



Universiteit
Leiden
The Netherlands

Cognitive control and binding in context-based decision-making : normal and dopamine deviant populations

Wouwe, N.C. van

Citation

Wouwe, N. C. van. (2009, December 3). *Cognitive control and binding in context-based decision-making : normal and dopamine deviant populations*. Retrieved from <https://hdl.handle.net/1887/14476>

Version: Not Applicable (or Unknown)

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/14476>

Note: To cite this publication please use the final published version (if applicable).

Discussion

The questions addressed by this thesis were the following: Do goal-driven processes as well as stimulus-driven processes contribute to context-based decision making? In what way do modulations of brain areas innervated by dopamine (DA), i.e. striatum or prefrontal cortex, affect control processes in context-based decision making and affect acquiring S-R associations? These questions were divided into subquestions and investigated in the previous chapters.

Do goal-driven and stimulus-driven processes account for variance in context-based decision making as measured by the AX-CPT?

The first two chapters investigated the contribution of goal-driven processes and stimulus-driven processes to context-based decision making, at behavioral as well as neural levels. Participants performed an adapted AX-continuous performance task (AX-CPT), using words (cues) followed by probe-pictures of male and female faces (chapter 2, experiment 1 and 2) or word-cues followed by face and house probe-pictures (chapter 3). The subjects' task was to respond to a target probe, for example a female face, given that it was preceded by a specific context-cue, for example a word in uppercase. Every other combination of context and probe called for a nontarget response.

In the *AX*-CPT task, the fact that *AX* target trials occur most frequently yields an error if a target probe appears in a nontarget context, or when a nontarget probe appears in a target context. According to goal-driven control theories of context-based decision making, context-cue information is used for preparation for a context-appropriate response (Braver & Barch, 2002; Braver et al., 2001; Miller & Cohen, 2001). Performance costs on nontarget trials are attributed to inefficient control. In contrast, decisions can also be biased by more stimulus-driven factors such as the automatic reactivation of features that previously accompanied the current event but that are currently irrelevant. Episodic bindings in the *AX*-CPT may arise between specific cue-words and probe-pictures. The binding account (Hommel, 2004, 1998; Logan, 1988) predicted that performance costs are caused by the need to overrule prepotent associations if features of the current event are partially incongruent with the previously created episodic binding.

The studies in chapter 2 and 3 revealed that although context-based decision making behavior is strongly affected by goal-driven preparation, learned associations between features of a previous event also play a role. This is consistent with other work demonstrating the role of episodic bindings in addition to effects of goal-driven biases (cf. Egner, 2007; Verguts & Notebaert, 2008).

Both behavioral experiments in chapter 2 demonstrated that performance costs induced by episodic bindings were stronger in trials where a cue-driven bias had to be overcome (when a nontarget probe appears in a target context) compared to trials where a probe-driven bias had to be overcome (when a target probe appears in a nontarget context). What does this convey about the relative contribution of goal-driven versus stimulus-driven biases to performance?

At first it seems that goal-driven preparation was more influential for *AX*-CPT performance than the effect of an association-driven bias; performance was highly impaired with *AY* trials (where the cue incorrectly indicates a target response) and binding costs are especially enhanced in these trials. On *BX* trials though, some performance costs were found due to the *X*-probe-bias, but oftentimes the *B*-cue overrides this bias. However, these findings might be biased by the task requirements. In an *AX*-CPT task *B*-cues are more predictable than *A*-cues in terms of response choice, which may have resulted in fast response times on all target probes that appear in a nontarget context (*BX* trials), without dissociating stimulus-specific *BX* performance

costs. Findings from a previous behavioral study (Van Wouwe, Band & Ridderinkhof, 2005) suggest that if the *B*-cue becomes less predictable, probe information will be more relevant to the decision and thereby affect performance. In the studies of chapter 2 however, the cue already provides a strong response tendency; therefore the priming effect of features associated with the cue may be enhanced compared to the probe-associated information.

Generally, conditions that determine the relative influence of proactive goal-driven or stimulus-driven biases on performance are not clearly defined yet.

Recent animal and modeling work indicate that overtraining certain stimulus-action relations reduces the ability to perform on a goal-based strategy (Adams, 1982; Daw et al., 2005). Likewise, Colzato, Raffone and Hommel (2006) showed that highly practiced associations (for example, yellow is highly associated with banana) affects the amount of attention that particular feature dimensions attract; that is, strongly associated features receive more attention than new unknown features, which can thereby influence which features are considered in binding.

Our task set-up did not explicitly require paying attention to stimulus-specific information: when a target probe appeared in a nontarget context, category information (stimulus category information, i.e. gender, defined whether a probe was target or nontarget) could be used to make the correct decision, instead of stimulus-specific information (i.e. individual faces). The task instructions lead to “intentional weighting” (e.g. Hommel et al., 2001; Wenke & Frensch, 2005) of stimulus features that are interpreted to be task-relevant. This would mean that category information would receive more attention than individual face information, which may have reduced performance costs that result from these stimulus-specific associations.

The fMRI study of chapter 3 demonstrated that the majority of the behavioral effects in the *AX*-CPT can indeed be explained in terms of a goal-driven account, although cue-probe bindings also explain unique variance in performance during decision making. Furthermore, the fMRI data shed light on the neural mechanisms underlying these processes. We found enhanced activation in ACC and DLPFC activation and impaired performance on trials that call for increased top-down control (*AY* and *BX*) compared to target trials (*AX*). This supports top-down control theories of context-based decision making (Braver et al., 2001; Durston et al., 2003; Kerns et al., 2004; MacDonald et al., 2000).

Episodic bindings were established by a behavioral training prior to the fMRI experiment and again presented during the experiment as complete repetitions or partial repetitions of trained cue and probe sequences. These bindings consisted of S-S associations, unlike in chapter 2 where S-R associations were investigated that were gradually built up throughout the experiment. With respect to the episodic-binding account, the following predictions were confirmed. If a specific cue stimulus was presented with a face probe on previous occasions, subsequent presentation of the associated cue and face probe increasingly reactivated the 'face area' in the brain (fusiform face area, FFA) with an improvement in performance, compared to face probes presented subsequent to cues previously associated with a house. We found similar results for cue stimuli followed by houses, which reactivated the 'house area' (parahippocampal place area, PPA) upon probe presentation, also with an improvement in performance. This pattern of activation and performance costs was predicted exclusively by the binding account (Hommel, 2004) and was in line with other imaging studies that investigated binding (Keizer et al., 2008; O'Craven, Downing & Kanwisher, 1999).

The reactivation of binding-related information in the AX-CPT seemed to take place with probe presentation (when a response is required) and not yet with cue-based preparation, while chapter 2 mainly provides evidence for cue-based priming effects. This may be the result of the experimental procedure of the fMRI study: during the experiment face and house probes appeared equally frequent subsequent to the cues; thus, participants may have anticipated equally for either a house or a face, based on frequency of presentation, which may have deferred binding-related reactivation.

The contribution of episodic bindings to the behavioral results of the fMRI study seemed to be smaller than its contribution to behavior in the experiments of chapter 2. This may be explained by the fact that in chapter 3, partial repetitions and complete repetitions were presented equally often, whereas in chapter 2 complete (but not partial) repetitions occurred as often as AX trials, which may have strengthened these associations throughout the task. Additionally, the S-S bindings in chapter 3 may have been less influential than S-R associations of chapter 2 in generating performance costs beyond the strong A-cue and X-probe induced response bias. Moreover, a recent study (Hommel, Colzato, van den Wildenberg, & Cellini, submitted) revealed that binding effects are severely reduced by scanner noise, which may have to do with subjects'

increase in motivation and a switch towards a more goal-based performance strategy when performing a task in the scanner.

Contrary to our expectations, we did not find increased activation in frontal control areas such as ACC or DLPFC induced by a partial repetition of cue and probe compared to complete repetitions. Future imaging studies should aim at increasing our understanding of the relative importance of goal-driven vis-à-vis stimulus-driven biases in adaptive behavior, for example whether and in what way trial-to-trial changes in top-down control affect attention and episodic binding.

Recently, Braver et al. (2007) suggested that episodic retrieval of information from WM may be subserved by the hippocampus. Additionally, it has been proposed that in WM tasks with a long delay interval, in tasks with an increase in WM load (Rissman, Gazzaley, & D'Esposito et al., 2008) or in participants with a dysfunctional PFC, such as in healthy aging (Paxton, Barch, Racine, & Braver, 2007), performance increasingly relies on a stimulus-driven reactivation of task-relevant information. Because it is currently unknown whether contribution of the hippocampus is crucial to explain episodic-binding effects as studied in our task it would be informative to study episodic binding after administration of the drug midazolam, which deactivates the hippocampus (Hirshman, Passannante, & Arndt, 2001; Kristiansen & Lambert, 1996; Poncer, Durr, Gahwiler, & Thompson, 1996).

The relation between top-down control and episodic binding may also be the other way around; unbinding or rebinding previously encountered episodic traces increases control on future trials, as suggested by Verguts and Notebaert (2008); that is, if both stimulus and action features remain equal between the current and the previous event, no conflict will be detected and thus no need to increase control on subsequent events.

Taken together, chapter 2 and 3 pointed out that in some decisions with rapidly changing environmental demands, goal-driven preparation is often beneficial but may also hamper performance which can be overcome by applying increased control. However, this top-down bias is regulated more efficiently when the specific stimulus is presented in the same context it was previously associated with, compared to when it is presented in a new and unusual context. Additionally, the fMRI study sheds light on the way in which these stimulus-driven performance changes may be represented in the brain.

What is the modulating influence of induced positive affect on control processes in context-based decision making?

Chapter 4 consisted of an ERP study that investigated the modulating influence of positive affect on proactive control (context maintenance and updating), reactive control (flexible adaptation to incoming task-relevant information), and evaluative control (performance monitoring) in an *AX-CPT* task.

The *AX-CPT* task in this study consisted of a classic letter-based CPT; during each *AX-CPT* trial participants were presented with a sequence of letters on the computer screen. Unknown to the participants, these letter sequences were constructed as trials of cue-probe pairs (types *AX*, *AY*, *BX*, *BY*). Subjects were instructed to respond to every letter with either a target or nontarget response. A target response was required only when the target *X*-probe was immediately preceded by an *A*-cue. In every other case participants had to respond with a nontarget response. Positive or neutral mood was induced (between subjects) before participants started the experimental task.

Changes in the cue- and probe-related ERPs enabled us to disentangle the modulating effect of positive affect on proactive (as reflected in the P3b and the Contingent Negative Variation components of the ERP), reactive (indexed by the N2 elicited by the target) and evaluative control mechanisms (indexed by the Error-Related Negativity, ERN, elicited after incorrect responses). Induced positive affect was found to enhance flexibility and to modulate the ability to evaluate control, whereas it did not change goal maintenance or preparation.

The modulating effect of positive affect on flexibility seems to be consistent with other studies that found improved flexibility with positive affect (Ashby, Turken, & Isen, 1999; Ashby, Valentin, & Turken, 2002) and *AX-CPT* performance with positive affect (Dreisbach, 2006). Reduced performance monitoring as a result of positive affect is in line with other studies that investigated the modulating effect of affect on the ERN: the amplitude of the ERN was found to be larger for individuals high in negative emotionality than for individuals low in negative emotionality (Hajcak, McDonald, &

Simons, 2004; Luu, Collins, & Tucker, 2000). Additionally, Wiswede, Münte, Goschke, & Rüsseler (2009) recently revealed that a phasic affect induction by means of IAPS pictures modulates the ERN; negative affect increased the ERN but, contrary to our results, positive affect did not modulate the ERN.

In our study, positive affect did not modulate maintenance as predicted by the goal-driven account (Braver et al., 2001; Braver et al., 2007; Cohen et al., 2002) and the maintenance-flexibility theory (Dreisbach & Goschke, 2004). Positive affect seemed to exclusively improve flexibility whereas maintenance remained unaffected. Increased flexibility may impair performance with more challenging task constraints, for example when task-irrelevant distracting information is presented during a delay interval which was not present in our study.

The improvement in behavioural flexibility as elicited by the positive-affect induction was attributed to a dopaminergic increase in the striatum; DA increases in the striatum are involved in some forms of cognitive flexibility (i.e. switching between relevant stimulus information), as substantiated by pharmacological fMRI studies with healthy controls and patients with Parkinson's disease (PD) on and off medication (Cools, Barker, Sahakian, & Robbins, 2001; Cools, Sheridan, Jacobs, & D'Esposito, 2007). These studies suggest that a DA enhancement in the striatum improves the efficacy of using incoming response-relevant stimulus information to control behavior.

Reduced performance monitoring (ERN) with positive affect also seems to correspond to results of studies on dopaminergic modulations of the striatum (Frank & O'Reilly, 2006; Ito & Kitagawa, 2006; Zirnheld et al., 2004): positive affect may have increased DA in the striatum, explaining why participants in the positive affect condition might not have been able to show the DA dips during error processing, as indicated by the reduced ERN in the positive compared to neutral affect conditions. The relation between phasic DA changes and the ERN has recently been subject of discussion; that is, there is some controversy about the order of DA signal and ERN signal. According to Holroyd and Coles (2002), a nigrostriatal dip in DA disinhibits the ACC which results in the ERN. Orthogonal to this it is argued that activity in the ACC (induced by errors or negative feedback) precedes the DA signal: projections from ACC to the ventral tegmental area (or indirectly via striatum) inhibit DA neurons (Botvinick, 2007; Frank et al., 2007; Jocham & Ullsperger, 2009). One of the arguments for this reasoning is that phasic DA signals are not fast enough to generate the ERN, therefore

the ERN may be present before a phasic DA change. The ERN might thus reflect an inhibitory teaching signal from the ACC to DA neurons in the BG, which again affects cortical activity on subsequent trials. It is also possible that both mechanisms are concurrently active: first, the BG might train the ACC to distinguish errors from correct trials while in later stages the ACC may induce DA dips in BG with erroneous decisions.

Holroyd and Coles (2008; Hewig et al., 2007) recently suggested that ACC integrates reinforcement signals over time and thereby guides decision making. The ERN amplitude reflects the subjective value attributed the response options as learned during previous responses and not an error per se. In line with this, Holroyd, Pakzad-Vaezi, and Krigolson (2008) put forward a suggestion that reconciles the ERN with the conflict-monitoring theory of the ACC (which predicts an increase in the N2 and ERN with enhanced response-conflict instead of errors; Yeung, Botvinick, & Cohen, 2004). Phasic increases in DA elicited by a correct response or positive feedback could inhibit the conflict-induced activity in the ACC and thereby reduce the amplitude of the N2. According to this view, DA signals train the ACC to perform a task in a way that minimizes response conflict.

Individual differences in tonic DA states might lead to different teaching signals to the ACC, which again affects the ACC response on future trials. High tonic levels of striatal DA (as induced by positive affect) may have resulted in an absence of teaching signals with errors, but an enhanced signal with correct trials, thereby leading to reduced error monitoring and response conflict on future trials.

Recently, AX-CPT performance has been studied in striatal DA deficient populations (Frank, Santamaria, O'Reilly & Willcutt, 2007; Moustafa et al., 2008) and in a pharmacological striatal DA manipulation in healthy subjects (Frank & O'Reilly, 2006).

In addition to the standard AX-CPT, these studies (Frank & O'Reilly, 2006; Frank, Santamaria, O'Reilly, & Willcutt, 2007; Moustafa et al., 2008) also included a distractor (to investigate updating of irrelevant information) and learning version of the AX-CPT, in which target sequences had to be learned based on feedback. The ability to acquire the correct target sequences depends on phasic bursts with positive feedback. However, although enhanced phasic DA bursts may lead to an overall Go bias and benefit learning, it may be maladaptive in the presence of distracting information.

Summarized, the studies by Frank and colleagues (2007; Frank & O'Reilly, 2006; Moustafa et al., 2008) conveyed that with an increase in DA in the BG, Go learning and

updating of WM improved, whereas NoGo learning and the ability to inhibit irrelevant or distracting information was reduced. Improved updating and maintenance was reflected in these studies by impaired *AY* but improved *BX* performance. Based on these findings one could argue that positive affect actually reduced DA, because *AY* performance improved with positive affect. However, since maintenance seemed to remain unaffected by positive affect, as reflected both in the ERPs and in behavior (*AX*, *BY* and *BX* trials), we suggest that positive affect may have improved learning task-relevant Y-probe-associated response information which enhanced performance on *AY* trials. In *BX* trials adequate maintenance of task-relevant information is sufficient to perform correctly and the ability to maintain information seemed unaffected by positive affect. Nevertheless, we expect that learning from negative feedback or inhibiting distracting information during a delay interval between cue and probe will be impaired with positive affect.

Our study provided some insight into the control processes that are modulated by positive affect and how this may be explained in terms of dopaminergic changes. However, we did not aim to provide conclusive results regarding the effect of positive affect on the different dopaminergic systems, as DA levels (let alone phasic changes therein) cannot be measured directly in humans. Recently developed methods for fMRI imaging of brain stem nuclei (D'Ardenne, McClure, Nystrom, & Cohen, 2008; Münte et al., 2008) offer the potential to shed light on the modulating influence of phasic affect changes compared to tonic affect changes on control processes and the currently assumed neural correlates like the DA system and the brain areas innervated by DA. Individual differences in DA level, for example in DA-deficient patient populations (PD, ADHD, schizophrenia) or genetic variation in DA polymorphisms, could also provide insight into the effect of short-term and long-term affect-related changes on cognitive performance, the presumed effect of affect on DA and individual differences in response to affect-related changes.

Based on a vast body of literature, the results of our studies were specifically interpreted in terms of the DA system and interactions between BG and PFC. However, we did not claim that this is an exclusive explanatory account. Other neurotransmitter systems may play a role as well; for example, the noradrenergic system has been shown to be engaged in action selection (Aston-Jones & Cohen, 2005 for a review) and explains behavior in a DA-deficient population like ADHD beyond what can be

explained by the dopaminergic system; that is, whereas DA enables learning and updating relevant information, norepinephrine effects may modulate the strength of input that enters the BG (Frank, Scheres, & Sherman, 2007).

What is the modulatory effect of DA and subthalamic nucleus stimulation on reward-based decision-learning processes related to specific striatal structures?

While chapter 2 and 3 sought to investigate the behavioral biases that affect context-based decision making, chapter 5 and 6 studied context formation, that is, the formation of stimulus-action-reward associations and the effect of BG modulations on this learning process. These studies were largely inspired by reward-based decision-learning studies in animals, DA-deficient populations and pharmacological DA manipulations.

Chapter 5 aimed to distinguish the effect of a DA modulation on reward-based decision-learning processes related to specific striatal structures (caudate and putamen). In order to investigate this, PD patients performed a probabilistic reward-based decision-learning task ON and OFF medication. Chapter 6 sought to examine the effect of STN modulation on reward-based decision-learning processes. Again, this was investigated by means of a probabilistic reward-based decision-learning task (Haruno & Kawato, 2006), performed by PD patients ON and OFF STN stimulation.

Subjects were instructed that the goal of the task was to make as much money as possible by pressing a left or a right button press to each picture stimulus that appeared on the computer screen. Each response provided the chance to either win or lose money, which was probabilistically determined. The probabilistic-learning task was designed to estimate reward-prediction errors (RPE), i.e. the difference between expected reward and actual reward which decreases with learning, and to measure the learning of stimulus-action-dependent reward prediction values predictions (SADRP), i.e., the association between visual input, action and reward which increases with

learning. RPE and SADRPs have been linked to caudate nucleus and putamen activity, respectively (Haruno & Kawato, 2006a, 2006b).

That is, in the initial stages of learning, SADRPs value is unknown. Thus, expectations about the potential reward of a decision are more often disappointed by the actual reward. The difference between expected and actual rewards, coined the reward-prediction error, is theorized to provide the feedback necessary to adjust decision-making strategies which corresponds to caudate activity. This global reward-related information (i.e. not fine-grained stimulus- and response-specific information) from the caudate is propagated to motor loops (which include the putamen) by means of a DA signal (subserved by reciprocal projections between the striatum and the substantia nigra; Haruno & Kawato, 2006b). At later stages of learning, activity in putamen increases to incorporate more specific motor information with the associated stimuli and expected reward, i.e. the reward associated with a specific stimulus and response becomes more predictable and learning is gradually fine-tuned (Haruno & Kawato, 2006b). As these SADRPs values are learned, the reward-prediction error is reduced as subjects more accurately anticipate the rewards associated with their actions.

A reinforcement model (Q-learning, Sutton & Barto, 1998) was used to estimate each participant's SADRPs and RPE during learning.

Similar to a previous study with this reward-based decision-learning task (Haruno & Kawato, 2006a, 2006b) participants' performance increased with increased stimulus-action-reward predictability. Moreover, dopaminergic medication in PD patients affected reward-prediction error and stimulus-action-reward-prediction value during probabilistic reward-based decision-learning. Both aspects of reward-based decision-learning, the evaluative component (RPE), measured in the first phase of learning, and SADRPs, measured at the end of learning, were improved by dopaminergic medication, but this depended on the amount of medication received. This suggests that their underlying neural structures caudate and putamen, and interaction between caudate and putamen during learning, benefit from dopaminergic modulations in PD patients.

There appears to be an optimal level of dopaminergic medication for these learning components; i.e. large daily doses may become suboptimal which suggests that there is a U-shaped curve of optimal performance (Schonberg, 2007; Cools et al., 2009) and this might also depend on individual baseline differences. According to the overdose hypothesis (Cools, Barker, Sahakian, & Robbins, 2001; Gotham, 1988), impairments on

reversal and extinction learning tasks (Cools et al., 2001; Czernecki et al., 2002; Swainson et al., 2000) in PD patients ON medication can be explained by the negative effects of an overdosed ventral striatum. Unlike in reversal learning, the learning functions in our study are not specifically associated with ventral striatum, but to putamen (dorsal) and caudate (ventral and dorsal), which may explain why we did not find impaired learning.

Frank's modeling work (2005) showed that PD patients OFF medication more effectively process negative feedback in comparison to positive feedback whereas PD patients ON medication show the opposite pattern. However, this model did not take individual differences into account (like medication dosage, DA genetics) or distinguish between feedback-based learning in different striatal structures like caudate and putamen. It remains to be investigated whether positive and negative feedback will differentially affect learning functions related to caudate and putamen in PD and on top of that, in what way genetic differences in PD patients modulate these effects.

STN stimulation differentially affected reward evaluation (RPE) and the degree of learning stimulus-action relations based on their outcomes (SADRP) during probabilistic reward-based decision-learning. Deep brain stimulation (DBS) of the STN in PD patients improved learning stimulus-action-reward relations. Without the stimulation, patients were less able to use outcomes of their actions adequately to change their behavior, which suggests that cognitive functions relying on the putamen benefit from DBS of the STN. Stimulating the motor area of the STN was predicted to affect cognitive and limbic loops in addition to the motor loop (Mallet et al., 2007), but it seemed to modulate the caudate related learning function (RPE) only when taking disease duration and age of the PD patients into account. The reduction in RPE early in the learning process with the stimulation ON compared to OFF was predicted by age and disease duration: the younger the patient and shorter the disease duration, the larger the beneficial effect of DBS STN on RPE.

Improved SADR learning with DBS of the STN is in line with findings of enhanced motor performance (Benabid et al., 1993; Kleiner-Fisman et al., 2003) and improved feedback-based learning (Funkiewiez, 2006) with STN stimulation in PD patients. It is also supported by STN lesions in rats indicating that STN lesions (note that STN stimulation effects are suggested to be comparable to lesioning effects) increases

'wanting' and thereby facilitates feedback-based learning (Uslaner & Robinson, 2006; Uslaner et al., 2008), particularly when the chance to receive positive reward is high.

In contrast with these findings, Frank (2007) did not reveal any effect of STN stimulation on either positive or negative feedback learning in PD patients as predicted by his model. In a DA depleted brain, like in PD, low levels of DA result in excessive activity in D2 NoGo striatal neurons which indirectly (via the GPe) removes the inhibition from the STN, resulting in an overactive STN and thereby stronger NoGo signal. Although there is an ongoing debate about the specific mechanisms underlying the therapeutic effect of STN stimulation (Benazzouz, & Hallett, 2000; Bergman, Wichmann, & DeLong, 1990; Liu, Postupna, Falkenberg, & Anderson, 2006; Meissner, et al., 2005), DBS is currently thought to inactivate the STN. This removes excessive activation of the GPi and thus disinhibits the thalamus, thereby facilitating thalamic excitation of the cortex. According to Frank's model (2007), the STN provides a global NoGo signal because projections from the STN to GPi are diffuse and not response specific. Thus STN stimulation is not predicted to affect learning specific stimulus-response associations.

However, the data in this study were not analyzed as a function of patient characteristics (like age or disease duration), which we showed to be important predictors of whether DBS STN leads to improvement. The task employed by Frank (2007) may also have been more difficult, or possibly less sensitive to effects on the later phases of learning.

Nevertheless, the exact mechanism that accounts for behavioral changes with DBS of the STN and the role of the STN remains unclear. In contrast with Frank's model, other BG models (e.g. Albin, Young & Penney, 1989) for example would have predicted that stimulating the STN in PD patients would impair NoGo learning but improve Go learning. Along the lines of these BG models, stimulating the STN could have reduced the excessive activity in the NoGo pathway in PD patients in our study and thereby improved SADRP learning.

The strongest effects of DBS in our study were found on the learning function associated with the putamen (SADRP) which might be explained by the placement of the stimulating electrodes in the motor areas of STN. Although stimulating the motor area of the STN may also affect cognitive and limbic loops, because STN output is not

sharply segregated (Mallet et al., 2007), the DBS effect may have been relatively stronger on regions within the corticostriatal motor loop.

The beneficial effect of STN stimulation on the putamen may have been established by STN influence on multiple sites within this motor loop. STN stimulation may have modulated the processing of motor input information from GPe (entering the GPe via the putamen). Moreover, STN is directly activated by projections from the motor cortex (hyperdirect pathway, Nambu et al., 2000). Thus, if several competing responses are active in the motor cortex, the STN becomes increasingly activated which leads to a global NoGo signal. Stimulating the STN may change the way these signals are processed, for example, if an already overactive STN in PD is excited by the motor cortex this leads to oscillatory activity and tremor, whereas stimulating or lesioning the STN normalizes this activity (Bergman et al., 1990).

Parametric modulation of STN stimulation in different functional STN areas might shed light on the modulating role of STN in reward-based decision-learning.

In sum, modulations of the BG, like dopaminergic medication and STN stimulation affect reward processing and associating an event with a response and an outcome. However, the improvement induced by dopaminergic medication largely depended on individual patient characteristics. DA modulation of the BG may be involved not only in gradual learning of SADRPs associations but may also be engaged in episodic binding of stimulus and response features which is suggested by preliminary results from an PD patients ON and OFF medication study performing an event-file task; in PD patients ON medications performance costs with partial repetitions of S-R associations were larger compared to OFF medication (Colzato, van Wouwe, Wylie, Band, Ridderinkhof, & Hommel, in preparation).

Taken together, the present thesis provides converging evidence that goal-driven biases as well as stimulus-driven biases both contribute to decision making in rapidly changing environmental demands. In a context-based decision making task like the AX-CPT, proactive goal-driven preparation seems to affect performance more than stimulus-driven preparation. Nevertheless, this top-down bias is modulated by previous context-probe associations. Goal-driven preparation is especially beneficial when a context is succeeded by a stimulus that was previously presented with this particular context compared to when a new or unusual stimulus appears. The beneficial effect of

this association-driven guidance of behavior was confirmed by the fMRI study; performance improved with increased activation in the brain areas representing the features previously associated with the current context.

In addition to the task demands that affect control processes in decision making, DA modulations alter decision making. DA increase by means of DA medication in PD patients and indirectly by means of positive affect in healthy controls seem to enable more flexible decision making either in terms of the ability to update currently task-relevant stimulus and response information (in healthy controls studies) or in terms of the ability to adequately associate stimuli and responses to gain reward on future decisions (in PD patients studies).

Moreover, different BG modulations (DBS of the STN and DA increase by medication) improve the capacity to associate stimulus action and reward; a function subserved by the putamen. Both BG treatments also improve the adequacy of evaluating reward early in the learning process, subserved by the caudate, although this effect depends on individual characteristics of the patients.

