might be a more automatic (bottom-up) form of cognitive control than stop-signal inhibition, which has been argued to be a more controlled (top-down) form of cognitive control (Logan, 1988; Verbruggen & Logan, 2008a). It might be that unconscious stimuli are able to trigger the more automatic version of control (Go/No-Go task), but not the “true” top-down version of control (stop task). However, an explanation that only takes into account the strength of S-R mapping cannot explain all the available evidence since we observed recently that the nature of an unconscious signal can vary on a trial-by-trial basis and still influence behaviour (Wokke et al., in preparation). Notwithstanding, there might still be (additional) differences in the complexity of both tasks and the demands on cognitive control that can explain the absence/presence of unconscious post-error adaptation.

A second potential explanation is related to the number and functional significance of response-errors in the stop-signal task vs. the Go/No-Go task. In the stop task participants are told explicitly that they will (almost by definition) fail to inhibit their response on ~50% of the trials. In many versions of the task a dynamic algorithm even adjusts the stop-signal delay in such a way that participants indeed inhibit their response in ~50% of the occasions. In contrast, in the Go/No-Go task, participants are instructed to respond as quickly as possible to Go trials and to inhibit their response when a No-Go trial is presented. Therefore, it is less evident that making errors is “part of the experiment”. Generally, in the Go/No-Go task participants fail to inhibit their response in approximately 30% of the occasions. Due to these differences (error rate and instruction), the functional significance of an error probably differs significantly between the stop-signal task and the Go/No-Go tasks, which could subsequently affect the impact of “unconscious errors” on subsequent post-error control adjustments (Overbeek, Nieuwenhuis, & Ridderinkhof, 2005). Indeed, it has been shown that the extent of (conscious) post-error slowing is modulated by the frequency of an error (Notebaert et al., 2009) and it seems that post-error slowing effects are generally less reliably found in the stop-signal task than in the Go/No-Go task (Emeric et al., 2007; Verbruggen et al., 2008).

Besides these two (and probably more) task-specific differences, it is interesting to note that not only the studies presented in this thesis have obtained somewhat contradicting results with respect to the workings of unconsciously triggered regulatory control, but that the same holds for studies performed by others. Whereas some authors have argued/observed that an unnoticed (“unconscious”) error/conflict signal triggers regulatory changes in cognitive control

(Cohen, van Gaal, Ridderinkhof, & Lamme, submitted; Hester et al., 2005; Praamstra & Seiss, 2005; Rabbitt, 2002; van Gaal, Lamme, & Ridderinkhof, submitted), others have not (Endrass, Reuter, & Kathmann, 2007; Klein et al., 2007; Kunde, 2003; Nieuwenhuis et al., 2001; O'Connell et al., 2007; van Gaal et al., 2009). Notably, the majority of these studies did not objectively evaluate the awareness-level of the committed errors or the stimulus of interest (through objective forced-choice discrimination measures), but instead used subjective measures to evaluate “error awareness”. These inter-experimental differences might be explained by the fact that many different paradigms have been used to study unconsciously triggered regulatory control, including the anti-saccade task (Endrass et al., 2007; Klein et al., 2007; Nieuwenhuis et al., 2001), the Go/No-Go task (Cohen et al., submitted; Hester et al., 2005; O'Connell et al., 2007), the stop task (van Gaal et al., 2009), simple choice RT tasks (Rabbitt, 2002) and different versions of the masked priming task (Kunde, 2003; Praamstra & Seiss, 2005; van Gaal et al., submitted). Together, this suggests that post-error control adjustments can be triggered unconsciously on some occasions, but not on others. At the moment, the exact nature of post-error adaptation is still ill-understood and some even argue that they are not related to cognitive control operations at all (Notebaert et al., 2009). Therefore, it seems too early to draw firm conclusions about the boundary conditions for unconscious regulatory control processes to occur.

Apart from the substantial inter-experimental differences, how can we explain the unconscious regulatory control effects that we observed in chapter 7 and 8? First of all, these results are rather difficult to explain in the framework of feedforward vs. feedback versions of control. In chapter 8 we have shown that unconsciously initiated regulatory control processes were associated with increased phase synchronization between prefrontal and visual cortex, a measure thought to reflect increased functional connectivity between brain areas. Especially these latter results suggest that the story is slightly more complicated than previously assumed. Apparently, under some circumstances unconscious stimuli are able to influence cognitive processes for relatively long periods of time (i.e. 1-2 seconds) and might even initiate long-range functional interactions between distant brain areas. Interestingly, recently, others have also observed relatively long-lasting neural activations elicited by unconscious words, up to approximately one second (Gaillard et al., 2009; Naccache et al., 2005) and even longer effects of unconscious priming (up to several minutes) have been reported, for example in “mere exposure” paradigms (Elliott & Dolan, 1998; Gaillard et al., 2007).

In combination, these results suggest that unconscious information is able to influence cognitive processes for longer periods of time than previously thought, which has important implications for the theoretical models that propose a rapid decay of the strength of unconscious activations in the human brain (Dehaene et al., 2006; Kouider & Dehaene, 2007; Rossetti, 1998). To further clarify the lifetime and long-term influence of unconscious information on our behavior, future studies are necessary to examine the specific characteristics of the tasks used that determine the existence/absence and magnitude of trial-by-trial adaptive control processes. To get a further grip on this controversial issue in the future it is important to design experimental tasks in which (exactly) the same stimulus parameters are used for conscious and unconscious conditions to overcome differences in bottom-up stimulus strength. A potential candidate for these studies might be the use of the attentional blink (AB) paradigm. In a typical AB task participants perceive a rapid stream of stimuli, including targets as well as distractors. When a target follows another target after ~200-300 ms, the second target is missed on many occasions. Thus, although all parameters are exactly the same on all trials, sometimes a stimulus is perceived consciously, whereas on other trials it is not. By combining the AB paradigm with traditional cognitive control tasks (e.g. the stop task) it might be possible to nicely equate the bottom-up stimulus strength between conditions, whereas the awareness-level of the stimulus is carefully manipulated.

General conclusion

Based on the experiments reported in this thesis the following conclusions can be drawn: 1) Contrary to common theorizing, high-level cognitive control processes can be initiated unconsciously. 2) Although the prefrontal cortex has traditionally been associated with conscious information processing only, the present results indicate that unconscious stimuli are also able to activate prefrontal control networks. 3) The extent and depth of processing of an unconscious (control-initiating) stimulus in the human brain depends on task instructions and the behavioral relevance of the stimulus. 4) On some occasions, regulatory control processes can be initiated unconsciously; however the underlying mechanisms and specific boundary conditions should be studied more extensively to clarify inter-experimental differences. More generally, the results presented in this thesis expand our understanding of the scope, depth and potential long-term effects of unconscious information processing in the human brain and directly impact the current debate about the neural correlates of consciousness and the intimate relationship between consciousness and cognitive control.