8

Chemical Neuromodulation of Goal-Directed Behavior

Roshan Cools

Abstract

Directing our behavior adequately to current goals requires a trade-off between cognitive flexibility and cognitive stability. In this chapter, empirical data and theories are reviewed which show that this trade-off depends on optimal modulation of frontostriatal circuitry by the major ascending neuromodulatory systems of dopamine, noradrenaline, and acetylcholine. Highlighted are the roles of dopamine in (a) the prefrontal cortex in the stabilization of goal-relevant representations and (b) in the basal ganglia in the flexible updating of those representations. The cognitive neurochemistry of cognitive flexibility is, however, complex, with different forms of flexibility implicating subcortical and/or cortical dopamine, noradrenaline, and/or acetylcholine. The review concludes with a number of open questions raised by attempts to reconcile the different, complementary theories about the neurochemistry of the flexibility-stability trade-off.

Introduction

Our environment changes constantly. The ability to adapt flexibly to these constant changes is unique in humans. We can persist with current behavioral strategies as long as these seem optimal for goal achievement, yet we can also update our strategies flexibly when the need for change becomes sufficiently salient. How do our minds achieve this flexibility? This is not a straightforward issue, because only some of the changes around us are relevant and require cognitive flexibility. Most other changes are irrelevant (i.e., they represent noise) and should be ignored. In the latter case, adaptive behavior depends on cognitive stability rather than cognitive flexibility. What we need is an ability to regulate dynamically the balance between cognitive flexibility and cognitive stability depending on current task demands.

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

The trade-off between cognitive flexibility and stability is related to that between divided and focused attention (Hasselmo and Sarter 2011) as well as exploration and exploitation (Daw et al. 2006). With regard to the latter trade-off, exploration generally refers to active cognitive search for new, potentially better alternatives, whereas exploitation generally refers to the pursuit of what is currently known to be the best option (Daw et al. 2006). Exploration or cognitive search has been proposed to be triggered by changes in overall utility; that is, reductions in the overall perceived costs and benefits of ongoing behavior (Aston-Jones and Cohen 2005b). However, it might also be elicited by a salient, novel, or unexpected stimulus, an effect that has been captured by the concept of an "exploration bonus" assigned to such stimuli. For instance, imagine sitting at your desk, engaged in an e-conversation with a colleague, when a fire breaks out in your building. How do our minds decide when the environmental change is sufficiently salient to trigger flexible attention shifting? And how do we make sure that we do not respond to every little distracting sensory event in our office? Setting the threshold adequately for such attention shifting (to external events in the environment or internal events in working memory) is critical for optimal goal-directed behavior and requires cognitive control.

The brain region that has been associated most commonly with cognitive control is the prefrontal cortex (PFC). We know that this region does not act in isolation to bias cognitive control, but rather interacts with a set of deep brain subcortical structures, in particular the striatum, in so-called frontostriatal circuits. Processing in these circuits is extremely sensitive to modulation by the major ascending neuromodulators—dopamine, noradrenaline, acetylcholine, and serotonin—which is not surprising given diffuse ascending inputs from the brainstem to both the PFC and various subcortical structures. The widely distributed and diffuse nature of these neuromodulatory projections has led many investigators to assume that they serve relatively nonspecific functions, such as arousal and sleep-wake cycle regulation. In this chapter, I review some current ideas about the role of these neuromodulators, in particular dopamine and to a lesser degree noradrenaline and acetylcholine, in cognitive flexibility and stability, which suggest that they serve more specific functions in goal-directed behavior. I begin by highlighting the role of dopamine in the PFC in the stabilization of goal-relevant representations. Then I describe evidence for a role of dopamine in the basal ganglia (BG) in a functionally opponent component process (i.e., the flexible updating of goal-relevant representations). Critically, I end by pointing out that this distinction is likely oversimplified, and that a full understanding of the neurochemistry of cognitive flexibility requires us to take into account the degree to which such flexible updating of goal-relevant representation involves top-down, goal-directed search, associated with the PFC, versus habitual control mechanisms, associated with the BG.

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

Neurochemical Modulation of the Prefrontal Cortex and the Stabilization of Goal-Relevant Representations

The neurochemical mechanisms of the stability component of the flexibility-stability trade-off are potentially somewhat better understood than are those of the flexibility component. Indeed, one of the best known functions of the PFC is the active stabilization of goal-relevant representations, an important component process of working memory (Baddeley 1986; Fuster 1989; Goldman-Rakic 1995). The importance of the PFC for working memory was first demonstrated by Jacobsen (1936), who showed that monkeys with frontal lobe lesions were impaired on the well-known delayed response task. Electrophysiological work with monkeys supported the primate lesion work by demonstrating that the firing of PFC neurons persists throughout the delay of delayed response tasks (Fuster and Alexander 1971), even in the face of distraction. Further, functional magnetic resonance imaging (fMRI) studies with human volunteers have revealed similarly persisting responses in the human PFC during delayed response tasks (Curtis and D'Esposito 2003). According to current ideas, these persistent responses during working memory tasks might correspond to the influence of excitatory top-down signals in the PFC, which bias the competition among brain regions in posterior sensory cortex. These PFC signals may increase the activity of brain regions processing goal-relevant representations and, by virtue of mutual inhibition, suppress activity of brain regions processing irrelevant representations (Miller and Cohen 2001).

In keeping with the pronounced sensitivity of the PFC to modulation by dopamine, there is extensive empirical support for an important role of dopamine, in particular D1 receptor (D1R) stimulation, in the PFC in these aspects of working memory (Goldman-Rakic 1995). Administration of the dopamine receptor agonist bromocriptine to healthy volunteers altered signal change in the PFC during distractor resistance in a working memory task (Cools et al. 2007b) (Figure 8.1). This paralleled effects of global dopamine depletion in the nonhuman primate PFC on task performance, which was more susceptible to distraction than that of control monkeys (Crofts et al. 2001). Although the actual mechanism by which dopamine alters stabilization of working memory representations requires further empirical study, hypotheses have been put forward based on *in vitro* electrophysiological and computational modeling work. Specifically, effects of D1R stimulation on cognitive stabilization might reflect dopamine-induced increases in the signal-to-noise ratio of neuronal firing in the PFC (Servan-Schreiber et al. 1990), leading to increased robustness of these representations in the face of intervening distractors (Durstewitz and Seamans 2008). For instance, recent neurophysiological data from monkeys (Vijayraghavan et al. 2007) have shown that D1 receptor stimulation in the nonhuman primate PFC improves the spatial tuning of cells during the performance of a spatial delayed response task by blocking task-irrelevant firing. The finding that dopamine-induced improvements of spatial tuning are

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

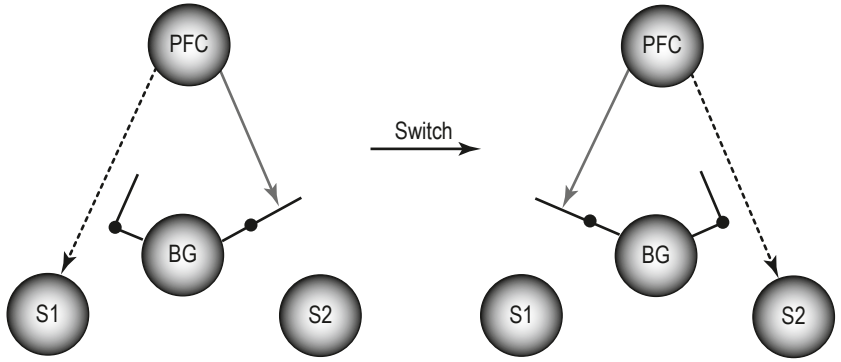

Figure 8.1 Schematic illustration of the working hypothesis that the basal ganglia (BG) control attention shifting by regulating top-down projections from prefrontal cortex (PFC) to posterior sensory areas. The PFC biases information processing in favor of posterior sensory regions that support currently goal-relevant representations (e.g., S1) away from regions that support currently goal-irrelevant representations (e.g., S2). In the model, this top-down control mechanism mediated by the PFC is in turn regulated by the BG, which implement a shift in attention (e.g., in response to novel salient stimuli) by closing the gate to one region (e.g., S1) while simultaneously opening the gate to another region (e.g., S2). Redrawn, with permission, after van Schouwenburg, Aarts, and Cools (2010a).

accompanied by suppressive effects on the firing of PFC cells concurs with the general observation from human neuroimaging that working memory improvement after dopamine-enhancing drug administration is accompanied by reductions in PFC activity.

Research indicates that the stabilization of goal-relevant representations depends not only on dopamine, but also on noradrenaline and acetylcholine transmission, possibly via modulation of attention (Arnsten 2009) and uncertainty signals (Yu and Dayan 2005), respectively. In the case of noradrenaline, for example, Arnsten (2009) has shown that the ability of a network of neurons to maintain firing over a delay period is strengthened by noradrenergic α2A receptor stimulation. According to her recent proposal (Arnsten 2009), dopamine and noradrenaline might subserve complementary roles in cognitive stabilization with α2A receptor stimulation enhancing network firing for shared inputs, thus increasing "signal," and D1 receptor stimulation sculpting neuronal firing by decreasing firing to nonpreferred inputs, thus decreasing "noise." In the case of acetylcholine, several cellular effects could contribute to the cholinergic enhancement of the stabilization of goal-relevant representations, including muscarinic receptor stimulation-induced persistence of spiking activity of PFC cells (Hasselmo and Sarter 2011).

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

Role of Dopamine in the Basal Ganglia in Cognitive Updating of Goal-Relevant Representations

The previous section highlighted the importance of dopamine, in particular, in the PFC for the stabilization of goal-relevant representations as well as for the filtering of new input that might be irrelevant to ongoing processing. One could say that the net effect of dopamine in the PFC is an elevation of the threshold for a new representation to be selected. Of course, this is adaptive when new input is irrelevant. However, it is maladaptive when new input is relevant. In this case, existing goal-relevant representations need to be flexibly updated rather than protected. Accumulating evidence indicates that dopamine is also implicated in this complementary updating aspect of cognitive control. Current theorizing suggests, however, that these effects of dopamine on updating might implicate not only the PFC but also, at least in some conditions, the BG (Frank 2005).

The proposal that dopamine in the BG subserves the flexible updating of goal-relevant representations fits with the traditional view of the BG as a selection or threshold-setting device, gating task-relevant representations to the PFC via the direct Go pathway, while simultaneously inhibiting competing task-irrelevant representations via the indirect NoGo pathway (Frank 2005; Mink 1996). Interestingly, dopamine has opposite effects on these two pathways, increasing activity in the direct BG pathway while suppressing activity in the indirect BG pathway. The net effect is a lowering of the threshold for a representation to be selected. This hypothesis is in line with suggestions that dopamine signals mediate the switching of attention to unexpected, behaviorally relevant stimuli (Redgrave et al. 1999) and more generally concurs with a rapidly growing body of data which shows BG involvement during updating of working memory representations (e.g., Dahlin et al. 2008). Furthermore, it is also consistent with empirical data that reveal effects of BG dopamine manipulations on set shifting (Haluk and Floresco 2009; Kellendonk et al. 2006). Furthermore, administration of the dopamine D2 receptor agonist bromocriptine to healthy volunteers altered signals related to set shifting in the BG, but not in the PFC (Cools et al. 2007b) (Figure 8.1). This finding paralleled later findings that behavioral effects of bromocriptine on set shifting could be predicted from baseline levels of dopamine in the BG (Cools et al. 2009) as well as selective set-shifting deficits in patients with BG dysfunction (Cools 2006).

As in the case of the modulation of the stabilization of working memory representations, the mechanism by which dopamine alters set shifting requires further empirical study. However, integration of ideas about the role of the PFC in top-down attention biasing and of the BG in selective gating raises the possibility that the BG facilitate set shifting by gating interactions between the PFC and posterior sensory cortex, thus controlling the top-down biasing of competition between goal-relevant and goal-irrelevant representations (Figure 8.2). This hypothesis is reminiscent of ideas that the (attentional or motor) output

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

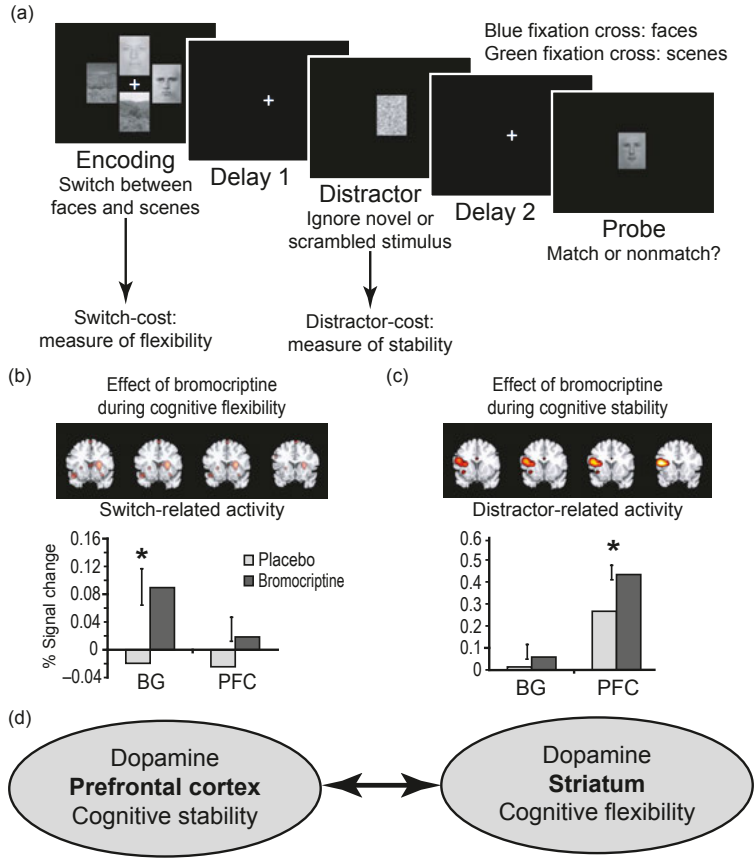

Figure 8.2 The effects of dopamine receptor stimulation depend on task demands and the neural site of modulation. (a) A delayed match-to-sample (DMS) task was used that provided a measure of cognitive flexibility (attention shifting during encoding) as well as a measure of cognitive stability (distractor resistance during the delay). Subjects memorized faces or scenes, depending on the color of the fixation cross. If the cross was blue, then subjects memorized the faces; if it was green, then they memorized the scenes. Subjects occasionally shifted during encoding between attending to faces and scenes. A distractor was presented during a delay. Subjects were instructed to ignore this distractor. (b) Top panel: effects of bromocriptine on basal ganglia (BG) activity during shifting, as a function of trait impulsivity. Whole-brain contrast values (>25) are overlaid on four coronal slices from the Montreal Neurological Institute high-resolution single-subject magnetic resonance image. Bottom panel: effects of bromocriptine on shift-related activity in the BG and left prefrontal cortex (PFC) in high-impulsive subjects only. (c) Top panel: effects of bromocriptine on PFC activity during distraction as a function of trait impulsivity (all contrast values >25 shown). Bottom panel: effects of bromocriptine on distractor-related activity in the BG and left PFC in high-impulsive subjects only. (d) Schematic representation of the hypothesis that dopamine modulates cognitive flexibility by acting at the level of the BG while modulating cognitive stability by acting at the level of the PFC. Reprinted with permission from Cools and D'Esposito (2011).

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

of the PFC can be gated by dopamine-dependent activity in the striatum (Hazy et al. 2007). Evidence for such output gating by dopamine in the BG came from a recent fMRI study, in which subjects shifted attention between the faces and the scenes of overlapping face/scene stimuli (van Schouwenburg et al. 2010b). The attention shifts were accompanied by potentiation of goal-relevant representations relative to goal-irrelevant representations in stimulus-specific posterior visual cortex (fusiform face area and parahippcampal place area), presumably reflecting top-down biases from the PFC. Effective connectivity analyses revealed that the BG indeed played a critical role in regulating these attention shifts by gating the top-down bias from the PFC on stimulus-specific posterior cortex. Dopamine could alter such top-down biasing of competition between goal-relevant and goal-irrelevant representations via stimulation of dopamine receptors on neurons in the BG, altering the balance between activity in the Go and NoGo pathways of the BG and lowering the threshold for gating top-down influences. Preliminary evidence concurs with this hypothesis, and showed that dopamine receptor stimulation with a dopamine receptor agonist in humans modulates activity in the BG, but not the PFC during attention shifting in this paradigm (van Schouwenburg et al., unpublished data). These data suggest that dopamine might modulate set shifting at the level of the BG (e.g., by modulating flow through frontostriatal circuits), and generally concur with empirical evidence from genetic and neurochemical imaging work, which reveals that variation in striatal dopamine function is associated with altered neural efficiency (Crofts et al. 2001) in the PFC and associated working memory updating and attention switching (Kellendonk et al. 2006; Landau et al. 2009; Nyberg et al. 2009; Stelzel et al. 2010).

Dopamine is not the only neuromodulator that modulates attention shifting. For example, drug-induced enhancement of noradrenaline activity has also been shown to potentiate attention shifting to motivationally significant stimuli in a manner fairly similar to dopamine (Sara 2009). Salient events are known to elicit both phasic noradrenaline and dopamine responses. As with dopamine, this orienting of attention to salient stimuli has also been compared with a temporary lowering of a decision threshold. Furthermore very similar ideas have been put forward to account for effects of acetylcholine on Posner target detection tasks (Hasselmo and Sarter 2011). Specifically, it has been argued that a salient target, which has been found to evoke phasic acetylcholine release in the PFC, may elicit an attentional shift akin to Posner's attentional orienting response, in order to align attention with a source of sensory input. Acetylcholine could do this by enhancing sensory input from the thalamus to the PFC and, at the same time, shutting down top-down suppression from the PFC. Interestingly, regions in the ventral parts of the BG that are strongly innervated by dopamine can selectively influence cholinergic modulation of thalamic sensory inputs to the PFC. Thus another mechanism by which BG dopamine might facilitate attention shifting is by gating acetylcholine-dependent

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

interactions between the PFC, the thalamus, and stimulus-specific sensory areas in posterior cortex.

Trading off Flexibility and Stability

Cognitive flexibility and stability might be conceptualized as representing functionally opposing processes. If we update too readily, then we are likely to get distracted, rendering our behavior unstable. Conversely, if our representations are overly persistent or stable, then there is a danger of inflexibility and unresponsiveness to new information. Empirical data support the hypothesis that these two opponent processes might be subserved by dopamine in the BG and the PFC, respectively. Roberts and colleagues (Robbins and Roberts 2007) injected the neurotoxin 6-OHDA into the BG or PFC of nonhuman primates and showed that, while dopamine lesions in the PFC *improved flexibility* (attentional set shifting), dopamine lesions in the BG actually *impaired flexibility* (attentional set shifting). Subsequent work showed that this modulation of flexibility during attentional set shifting may have resulted from effects on performance during the preceding set-maintenance stages of the task (Crofts et al. 2001). Specifically, that subsequent study revealed that dopamine lesions in the PFC led to *enhanced distractibility* (poor attentional set maintenance), whereas dopamine lesions in the BG actually *reduced distractibility* (enhanced attentional set maintenance). Thus the contrasting effects on set maintenance may well underlie the contrasting changes measured in the subsequent attentional set-shifting stages of the task. Interestingly, an analogous observation was recently made in Parkinson's disease patients, who exhibit relatively selective dopamine depletion in the BG. These patients exhibit not only impaired set shifting on a variety of tasks but also enhanced distractor resistance (Cools et al. 2010a). Overall, the opposing effects of BG and frontal dopamine lesions suggest that a dynamic balance between cognitive stability and flexibility may depend on precisely balanced dopamine transmission within the PFC and the BG, respectively. The functional opponency between stability and flexibility maps well onto the neurochemical reciprocity between dopamine in the PFC and the BG. Increases and decreases in PFC dopamine lead to decreases and increases in BG dopamine, respectively (Kellendonk et al. 2006; Pycock et al. 1980).

This working hypothesis is reminiscent of the dual-state theory put forward recently by Durstewitz and Seamans (2008), which is grounded in *in vitro* neurophysiology, biophysically realistic computational modeling work, as well as empirical pharmacological work (Floresco et al. 2006). According to this theory, PFC networks can be either in a D1-dominated state, which is characterized by a high-energy barrier that favors robust stabilization of representations, or in a D2-dominated state, characterized by a low-energy barrier favoring fast, flexible shifting between representations. This alternative receptor-based

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

theory is not necessarily inconsistent with the presented working hypothesis, according to which dopamine in the BG and the PFC subserve the distinct roles of flexibility and stability, respectively, particularly given the observation that D2 receptors are more abundant in the BG than in the PFC, which contains fewer D2 than D1 receptors.

Neurochemical Modulation of Cognitive Flexibility and Exploration

At first glance, the hypothesis that dopamine modulates certain forms of set shifting by acting at the level of the BG rather than the PFC is perhaps incompatible with traditional notions that effects of dopamine on high-level cognitive control are mediated by the PFC. In fact, not all forms of set shifting depend on dopamine in the BG. For example, although several studies have observed sensitivity of the set-shifting deficit in Parkinson's patients to withdrawal of dopaminergic medication (Cools 2006), other studies have failed to reveal such dependency on dopamine in the BG (Kehagia et al. 2010). Similarly, while a range of pharmacological neuroimaging studies has revealed selective modulation of BG signals by dopamine during set shifting, several other pharmacological studies have revealed effects of dopamine in the PFC during set shifting.

One possible explanation is that the extent to which flexible behavior implicates (neuromodulation of) the BG or the PFC depends on the degree of exploration, or cognitive search, required for the type of set shifting assessed. This observation concurs with recent evidence which indicates that the catecholamine-*O*-transferase gene, which primarily controls dopamine in the PFC, affects exploratory decisions during a learning task (Frank et al. 2009). Furthermore, a recent microdialysis study (van der Meulen et al. 2007) demonstrated increased catecholamine release in the PFC during serial reversal learning, an effect that was particularly pronounced in the early stages of the task, when reversals presumably required a relatively greater degree of exploration than during the late stages of the task. Conversely, a task that is disproportionally sensitive to dopaminergic medication in Parkinson's disease, associated with BG dopamine depletion, is the task-switching paradigm, where switches are externally cued, thus requiring little to no cognitive search. Demands for cognitive search are particularly low in some versions of this paradigm (e.g., those requiring switches to naming the direction of the arrow of an arrow/word stimulus), in which task sets are well established. It is these "habitual" shifts that are sensitive to dopaminergic medication in Parkinson's disease. The same medication in Parkinson's disease, however, has no effect on other versions of this paradigm, such as those requiring switches to poorly established task sets (e.g., classifying digits as odd or even, versus high or low), when demands for cognitive search might be enhanced (Kehagia et al. 2010). A similar argument might be put forward when considering the insensitivity to BG dopamine

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

of performance on Wisconson card sort-like tasks, such as extra-dimensional set shifting (EDS), which requires cognitive search for a newly rewarded stimulus according to changes in the relevance of stimulus dimensions. Both dopaminergic medication in Parkinson's disease and BG dopamine lesions in nonhuman primates leave unaffected performance on an initial EDS (Lewis et al. 2005). By contrast, a subsequent EDS back to the originally relevant attentional set is severely impaired by dopamine lesions in the BG (Collins et al. 2000), consistent with the dopamine-dependent deficit seen in Parkinson's patients during task switching between well-established sets. Another form of set shifting that seems to critically involve dopamine in the BG is reversal learning (Clatworthy et al. 2009; Cools et al. 2007a, 2009). In the traditional version of this task, a negative prediction error encountered upon contingency reversal, due to choice of the previously rewarded stimulus, also implies that the nonchosen stimulus is now rewarded. Accordingly, demands for exploration, or search, in traditional tasks of reversal learning are relatively low. Instead, adequate reversal learning depends on the optimal pursuit of what is currently known, based on experience, to be the best option (i.e., exploitation).

The hypothesis that BG dopamine is concerned with forms of set shifting that do not involve exploration or cognitive search, but rather only exploitation of learned information, concurs with the well-known implication of dopamine and the BG in "model-free" reinforcement learning (i.e., trial-and-error learning to maximize rewards). Conversely, the hypothesis that the PFC (and its neuromodulation) is concerned with forms of set shifting that implicate exploration concurs with empirical neuroimaging data (Daw et al. 2006) as well as with current theories about the role of prefrontal neuromodulation in exploration (Aston-Jones and Cohen 2005b). In particular, Aston-Jones and Cohen have invoked the adaptive gain theory, according to which different modes of noradrenaline transmission regulate the trade-off between exploitation and exploration. In this model, a high phasic mode promotes exploitative behavior and focused attention by facilitating processing of task-relevant information, whereas a low tonic noradrenaline mode ensures that irrelevant stimuli are filtered. Increasing the tonic mode promotes behavioral disengagement and divided attention, thus allowing potentially new and more rewarding behaviors to be explored. The transition from the phasic to the tonic noradrenaline mode is controlled by specific regions in the PFC (i.e., the orbitofrontal cortex and the anterior cingulate cortex), which in turn control the firing of noradrenaline neurons in the brainstem in a top-down manner.

The notion that (tonic) cortical noradrenaline is particularly important for explorative modes of behavior concurs with empirical findings from work with experimental animals as well as humans, which show that EDS is sensitive to manipulation of (tonic) noradrenaline transmission (Robbins and Roberts 2007). This series of findings also raises the possibility that the effect of nonspecific catecholamine modulation of EDS reflects modulation by noradrenaline rather than dopamine. Furthermore, the dopamine-insensitive EDS deficit

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

in Parkinson's patients, which is restricted to conditions that require shifting to a dimension that is not very salient (thus maximizing demands for cognitive search; Cools et al. 2010b), might also be mediated by frontoparietal cortical abnormalities in catecholamine (e.g., noradrenaline) neurotransmitter systems rather than BG dopamine dysfunction.

The adaptive gain theory emphasizes the importance of noradrenaline for exploration and is complementary to a different influential proposal that tonic noradrenaline activity serves a neural interrupt or network reset function, thus enabling the interruption of ongoing activity, or revision of internal representations, based on new sensory input (Yu and Dayan 2005). A unique feature of the model by Yu and Dayan is that it predicts noradrenaline to be involved predominantly when changes in the environment are unexpected (as opposed to expected). In their conceptualization, unexpected uncertainty is induced by gross changes in the environment that produce sensory observations strongly violating top-down expectations, as in the case of EDS. This is contrasted with expected uncertainty, which arises from known unreliability of predictive relationships within a familiar environment (Yu and Dayan 2005). Critically, they argue that expected uncertainty is signaled by acetylcholine, a stance that is consistent with observations, mentioned earlier, that cholinergic changes are associated with attentional shifts in Posner-like attention-orienting paradigms where subjects are aware of cue invalidity (Hasselmo and Sarter 2011). By contrast, cholinergic manipulations generally leave EDS unaffected. Thus according to these ideas, both increases in (tonic) noradrenaline and acetylcholine align attention with a source of sensory input, by enhancing sensory input from the thalamus to the PFC and by shutting down top-down internal models held online by the PFC. However, the signals that trigger this noradrenaline- and acetylcholine-mediated flexibility might differ. The theory is generally consistent with observed sensitivity of EDS to noradrenaline, but not acetylcholine. Furthermore, it also concurs with observed sensitivity to acetylcholine, but not noradrenaline, of (late but not early) reversal learning (Chamberlain et al. 2006; Robbins and Roberts 2007).

Conclusion and Open Questions

The empirical data and theories reviewed in this chapter indicate that the balance between cognitive flexibility and stability depends critically on modulation by the major ascending neuromodulatory systems. I have focused on the roles of dopamine, but also mentioned those of noradrenaline and acetylcholine. While cognitive stabilization is well established to depend critically on D1R stimulation in the PFC, the literature on the cognitive neurochemistry of cognitive flexibility is more complex, with striatal dopamine, and frontal noradrenaline and acetylcholine being important for different forms of shifting. An understanding of these apparent discrepancies requires us to recognize that

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

cognitive flexibility is not a unitary phenomenon, with distinct forms of flexibility implicating different cortical and subcortical neurochemical mechanisms.

One factor that might be taken into account when assessing the neurochemical mechanisms of flexibility is the degree of exploration, or search, for new, potentially better alternatives as opposed to the exploitative pursuit of what is currently known to be the best option. Explorative forms of shifting that involve cognitive search, such as EDS, seem more sensitive to catecholaminergic modulation of the PFC, in particular by noradrenaline, whereas exploitative (or habitual) forms of shifting that do not involve cognitive search (e.g., certain forms of task switching and reversal learning) seem more sensitive to dopaminergic modulation of the BG. Future work should address the further question of whether, and if so how, issues of unexpected versus expected uncertainty relate to issues of explorative versus exploitative shifting. For instance, the disproportionate sensitivity to cholinergic manipulations of late versus early reversals (Robbins and Roberts 2007) might be interpreted to reflect the reduced degree of exploration required for late versus early reversals. However, it might also reflect the fact that late reversals are more expected than are early reversals. Similarly, the disproportionate catecholamine release in the PFC during early versus late reversals (van der Meulen et al. 2007) might reflect the greater degree of exploration required for early versus late reversals, but it might also reflect the fact that early shifts are less expected than are late reversals. Finally, along the same lines, one might also raise the question whether "habitual" shifting, such as task-set switching and repeated EDSs, are disproportionately sensitive to dopamine in the BG due to the fact that such paradigms involve relatively little cognitive search, or rather because the uncertainty that triggers these "habitual" shifts is more expected than unexpected.

A further factor that should be taken into account in future cognitive neurochemical work concerns the hierarchical nature of cognitive search. Search goals can be defined at different levels of abstraction, something that is well illustrated by the difference between intra-dimensional shifting (IDS) and EDS. Both types of shift have relatively high demands for cognitive search and both are triggered by relatively unexpected uncertainty. However, IDS involves changes within a stimulus dimension (novel exemplars, e.g., yellow or blue), whereas EDS involves changes between stimulus dimensions (e.g., shape or color). This factor of hierarchy may become relevant when considering findings that (tonic) noradrenaline manipulations affect EDS, but not exploration of changes along one and the same stimulus dimension in a four-arm bandit task (Jepma et al. 2010).

More generally, it is clear from the above that cognitive approaches to neurochemistry have revealed that dopamine, noradrenaline and acetylcholine likely serve more specific functions in goal-directed behavior than has been traditionally assumed. This specificity arises in part from the different computations that are carried by the targeted regions, which differ in receptor distribution, but also reflects most likely a number of other factors that were not

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

addressed explicitly in this chapter. These factors include the computations carried out by brain structures that control the ascending systems in a top-down manner, the baseline dependency of the neuromodulatory effects, and the (phasic versus tonic) timescale of neurotransmitter effects. The particular importance of considering the timescale of neuromodulatory effects is illustrated by the adaptive gain theory of Aston-Jones and Cohen (2005b), which attributes distinct exploratory and exploitative functions to the tonic and phasic modes of noradrenaline transmission. However, the timescale of neurotransmission also plays a central role in current thinking about dopamine (Niv et al. 2007), acetylcholine as well as serotonin (Cools et al. 2011). These modes may serve partly antagonistic and partly synergistic roles, the latter possibly realized by synaptic overflow from phasic events followed by slower reuptake. For example, the reward-and-punishment prediction error signals that reinforcement learning theories hypothesize to be carried by phasic dopamine and serotonin responses, respectively, might also contribute, when averaged slowly over time, to response vigor or action threshold setting by measuring average reward and punishment rate (Cools et al. 2011). Clearly, it will be crucial to obtain better insights in the degree to which commonly used neurochemical manipulations affect phasic versus tonic transmission.

In conclusion, future work will benefit from adopting a cognitive mechanistic approach to neurochemistry, to allow us to move beyond apparent discrepancies between theories of dopamine, noradrenaline and acetylcholine in terms of cognitive control, attention, working memory, or learning. This is pertinent given the implication of most neuromodulators in all of these processes and will help us further define the computational nature of the flexibility-stability paradox.

Acknowledgments

Supported by the Netherlands Organization for Scientific Research (Vidi grant 016.095.340), the Dutch Brain Foundation, a Hersenstiching fellowship F2008(1)-01), the National Institute of Health (R01 DA020600), and the Human Frontiers Science Program (RGP0036/2009-C).

First column (top to bottom): Trevor W. Robbins, Roshan Cools, A. David Redish, Cyriel M. A. Pennartz, John P. O'Doherty, Christian Büchel
Second column: Cyriel M. A. Pennartz, Trevor W. Robbins, A. David Redish, Jeremy K. Seamans, Daniel Durstewitz, Joshua W. Brown, Catharine A. Winstanley
Third column: Catharine A. Winstanley, Bernard W. Balleine, Trevor W. Robbins, Bernard W. Balleine, Joshua W. Brown, Roshan Cools, Jeremy K. Seamans

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

From “Cognitive Search: Evolution, Algorithms, and the Brain,” edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.