



Universiteit
Leiden
The Netherlands

Arousal, exploration and the locus coeruleus-norepinephrine system
Jepma, M.

Citation

Jepma, M. (2011, May 12). *Arousal, exploration and the locus coeruleus-norepinephrine system*. Retrieved from <https://hdl.handle.net/1887/17635>

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/1887/17635>

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

Arousal, exploration and the locus coeruleus-norepinephrine system

Marieke Jepma

ISBN 978-94-91211-30-0

Copyright © Marieke Jepma, 2011

Cover design: Hanneke and Marieke Jepma

Printed by Ipskamp Drukkers B.V. Amsterdam

The research presented in this thesis was supported by a VIDI grant from the Netherlands Organization for Scientific Research (NWO) to Sander Nieuwenhuis.

Printing of this thesis was supported by BioSemi.

Arousal, exploration and the locus coeruleus-norepinephrine system

Proefschrift ter verkrijging van
de graad van Doctor aan de Universiteit Leiden,
op gezag van de Rector Magnificus Prof. mr. P. F. van der Heijden,
volgens besluit van het College voor Promoties
te verdedigen op donderdag 12 mei 2011
klokke 11:15 uur

Door
Marieke Jepma
geboren te Utrecht
in 1983

Promotiecommissie

Promotor Prof. Dr. B. Hommel

Co-promotor Dr. S. Nieuwenhuis

Overige leden Dr. G. P. H. Band

Dr. R. Cools, Radboud Universiteit Nijmegen

Prof. Dr. R. de Jong, Rijksuniversiteit Groningen

Prof. Dr. L. J. Kenemans, Universiteit Utrecht

Prof. Dr. K. R. Ridderinkhof, Universiteit van Amsterdam

Prof. Dr. S. A. R. B. Rombouts

Contents

Chapter 1:	General introduction	7
Chapter 2:	Pupil diameter predicts changes in the exploration-exploitation trade-off: Evidence for the adaptive gain theory	17
Chapter 3:	The role of the noradrenergic system in the exploration-exploitation trade-off: A psychopharmacological study	37
Chapter 4:	Neurocognitive function in dopamine- β -hydroxylase deficiency	63
Chapter 5:	Neural mechanisms underlying the induction and relief of perceptual curiosity	97
Chapter 6:	The effects of accessory stimuli on information processing: Evidence from electrophysiology and a diffusion-model analysis	113
Chapter 7:	Temporal expectation and information processing: A model-based analysis	141
Chapter 8:	Summary	167
	References	170
	Nederlandse samenvatting	191
	Dankwoord	199
	Curriculum Vitae	200

Chapter 1

General introduction

Part of this chapter is based on: Nieuwenhuis, S., & Jepma, M. (2010). Investigating the role of the noradrenergic system in human cognition. In T. Robbins, M. Delgado, & E. Phelps (Eds.), *Decision making. Attention & Performance, Vol. XXIII*. Oxford: Oxford University Press.

The locus coeruleus-norepinephrine system

As their name suggests, neuromodulators such as dopamine, acetylcholine and norepinephrine modify the effects of neurotransmitters—the molecules that enable communication between neurons. Neuromodulatory systems are involved in almost every mental function, including attention, learning and emotion (Robbins, 1997), and they are disturbed in many neurological and psychiatric disorders, such as attention-deficit/hyperactivity disorder (ADHD), post-traumatic stress disorder, and schizophrenia. This thesis focuses specifically on the role of the noradrenergic system in human cognition and brain function.

The locus coeruleus (LC) is the brainstem neuromodulatory nucleus responsible for most of the norepinephrine (NE) released in the brain. The LC has widespread projections throughout the neocortex, thalamus, midbrain, cerebellum and spinal cord (Aston-Jones, Foote, & Bloom, 1984; Berridge & Waterhouse, 2003). The LC-mediated noradrenergic innervation increases the responsiveness of efferent target neurons (Berridge & Waterhouse, 2003), which can be modeled as a change in the gain (steepness) of the neurons' activation function (Servan-Schreiber, Printz, & Cohen, 1990). Although cell recordings in non-human primates have yielded a wealth of information regarding the dynamics of the noradrenergic system, to date there has been very little empirical research on the activation dynamics and function of this system in humans. This is not so surprising since the study of the noradrenergic system in humans poses considerable methodological challenges. For example, it is not possible to directly measure the neurophysiological effects of NE in the human brain. The study of these effects requires the development of indirect measures, or the measurement of changes in behavior and brain activity brought about by pharmacological manipulations of the noradrenergic system.

The adaptive gain theory of LC-NE function

For a long time researchers have associated the LC-NE system with basic, nonspecific functions such as regulating arousal and the sleep-wake cycle (Aston-Jones et al., 1984; Jouvet, 1969). But recent research has shown that neuromodulators have more specific functions in the control of behavior (e.g., Aston-Jones & Cohen, 2005; Sara, 2009). According to an influential recent theory, the adaptive gain theory (Aston-Jones & Cohen, 2005), the LC-NE system has a critical role in the optimization of behavioral performance—by facilitating responses to motivationally significant stimuli and regulating the tradeoff between exploitative and exploratory behaviors. The adaptive gain theory is largely based on neurophysiological observations in behaving animals, which will be described in the following sections.

The function of the phasic LC response

When an animal is actively engaged in performing a task, LC neurons exhibit a rapid, phasic increase in discharge rate to task-relevant and otherwise motivationally salient stimuli. For example,

such *LC phasic responses* are observed for target stimuli in a simple target-detection task in which monkeys are required to respond to rare target stimuli presented at random intervals embedded in a train of distractor stimuli. Provided that the animal is engaged in the task, these target stimuli cause a phasic increase in LC firing rate that peaks approximately 100-150 ms post-target and approximately 200 ms prior to the response (e.g., Aston-Jones, Rajkowski, Kubiak, & Alexinsky, 1994; Clayton, Rajkowski, Cohen, & Aston-Jones, 2004). Importantly, the LC does not exhibit this type of phasic response to distractor stimuli, nor is the phasic response associated with any other task-related events once training is complete (reward delivery, fixation point, response movements, etc.). However, similar phasic responses are elicited by unexpected, intense, threatening, or otherwise salient stimuli that demand effective processing and action (Aston-Jones, Rajkowski, & Cohen, 1999). The ensuing release of NE in cortical areas temporarily increases the responsivity of these areas to their afferent input (Berridge & Waterhouse, 2003). When applied in a temporally strategic manner (e.g., when driven by the identification and evaluation of motivationally relevant stimuli), increases in responsivity produce an increase in the signal-to-noise ratio of subsequent processing and a concomitant improvement in the efficiency and reliability of behavioral responses (Servan-Schreiber et al., 1990). Accordingly, it has been found that LC phasic activation reliably precedes and is temporally linked to behavioral responses to task-relevant stimuli (Bouret & Sara, 2004; Clayton et al., 2004). In addition, studies have reported a direct relation between the strength of LC activity and response accuracy in choice-reaction time tasks (Rajkowski, Majczynski, Clayton, & Aston-Jones, 2004). Together, these findings suggest that phasic noradrenergic signals play an important role in optimizing responses to motivationally significant stimuli.

Phasic versus tonic LC firing mode and corresponding control states

Besides the phasic increases in activity following motivationally significant stimuli, there are also tonic (baseline) changes in LC activity (i.e. changes happening over the course of multiple seconds or minutes). Levels of LC tonic activity vary systematically in relation to measures of task performance (Figure 1). Aston-Jones and colleagues (1994) recorded LC activity in monkeys during performance of a target-detection task. Periods of intermediate tonic LC activity were accompanied by large LC phasic responses to target stimuli, and rapid and accurate responding. In contrast, periods of elevated tonic LC activity were consistently accompanied by relatively poor task performance, and distractible, restless behavior. Such phases were also consistently associated with a diminution or absence of the target-evoked LC phasic responses observed during periods of good performance. These findings have led to the proposal that in the waking state there are two distinguishable modes of LC activity (Aston-Jones et al., 1999; Figure 1): In the *phasic mode*, bursts of LC activity are observed in association with the outcome of task-related decision processes, and are closely associated with goal-directed behavior. In the *tonic mode*, LC baseline activity is elevated but phasic bursts of activity are absent and behavior is more distractible.

According to the adaptive gain theory (Aston-Jones & Cohen, 2005; Cohen, Aston-Jones, & Gilzenrat, 2004), the different modes of LC activity serve to regulate a fundamental tradeoff

between two control states: exploitation and exploration. The LC phasic mode promotes exploitative behavior by facilitating processing of task-relevant information (via the phasic response), while filtering out irrelevant stimuli (through low tonic responsivity). By increasing the phasic character of LC firing, the cognitive system is better able to engage in the task at hand, and maximize rewards harvested from this task. In contrast, the LC tonic mode promotes behavioral disengagement by producing a more enduring and less discriminative increase in responsivity. Although this degrades performance within the current task, it facilitates the disengagement of attention from this task, thus allowing potentially new and more rewarding behaviors to be emitted. Thus, the transition between the two LC modes can serve to optimize the trade-off between exploitation and exploration of opportunities for reward, and thereby maximizes overall utility.




Figure 1. Inverted-U relationship between tonic LC activity and performance in tasks that require focused attention. Moderate LC tonic activity is associated with optimal performance and prominent phasic LC activation following task-relevant stimuli (phasic LC mode). High levels of tonic LC activity are associated with poor performance and the absence of phasic LC activity (tonic LC mode). According to Aston-Jones and Cohen (2005), shifts along the continuum between the phasic and tonic LC modes drive corresponding changes in the exploitation-exploration tradeoff. Figure adapted from Aston-Jones and Cohen (2005).

The adaptive gain theory further holds that the transition between phasic and tonic LC firing modes and the corresponding control states are driven by online assessments of utility by the frontal structures that provide a major input to the LC, the anterior cingulate and the orbitofrontal cortex. According to the theory, the utility signals in these brain areas are integrated over different timescales and then used to regulate LC mode (Aston-Jones & Cohen, 2005). Brief lapses in performance, in the context of otherwise high utility, augment the LC phasic mode, resulting in improved task performance. In contrast, enduring decreases in utility drive transitions to the LC

tonic mode, promoting disengagement from the current task and facilitating exploration of behavioral alternatives.

Most of the evidence for the hypothesized link between utility, LC firing mode and exploitative vs. exploratory behavior comes from animal studies, but even that evidence is sparse. Importantly, crucial empirical tests of the theory in humans have been lacking. To fill this gap, we have used noninvasive methods to test the main assumptions of the adaptive gain theory in human participants (Chapters 2 and 3).

Other recent theories on the role of the LC-NE system in cognition

Since the publication of the adaptive gain theory, researchers have proposed several new accounts of the role of the LC-NE system in cognitive function. Yu and Dayan (2005), for example, proposed that tonic NE activity signals *unexpected uncertainty* arising from unanticipated changes in the nature of a task or behavioral context. According to Yu and Dayan, this elevated tonic NE activity in turn promotes bottom-up relative to top-down processing which facilitates learning about the external environment. As a complementary extension of this idea, Dayan and Yu (2006) proposed that phasic increases in LC/NE activity encode unexpected uncertainty arising from unexpected events or state changes *within* a task, and serve to interrupt ongoing cognitive processing associated with the default task state. In a similar vein, Bouret and Sara (2005) conceptualized the phasic LC response as a “network reset” signal that allows rapid stimulus-induced cognitive shifts and behavioral adaptation by facilitating the reorganization of target neural networks.

Whereas the adaptive gain theory mainly focuses on the regulation of attention and performance, these other accounts address the role of the LC-NE system in learning-related processes, and hence can be seen as complementary to the adaptive gain theory. The functions of the LC-NE system proposed by these accounts are broadly consistent with the adaptive gain theory. The adaptive gain theory’s assumption that the LC tonic mode promotes an exploratory control state, for example, implicitly suggests that this will facilitate learning about the external environment, consistent with Yu and Dayan’s (2005) account.

The role of the LC-NE system in neuropsychiatric disorders

Given the important role of the LC-NE system in cognition and behavior (e.g., Sara, 2009), it is not surprising that dysfunctions of this system have been associated with several neuropsychiatric disorders (e.g., Siever & Davis, 1985). Aston-Jones, Iba, Clayton, Rajkowski, and Cohen (2007) have proposed that dysregulation of the tonic and phasic components of LC activity may give rise to a variety of psychiatric conditions. For example, they hypothesized that a “hypertonic” LC mode may underlie some symptoms of attention-deficit/hyperactivity disorder (ADHD), post-traumatic stress disorder, and manic-depressive disorder. These disorders are

associated with concentration problems, sleeplessness and impulsivity—symptoms that resemble the distractible behaviors of monkeys in the tonic LC mode. Conversely, a chronically “hypotonic” LC mode may give rise to the limited emotionality and flat affect that are common symptoms in depressed patients. The idea that LC dysfunction is implicated in depression is supported by findings of LC cell loss and depleted NE levels in the brains of suicide victims (e.g., Arango, Underwood, & Mann, 1996; Ordway, Schenk, Stockmeier, May, & Klimek, 2003). In addition, Aston-Jones et al. (2007) speculated that a “hyperphasic” LC mode may be responsible for the extremely focused attentive state and impaired ability to shift attention to new stimuli that are observed in autistic patients (Mann & Walker, 2003). It is important to note that these ideas are still very speculative. Thus, although there is substantial evidence that the noradrenergic system is involved in various neuropsychiatric conditions, the exact etiology underlying the relationship between LC/NE dysfunction and neuropsychiatric disorders remains to be determined.

Chapter 4 of this thesis focuses on a very special case of noradrenergic dysfunction: dopamine- β -hydroxylase (D β H) deficiency. D β H deficiency is a rare genetic disorder characterized by a complete lack of NE in both the peripheral and central nervous system. Thus, patients with D β H deficiency may be seen as having a selective and complete lesion of the noradrenergic system. Informal clinical observations suggest that D β H-deficient patients do not have obvious cognitive impairments, which is remarkable given the important role of the LC-NE system in normal cognitive function and in neuropsychiatric disorders. This suggests that D β H-deficient patients may have subtle neurocognitive deficits that have remained unnoticed in informal observations. We tested five D β H-deficient patients and a healthy control group on a comprehensive neurocognitive test battery to provide a systematic evaluation of neurocognitive function in D β H deficiency (Chapter 4).

Curiosity and exploration

As described above, the adaptive gain theory proposes that the LC-mediated trade-off between exploitative and exploratory behaviors is driven by assessments of task-related utility. However, there are also many examples of exploratory behaviors that are not directly related to task utility but seem to be driven by the innate desire to learn or experience something that is unknown. This drive to know or experience new things is typically referred to as curiosity. In many circumstances, both animals and humans have a natural tendency to explore novel, unexpected or uncertainty-inducing stimuli (Berlyne, 1960; Daffner, Mesulam, Scinto, Cohen, Kennedy, et al., 1998; Ennaceur & Delacour, 1988; Hughes, 2007; Wittmann, Daw, Seymour, & Dolan, 2008), which suggests that the exploration of curiosity-inducing stimuli is intrinsically rewarding. In the reinforcement-learning literature, the bias towards the exploration of novel or uncertain options is captured by the concept of an “exploration bonus” that is assigned to novel or uncertain stimuli to increase their expected value and promote their exploration (e.g., Kakade & Dayan, 2002; Sutton & Barto, 1998).

Pharmacological studies in rats have suggested that curiosity-related exploratory behavior is mediated by the LC-NE system (Devauges & Sara, 1990; Sara, Dyon-Laurent & Hervé, 1995; Mansour, Babstock, Penney, Martin, McLean, et al., 2003). These studies found that drug-induced enhancements of phasic LC/NE activity resulted in increased exploration of novel and unexpected objects (i.e. specific exploration), but did not increase general exploratory activity (Devauges & Sara, 1990; Mansour et al., 2003). In contrast, pharmacological and environmental manipulations that enhance tonic LC/NE activity have been found to result in increased spontaneous sampling of random environmental stimuli, and in wider-ranging and more varied movement patterns (i.e. diversive exploration; Flicker & Geyer, 1982; Mansour et al., 2003). These findings are consistent with the assumptions of the adaptive gain theory that the phasic and tonic modes of LC activity promote, respectively, focused and divided attention.

The distinction between specific and diversive exploration resembles the distinction that has been proposed between specific and diversive curiosity, referring to the desire for a particular piece of information versus the more general stimulation-seeking motive that is closely related to boredom (Berlyne, 1960). A second, orthogonal, distinction has been made between perceptual curiosity, which is evoked by novel, strange or ambiguous perceptual stimuli, and epistemic curiosity, which refers to the desire for intellectual knowledge which applies mainly to humans (Berlyne, 1954).

In the 1960's and 70's, curiosity was a topic of intense investigation among experimental psychologists, resulting in an extensive theoretical framework for understanding curiosity and related behaviors. According to a classic psychological theory, curiosity evoked by ambiguous or conflict-inducing stimuli produces increased levels of arousal and is experienced as an aversive state, due to lack of information (e.g., Berlyne, 1966). The theory further proposes that termination of this condition, through access to relevant information, is rewarding and promotes learning. Although curiosity is one of the most basic biological drives in both animals and humans, and has been identified as one of the key motives for learning and discovery, the topic has been largely neglected in cognitive neuroscience; hence the neural mechanisms underlying curiosity are still poorly understood. Chapter 5 of this thesis describes a study in which we investigated the neural correlates of human perceptual curiosity.

Arousal, accessory stimuli and temporal uncertainty

Most of the topics discussed in this thesis are closely linked to arousal, a fundamental property of behavior. The concept of arousal is strongly related to attention, anxiety, stress and motivation, but has proven difficult to define. The LC-NE system is often associated with arousal, based on classical findings that tonic LC activity covaries with stages of the sleep-wake cycle (e.g., Aston-Jones & Bloom, 1981a; Hobson, McCarley, & Wyzinski, 1975) and that LC neurons exhibit strong phasic responses to salient and arousing stimuli (e.g., Aston-Jones & Bloom, 1981b; Grant, Aston-Jones, & Redmond, 1988). In addition, the inverted-U relationship between tonic LC activity

and performance (Figure 1) resembles the Yerkes-Dodson relationship between arousal and performance, one of the most important components of arousal theory (Duffy, 1957; Yerkes & Dodson, 1908). Recent studies have corroborated the notion that the LC-NE system plays a crucial role in the regulation of arousal (e.g., Gompf, Mathai, Fuller, Wood, Pedersen, et al., 2010; Carter, Yizhar, Chikahisa, Nguyen, Adamantidis, et al., 2010).

Obviously, the LC-NE system is not the only arousal-related system. It is generally accepted that arousal is a multifaceted construct which comprises a constellation of brain and somatic systems that subservise distinct but often overlapping functions (Neiss, 1988; Pribram & McGuinness, 1975; Robbins, 1997). One of these systems is the peripheral sympathetic nervous system. Motivationally significant stimuli or events typically elicit both a phasic LC response and a phasic response of the peripheral sympathetic nervous system that is often referred to as the *orienting response* (Lynn, 1966; Pavlov, 1927; Sokolov, 1963). The orienting response entails a collection of physiological changes, including a temporary dilation of the pupils, a rise in skin conductance, and a momentary change in heart rate, and is typically accompanied by a shift of attention toward the eliciting event. Anatomical considerations suggest that the parallel activation of the peripheral sympathetic nervous system and the LC-NE system following motivationally significant events reflects co-activation of these two systems by a common afferent source in the medulla (Aston-Jones, Ennis, Pieribone, Nickell, & Shipley, 1986; Nieuwenhuis, De Geus, & Aston-Jones, 2011). Nieuwenhuis et al. (2011) hypothesized that the co-activation of the LC-NE system and the peripheral sympathetic nervous system allows efficient mobilization for action in response to motivationally significant events: the LC-NE system facilitates the execution of cognitive decisions concerning proper behaviors in the face of urgent stimulus demand while, at the same time, the peripheral sympathetic nervous system facilitates physical execution of the chosen behaviors.

As described above, the orienting response and the phasic LC response are driven by motivationally significant task-relevant stimuli, but also by novel or intense task-irrelevant stimuli, such as unexpected loud sounds. The automatic orienting of attention towards salient task-irrelevant stimuli generally disrupts performance on the concomitant task (e.g., Parmentier, Elford, Escera, Andrés, & San Miguel, 2008; Schröger and Wolff, 1998). However, there are also instances where the occurrence of a task-irrelevant sound leads to faster responses to a simultaneously presented imperative stimulus in another modality (e.g., Bernstein, Clark, & Edelstein, 1969a,b; Hackley & Valle-Inclan, 1998, 1999; Valls-Solé, Solé, Valdeoriola, Muñoz, Gonzalez, et al., 1995). This phenomenon has been referred to as the *accessory-stimulus effect*, and is generally attributed to a temporary increase in arousal. Besides their effect on reaction times, accessory stimuli have been found to elicit an increase in response force (Miller, Franz, & Ulrich, 1999; Stahl & Rammsayer, 2005). Pharmacological manipulations in cats have shown that the availability of NE is critical for accessory-stimulus induced increases in motor activity, at least in the case of reflexive responses (Stafford & Jacobs, 1990). A possibility that remains to be explored is that an NE-mediated temporary increase in neuronal responsivity (or gain) also underlies the accessory-stimulus induced

speeding of reaction times. It is interesting to note in this regard that changes in gain are closely related to, and under certain conditions can be equivalent to, changes in decision threshold (Servan-Schreiber, Printz, & Cohen, 1990). Thus, one possible mechanism underlying the speeding of responses by accessory stimuli is a temporary lowering of the decision threshold. Despite a substantial empirical database, there is no general agreement among researchers regarding the neurocognitive mechanisms underlying the facilitatory effect of accessory stimuli. Chapter 6 of this thesis describes two experiments that aimed to shed more light on the effects of accessory stimuli on different components of information processing.

The effects of task-irrelevant accessory stimuli on information processing are exogenously-driven (i.e. automatic). A possibly related endogenously-driven phenomenon is the speed-up of reaction times to an imperative stimulus when its timing is highly predictable. This phenomenon, referred to as the warning effect or temporal-preparation effect, is typically investigated by means of paradigms in which participants use temporal cues to anticipate the onset of an imperative stimulus. In contrast to the accessory-stimulus paradigm, the interval between the temporal cue and the imperative stimulus is long enough to enable deliberate preparation. Like the accessory-stimulus effect, temporal-preparation effects have been attributed to NE-mediated changes in alertness (Coull, Nobre, & Frith, 2001; Fernandez-Duque & Posner, 1997; Witte & Marrocco, 1997). Furthermore, it has been found that the firing rate of LC neurons increases during the interval between the temporal cue and the imperative stimulus (Yamamoto & Ozawa, 1989). This raises the possibility that the temporal-preparation effect and the accessory-stimulus effect may correspond to endogenous and exogenous instances of the same underlying process: whereas accessory stimuli, by virtue of their salience, may elicit an automatic NE-mediated increase in gain, temporal preparation may allow controlled gain modulations resulting in the optimization of system parameters at the expected onset of the imperative stimulus. Chapter 7 of this thesis describes two experiments in which we investigated temporal-preparation effects on information processing.

Chapter 2

Pupil diameter predicts changes in the exploration-exploitation trade-off: Evidence for the adaptive gain theory

This chapter is published as: Jepma, M., & Nieuwenhuis, S. (in press). Pupil diameter predicts changes in the exploration-exploitation tradeoff: Evidence for the adaptive gain theory of locus coeruleus function. *Journal of Cognitive Neuroscience*.

Abstract

The adaptive regulation of the balance between exploitation and exploration is critical for the optimization of behavioral performance. Animal research and computational modeling have suggested that changes in exploitative vs. exploratory control state in response to changes in task utility are mediated by the neuromodulatory locus coeruleus-norepinephrine (LC-NE) system. Recent studies have suggested that utility-driven changes in control state correlate with pupil diameter, and that pupil diameter can be used as an indirect marker of LC activity. We measured participants' pupil diameter while they performed a gambling task with a gradually changing pay-off structure. Each choice in this task can be classified as exploitative or exploratory, using a computational model of reinforcement learning. We examined the relationship between pupil diameter, task utility and choice strategy (exploitation vs. exploration), and found that (i) exploratory choices were preceded by a larger baseline pupil diameter than exploitative choices; (ii) individual differences in baseline pupil diameter were predictive of an individual's tendency to explore; and (iii) changes in pupil diameter surrounding the transition between exploitative and exploratory choices correlated with changes in task utility. These findings provide novel evidence that pupil diameter correlates closely with control state, and are consistent with a role for the LC-NE system in the regulation of the exploration-exploitation trade-off in humans.

Introduction

Imagine you are in a restaurant, and are faced with the decision what food to order. One option is to choose a familiar dish that you know and like. Alternatively, you could try an unfamiliar dish, and take the risk that you might not like it. However, it is also possible that the unfamiliar dish turns out to become your new favorite, which you would never have discovered when sticking to the familiar dish. This example illustrates the dilemma between exploiting well-known options and exploring new ones. The trade-off between exploitation and exploration plays an important role in all kinds of decisions, especially in unfamiliar or changing environments. Although there has been a recent rise in studies investigating the strategies that are used to handle this trade-off and the neural mechanisms involved (for a review see Cohen, McClure, & Yu, 2007), these issues are still poorly understood.

One relevant line of research that addresses this issue suggests that the locus coeruleus-norepinephrine (LC-NE) neuromodulatory system plays an important role in regulating the balance between exploitation and exploration (Aston-Jones & Cohen, 2005; Usher, Cohen, Servan-Schreiber, Rajkowski, & Aston-Jones, 1999). Aston-Jones and Cohen have proposed that exploitative and exploratory *control states* are mediated by two modes of LC activity, called the ‘phasic’ and the ‘tonic mode’, respectively. The phasic LC mode is characterized by an intermediate level of LC baseline activity and large phasic increases in activity in response to task-relevant stimuli. The ensuing phasic release of NE in cortical areas temporarily increases the responsivity (or gain) of these areas to their afferent input, selectively potentiating the processing of these task-relevant stimuli (Berridge & Waterhouse, 2003; Doya, 2002; Servan-Schreiber, Printz, & Cohen, 1990). Conversely, the tonic LC mode is characterized by an elevated level of LC baseline activity and tonic NE release, and the absence of phasic responses¹.

According to the adaptive gain theory (Aston-Jones & Cohen, 2005), the two LC modes promote exploitation and exploration by adaptively adjusting the responsivity of cortical neurons: the phasic mode produces selective increases in neuronal responsivity in response to task-related stimuli, thereby optimizing performance in the current task (i.e. exploitation). In contrast, the tonic mode produces a more enduring and less discriminative increase in neuronal responsivity. Although this degrades performance within the current task, it facilitates the disengagement of attention from this task and the processing of other non-task related stimuli and/or behaviors (i.e. exploration). A second assumption of the theory is that transitions between the phasic and tonic LC modes and corresponding control states are driven by online assessments of task-related utility carried out in ventral and medial frontal structures (Aston-Jones & Cohen, 2005). Consistent with this hypothesis, anatomical studies have shown that the primary neocortical projections to LC come from

¹ Whereas we discuss the phasic and tonic LC modes as distinct, they likely represent the extremes of a continuum of function. When we refer to the phasic or tonic LC mode, we mean a *more* phasic or tonic LC mode, not necessarily the extremes of the continuum.

orbitofrontal and anterior cingulate cortex (Aston-Jones et al., 2002; Rajkowski, Lu, Zhu, Cohen, & Aston-Jones, 2000; Zhu, Iba, Rajkowski, & Aston-Jones, 2004)—areas known to be responsive to task-related rewards and costs of performance (Botvinick, 2007; Ridderinkhof, Ullsperger, Crone, & Nieuwenhuis, 2004). In order to adaptively regulate the balance between exploitation and exploration, utility assessments are integrated over both short (e.g., seconds) and longer (e.g., tens of seconds) timescales. If long-term utility is high, temporary decreases in utility augment the phasic LC mode, in order to restore task performance. Conversely, long-term decreases in utility drive the LC toward the tonic mode, which facilitates disengagement from the current task and exploration of alternative behaviors.

The adaptive gain theory has been supported mainly by computational modeling studies (Usher et al., 1999) and neurophysiological studies in monkeys that have used relatively simple tasks (Aston-Jones & Cohen, 2005). In contrast, with one notable exception (Gilzenrat, Nieuwenhuis, Jepma, & Cohen, 2010), there have been no tests of this theory in humans yet. In order to test the theory in humans, a non-invasive measure of LC activity is required. There is preliminary evidence that pupil diameter can provide such a measure: although it does not appear to be under direct control of the LC, pupil diameter is correlated with LC activity and thus may be useful as a “reporter variable” (Nieuwenhuis, de Geus, & Aston-Jones, in press). Rajkowski, Kubiak, and Aston-Jones (1993), for example, found a strong correlation in monkeys between baseline pupil diameter and tonic LC firing rate over the course of 90 minutes of performance in a target-detection task. Furthermore, a recent study that investigated how pupil diameter is related to experimental manipulations of task-related utility and behavioral indices of task (dis)engagement showed that pupil diameter varied in a way consistent with predicted LC dynamics (Gilzenrat et al., 2010). Specifically, this study showed that decreases in long-term utility and behavioral indices of task disengagement were associated with increased baseline pupil diameters and decreased pupil dilations, mirroring the high tonic and low phasic activity associated with the tonic LC mode. However, although this study assessed pupil effects associated with task (dis)engagement, it did not explicitly investigate the exploitation-exploration trade-off since participants were not given the opportunity to explore different task options.

Inspired by the recent evidence that pupil diameter might be used as an indirect index of LC activity, we measured participants’ pupil diameter while they performed a ‘four-armed bandit’ task with a gradually changing pay-off structure in which the trade-off between exploitation and exploration is a central component (Daw, O’Doherty, Dayan, Seymour, & Dolan, 2006; Figure 1; Appendix). Optimal performance in this task requires a delicate balance between exploitative and exploratory choices. We examined whether the relationship between pupil diameter, control state and task-related utility was consistent with the two main assumptions of the adaptive gain theory, namely that LC mode regulates the trade-off between exploitative and exploratory control states, and that transitions between LC modes are driven by assessments of task-related utility. The first assumption predicts that exploratory choices will be associated with a larger baseline pupil diameter, possibly reflecting a more tonic LC mode, than exploitative choices. In addition, this

assumption suggests that individual differences in overall pupil diameter might be correlated with individual differences in exploratory choice behavior: participants with larger overall pupil diameters, perhaps suggestive of a more tonic LC mode, may make more exploratory choices. The second assumption predicts that changes in utility surrounding the transition between control states will be accompanied by specific changes in baseline pupil diameter: a steady increase in baseline pupil diameter as decreasing utility drives the participant toward exploration; a monotonic decrease in baseline pupil diameter as utility increases after the participant has started a new series of exploitative choices.




Figure 1. The four-armed bandit task. Participants made repeated choices between four slot machines. Unlike standard slots, the mean pay-offs of the four machines changed gradually and independently from trial to trial (four colored lines). Participants were encouraged to earn as many points as possible during the experiment. After the experiment, each choice was classified as exploitative or exploratory, using a computational model of reinforcement learning.

Materials and Methods

Participants

Seventeen volunteers participated (11 women; aged 18-33 years; mean age = 22.4). The experiment was approved by the local ethics review board and conducted according to the principles expressed in the Declaration of Helsinki. Informed consent was obtained from all participants.

Stimuli and Procedure

Participants performed a ‘four-armed bandit’ task, while their pupil diameters were continuously measured. The task was a slightly modified version of the task used by Daw et al. (2006). Participants were presented with pictures of four different colored slot machines (of equal luminance) on a medium gray background. The slot machines stayed on the screen during the entire experiment. Each trial started with a 4 s interval during which the slot machines were displayed, but

participants could not select a machine yet. After this, a black fixation cross appeared in the center of the screen, indicating that participants could select one of the slot machines, by pressing the ‘q’-, ‘w’-, ‘a’- or ‘s’- key. Participants had a maximum of 1.5 s in which to make their choice; if no choice was made during that interval, a ‘TIME OUT’ message appeared in the center of the screen for 3 s to signal a missed trial (average number of missed trials = 1.7). If participants responded within 1.5 s, the lever of the chosen slot machine was lowered and the number of points earned was displayed in the chosen machine. These points were displayed until the end of the trial, which was 7 s after trial onset. Importantly, the number of points paid off by the four slot machines gradually and independently changed from trial to trial (Figure 1; Appendix).

The experiment was conducted at a slightly dimmed illumination level (room illumination 100 lux). We recorded pupil diameter at 60 Hz using a Tobii T120 eye tracker, which is integrated into a 17-inch TFT monitor (Tobii Technology, Stockholm, Sweden). Participants were seated at a distance of approximately 60 cm from the monitor. Prior to the start of the experimental session, participants viewed visually presented instructions, including an instruction that the pay-offs of the machines would change throughout the experiment, and were given 24 practice trials to familiarize them with the task. After the practice trials, participants were instructed that the machines had been reset for the experimental session. The experimental session consisted of 180 trials, and lasted about 20 minutes. We instructed the participants that they would be paid according to how many points they had earned during the experimental session. We also instructed them that on average participants earned 2.50 euros in this experiment. However, we did not tell participants how the number of points was converted into euros, or what their cumulative point total was. At the end of the experiment, each participant was paid 3 euros.

Data Analysis

Behavioral Analysis. In order to classify each choice as exploitative or exploratory, we fitted a reinforcement-learning model to the data of each participant. We used the same model as used by Daw et al. (2006). This model consists of a mean-tracking rule that estimates the mean pay-off of each machine, and a choice rule that selects a machine based on these estimations (Appendix). The choice rule we used was the ‘softmax’ rule. This rule assumes that choices between different options are made in a probabilistic manner, such that the probability that a particular machine is chosen depends on its relative estimated pay-off. The exploitation-exploration balance is adjusted by a parameter referred to as gain, or inverse temperature: with higher gain, action selection is determined more by the relative estimated pay-offs of the different options, whereas with lower gain, action-selection is more evenly distributed across the different options. We classified each choice as exploitative or exploratory according to whether the chosen slot machine was the one with the maximum estimated pay-off (exploitation) or not (exploration).

Pupil Analysis. Pupil data were processed and analyzed using Brain Vision Analyzer (Brain Products, Gilching, Germany). Artifacts and blinks were removed using a linear interpolation algorithm. We assessed the baseline pupil diameter prior to the selection of a slot machine, as well

as the magnitude of the pupil dilation following the selection of a slot machine. To determine baseline pupil diameter, we averaged the pupil data in the period from 2.5 s to 0.5 s before the key-press. The pupil data during the 0.5 s immediately preceding the key-press were not included in the baseline period because most participants showed an anticipatory increase in pupil diameter starting about 0.5 s before their key-press response. The pupil dilation evoked by choosing a machine and perceiving the received pay-off was measured as the highest deviation from the baseline in the 3 s following the key-press response.

We compared the average baseline pupil diameter and pupil dilation on exploitation versus exploration trials. In addition, we calculated the degree of exploration for each exploratory choice, by subtracting the estimated pay-off of the chosen machine from the maximum estimated pay-off. We divided all exploration trials into three equally sized bins based on the degree of exploration (low, medium and high), and assessed the average baseline pupil diameter for these three exploration bins. Since the number of points earned was displayed immediately after the selection of a slot machine, the pupil dilation on each trial reflected both the selection of a machine and the processing of the received pay-off. Due to this confound, we could not unequivocally interpret differences in pupil dilation between exploitation and exploration trials, and focused our analyses on the baseline pupil diameter.

Compared to exploitative choices, exploratory choices were more often preceded by other exploratory choices. In addition, exploratory choices were associated with a lower pay-off and more negative prediction error on the previous trial, and a lower expected pay-off and higher entropy on the current trial (see Results). Entropy is an index of the similarity of the four slot machines' expected pay-offs; it increases as the expected pay-offs of the four slot machines become more similar. Entropy thus provides an estimate of the level of uncertainty, or conflict, associated with figuring out which slot machine is the most valuable. The entropy $H(X)$ on each trial was calculated as:

$$H(X) = -\sum_i P(x_i) \log_2 P(x_i)$$

where $P(x_i)$ is the probability of choosing slot machine x_i . To assess whether these potential confounds could account for the differences in baseline pupil diameter on exploration and exploitation trials, we subjected the single-trial baseline pupil diameter values to a multiple linear regression analysis, separately for each participant. Choice strategy (explore vs exploit) and the five above-mentioned nuisance variables (expected pay-off, entropy, and the pay-off, prediction error and strategy on the previous trial) as well as a constant were included as explanatory factors. For choice strategy and choice strategy on the previous trial, we used binary factors that have a value of 1 on exploit trials and 0 on explore trials. To assess which variables were significant predictors of baseline pupil diameter, we conducted a one-sample t-test on the regression coefficients of each explanatory factor (Lorch & Myers, 1990).

We also assessed whether individual differences in pupil diameter predicted individual differences in exploratory behavior. In this analysis, we computed the between-subjects correlation

between the average baseline pupil diameter and the proportion of exploratory choices, and between the average baseline pupil diameter and the value of the gain/inverse temperature parameter of the reinforcement learning model.

To assess the development of our utility measures (pay-off, expected pay-off and entropy) and baseline pupil diameter surrounding the transition between exploitative and exploratory choice strategies, we averaged trials as a function of their position relative to the transition from an exploitative to an exploratory choice strategy, and vice versa. For this analysis, we only considered the exploration trials that were preceded or followed by a minimum of three exploitation trials.

Results

Participants alternated between choosing the slot machine with the highest estimated current pay-off (exploitation) and choosing slot machines with a lower expected pay-off (exploration). In comparison to the exploitation trials, exploration trials were more often preceded by other exploration trials (Table 1), indicating that participants tended to explore for several successive trials before settling on a new slot machine. The main characteristics of the exploitation and exploration trials are summarized in Table 1.

Table 1. Characteristics of exploration and exploitation trials (standard deviation in parentheses)

	exploration	exploitation	<i>p</i> -value
Proportion of total number of trials	0.31 (0.10)	0.69 (0.10)	< 0.001
Proportion preceded by exploration trial	0.41 (0.07)	0.28 (0.13)	0.001
RT (ms)	492 (82)	508 (75)	0.15
RT variability (SD of RTs)	150 (45.5)	151 (40.0)	0.912
RT trial N-1 (ms)	498 (72)	504 (79)	0.36
Pay-off (points)	48 (1.6)	63 (1.9)	< 0.001
Prediction error (points)	-2.8 (6.5)	-1.0 (5.1)	0.07
Expected pay-off (points)	51 (6.4)	64 (4.0)	< 0.001
Entropy (bits)	1.5 (0.14)	1.2 (0.33)	< 0.001
Pay-off preceding trial (points)	54 (2.4)	60 (3.1)	< 0.001
Prediction error preceding trial (points)	-3.6 (4.4)	-1.0 (5.9)	0.001

Pupil Diameter on Exploitation versus Exploration Trials

First, we compared the baseline pupil diameter preceding exploitative and exploratory choices. Baseline pupil diameters preceding exploratory choices were larger than those preceding exploitative choices [3.93 vs. 3.88 mm, $t(16) = 3.0$, $p = 0.008$; Figure 2, left panel]. Furthermore, within the exploration trials, baseline pupil diameter increased as a function of the degree of exploration (Materials and Methods), as revealed by a repeated-measures linear-trend analysis

[$F(1,16) = 15.3, p = 0.001$; Figure 2, right panel]. We also examined the pupil dilations evoked by exploratory and exploitative choices. There was a trend towards larger dilations on exploration than exploitation trials [0.17 vs. 0.13 mm; $t(16) = 2.1, p = 0.051$]. This was probably due to the higher incidence of negative prediction errors on exploration trials (Satterthwaite et al., 2007), since the effect disappeared when only the trials with positive prediction errors were included ($p = 0.14$).




Figure 2. Pupil diameter on exploration and exploitation trials. (A) Time course of grand-average pupil diameter aligned to the key-pres indicating the selection of a slot machine, for exploratory and exploitative choices. (B) Average baseline pupil diameter for exploitative choices (black bar), and exploratory choices with a low, medium and high degree of exploration (striped bars).

The difference in baseline pupil size between exploitation and exploration trials already started to develop during the pupil response on the preceding trial (Figure 3): trials immediately preceding exploration trials were associated with a larger pupil dilation than trials immediately preceding exploitation trials [0.17 vs. 0.13 mm, $t(16) = 3.2, p = 0.006$]. However, this effect on the preceding trial could not (fully) explain the difference in baseline pupil diameters between exploitation and exploration trials, because the difference remained significant when pupil dilation on the previous trial was included as a covariate in the analysis [$F(1, 15) = 4.69, p = 0.047$].




Figure 3. Time course of grand-average post-choice pupil dilation for the trials preceding exploration and exploitation trials.

Exploitation and exploration trials differed in several aspects other than choice strategy (Table 1). Trials preceding exploration trials were characterized by a larger proportion of exploratory choices, a lower pay-off and a more negative prediction error than trials preceding exploitation trials. In addition, exploration trials were characterized by a lower model-estimated expected pay-off (of the chosen slot machine) and higher entropy than exploitation trials. We investigated whether choice strategy (explore vs. exploit) could predict baseline pupil diameter independently of these potential nuisance variables by means of a linear multiple regression analysis (Materials and Methods). Importantly, when adjusted for all other variables, choice strategy made a unique contribution to the prediction of baseline pupil diameter [$t(16) = 3.43, p = 0.003$]. The only other significant predictor of baseline pupil diameter was the strategy on the previous trial [$t(16) = 2.98, p = 0.009$]. Additional control analyses that yielded similar results are reported in the Appendix.

Together, these findings confirm our first prediction that exploratory choices are associated with a larger baseline pupil diameter, while excluding a range of alternative interpretations for the observed pupil effect.

Individual Differences in Pupil Diameter and Exploratory Choice Behavior

Sofar we have examined pupil diameter as a function of the within-subject factor choice strategy. We next assessed whether individual differences in overall pupil diameter were predictive of individual differences in exploratory choice behavior. There was a positive correlation, across participants, between the average pupil diameter over all trials and the proportion of exploratory choices ($r = 0.50, p = 0.04$; Figure 4, left panel). Similarly, there was a negative correlation between the average pupil diameter and the value of the gain parameter of the reinforcement learning model ($r = -0.53, p = 0.03$; Figure 4, right panel). These correlations were also present when the baseline pupil diameters on exploitation and exploration trials were considered separately (pupil diameter on exploitation trials and proportion exploratory choices: $r = 0.49, p = 0.04$; pupil diameter on exploitation trials and gain parameter: $r = -0.52, p = 0.03$; pupil diameter on exploration trials and proportion exploratory choices: $r = 0.48, p = 0.05$; pupil diameter on exploration trials and gain parameter: $r = -0.53, p = 0.03$). Unlike the gain parameter, the other model parameters did not correlate with pupil diameter (decay parameter: $r = -0.24, p = 0.36$; decay center: $r = 0.07, p = 0.78$).

Obviously, individual differences in pupil diameter relate to many factors other than control state, such as age, personality and intelligence (Janisse, 1977). Importantly, these factors presumably increased the between-subjects error variance in our data, which decreased the power for detecting a correlation. Thus, the fact that we found a correlation in spite of a presumably large error variance in the between-subjects pupil data affirms the existence of the correlation. However, it is also possible that individual differences in pupil diameter reflect individual differences in motivation or the amount of attention paid to the task. Such motivational factors might influence

choice strategy, which could provide an alternative explanation for the correlations between pupil diameter and exploratory behavior across participants.

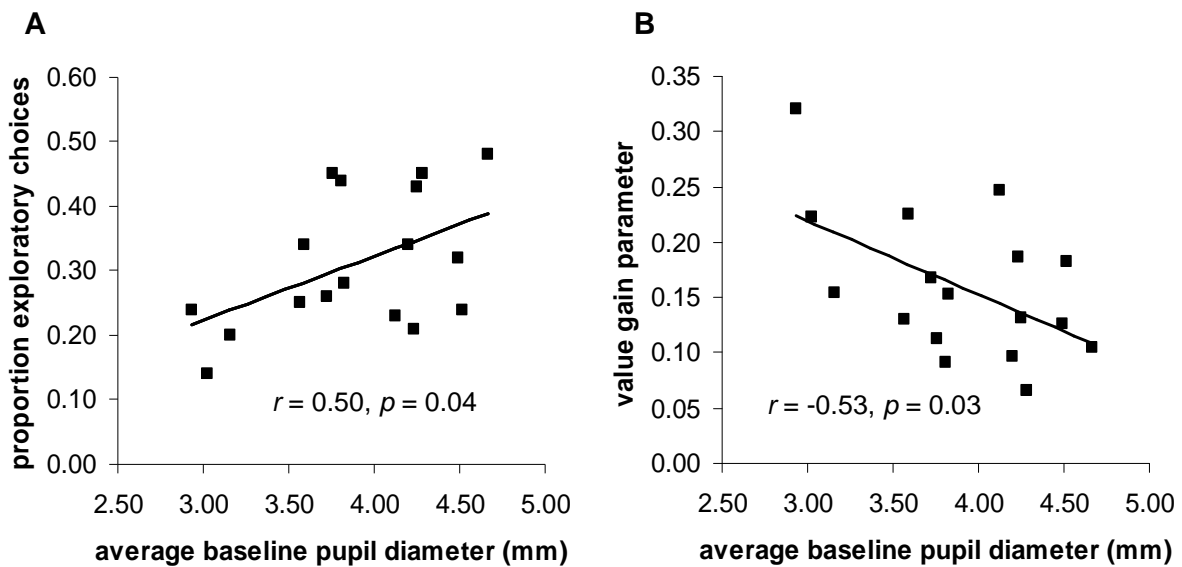


Figure 4. Individual differences in pupil diameter and exploratory choice behavior. (A) Scatter plot of the between-subjects correlation between average baseline pupil diameter and the proportion of exploratory choices. (B) Scatter plot of the between-subjects correlation between average baseline pupil diameter and the value of the gain or inverse-temperature parameter of the reinforcement-learning model. A lower value of this parameter indicates a more exploratory choice strategy.

Changes in Utility and Pupil Diameter Surrounding a Transition between Choice Strategies

Sofar we have examined the difference in pupil diameter between exploitation and exploration trials. We next examined the changes in utility measures surrounding the transition between exploitative and exploratory choice strategies. As measures of utility, we used the model-estimated expected pay-off of the chosen machine, the received pay-off, and the entropy (Materials and Methods). Subsequently, we tested whether such changes in utility were accompanied by changes in pupil diameter.

Figure 5 (upper panel) shows the expected pay-off, received pay-off and entropy for the first and the last of a series of exploration trials and the three preceding and following exploitation trials. During the three exploitation trials that preceded the switch to an exploratory choice strategy, entropy gradually increased [$F(1, 16) = 10.16, p = 0.006$] and pay-off gradually decreased [$F(1, 16) = 50.72, p < 0.001$], as revealed by a repeated-measures linear-trend analysis. Expected pay-off also showed a decrease over the three trials preceding the first explore trial, but this effect missed significance [$F(1, 16) = 2.85, p = 0.11$]. Thus, there was a gradual decrease in utility preceding the switch from an exploitative to an exploratory choice strategy, suggesting that, on average, participants began exploring when task utility was at a minimum. In addition, during the three exploitation trials following the last exploration trial, entropy gradually decreased [$F(1, 16) = 9.74, p = 0.007$] and expected pay-off gradually increased [$F(1, 16) = 13.72, p = 0.002$]. Thus, there was

a gradual increase in utility following the switch from an exploratory to an exploitative choice strategy.

We next examined the development of baseline pupil diameter over the trials surrounding the switch between exploitative and exploratory choice strategies (Figure 5, lower panel). Baseline pupil diameter did not differ significantly across the three exploitation trials preceding the first exploration trial [$F(2, 32) = 1.30, p = 0.29$]. However, baseline pupil diameter showed a gradual decrease over the three exploitation trials following the last exploration trial [$F(1, 16) = 6.18, p = 0.024$], resembling the gradual decrease in entropy and increase in expected pay-off during these trials. As predicted, baseline pupil diameter correlated negatively with expected pay-off [$r = -0.72, p(1\text{-tailed}) = 0.023$] and positively with entropy [$r = 0.68, p(1\text{-tailed}) = 0.032$] across the eight trial positions in Figure 5. These findings provide some evidence for our second prediction, that changes in utility surrounding the transition between control states would be systematically correlated with changes in baseline pupil diameter.

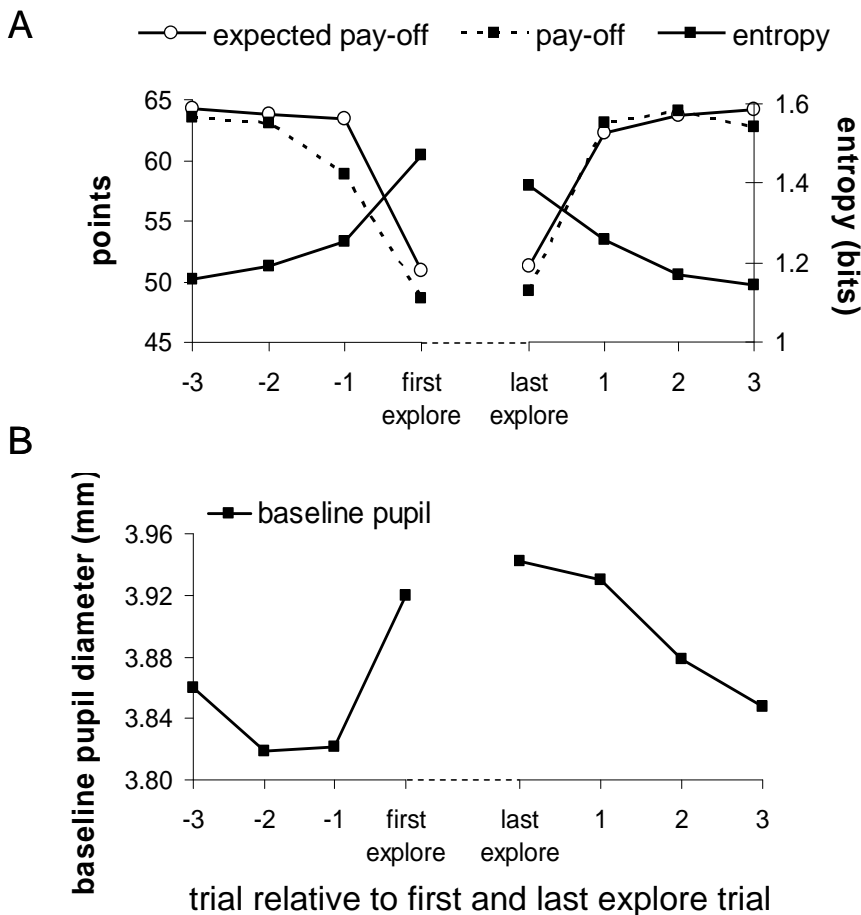


Figure 5. Grand-average dependent measures for the first and last of a series of exploration trials, and the three preceding and following exploitation trials. (A) Our measures of utility: expected pay-off, received pay-off and entropy. (B) Baseline pupil diameter.

Discussion

We investigated the relationship between pupil diameter, choice strategy (exploitation vs. exploration) and task utility, in order to test predictions of the adaptive gain theory of LC function in humans. This study was inspired by recent observations that pupil diameter might be used as a reliable index of LC activity. Our main findings can be summarized as follows: (i) exploratory choices were associated with a larger baseline pupil diameter than exploitative choices; (ii) individual differences in baseline pupil diameter predicted individual differences in exploratory choice behavior: participants with a larger pupil diameter made more exploratory choices and were characterized by a smaller gain parameter of the reinforcement-learning model; and (iii) trial-to-trial changes in baseline pupil diameter surrounding the transition between choice strategies correlated systematically with changes in utility, at least during the transition from exploration to exploitation. At the least, these findings provide novel evidence for a close relationship between pattern of pupillary response and control state. More tentatively, these findings provide indirect support for the two main assumptions of the adaptive gain theory, namely that LC firing mode regulates the trade-off between exploitative and exploratory control states, and that changes in LC mode are driven by online assessments of task-related utility (Aston-Jones & Cohen, 2005).

Our finding that pupil diameter is predictive of choice strategy, in a manner consistent with the adaptive gain theory, corroborates recent findings by Gilzenrat et al. (2010) that pupil diameter is related to behavioral indications of the tonic and phasic LC mode. Gilzenrat et al. found that large baseline pupils were associated with slower, more variable reaction times and less accurate performance in a target-detection task, and with task disengagement in a task in which participants were given the opportunity to disengage from the current task context when utility decreased. Furthermore, several pharmacological studies have shown that drug-induced activation of the LC-NE system increases cognitive flexibility and behavioral disengagement. For example, drugs that increase tonic NE levels (i.e. mimic the effects of elevated NE release that characterize the tonic LC mode) have been found to improve attentional-set shifting and reversal learning in rats and monkeys (Devauges & Sara, 1990; Lapid & Morilak, 2006; Lapid, Bondi, & Morilak, 2007; Seu, Lang, Rivera, & Jentsch, 2008; Steere & Arnsten, 1997; but see Chamberlain et al., 2006). In humans, increased NE levels induced by the selective NE reuptake inhibitor atomoxetine have been found to improve the ability to stop an ongoing motor response when cued to do so (Chamberlain et al., 2006). A possible explanation for this finding is that the drug-related increase in cognitive flexibility facilitates disengaging from one task (responding) and switching to a new task (stopping the response). In addition, increased NE levels induced by the selective NE reuptake inhibitor reboxetine have been found to enhance social flexibility in human participants, as indicated by increased social engagement and cooperation and a reduction in self-focus (Tse & Bond, 2002). Although none of these studies directly investigated exploitative versus exploratory behaviors, their findings support the idea that the tonic LC mode produces an enduring and largely nonspecific increase in responsivity, which promotes a flexible, exploratory control state.

Modeling studies have started to investigate the relationship between LC mode and task-related utility, integrated over different timescales (Aston-Jones & Cohen, 2005, Figure 10; McClure, Gilzenrat, & Cohen, 2005). However, to date there has been hardly any empirical research on the temporal dynamics of utility-driven changes in LC mode. We addressed this issue by assessing the trial-to-trial changes in utility and baseline pupil diameter surrounding the switch between exploitative and exploratory choice strategies. The switch to an exploratory choice strategy was preceded by a gradual decrease in utility, but an abrupt increase in baseline pupil diameter. When participants started to exploit a new machine after a period of exploration, utility gradually increased and baseline pupil diameter gradually decreased again. This pattern suggests that the transition from the tonic to the phasic mode is rather gradual, whereas the transition from the phasic to the tonic LC mode is more abrupt. A somewhat similar pattern was found by Gilzenrat and colleagues: Baseline pupil diameter showed a marked gradual decrease when participants started to engage in a new task; the increase in baseline pupil diameter leading up to task disengagement was less gradual and less pronounced. The implications of these data for our understanding of the specific mechanisms by which changes in short- and long-term utility control LC mode remain a matter for further research. One possibility is that LC baseline activity abruptly increases when long-term utility falls below a certain value. Consistent with this possibility, there is some evidence that tonic LC activity in monkeys can increase abruptly after a change in task contingency (Aston-Jones, Rajkowski, & Kubiak, 1997) or during the transition from a drowsy to an alert behavioral state (Rajkowski, Kubiak, & Aston-Jones, 1994). In any case, more empirical data is needed to determine how different measures of utility are integrated over different timescales and to specify the function relating overall utility to changes in LC mode. Such knowledge will also inform the implementation of a utility-sensitive adaptive gain mechanism in reinforcement-learning models. This will present a significant advance compared to current models, such as the model used here, in which the gain parameter is estimated for each participant but fixed across the experiment.

The abrupt increase in baseline pupil diameter prior to an exploratory choice might also be related to the specific task that we used. An aspect of the task that could be important in this respect is the high learning rate (see Appendix). A comparably high learning rate was found in a previous study using this task (Daw et al., 2006), so it seems to be characteristic of participants' choice behavior in this task. Such high learning rates imply that participants base their expectations regarding the slot machines' pay-offs primarily on their most recent experience with each machine. Accordingly, a single bad outcome on a certain trial is likely to be experienced as a substantial decrease in utility and to promote the exploration of another machine. This possibly explains the abrupt increase in baseline pupil diameter we observed immediately preceding the first of a series of exploratory choices. Thus, it will be important to assess in future studies whether tasks that are associated with lower learning rates will result in a more gradual increase in pupil diameter preceding the switch to an exploratory choice strategy.

Because the evidence for a close relationship between pupil diameter and LC activity is currently limited (Gilzenrat et al., 2010; Nieuwenhuis et al., in press; Rajkowski et al., 1993), more

neurophysiological studies are needed to further establish this relationship. In addition, the neural mechanism underlying this putative relationship remains to be determined. To date, there are no known direct connections from the LC to the autonomic centers that regulate pupil size. It is more likely that pupil diameter and LC activity are closely linked because they receive downstream influences from a common afferent source. This common afferent might be the paragigantocellularis (PGi) nucleus of the ventral medulla, which plays a pivotal role in controlling both the LC and the sympathetic axis of the autonomic nervous system (Aston-Jones, Ennis, Pieribone, Nickell, & Shipley, 1986; Nieuwenhuis et al., in press). The notion that the LC and the autonomic nervous system receive their major input from a common source is consistent with several findings that suggest a strong temporal correlation between LC-NE activity and sympathetic nervous system activity (Elam, Svensson, & Thorén, 1986; Abercrombie & Jacobs, 1987; Reiner, 1986). Anatomical studies have revealed widespread afferents to the PGi from numerous brain areas, including the medial prefrontal cortex, insula, hypothalamus and periaqueductal grey, suggesting that activity in these areas might influence pupil diameter by way of the PGi (Aston-Jones et al., 1986). Consistent with this possibility, fMRI studies in humans and single-cell stimulation/recording studies in animals have shown that activity in this afferent network (including prefrontal cortex) is related to changes in pupil diameter (Critchley, Tang, Glaser, Butterworth, & Dolan, 2005; Loewenfeld, 1993; Siegle, Steinhauer, Stenger, Konecky, & Carter, 2003).

Although our study focused on a possible role of the LC-NE system, it is unlikely that this is the only brain system involved in regulating the exploration-exploitation tradeoff. There is some evidence that the dopamine system also influences levels of exploration or task (dis)engagement (Dreisbach et al., 2005; Frank, Doll, Oas-Terpstra, & Moreno, 2009). For example, in one study, individuals with high spontaneous eyeblink rates (a marker of central dopaminergic activity) showed enhanced cognitive flexibility, as measured by the tendency to disengage from previously task-relevant stimuli and orient to novel stimuli (Dreisbach et al., 2005). Furthermore, this effect was modulated by the D4 dopamine receptor gene polymorphism. Another study reported that the val158met polymorphism of COMT, a gene that substantially affects prefrontal dopamine levels, could account for individual differences in uncertainty-based exploration (Frank et al., 2009). In addition to other neuromodulator systems, recent studies have implicated the frontopolar cortex in the control of exploratory behaviors (Daw et al., 2006; Bourdaud, Chavarriaga, Galan, & Millán, 2008), although the specific computations performed by the frontopolar cortex in this context are a topic of ongoing debate (Boorman, Behrens, Woolrich, & Rushworth, 2009). A key objective for future research is to specify the distinct contributions and interactions of the dopamine and LC-NE systems and the prefrontal cortex in the regulation of the exploration-exploitation tradeoff.

Our experimental design enabled examination of the baseline pupil diameter but, due to the overlap of the decision and outcome processing, did not allow examination of the decision-induced pupil dilation. Hence, the hypotheses we tested were restricted to the adaptive gain theory's assumptions about tonic LC activity. To provide complementary data with regard to phasic LC responses, an important aim for future studies is to use a task in which the decision and outcome

presentation are separated in time such that the pupil dilations associated with these two processes can be isolated.

The present study tested specific predictions regarding the relationship between pupil diameter, utility measures and choice strategy based on a mechanistic theory about the role of the LC-NE system in regulating control state, and preliminary evidence from previous studies for a close relationship between LC activity and pupil diameter. Given the specificity, and therefore the intrinsic unlikelihood, of our predictions, the fact that the predicted effects were observed lends provisional support to the hypotheses that drove the predictions. However, since this is an inductive argument, it is important to note that we cannot rule out the possibility that the observed relationships were not related to LC-mediated modulation of control state. Thus, future studies using more direct measures or manipulations of the LC-NE system are needed to either confirm or invalidate the conclusions from the present study.

For a long time, the LC-NE system has been associated with basic functions such as arousal and the sleep-wake cycle. Only recently, researchers have begun to examine its involvement in more specific cognitive functions, such as attention, memory, perceptual selection and the signaling of unexpected uncertainty (Einhäuser, Stout, Koch, & Carter, 2008; Robbins, 1997; Sara, 2009; Yu and Dayan, 2005). The present study contributes to this work by addressing, albeit indirectly, the role of the LC-NE system in the control of human behavior. Specifically, the findings reported here support the adaptive gain theory (Aston-Jones & Cohen, 2005), which posits an important role for the LC-NE system in the optimization of behavioral performance by regulating the balance between exploitative and exploratory control states.

Appendix

Pay-off structure of the gambling task

The number of points paid off by slot machine i on trial t ranged from 1 to 100, drawn from a Gaussian distribution (standard deviation $\sigma_o = 4$) around a mean $\mu_{i,t}$ and rounded to the nearest integer. On each trial, the means diffused in a decaying Gaussian random walk:

$$\mu_{i,t+1} = \lambda\mu_{i,t} + (1-\lambda)\theta + \nu.$$

The decay parameter λ was 0.9836, the decay center θ was 50, and the diffusion noise ν was zero-mean Gaussian (standard deviation $\sigma_d = 2.8$). We used one instantiation of this process (Figure 1).

Reinforcement-learning model

We used a Bayesian mean-tracking rule (i.e. a Kalman filter) that tracked the mean expected pay-off of each machine ($\hat{\mu}_{i,t}$) and the variance of these pay-offs ($\hat{\sigma}_{i,t}^2$). On the first trial of the task, all four machines had the same prior mean $\hat{\mu}_{i,1}^{pre}$ and variance $\hat{\sigma}_{i,1}^{2pre}$. These start values were based on the pay-offs received during the practice block, and were determined separately for each participant (mean $\hat{\mu}_{i,1}^{pre} = 51.9$, SD = 2.7; mean $\hat{\sigma}_{i,1}^{2pre} = 52.3$, SD = 14.9). When a participant chose machine c on trial t and received pay-off r , the estimated pay-off distribution ($\hat{\mu}_{c,t}^{post}, \hat{\sigma}_{c,t}^{2post}$) was updated according to:

$$\hat{\mu}_{c,t}^{post} = \hat{\mu}_{c,t}^{pre} + \kappa_t \delta_t$$

$$\hat{\sigma}_{c,t}^{2post} = (1 - \kappa_t) \hat{\sigma}_{c,t}^{2pre}$$

with prediction error $\delta_t = r_t - \hat{\mu}_{c,t}^{pre}$ and learning rate $\kappa_t = \hat{\sigma}_{c,t}^{2pre} / (\hat{\sigma}_{c,t}^{2pre} + \hat{\sigma}_o^2)$.

The estimated pay-off distributions for the unchosen machines did not change.

Then, the estimated prior pay-off distributions on the subsequent trial (trial $t+1$) were updated in time according to:

$$\hat{\mu}_{i,t+1}^{pre} = \hat{\lambda} \hat{\mu}_{i,t}^{post} + (1 - \hat{\lambda}) \hat{\theta}$$

$$\hat{\sigma}_{i,t+1}^{2pre} = \hat{\lambda}^2 \hat{\sigma}_{i,t}^{2post} + \hat{\sigma}_d^2.$$

We modeled the choice of the participants by a softmax rule. The probability $P_{i,t}$ of choosing machine i on trial t was given by:

$$P_{i,t} = \frac{\exp(\beta \hat{\mu}_{i,t}^{pre})}{\sum_j \exp(\beta \hat{\mu}_{j,t}^{pre})}$$

with exploration parameter β (often referred to as gain, or inverse temperature).

For a discussion of the Kalman filter and the softmax rule, we refer the reader to Anderson and Moore (1979), and Sutton and Barto (1998), respectively.

We fitted the model to each individual participant’s choice data. The trials in which no response was made within the 1.5-s time limit were omitted. The parameters $\hat{\lambda}$, $\hat{\theta}$ and β were estimated per participant by maximizing the log-likelihood of the observed choices (Supplemental Table 1). Parameter $\hat{\sigma}_o$ was fixed at 4. Estimation of parameter $\hat{\sigma}_d$ resulted in extreme values for most of the participants (values larger than 1000 for ten of the seventeen participants), suggesting unreliable fits. Therefore, we fixed this parameter at 50, which is similar to the best fitting $\hat{\sigma}_d$ parameter found in a previous study (Daw, O’Doherty, Dayan, Seymour, and Dolan, 2006). This large value of $\hat{\sigma}_d$ implies that participants overestimate the speed of diffusion in the pay-offs. Large values of $\hat{\sigma}_d$ induce high learning rates, indicating that the expected pay-offs are determined primarily by the most recent experience with each machine.

Supplemental Table 1. Mean parameter estimates and negative log likelihood for the fit of the softmax model to the choice data of each participant. The parameter values used to generate the pay-offs, and the negative log likelihood of a model in which choices are made randomly are also shown.

	Estimated values	Generative values
β	0.160 (0.066)	
λ	0.894 (0.083)	0.9836
θ	56.9 (17.6)	50
σ_d	50 (fixed)	2.8
σ_o	4 (fixed)	4
-LL	153.1 (34.8)	
-LL randomly choosing model	247.2 (2.0)	

Note: SD in parentheses; -LL = negative log likelihood

Additional control analysis

Besides the multiple regression analyses, we performed a second set of control analyses to investigate whether differences in each of the potential confound variables could account for the different baseline pupil diameter on exploration and exploitation trials (and hence might provide an alternative interpretation of the effect). We repeated the comparison of baseline pupil diameter on exploitation and exploration trials while, in separate analyses, controlling for differences in each of the potential confound variables (pay-off on the previous trial, prediction error on the previous trial, expected pay-off on the current trial and entropy on the current trial), by matching the values of these variables across exploration and exploitation trials (Bernstein, Scheffers, & Coles, 1995). We sorted each participant’s exploitation and exploration trials by one of these variables, and then successively removed the most extreme exploitation and exploration trials, thereby reducing the difference between the mean value of this confound variable on exploitation and exploration trials. After each trial removal, we calculated the difference between the mean values of the confound

variable on exploitation and exploration trials, and we stopped the removal process when this difference was not further decreased by removal of a subsequent trial (Supplemental Table 2). We also controlled for choice strategy on the previous trial, by including only the trials that were preceded by an exploitation trial. Finally, in order to control more explicitly for the possibility that the higher incidence of negative prediction errors preceding exploratory choices was driving the effect, we repeated the analysis while only including the trials that were preceded by a positive prediction error.²

Importantly, none of the potential confound variables could account for the larger baseline pupils preceding exploratory compared to exploitative choices: the critical effect remained significant after correction for choice strategy on the previous trial [$t(16) = 2.5, p = 0.026$]; pay-off on the previous trial [$t(16) = 2.9, p = 0.009$]; prediction error on the previous trial [$t(16) = 2.5, p = 0.025$]; expected pay-off [$t(12) = 3.1, p = 0.010$]; and entropy [$t(16) = 3.5, p = 0.003$]. Furthermore, the effect remained significant when only the trials that were preceded by a positive prediction error were considered ($t(12) = 2.3, p = 0.037$), suggesting that the larger baseline pupil on explore compared to exploit trials was not due to the larger incidence of negative prediction errors preceding explore trials.

Supplemental Table 2. The number of excluded trials and the values of the potential confound variables on exploration and exploitation trials after correction.

	# excluded trials	Exploration	Exploitation	<i>p</i> -value
Expected pay-off	83.8 (13.5)	60.0 (4.1)	60.1 (4.0)	.29
Entropy	27.0 (14.0)	1.25 (0.30)	1.26 (0.29)	.02
Pay-off preceding trial	23.2 (15.9)	57.8 (2.0)	57.9 (2.0)	.13
Prediction error preceding trial	14.3 (9.6)	-1.81 (5.24)	-1.83 (5.28)	.40

Note: SD in parentheses. The difference in entropy after correction is in the opposite direction (larger entropy on exploitation trials) compared to the original effect.

Uncertainty-driven exploration and pupil diameter

In the softmax rule described above, the probability that a particular machine is chosen is determined by its relative mean estimated pay-off (and the value of the gain parameter), but not by the uncertainty about its potential pay-offs (i.e. the variance of the estimated pay-off distribution $\hat{\sigma}_i^{2pre}$). On the other hand, modeling studies have suggested that exploration might be directed towards particular choices in proportion to the uncertainty about their outcomes, which can be implemented by adding an ‘uncertainty bonus’ to the expected value of options with uncertain outcomes (e.g., Sutton, 1990). It has recently been shown that individual differences in uncertainty-

² Four participants had to be excluded from the analysis that corrected for expected pay-off, because the difference in expected pay-off between their exploration and exploitation trials was so large that no exploration trials were left using this procedure. Similarly, four participants were excluded from the analysis in which only the trials preceded by a positive prediction error were considered, since less than ten explore and/or exploit trials were left for these participants.

based exploration are associated with the val158met polymorphism of the COMT gene, which substantially affects prefrontal dopamine levels (Frank, Doll, Oas-Terpstra, & Moreno, 2009). According to the adaptive gain theory, the increased NE level in the tonic LC mode indiscriminately facilitates processing of all stimuli and/or behaviors, which promotes a nonspecific type of exploration. Hence, the theory predicts that individual differences in tonic LC activity (as indexed by baseline pupil diameter in this study) will be related to individual differences in exploratory behavior (Results section), but not to individual differences in uncertainty-specific exploration.

To assess this last prediction, we considered a softmax rule in which an ‘uncertainty bonus’ of ϕ standard deviations was added to the estimated mean pay-offs:

$$P_{i,t} = \frac{\exp(\beta[\hat{\mu}_{i,t}^{pre} + \phi\hat{\sigma}_{i,t}^{pre}])}{\sum_j \exp(\beta[\hat{\mu}_{j,t}^{pre} + \phi\hat{\sigma}_{j,t}^{pre}])}$$

The best fitting uncertainty bonus parameter in this model varied across participants: four participants had a positive bonus and thirteen participants had a negative bonus (mean bonus = -0.117, SD = 0.336). Thus, for the majority of the participants, uncertainty about the potential outcomes of a machine *discouraged* exploration of that machine. Importantly, the value of the uncertainty bonus parameter did not correlate with baseline pupil diameter ($r = 0.05$, $p = 0.86$), consistent with the assumption that the tonic LC mode is not associated with uncertainty-specific exploration.

Chapter 3

The role of the noradrenergic system in the exploration-exploitation trade-off: A psychopharmacological study

This chapter is published as: Jepma, M., te Beek, E.T., Wagenmakers, E.-J., van Gerven, J.M.A, & Nieuwenhuis, S. (2010). The role of the noradrenergic system in the exploration-exploitation trade-off: a psychopharmacological study. *Frontiers in Human Neuroscience*, 4, 170.

Abstract

Animal research and computational modeling have indicated an important role for the neuromodulatory locus coeruleus-norepinephrine (LC-NE) system in the control of behavior. According to the adaptive gain theory, the LC-NE system is critical for optimizing behavioral performance by regulating the balance between exploitative and exploratory control states. However, crucial direct empirical tests of this theory in human subjects have been lacking. We used a pharmacological manipulation of the LC-NE system to test predictions of this theory in humans. In a double-blind parallel-groups design (N = 52), participants received 4 mg reboxetine (a selective norepinephrine reuptake inhibitor), 30 mg citalopram (a selective serotonin reuptake inhibitor) or placebo. The adaptive gain theory predicted that the increased tonic NE levels induced by reboxetine would promote task disengagement and exploratory behavior. We assessed the effects of reboxetine on performance in two cognitive tasks designed to examine task (dis)engagement and exploitative versus exploratory behavior: a diminishing-utility task and a gambling task with a non-stationary pay-off structure. In contrast to predictions of the adaptive gain theory, we did not find differences in task (dis)engagement or exploratory behavior between the three experimental groups, despite demonstrable effects of the two drugs on non-specific central and autonomic nervous system parameters. Our findings suggest that the LC-NE system may not be involved in the regulation of the exploration-exploitation trade-off in humans, at least not within the context of a single task. It remains to be examined whether the LC-NE system is involved in random exploration exceeding the current task context.

Introduction

The locus coeruleus (LC) is one of the major brainstem neuromodulatory nuclei, with widely distributed, ascending projections throughout the neocortex. LC activation results in the release of norepinephrine (NE) in cortical areas, which increases the responsivity of these areas to their afferent input (Berridge and Waterhouse, 2003; Servan-Schreiber *et al.*, 1990). Traditionally, the LC-NE system has been associated with basic functions such as arousal and the sleep-wake cycle (Aston-Jones *et al.*, 1984; Jouvet, 1969), but recent studies have suggested that this system also plays a more specific role in the control of behavior (Aston-Jones *et al.*, 1997; Clayton *et al.*, 2004; Usher *et al.*, 1999). According to an influential recent theory of LC function, the adaptive gain theory (Aston-Jones and Cohen, 2005), the LC-NE system plays an important role in regulating the balance between exploiting known sources of reward versus exploring alternative options.

Neurophysiological studies in monkeys have revealed spontaneous fluctuations of tonic (baseline) LC activity over the course of a test session (Aston-Jones *et al.*, 1996; Kubiak *et al.*, 1992). Interestingly, these variations in tonic LC activity were closely related to the monkeys' control state: periods of moderate tonic LC activity were consistently associated with task engagement and accurate task performance, whereas periods of elevated tonic LC activity were associated with distractible behavior and poor task performance. Periods of very low or absent tonic LC activity were associated with drowsiness and inattention. Furthermore, periods of moderate tonic LC activity were accompanied by large phasic increases in LC activity following task-relevant stimuli, whereas such phasic LC responses were diminished during periods of elevated or low tonic LC activity. Thus, during alert task performance, the pattern of LC activity varied between moderate tonic/large phasic activity and elevated tonic/small phasic activity, which are referred to as the phasic and the tonic LC mode, respectively.

According to the adaptive gain theory (Aston-Jones and Cohen, 2005), the phasic and tonic LC modes promote, respectively, exploitative and exploratory control states. In the phasic mode, NE is released selectively in response to task-relevant events, which promotes task engagement and the optimization of performance in the current task (exploitation). In the tonic mode the sustained release of NE indiscriminately facilitates processing of all events, including non-task-related events, which promotes task disengagement and exploration. The theory further proposes that transitions between the phasic and tonic LC modes are driven by assessments of task-related costs and rewards (task utility), carried out in ventral and medial frontal structures.

The adaptive gain theory has been supported by computational modeling and neurophysiological studies in monkeys (Aston-Jones and Cohen, 2005; Usher *et al.*, 1999) and, indirectly, by recent pupillometry studies in humans (Gilzenrat *et al.*, 2010; Jepma and Nieuwenhuis, in press). However, crucial direct empirical tests of the theory in human participants have been lacking.

In the present study, we used a pharmacological manipulation to test in humans one of the central tenets of the adaptive gain theory, namely the assumption that the tonic LC mode promotes an exploratory control state. Participants received a single dose of reboxetine (a selective NE reuptake inhibitor), citalopram (a selective serotonin reuptake inhibitor) or placebo. Acute administration of reboxetine has opposing effects in the forebrain (increased NE levels via the inhibition of NE reuptake) and in the LC (reduction of firing activity via the increased activation of inhibitory $\alpha 2$ -autoreceptors; Szabo and Blier, 2001). However, microdialysis studies have shown that the net effect of these two actions is an increase in NE levels in various regions of the brain (for a wide range of reboxetine doses; Invernizzi and Garattini, 2004; Page and Lucki, 2002), which supposedly resembles the effects of elevated NE release in the tonic LC mode. To determine whether potential effects were selective for manipulations of the LC-NE system, we used citalopram as a control drug; it increases serotonin but not NE levels (Bymaster *et al.*, 2002). To confirm that these drugs at the doses employed in this study were pharmacologically active, we determined pupil size and several of the most drug-sensitive central nervous system (CNS) effects, including adaptive-tracking performance (index of visuomotor coordination and vigilance; Van Steveninck *et al.*, 1991, 1993) and saccadic peak velocity (index of alertness; Van Steveninck *et al.*, 1991, 1999).

The adaptive gain theory predicted that the increased tonic NE levels that were presumably induced by reboxetine would result in more task disengagement and exploratory behavior in the reboxetine group compared to the citalopram and placebo groups. We used two cognitive tasks to test these predictions. We measured task (dis)engagement using a diminishing-utility task (Gilzenrat *et al.*, 2010), in which task difficulty and potential reward—two determinants of task utility—increased over time. Importantly, participants had the opportunity to reset the level of task difficulty and reward, and hence disengage from the current task set. We measured exploratory behavior using a gambling task with a gradually changing pay-off structure (Daw *et al.*, 2006; Figure 2), in which optimal performance required a delicate balance between exploitative and exploratory choices.

Materials and methods

Participants

Fifty-two healthy university students, aged 18–25 years, took part in a single experimental session in return for €100,-. After signing an informed consent, participants were medically screened within 3 weeks before study participation. Exclusion criteria included history or presence of psychiatric disease and evidence of relevant clinical abnormalities.

Participants received a single oral dose of 4 mg reboxetine, 30 mg citalopram or placebo in a double-blind, parallel-groups design. The doses of reboxetine and citalopram were based on previous studies that have found significant behavioral effects using these doses of reboxetine (e.g., De Martino *et al.*, 2008; Miskowiak *et al.*, 2007; Tse and Bond, 2002) and citalopram (e.g., Chamberlain *et al.*, 2006). Unfortunately, the random-block design intended to produce equal numbers of men and women in each treatment group was thwarted by early dropouts and planning

problems, causing a somewhat unbalanced sex distribution. The reboxetine group (8 men, 10 women, mean age = 20.6), the citalopram group (8 men, 8 women, mean age = 21.6) and the placebo group (10 men, 8 women, mean age = 21.5) had similar mean ages ($F(2,49) = 1.66, p = 0.20$). The study was approved by the medical ethics committee of the Leiden University Medical Center and conducted according to the Declaration of Helsinki.

Procedure

All participants came to the research centre at 8AM after an overnight fast (except from water). We instructed participants to abstain from caffeine, nicotine, alcohol and other psycho-active substances from 10PM the night prior to the study day. On arrival, participants underwent a medical screening. Approximately one hour after arrival, participants in the citalopram group received a capsule with 2 mg granisetron, to prevent nausea as a potential side effect of citalopram. Participants in the reboxetine and placebo groups received a placebo capsule instead of granisetron. Sixty minutes later, participants received a capsule with reboxetine, citalopram or placebo.

Peak plasma concentrations of reboxetine and citalopram occur, respectively, 2 and 2-4 hours after drug administration (Dostert *et al.*, 1997; Edwards *et al.*, 1995; Hyttel, 1994; Noble and Benfield, 1997). Accordingly, the experimental tasks designed to measure task (dis)engagement and exploratory behavior were performed between 2 and 3 h post-treatment. All participants started with the diminishing-utility task, followed by the gambling task³. We measured participants' pupil-iris ratio (Twa *et al.*, 2004) and subjective state at several time points during the study day. Subjective state was assessed by means of sixteen 100-mm visual analogue scales measuring alertness, calmness and contentment (Bond and Lader, 1974). In addition, at several time points during the study day, we measured participants' adaptive-tracking performance (Borland and Nicholson, 1984; see Appendix for a description of the task) and saccadic eye movements (Van Steveninck *et al.*, 1989). These measures were part of a more extensive CNS test battery, the results of which will be reported more comprehensively elsewhere (te Beek *et al.*, in preparation). To assess drug-related effects on subjective state, pupil size, adaptive-tracking performance and saccadic eye movements, we compared the pre-treatment values with the average values from the time points surrounding performance of the diminishing-utility task and the gambling task (i.e., 2-3 h post-treatment). The complete time courses of these measures will be reported elsewhere (te Beek *et al.*, in preparation).

Diminishing-utility task

Participants performed an auditory pitch-discrimination task (Gilzenrat *et al.*, 2010). Each trial began with a sequence of two 250-ms sinusoidal tones: a reference tone, followed 3 s later by a comparison tone. Participants were instructed to indicate whether the comparison tone was higher

³ Due to technical problems, three participants did not complete one of the tasks and were excluded from the corresponding analyses. For the diminishing-utility task this was the case for one female participant in the citalopram group and one male participant in the placebo group, and for the four-armed bandit task this was the case for one male participant in the placebo group.

or lower in pitch than the reference tone, and earned points for each correct response. If participants responded correctly on a particular trial, the value of that trial was added to the participant's total score. In addition, in the next trial, the reward that could be earned increased by 5 points, and the pitch discrimination was made more difficult by halving the difference in pitch between the reference and comparison tones. Following an incorrect response, the reward value of the subsequent trial decreased by 10 points (but with a floor value of 0 points), and the level of task difficulty remained the same. Importantly, prior to each trial, participants had the opportunity to "escape" from the current series of discriminations without score penalty and receive a new discrimination task (i.e., comparison against a new reference tone), with the point value reset to 5 points and the easiest pitch discriminability. Participants were instructed to maximize their total score over the 20 minutes of the experiment.

The task procedure is illustrated in Figure 1. At the start of each trial participants were shown a score/value screen that displayed the total score accumulated thus far and the point value of the next trial. Participants then indicated with a key press whether they wanted to "accept" this trial or "escape". If the participant accepted the trial, a reference/comparison tone pair followed after a delay of one second. Participants were instructed to indicate as quickly and accurately as possible whether the comparison tone was lower or higher in pitch than the reference tone. After a delay of one second, the accuracy of the participant's response was indicated by a 250-ms feedback sound: a bell sound for correct responses and a buzzer sound for incorrect responses. Two seconds after the feedback sound, the next trial started. If participants pressed the "escape" button at the score/value screen, a 250-ms "escape sound" was played, immediately followed by a new score/value screen. We refer to a series of trials accepted by a participant as an "epoch" of play. Electing to escape begins a new epoch. We considered the average number of trials in an epoch as an index of task (dis)engagement.

In the first trial of each epoch, the difference in pitch between the two tones was 64 Hz. As noted above, this difference was halved following each correct response. If participants correctly discriminated a $\frac{1}{4}$ -Hz difference, the tones presented in the next trial were impossible to discriminate (i.e., 0 Hz difference), and impossible discrimination trials continued to be presented until the participant elected to escape. Accordingly, participants would exhaust any real discriminable differences between reference and comparison tone after nine correct trials; the tenth and subsequent trials within an epoch were impossible to discriminate. The feedback signal on impossible-discrimination trials was randomly picked. The same reference tone was presented on each trial within a given epoch. After an escape, a new reference tone was selected randomly without replacement from the set [400, 550, 700, and 850 Hz]. The set was replenished if all reference tones were exhausted. On 50% of the trials, the comparison tone was higher in pitch and on the remaining trials it was lower in pitch than the reference tone.




Figure 1. Illustration of a sample trial in the diminishing-utility task. See text for further details.

Gambling task

Participants performed a ‘four-armed bandit’ task (Daw *et al.*, 2006). On each trial, participants were presented with pictures of four different-colored slot machines, and selected one by pressing the ‘q’-, ‘w’-, ‘a’- or ‘s’- key. Participants had a maximum of 1.5 s in which to make their choice; if no choice was made during that interval, a red X appeared in the center of the screen for 4.2 s to signal a missed trial (average number = 2.5). If participants responded within 1.5 s, the lever of the chosen slot machine was lowered and the number of points earned was displayed in the chosen machine for 1 s after which the next trial started. The task consisted of 300 trials. Importantly, the number of points paid off by the four slot machines gradually and independently changed from trial to trial (Figure 2; Appendix).

Before the start of the experimental session, participants were given 24 practice trials. We instructed the participants that, on top of the standard payment for participation in the study, they would receive a bonus sum of money that depended on the number of points they would obtain in this task, and that the average bonus earned in this task was 9 euros. However, we did not tell participants how the number of points was converted into euros, or what their cumulative point total was. After completion of the study, each participant received a bonus of 10 euros.

Analysis. We fitted three reinforcement-learning models to the data. All models estimated the pay-offs of each machine on each trial, and selected a machine based on these estimations. The models differed in how they calculated the estimated pay-offs (Appendix). All models selected a machine according to the ‘softmax’ rule. This rule assumes that choices between different options are made in a probabilistic manner, such that the probability that a particular machine is chosen depends on its relative estimated pay-off. The exploitation-exploration balance is adjusted by a parameter referred to as gain, or inverse temperature: with higher gain, action selection is determined more by the relative estimated pay-offs of the different options (exploitation), whereas with lower gain, action-selection is more evenly distributed across the different options

(exploration). We classified each choice as exploitative or exploratory according to whether the chosen slot machine was the one with the maximum estimated pay-off (exploitation) or not (exploration). In addition, we calculated the degree of exploration for each exploratory choice, by subtracting the estimated pay-off of the chosen machine from the maximum estimated pay-off. We assessed the value of the gain parameter and the proportion of exploratory choices as a function of pharmacological treatment. Only the results from the best-fitting model are reported, although the other models yielded similar results.




Figure 2. The four-armed bandit task. Participants made repeated choices between four slot machines. Unlike standard slots, the mean pay-offs of the four machines changed gradually and independently from trial to trial (four colored lines). Participants were encouraged to earn as many points as possible during the task. Each choice was classified as exploitative or exploratory, using a computational model of reinforcement learning.

Results

Subjective state

The participants assigned to the three treatment groups did not differ in their pre-treatment ratings of alertness, calmness or contentment (all $ps > 0.7$; Table 1). To assess the effects of reboxetine and citalopram on subjective state we conducted analyses of covariance (ANCOVAs) on the subjective ratings of alertness, calmness and contentment, with treatment and sex as between-subject factors and the pre-treatment ratings as covariate. There were no main effects of treatment or sex, and no treatment by sex interactions on any of these ratings (all $ps > 0.16$), suggesting that reboxetine and citalopram did not affect subjective state.

Table 1. Pre- and post-treatment ratings of alertness, calmness and contentment in the placebo, citalopram and reboxetine group (SD in parentheses)

	Time of measurement	Placebo	Citalopram	Reboxetine
Alertness (mm)	Pre-treatment	51.2 (7.9)	52.2 (5.3)	50.6 (4.4)
	Post-treatment	50.2 (8.9)	52.4 (6.4)	48.6 (5.5)
Calmness (mm)	Pre-treatment	57.5 (9.9)	57.9 (10.2)	56.2 (4.4)
	Post-treatment	59.2 (10.7)	54.9 (9.4)	56.3 (6.1)
Contentment (mm)	Pre-treatment	55.9 (7.4)	56.7 (9.1)	55.9 (4.1)
	Post-treatment	57.5 (8.3)	56.4 (8.6)	56.9 (5.2)

Non-specific central and autonomic nervous system effects

Figure 3 (left panel) shows the adaptive-tracking performance pre-treatment (averaged across 1.5 and 0.5 h pre-treatment) and post-treatment (averaged across 2 and 3 h post-treatment) for each treatment group. We conducted an ANCOVA on the post-treatment adaptive-tracking performance with treatment and sex as between-subjects factors and pre-treatment performance as covariate. This analysis revealed a main effect of treatment [$F(2, 45) = 5.2, p = 0.009$]. There was no main effect of sex [$F(1, 45) = 0.8, p = 0.4$] and no interaction between treatment and sex [$F(2, 45) = 1.1, p = 0.3$]. Follow-up comparisons indicated that the reboxetine group showed worse post-treatment adaptive-tracking performance than the placebo group [$F(1, 31) = 12.0, p = 0.02$], whereas there was no difference between the citalopram and the placebo group [$F(1, 29) = 0.5, p = 0.5$]. The difference in post-treatment adaptive-tracking performance between the reboxetine and the citalopram group just failed to reach significance [$F(1, 29) = 3.8, p = 0.06$]. These results suggest that reboxetine led to a decrease in adaptive-tracking performance.

Figure 3 (middle panel) shows the saccadic peak velocity measured pre-treatment (averaged across 1.5 and 0.5 h pre-treatment) and post-treatment (averaged across 2 and 3 h post-treatment) for each treatment group. An ANCOVA on the post-treatment saccadic peak velocity with treatment and sex as between-subjects factors and pre-treatment saccadic peak velocity as covariate revealed a main effect of treatment [$F(2, 45) = 15.3, p < 0.001$]. There was no main effect of sex [$F(1, 45) = 1.8, p = 0.2$] and no significant interaction between treatment and sex [$F(2, 45) = 0.6, p = 0.6$]. Follow-up comparisons indicated that the reboxetine group showed smaller post-treatment saccadic peak velocity than the placebo group [$F(1, 31) = 5.1, p = 0.03$], whereas the citalopram group showed larger post-treatment saccadic peak velocity than the placebo group [$F(1, 29) = 8.6, p = 0.007$]. Thus, both reboxetine and citalopram affected saccadic eye movements, but the effects were in opposite directions. The time courses of saccadic peak velocity and adaptive-tracking performance showed that the effects of reboxetine and citalopram on these measures were maximal at the time points surrounding performance of the diminishing-utility task and the gambling task (te Beek *et al.*, in preparation), suggesting that the drug-related CNS effects were maximal during performance of these tasks.




Figure 3. Adaptive-tracking performance, saccadic peak velocity and pupil-iris ratio pre-treatment and post-treatment, separately for each treatment group (error bars indicate standard errors of the mean). PLA = placebo, CIT = citalopram, RBX = reboxetine.

Figure 3 (right panel) shows the pupil-iris ratio measured pre-treatment (averaged across 1.5 and 0.5 h pre-treatment) and post-treatment (averaged across 2, 2.5 and 3 h post-treatment) for each treatment group. An ANCOVA on the post-treatment pupil-iris ratio with treatment and sex as between-subjects factors and pre-treatment pupil-iris ratio as covariate revealed a main effect of treatment [$F(2, 45) = 22.1, p < 0.001$]. There was no main effect of sex [$F(1, 45) = 0.1, p = 0.7$] and no significant interaction between treatment and sex [$F(2, 45) = 2.8, p = 0.07$]. Follow-up comparisons indicated that both the reboxetine group and the citalopram group had larger post-treatment pupil-iris ratios than the placebo group [$F(1, 31) = 7.1, p = 0.01$ and $F(1, 29) = 44.4, p < 0.001$, respectively]. In addition, post-treatment pupil-iris ratio was larger in the citalopram group than the reboxetine group [$F(1, 29) = 13.7, p = 0.001$]. Thus, consistent with previous studies (Phillips *et al.*, 2000; Schmitt *et al.*, 2002), both citalopram and reboxetine led to an increase in pupil diameter, and this effect was more pronounced in the citalopram group. There is no reliable evidence for direct projections from the LC to the autonomic nuclei that control the pupil (Aston-Jones, 2004), but there are a number of possible indirect pathways by which LC manipulation could affect the sympathetic nervous system (cf. Berntson *et al.*, 1998). Therefore, it is possible that the increase in pupil diameter in the reboxetine group reflects drug-induced changes in LC activity. However, it is also possible that the pharmacological effects on pupil diameter were produced at the level of the autonomic nuclei controlling the pupil, and thus reflect other drug actions than changes in LC activity.

Diminishing-utility task

The progressive increase in both task difficulty and potential reward during each series of tone discriminations produces a nonlinear development of task-related utility. Initially, the increases

in reward value for correct performance outpace the increases in difficulty, such that the expected value (utility) of task performance progressively increases. However, after several trials, the increases in difficulty will lead to sufficient number of errors as to reduce the expected value of performance, even in the face of increasing reward value for correct responses.

To examine changes in performance and task-related utility leading up to and following participants' choice to 'escape' (i.e., abandon the current series and start a new one), we averaged trials as a function of their position relative to the escape events. For this analysis, we considered only escape events that were preceded and followed by a minimum of four regular (i.e., non-escape) trials. As a measure of task utility, we calculated an estimate of expected value for each trial. For a given trial, expected value was computed individually for each participant by multiplying the point value of the trial (representing the potential reward value if the trial was accepted) by the expected accuracy on that trial for that participant. Expected accuracy was defined as the probability that the participant would give a correct response, given the level of difficulty of the required pitch discrimination. To determine this, we averaged the accuracy of all other trials for that participant with the same frequency difference between reference and comparison tones.

Figure 4 (left panels) shows the average accuracy and RT on the trials flanking an escape for each treatment group. All treatment groups showed a sharp decrease in accuracy and an increase in RT over the trials leading up to an escape, which was confirmed by significant linear trends [$F(1,44) = 462.5, p < 0.001$ and $F(1,44) = 14.3, p < 0.001$, respectively]. As expected, performance was best on the first trial following an escape, after which accuracy gradually decreased and RT increased again [$F(1,44) = 54.5, p < 0.001$ and $F(1,44) = 35.1, p < 0.001$, respectively]. Figure 4 (right panels) shows how our measure of expected value and the actual point value varied across the trials surrounding an escape. In all treatment groups, participants on average selected to escape when expected value approached the start value of a new series of discriminations. Both expected value and point value gradually decreased over the trials leading up to an escape [$F(1,44) = 100.1, p < 0.001$ and $F(1,44) = 30.5, p < 0.001$, respectively], and gradually increased again over the trials following an escape [$F(1,44) = 422.1, p < 0.001$ and $F(1,44) = 1079.0, p < 0.001$, respectively]. Importantly, the effects of peri-escape trial position on performance and task utility did not interact with treatment or sex (all $ps > 0.3$).

We next examined the average number of accepted trials in an epoch. The average number of trials in an epoch did not differ between the three treatment groups [$F(2,44) = 0.26, p = 0.77$]. There was no main effect of sex either [$F(1,44) = 1.08, p = 0.30$], and no interaction between treatment and sex [$F(2,44) = 0.33, p = 0.72$]. Furthermore, there was no significant across-subject correlation between the mean epoch length and the reboxetine-related change in adaptive-tracking performance [$r = 0.43, p = 0.08$]. Note that, if anything, this correlation showed a trend in the opposite direction than predicted by the adaptive gain theory. Mean epoch length was not significantly correlated with the drug-related increase in pupil diameter either [$r = -0.13, p = 0.62$ in the reboxetine group; $r = 0.24, p = 0.38$ in the citalopram group].

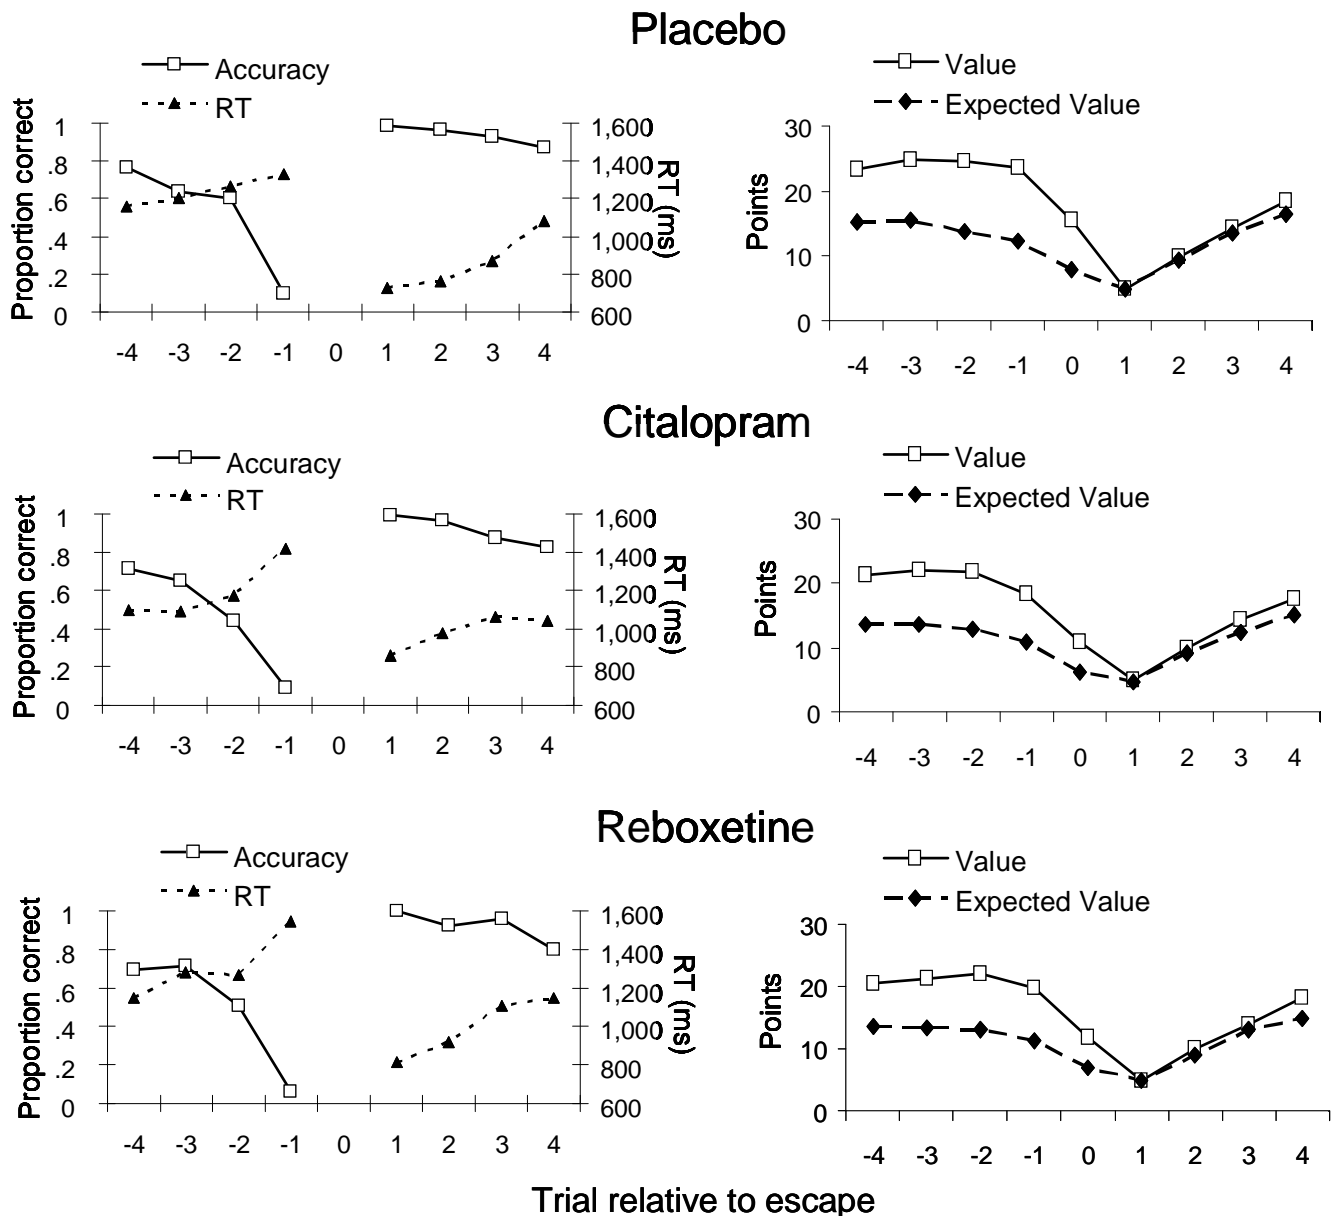


Figure 4. Dependent measures for peri-escape trials in the three treatment groups. Trial number “0” indicates the escape trial. Left panels: accuracy and response time (RT). Right panels: Trial value and its computed expected value. Note that no measures of accuracy and RT are available for escape trials, because, on these trials, no comparison tone was presented.

There were no effects of treatment or sex on the total number of trials completed or total number of points obtained (all $ps > 0.3$), except for a significant interaction between treatment and sex on the total number of point obtained [$F(2,44) = 3.68, p = 0.03$]. Follow-up contrasts indicated that the male participants obtained significantly more points than the female participants in the reboxetine group [$t(16) = 3.08, p = 0.007$], whereas there were no significant sex effects in the placebo and citalopram groups ($ps > 0.48$). An overview of the dependent variables in this task as a function of treatment and sex is shown in Table 2. An analysis of the improvement in tone-

discrimination performance over the course of the task (i.e., learning curve) is reported in the Appendix.

Table 2. Overview of the dependent variables in the diminishing utility task, as a function of treatment and sex (SD in parentheses).

	Placebo		Citalopram		Reboxetine	
	men	women	men	women	men	women
Mean epoch length (trials)	10.3 (2.3)	12.1 (4.3)	9.9 (2.5)	10.9 (4.1)	11.0 (3.8)	11.0 (2.3)
Number of escapes	12.8 (3.1)	11.5 (4.0)	13.4 (3.7)	12.9 (5.1)	13.3 (5.9)	11.8 (4.7)
Total score	1694 (380)	1749 (418)	1496 (537)	1674 (404)	1904 (353)	1356 (391)
Total number of trials	136 (3)	136 (3)	135 (3)	136 (3)	138 (3)	132 (3)

Gambling task

Each participant's tendency to make exploratory choices is reflected in the estimated gain parameter of the reinforcement-learning model: a lower value of the gain parameter indicates a more exploratory choice strategy (Materials and Methods; Appendix). The value of the gain parameter did not differ between the three treatment groups [$F(2,45) = 0.70, p = 0.51$; Supplemental Table 1] or between the male and female participants [$F(2,45) = 2.50, p = 0.12$]. In addition, we classified each choice as exploitative or exploratory according to whether the chosen slot machine was the one with the maximum estimated pay-off (exploitation) or not (exploration). The proportion of exploratory choices did not differ between the three treatment groups [28%, 32% and 27% in the placebo, citalopram and reboxetine group, respectively; $F(2,45) = 0.92, p = 0.41$] or between male and female participants [26% vs. 31%; $F(2,45) = 2.43, p = 0.13$]. The three treatment groups did not differ in the degree of exploration of the exploratory choices either (section 2.4.1); the degrees of exploration in the placebo, citalopram and reboxetine groups were 0.39, 0.37 and 0.37, respectively ($F(2,45) = 0.43, p = 0.65$).

Neither the value of the gain parameter nor the proportion of exploratory decisions was significantly correlated with the reboxetine-related change in adaptive-tracking performance [gain parameter: $r = 0.41, p = 0.09$; proportion exploration: $r = -0.25, p = 0.32$]. Our measures of exploration were not significantly correlated with the drug-related increase in pupil diameter either ($ps > 0.15$ in the reboxetine group; $ps > 0.35$ in the citalopram group).

There were no across-subject correlations between our measure of task disengagement in the diminishing-utility task (mean epoch length) and our measures of exploration in the gambling task (value gain parameter and proportion of exploratory choices; $ps > 0.8$). This suggests that the disengagement and exploration measures in these tasks reflect separate aspects of the exploratory control state hypothesized to be mediated by the tonic LC mode.

Discussion

The present study provided the first direct test in humans of one of the central tenets of the adaptive gain theory of LC function (Aston-Jones and Cohen, 2005), namely the assumption that an elevated level of tonic LC/NE activity (tonic LC mode) promotes a more exploratory control state. Contrary to predictions of the adaptive gain theory, we found no evidence that the increased NE levels induced by reboxetine were associated with task disengagement or exploratory behavior in our experimental tasks.

Our null effects cannot be explained by a general ineffectiveness of our pharmacological manipulations, since there were significant drug effects on several central and autonomic nervous system parameters. Reboxetine caused reductions in adaptive-tracking performance and in saccadic peak velocity, which corroborates previous findings suggesting the involvement of the noradrenergic system in visuomotor control of movements (Wang *et al.*, 2009). Citalopram increased saccadic peak velocity, which is in line with the mild stimulating properties of the SSRI on the electroencephalogram (Itil *et al.*, 1984; Saletu *et al.*, 2002). The time course of the effects suggests that reboxetine was maximally effective during performance of the diminishing-utility task and gambling task (te Beek *et al.*, in preparation). In addition, both citalopram and reboxetine resulted in an increase in pupil diameter, but it is unknown whether these pupil modulations were produced by changes in LC activity or by other drug influences peripheral to the LC (e.g., on lower medullary NE cell groups or autonomic nervous system). Furthermore, previous studies using the same dose of reboxetine, between-subject designs and similar group sizes have found significant group differences in behavioral measures (De Martino *et al.*, 2008; Miskowiak *et al.*, 2007; Tse and Bond, 2002). The absence of significant across-subject correlations between our measures of disengagement/exploration and the reboxetine-related effects on adaptive-tracking performance suggests that the effectiveness of the reboxetine manipulation in individual participants did not predict their tendency to disengage or explore.

The two experimental tasks we used to measure exploratory behavior and task (dis)engagement seem well suited for detecting individual differences in control state. The n -armed bandit task with non-stationary pay-off structure is the most commonly used paradigm for studying the exploration-exploitation trade-off in reinforcement-learning research (Sutton and Barto, 1998). Combined with computational modeling, it allows a formal description of participants' choice behavior and provides an index of their tendency to explore. The diminishing-utility task is a more novel paradigm in which task engagement is modulated by means of dynamic changes in task-related utility. Importantly, the opportunity to "escape" from the current task set provides an overt behavioral index of disengagement. In line with a previous study using this task (Gilzenrat *et al.*, 2010), we found that participants behaved optimally on average, and chose to disengage from the current task set when estimated task utility approached the baseline utility of a new task set. In addition, in a recent study using the same gambling task as used here (Jepma and Nieuwenhuis, in press) we have found that changes in utility measures and pupil diameter leading up to the switch

from an exploitative to an exploratory choice strategy were similar to those leading up to an “escape” in the diminishing-utility task (Gilzenrat *et al.*, 2010). This suggests that disengagement in the diminishing-utility task and exploration in the gambling task are both driven by decreases in task utility. That said, optimal exploration strategies in our experimental tasks may differ from those needed in the real world; the changes in pay-offs and task-related utility in our tasks developed gradually and relatively slowly over time, which may not correspond to the dynamics of utility changes in real-world environments (Cohen *et al.*, 2007).

Although disengagement and exploration are both considered behaviors indicative of an exploratory control state associated with the tonic LC mode, it is important to note that disengagement in the diminishing utility task (i.e., choosing to “escape” from the current series of tone discriminations) is not equivalent to exploration in the gambling task, which may explain the absence of a correlation between our measures of disengagement and exploratory behavior. The development of a computational model for the diminishing-utility task is an important objective for future studies, as this will allow a more formal description of participants’ behavior in this task and a better comparison with exploratory behavior in other tasks.

One possible explanation for the absence of reboxetine effects on our measures of task disengagement and exploratory behavior is that the LC-NE system is not involved in regulating the balance between exploitative and exploratory control states in humans. The adaptive gain theory is based on findings from neurophysiological studies in monkeys using relatively simple target-detection tasks, and it is possible that the results from these studies cannot be generalized to the regulation of control state in humans. Moreover, although it is intuitively appealing to interpret the observations of increased distractibility, labile attention and impaired focused performance during elevated tonic LC/NE activity in animals as reflections of an exploratory control state (Aston-Jones and Cohen, 2005), it is important to note that the neurophysiological studies did not explicitly investigate the exploration-exploitation trade-off; the proposed link between the tonic LC mode and an exploratory control state is an assumption. Because we did not find evidence for this assumption, it seems appropriate to consider alternative explanations for the distractible behavior associated with the tonic LC mode. When taking a reinforcement-learning model perspective, it may be possible to explain the behaviors observed in the tonic LC mode by changes in reinforcement-learning parameters other than the exploration parameter. One possibility is that high LC/NE activity increases the rate at which action values are updated based on new information (i.e., the learning rate parameter). This hypothesis would be compatible with a recent proposal that increased NE levels boost the learning of new task contingencies (Yu and Dayan, 2005). In line with this hypothesis, the estimated learning rate of the reinforcement-learning model that we fit to the choice data of the gambling task was somewhat larger in the reboxetine group than in the other treatment groups (Appendix, Supplemental Table 2; Supplemental Figure 2). However, because of the very high learning rates associated with this task, this result must be interpreted with caution. Alternatively, high LC/NE activity may increase the importance attached to immediate vs. delayed rewards (i.e., the future-reward discount factor). Support for this hypothesis comes from findings

from a recent study in mice that suggest that drug-induced increases in NE levels impair the ability to take future rewards into account, which would lead to the impulsive selection of options with short-term rewards (Luksys et al., 2009). Luksys et al. suggested that the distractible behavior observed in animals with elevated LC/NE activity can be produced by an increased devaluation of future, relative to immediate, rewards combined with high *exploitation* (as opposed to exploration; see Doya, 2002, for a similar proposal). Thus, the behaviors associated with the tonic LC mode that have been interpreted as indices of an exploratory control state by the adaptive gain theory may also be explained by modulations of other reinforcement-learning parameters. To further address this issue, future studies need to dissociate the role of the LC-NE system and other neuromodulatory systems in the regulation of different components of reinforcement learning and decision making.

Another possibility is that the tonic LC mode promotes a type of exploratory behavior and disengagement that was not measured in the present study. It is likely that exploration is not a single process but comprises several distinct functions involving different neural mechanisms. An important aspect may be whether exploration is driven by top-down motives or by bottom-up stimulation. Exploratory behavior in the four-armed bandit task may be referred to as ‘controlled’ or ‘systematic’ exploration, since it is aimed at obtaining information in order to optimize performance in the current task. Similarly, disengaging from the current task set in the diminishing-utility task serves the higher-level goal of maximizing the total score obtained in the task. Such controlled, top-down driven exploration and disengagement *within the current task context* might be mediated by different neural mechanisms and/or neuromodulatory systems than random, bottom-up driven exploration *exceeding the current task context*. Controlled exploration presumably requires cognitive control functions that rely on the prefrontal cortex (PFC), which is supported by the finding of PFC activation during exploratory decisions in the four-armed bandit task (Daw *et al.*, 2006). There is also some evidence that the dopamine system plays a role in the regulation of a particular type of controlled exploration (Frank *et al.*, 2009). Our findings suggest that the LC-NE system may not be involved in controlled exploration. However, our study leaves open the possibility that the LC-NE system is involved in random exploration exceeding the current task context. Random exploration is likely to be associated with an increased sensitivity to bottom-up activation, resulting from a global increase in neuronal responsiveness. The widespread projection system of the LC and the neuromodulatory effects of NE on cortical neurons suggest that the LC-NE system is well suited to produce such global changes in responsiveness.

The idea that the tonic LC mode promotes a more random type of exploration outside the current task context is supported by findings that drug-related increases in tonic NE levels improve attentional-set shifting and reversal learning in rats and monkeys (Devauges and Sara, 1990; Lapid and Morilak, 2006; Lapid *et al.*, 2007; Seu *et al.*, 2008), whereas noradrenergic lesions impair attentional-set shifting (McGaughy *et al.*, 2008; Newman *et al.*, 2008; Tait *et al.*, 2007). These functions require the adaptation of behavior according to unexpected changes in the task environment, which depends on a shift of attention to previously irrelevant stimulus dimensions. These types of attention shifts are likely to be facilitated by random exploration (although an

increased learning rate may provide an alternative explanation). Investigating the noradrenergic modulation of random exploration outside the current task context in humans is an important objective for future studies.

The distinction between controlled and random exploration might be related to the proposed distinction between expected and unexpected uncertainty (Yu and Dayan, 2005). Yu and Dayan have proposed that acetylcholine signals expected uncertainty (i.e., anticipated variation in task outcome), whereas NE signals unexpected uncertainty (i.e., unanticipated changes in the task context resulting in strong violations of top-down expectations; see Bouret and Sara, 2005, for a similar account). Yu and Dayan have also proposed that the NE-related signaling of unexpected uncertainty facilitates the learning of predictive relationships within a behavioral context, and therefore accelerates the detection of a change in task contingencies, which could explain the improvements in attentional-set shifting associated with increased tonic NE levels. Yu and Dayan's account thus suggests that the tonic LC mode boosts learning about new predictive relationships in noisy and changing environments. This account is closely related to the adaptive gain theory's assumption that the tonic LC mode promotes exploration, at least when applied to random exploration exceeding the current task context, since this type of exploration is likely to facilitate the learning of contextual changes. The detection of unexpected uncertainty might be an important factor in driving the LC towards a more tonic LC mode. However, how assessments of unexpected uncertainty interact with assessments of task-related utility on different timescales to regulate LC mode and control state remains to be investigated. An interesting speculation is that the degree of unexpected uncertainty determines how much weight is given to assessments of long versus short-term utility, such that long-term utility has relatively less influence in situations of high unexpected uncertainty. In terms of reinforcement-learning models, this would be similar to the suggested modulation of the learning rate parameter by the volatility of the environment (Behrens *et al.*, 2007).

Finally, it is important to note that although microdialysis studies have shown that a single dose of reboxetine increases NE concentrations, these studies, due to their limited temporal resolution, do not provide unequivocal evidence that this reflects purely an increase in tonic NE levels. Since the effects of selective NE reuptake inhibitors on the phasic LC response in awake animals are not known, we cannot exclude the possibility that our reboxetine manipulation also affected phasic LC activity and NE release, for example via modulations of the electrotonic coupling strength between LC neurons (Alvarez *et al.*, 2002). Thus, determining the exact effects of selective NE reuptake inhibitors on the phasic and tonic components of LC/NE activity will be important for a better understanding of their effects on cognition. In addition, the effects of pharmacologically increasing NE levels on control state might depend on individual differences in baseline (pre-treatment) NE level. Accordingly, individual differences in baseline NE level could have been partly responsible for the absence of group differences on our measures of disengagement and exploration. Consistent with this possibility, a recent study in mice has shown that pharmacological manipulations of the LC-NE system interact with several other factors, such as

individual differences in genotype and trait anxiety, stress and motivation, in modulating the exploration-exploitation trade-off (Luksys *et al.*, 2009). Thus, it seems that multiple factors need to be taken into account to enable predictions of exploratory behavior and its modulation by NE.

To conclude, our findings suggest that the acute induction of an elevated tonic NE level does not affect people's tendency to explore or disengage, at least not within the current task context. These findings challenge the adaptive gain theory's claim that the LC-NE system regulates the balance between exploitative and exploratory control states (Aston-Jones and Cohen, 2005). It remains to be examined whether the LC-NE system is involved in random exploration outside the current task context, possibly driven by the detection of unexpected uncertainty. The present study contributes to our understanding of the noradrenergic modulation of human control state, and hopefully encourages further investigation of this topic.

Appendix

Adaptive-tracking task

The adaptive-tracking task is a pursuit-tracking task (Borland and Nicholson, 1984). A target circle moves randomly on a computer screen, and the participant must try to keep a marker dot inside the moving circle by operating a joystick. The mean velocity of the moving circle is automatically adjusted to match the participant's skill. If the participant is successful in maintaining the dot inside the circle, the velocity of the moving circle gradually increases. Conversely, if the participant cannot maintain the dot inside the circle, the velocity is reduced. The task lasts 3.5 minutes, including a run-in period of 0.5 minute during which no data is recorded. Performance is measured as the percentage of time that the participant is able to keep the dot in the circle. The adaptive-tracking task has proved to be useful for measurement of CNS effects of alcohol, various psychoactive drugs and sleep deprivation (Cohen et al. 1985; Van Steveninck et al., 1991, 1999).

Pay-off structure of the gambling task

The number of points paid off by slot machine i on trial t ranged from 1 to 100, drawn from a Gaussian distribution (standard deviation $\sigma_o = 4$) around a mean $\mu_{i,t}$ and rounded to the nearest integer. At each trial, the means diffused in a decaying Gaussian random walk:

$$\mu_{i,t+1} = \lambda\mu_{i,t} + (1 - \lambda)\theta + \nu$$

The decay parameter λ was 0.9836, the decay center θ was 50, and the diffusion noise ν was zero-mean Gaussian (standard deviation $\sigma_d = 2.8$). We used three instantiations of this process; one is illustrated in Figure 2.

Description of the reinforcement-learning models

We fitted three reinforcement-learning models to the choice data of the gambling task. All models consisted of a mean-tracking rule that tracked the expected pay-offs of each machine ($\hat{\mu}_{i,t}$), and a choice rule that selected a machine based on these estimations. The estimated pay-offs were calculated as follows:

Model 1 (mean pay-off estimation without decay; Dayan and Abbott, 2001)

When a participant chooses machine c on trial t and receives pay-off r , the estimated pay-off of the chosen machine is updated according to:

$$\hat{\mu}_{c,t}^{post} = \hat{\mu}_{c,t}^{pre} + \hat{\kappa}\delta_t$$

with prediction error $\delta_t = r_t - \hat{\mu}_{c,t}$ and learning rate parameter $\hat{\kappa}$. The estimated pay-offs of the unchosen machines do not change.

Model 2 (mean pay-off estimation with decay)

The chosen machine's estimated pay-off is updated as in model 1:

$$\hat{\mu}_{c,t}^{post} = \hat{\mu}_{c,t}^{pre} + \hat{\kappa} \delta_t$$

In addition, the estimated pay-offs of all machines, regardless of choice, are updated in time according to:

$$\hat{\mu}_{i,t+1}^{pre} = \hat{\lambda} \hat{\mu}_{i,t}^{post} + (1 - \hat{\lambda}) \hat{\theta}$$

in which $\hat{\lambda}$ is the decay parameter (a smaller value of $\hat{\lambda}$ indicates a faster decay rate) and $\hat{\theta}$ is the decay-center parameter.

Model 3 (Kalman filter; Daw et al., 2006)

The pay-offs of the machines are updated as in model 2. In addition to tracking the mean pay-offs ($\hat{\mu}_{i,t}$), this model also tracks the uncertainties about these pay-offs ($\hat{\sigma}_{i,t}^2$, i.e., the variance of the expected pay-off distributions) which determine the trial-specific learning rates κ_t . When a participant chooses machine c on trial t and receives pay-off r , the estimated pay-off distribution of the chosen machine ($\hat{\mu}_{c,t}^{post}, \hat{\sigma}_{c,t}^{2post}$) is updated according to:

$$\hat{\mu}_{c,t}^{post} = \hat{\mu}_{c,t}^{pre} + \kappa_t \delta_t$$

$$\hat{\sigma}_{c,t}^{2post} = (1 - \kappa_t) \hat{\sigma}_{c,t}^{2pre}$$

with prediction error $\delta_t = r_t - \hat{\mu}_{c,t}^{pre}$ and learning rate $\kappa_t = \hat{\sigma}_{c,t}^{2pre} / (\hat{\sigma}_{c,t}^{2pre} + \hat{\sigma}_o^2)$.

Then, the estimated prior pay-off distributions of all machines on the subsequent trial (trial $t+1$) are updated in time according to:

$$\hat{\mu}_{i,t+1}^{pre} = \hat{\lambda} \hat{\mu}_{i,t}^{post} + (1 - \hat{\lambda}) \hat{\theta}$$

$$\hat{\sigma}_{i,t+1}^{2pre} = \hat{\lambda}^2 \hat{\sigma}_{i,t}^{2post} + \hat{\sigma}_d^2$$

In all models, the selection of a machine on each trial was determined by a softmax rule; the probability $P_{i,t}$ of choosing machine i on trial t as the function of the estimated pay-offs was:

$$P_{i,t} = \frac{\exp(\beta \hat{\mu}_{i,t}^{pre})}{\sum_j \exp(\beta \hat{\mu}_{j,t}^{pre})}$$

with exploration parameter β (referred to as the gain, or inverse temperature).

We fitted each model to the participants' choice data by maximizing the log-likelihood of the observed choices. To optimize the parameter fits, we used a nonlinear optimization algorithm (Matlab's `fminsearch` function; Lagarias et al., 1998), together with a search of different starting values. The trials in which no response was made within the 1.5-s time limit were omitted. The pay-off tracking parameters ($\hat{\kappa}$, $\hat{\lambda}$ and $\hat{\theta}$) were shared by all participants that had received the same

pharmacological treatment, whereas the exploration parameter (β) was estimated separately for each participant. Parameter $\hat{\sigma}_o$ in model 3 was fixed at 4. Estimation of parameter $\hat{\sigma}_d$ in model 3 resulted in extreme values for most of the participants, suggesting unreliable fits. Therefore, we fixed this parameter at 50, which is similar to the best fitting $\hat{\sigma}_d$ parameter found in a previous study (Daw et al., 2006). Large values of $\hat{\sigma}_d$ induce high learning rates, indicating that the expected pay-offs are determined primarily by the most recent experience with each machine. Given that the estimated learning rate parameters in models 1 and 2 were very near or even slightly above 1 as well (Supplemental Table 1), and that previous studies have also associated this task with high learning rates (Daw et al., 2006; Jepma and Nieuwenhuis, in press), the oversensitivity to the most recent pay-off of each machine seems to be characteristic of participants' choice behavior in this task.

To compare the adequacy of the three models in explaining the observed data we used the Bayesian Information Criterion (BIC; Raftery, 1996), a statistical criterion for model selection. The BIC is an increasing function of the residual sum of squares from the estimated model, and an increasing function of the number of free parameters to be estimated. Thus, the best model is the model with the lowest BIC value. In addition, the raw BIC values were transformed to a probability scale (BIC model weights or "Schwarz weights"), enabling a more intuitive comparison of the probabilities of each model being the best model, given the data and the set of candidate models (Wagenmakers & Farrell, 2004). Supplemental Table 1 shows the estimated parameter values and the BIC values and model weights of each model. Model 2 (mean pay-off estimation with decay) provided by far the best fit to the choice data.

Supplemental Table 1. Mean parameter estimates and fit information for the three models, separately for each treatment group (SD in parentheses). Model 2 provided the best fit to the data.


		Model 1	Model 2	Model 3
β	Placebo	0.095 (0.028)	0.137 (0.042)	0.197 (0.058)
	Reboxetine	0.105 (0.039)	0.152 (0.081)	0.245 (0.129)
	Citalopram	0.093 (0.035)	0.135 (0.053)	0.157 (0.061)
$\hat{\lambda}$	Placebo	-	0.73	0.70
	Reboxetine	-	0.73	0.65
	Citalopram	-	0.85	0.84
$\hat{\theta}$	Placebo	-	45.9	45.6
	Reboxetine	-	45.6	45.3
	Citalopram	-	49.7	49.5
$\hat{\kappa}$	Placebo	0.93	1.07	-
	Reboxetine	1.03	1.17	-
	Citalopram	0.86	1.01	-
-LL	Placebo	4380	3789	3821
	Reboxetine	4415	3751	3780
	Citalopram	4349	3858	3901
BIC	Placebo	8913	7757	7804
	Reboxetine	8994	7691	7732
	Citalopram	8842	7885	7954
p	Placebo	< 0.001	> 0.999	< 0.001
	Reboxetine	< 0.001	> 0.999	< 0.001
	Citalopram	< 0.001	> 0.999	< 0.001

Note: Model 1 = mean pay-off estimation without decay; Model 2 = mean pay-off estimation with decay; Model 3 = pay-off distribution estimation with decay; -LL = negative log likelihood (smaller values indicate better fit); BIC = Bayesian information criterion; p = BIC model weight.

Tone discrimination learning curves in the diminishing-utility task

To examine whether the three treatment groups showed different rates of improvement in tone-discrimination performance over the course of the task (i.e., different learning curves), we divided all trials in four equally sized consecutive trial bins, separately for each participant and each level of task difficulty, and assessed the mean percentage of correct tone discriminations in each trial bin (Supplemental Figure 1). The trials with impossible discriminations (i.e., 0 Hz tone differences) were excluded from the analysis. There was a significant main effect of trial bin on tone-discrimination performance [$F(3,132) = 10.1, p < 0.001$], which was best described by a linear improvement over the four sequential bins [$F(1,44) = 15.9, p < 0.001$]. This learning effect interacted with treatment at a trend level [$F(6,132) = 2.1, p = 0.057$], but did not differ between the male and female participants ($p = 0.48$). Follow-up comparisons indicated that the learning curve in

the reboxetine group differed from those in the placebo and citalopram groups [$F(3,93) = 2.5, p = 0.07$ and $F(3,87) = 2.7, p = 0.05$, respectively]; whereas the placebo and citalopram groups showed a significant linear improvement over the four consecutive bins (linear trend effect $ps < 0.002$ for both groups), the effect of trial bin in the reboxetine group was best described by a cubic trend [$F(1,17) = 11.8, p = 0.003$] reflecting the initial decrease in performance in trial bins 2 and 3 followed by an increase in performance in the last bin.



Supplemental Figure 1. Learning curves illustrating the change in tone discrimination performance over the four consecutive trial bins in the diminishing-utility task, separately for each treatment group (error bars indicate standard errors of the mean).

Bootstrap analysis of the shared parameters in the reinforcement-learning model




To approximate the distribution of the shared parameters, $(\hat{\lambda}, \hat{\theta}$ and $\hat{\kappa})$, we conducted a bootstrap analysis (Efron & Tibshirani, 1993). For each treatment group, the computer generated 2162 bootstrap sets by sampling with replacement from the original group of participants; each bootstrap set had the same number of "participants" as the original data set. Model 2 was fitted to the choice data from each bootstrap set, which resulted in a bootstrap sampling distribution for each parameter in each treatment group (Supplemental Figure 2).

To assess whether the $\hat{\lambda}$, $\hat{\theta}$ and $\hat{\kappa}$ parameter values differed between the three treatment groups we determined the 95% confidence interval of each parameter in each group (Supplemental Table 2). The distributions of the $\hat{\lambda}$ parameter suggest that $\hat{\lambda}$ is larger in the citalopram group than in the other two groups, indicating a slower decay rate (i.e., slower forgetting of the estimated values) in the citalopram group. However, the bootstrap-based 95% confidence interval of the citalopram group partly overlaps with that of the other treatment groups, hence the difference misses significance. The trend for a slower decay rate in the citalopram group may be consistent with findings that serotonin manipulations affect the sensitivity for short- vs. long-term consequences of

actions (e.g., Schweighofer et al., 2008). The values of $\hat{\theta}$ and $\hat{\kappa}$ did not differ significantly between the three groups, although there was a trend for a somewhat higher learning rate in the reboxetine group.

Supplemental Table 2. The 2.5, 50 and 97.5 percentile of the bootstrap sampling distributions of the $\hat{\lambda}$, $\hat{\theta}$ and $\hat{\kappa}$ parameters. The 2.5 and 97.5 percentiles indicate the lower and upper bound of the 95% confidence interval.

		percentile		
		2.5	50	97.5
$\hat{\lambda}$	Placebo	0.31	0.73	0.85
	Reboxetine	0.20	0.73	0.84
	Citalopram	0.78	0.85	0.89
$\hat{\theta}$	Placebo	42.4	45.9	49.6
	Reboxetine	40.1	45.7	49.9
	Citalopram	45.1	49.7	53.5
$\hat{\kappa}$	Placebo	0.17	1.05	1.22
	Reboxetine	0.49	1.16	1.30
	Citalopram	0.89	1.00	1.09



Supplemental Figure 2. Bootstrap distributions of the $\hat{\lambda}$ (decay parameter; larger values indicate slower decay rate), $\hat{\theta}$ (decay center) and \hat{K} (learning rate) parameters in each treatment group

Chapter 4

Neurocognitive function in dopamine- β -hydroxylase deficiency

This chapter is published as: Jepma, M., Deinum, J., Asplund, C.L., Rombouts, S.A., Tamsma, J.T., Tjeerdema, N., Spapé, M.M., Garland, E.M., Robertson, D., Lenders, J.W., & Nieuwenhuis, S. (in press). Neurocognitive function in dopamine- β -hydroxylase deficiency. *Neuropsychopharmacology*

Abstract

Dopamine- β -hydroxylase (D β H) deficiency is a rare genetic syndrome characterized by the complete absence of norepinephrine in the peripheral and the central nervous system. D β H-deficient patients suffer from several physical symptoms, which can be treated successfully with L-threo-3,4-dihydroxyphenylserine, a synthetic precursor of norepinephrine. Informal clinical observations suggest that D β H-deficient patients do not have obvious cognitive impairments, even when they are not medicated, which is remarkable given the important role of norepinephrine in normal neurocognitive function. The present study provided the first systematic investigation of neurocognitive function in human D β H deficiency. We tested five D β H-deficient patients and ten matched healthy control participants on a comprehensive cognitive task battery, and examined their pupil dynamics, brain structure and the P3 component of the electroencephalogram. All participants were tested twice; the patients were tested once ON and once OFF medication. Magnetic resonance imaging scans of the brain revealed that the patients had a smaller total brain volume than the control group, which is in line with the recent hypothesis that norepinephrine has a neurotrophic effect. In addition, the patients showed an abnormally small or absent task-evoked pupil dilation. However, we found no substantial differences in cognitive performance or P3 amplitude between the patients and the control participants, with the exception of a temporal-attention deficit in the patients OFF medication. The largely spared neurocognitive function in D β H-deficient patients suggests that other neuromodulators have taken over the function of norepinephrine in the brains of these patients.

Introduction

The locus coeruleus–norepinephrine (LC–NE) system is one of the major neuromodulatory systems in the brain. For a long time, investigators have associated this system with basic functions such as arousal and the sleep-wake cycle (Aston-Jones et al., 1984; Jouvet, 1969), and with various neuropsychiatric disorders such as depression and attention-deficit hyperactivity disorder (Ressler and Nemeroff, 2001; Siever and Davis, 1985). In addition, recent studies have shown that the LC–NE system is involved in more specific cognitive functions, such as memory, attention, perception, and decision making (Aston-Jones and Cohen, 2005; Robbins, 1997; Sara, 2009). These findings suggest that norepinephrine (NE) is essential for normal cognitive function in humans.

D β H deficiency is a rare genetic syndrome that is characterized by the congenital absence of the enzyme dopamine- β -hydroxylase (D β H), which is responsible for the conversion of dopamine (DA) to NE (Man in't Veld et al., 1987a; Robertson et al., 1986). As a result, D β H deficiency is characterized by a complete lack of NE and epinephrine in both the central and the peripheral nervous system (Man in't Veld et al., 1987a). There are currently approximately 15 patients with D β H deficiency known worldwide. These patients suffer from several physical symptoms, including severe orthostatic hypotension, fatigue and impaired exercise tolerance (Robertson and Garland, 2010). The only effective treatment of D β H deficiency involves administration of the drug L-threo-3,4-dihydroxyphenylserine (DOPS, droxidopa), which is converted directly into NE via L-aromatic-amino-acid decarboxylase, thereby bypassing D β H (Biaggioni and Robertson, 1987; Goldstein, 2006; Man in 't Veld et al., 1987b). Studies in rats and mice have shown that DOPS crosses the blood-brain barrier, and activates the production of NE in the central nervous system as well as the peripheral nervous system (Ishikawa et al., 1987; Kato et al., 1987a,b; Semba et al., 1985; Thomas et al., 1998). Treatment with DOPS results in a dramatic relief of physical symptoms and a substantial improvement of the quality of life of D β H-deficient patients.

The biochemical features, autonomic physiology and physical symptoms associated with human D β H deficiency have already been described in several studies (e.g., Mathias et al., 1990; Robertson et al., 1991; Thompson et al., 1995; Timmers et al., 2004). In addition, a post-mortem microscopic examination of the brain of one D β H-deficiency patient has revealed no histological abnormalities and no evidence for neuronal loss (Cheshire et al., 2006). However, to date there have been no systematic studies on cognitive and brain function in D β H deficiency. Informal clinical observations suggest that even before starting treatment, D β H-deficient patients do not have obvious cognitive impairments, which is striking given the large amount of evidence that NE plays an important role in normal cognitive function (Sara, 2009). This suggests that more carefully controlled laboratory tests may reveal subtle neurocognitive deficits in D β H-deficient patients that have remained unnoticed in informal observations.

The present study provides the first systematic evaluation of neurocognitive function in D β H deficiency. We tested 5 patients with D β H deficiency on a battery of cognitive tasks that have been proposed to depend on normal noradrenergic function, including an emotional working-

memory task (Chamberlain et al., 2006; Oei et al., 2010) and a temporal-attention task (attentional-blink task; De Martino et al., 2007; Nieuwenhuis et al., 2005a; Warren et al., 2009), expecting that these tasks would reveal possible abnormalities in the D β H-deficient patients. In addition, we examined task-evoked changes in pupil diameter, and recorded the electroencephalogram (EEG) during a target-detection task to examine event-related potential (ERP) correlates of noradrenergic activity (Liu et al., 2009; Nieuwenhuis et al., 2005b; Pineda et al., 1989). To assess whether potential abnormalities in performance were restricted to NE-mediated tasks, we also tested the patients on a spatial-attention task that does not probe noradrenergic function (Greenwood et al., 2005; Nieuwenhuis et al., 2007). Finally, we acquired an MRI scan of the patients' brain to assess possible abnormalities in brain volume and structure. We tested the patients once ON and once OFF DOPS medication, and compared their results with those of a matched healthy control group.

Materials and methods

Participants

We tested five D β H-deficient patients (two Dutch, two American, and one Canadian) and ten healthy controls (all Dutch). The two American patients were brothers, and the other patients were unrelated (see Supplementary Table 1 for the patients' demographic and clinical details). The genetic mutations in the DBH gene have been identified for all patients. Patient 1 is homozygous for the IVS1 +2T>C mutation, a mutation of the 5' splice site in the first intron which leads to abnormal splicing and hence a dysfunctional protein. Patient 2 is homozygous for a missense mutation in 764G>T (C255F; Deinum et al., 2004). Patients 3 and 4 are heterozygous for both the IVS1 +2T>C mutation and the 991G>A (D331N) missense mutation. Patient 5 is homozygous for two missense mutations in 259G>A (V87M) and 991G>A (D331N). Patient 5 also has a rare mosaic deletion at chromosome 11p13 [46,XX,del(11)(p12p14)/46,XX] which is unrelated to her D β H-deficiency (Erez et al., 2010).

The patient and control group were matched for age, sex and IQ (Table 1). We used the Vocabulary subtest of the Wechsler Adult Intelligence Scale (WAIS III, Wechsler, 1997) and the Raven's Standard Progressive Matrices test (SPM; Raven et al., 1988) to estimate IQ. The Dutch patients and their controls were matched for educational level as well. Given the different educational systems in the US and the Netherlands it was not possible to match the American patients and their Dutch control participants in terms of educational level; hence we matched for estimated IQ instead of educational level. Participants gave written informed consent before participation, and the study was approved by the medical ethics committee of the Leiden University Medical Center and the Institutional Review Board of Vanderbilt University.

Table 1. Demographic details of the control group and the patient group (means \pm standard deviations)

	Control group (N = 10)	Patient group (N = 5)
Age (years)	24.6 \pm 11.0	24.4 \pm 10.0
Sex (proportion female)	6/10	3/5
Interval between test sessions (days)	7.5 \pm 3.2	7.6 \pm 2.7
Scaled WAIS-III vocabulary score	8.6 \pm 2.3	11.4 \pm 3.4
Raven's SPM score	44.5 \pm 6.9	45.6 \pm 4.6
Estimated IQ (based on SPM score)	106.5 \pm 11.1	107.2 \pm 8.6

Notes: WAIS = Wechsler Adult Intelligence Scale, highest possible scaled vocabulary score = 19; SPM = standard progressive matrices, highest possible score = 66.

General procedure

All participants were tested twice on the same cognitive-task battery, with an intervening period of six to thirteen days. The patient and control groups had similar intervening periods (Table 1). Two patients were tested ON medication on the first test day and OFF medication on the second test day, and the other three patients were tested in the opposite order. Two of these patients had never been on DOPS medication before and started taking medication at least two days before the second test day. The other patients stopped taking their daily medication four to thirteen days before the OFF-medication test day and stayed off medication up to and including this day. Preceding and during the ON-medication test day, the patients took their DOPS medication as usual (see Supplementary Table 1 for the patients' demographic and clinical details).

The task battery included five cognitive tasks, described below and, in more detail, in Appendix I. At the beginning and end of each test day, participants completed the Positive Affect and Negative Affect Schedule (PANAS; Watson et al., 1988; translated into Dutch by Peeters et al., 1996). To measure catecholamine levels, we collected blood and 24-hour urine samples from the patients, prior to each test session (Table 1). Blood samples were taken after fifteen minutes of supine rest. We also collected blood samples from most control participants. Since we expected no differences in catecholamine levels between the two sessions for the control participants, their blood samples were collected only once. Finally, on one of the test days a structural T1-weighted MRI brain scan was acquired (see Appendix I for details of acquisition and analysis).

Emotional working-memory task

NE plays an important role in emotional memory (i.e., Chamberlain et al., 2006). The well-known phenomenon that emotional events are memorized better than neutral events (e.g., Cahill and McGaugh, 1998), for example, is associated with β -adrenergic-dependent modulations of amygdala-hippocampus interactions (Strange et al., 2003; Strange and Dolan, 2004). In addition, emotional distractor stimuli impair working-memory performance to a higher degree than neutral distractor stimuli (e.g., Buchner et al., 2004; Dolcos and McCarthy, 2006; Oei et al., 2009, 2010), an effect that is reduced by administration of the β -adrenergic antagonist propranolol (Oei et al.,

2010). We examined the effects of emotional and neutral distractor stimuli on performance in the working-memory task used by Oei et al. (2009, 2010).

Each trial of this task started with the presentation of either one or four letters (the target set), which had to be held in memory for later recognition. The target set was followed by a 1,500-ms delay period during which either a neutral picture or a negatively arousing picture was presented. After this, four letters (the probe set) were presented and participants had to indicate, as quickly and accurately as possible, whether or not the probe set contained a letter from the preceding target set.

Attentional-blink task

The attentional-blink paradigm is the most commonly used paradigm for investigating attentional selection in the temporal domain (for a review see Martens and Wyble, 2010). The attentional blink refers to a deficit in processing the second of two target stimuli that are presented in close temporal succession. This deficit is most severe when the second target is presented within 200-400 ms after the first target (Raymond et al., 1992), and is thought to result from competition between the two target stimuli for limited attentional resources (Shapiro et al., 1997). When the two targets are presented within approximately 200 ms, performance is often spared (e.g., Hommel and Akyürek, 2005), a phenomenon termed “lag-1 sparing”.

The temporal dynamics of the LC-NE system suggest that the LC-NE system mediates attentional selection in the temporal domain (Cohen et al., 2004; Dayan and Yu, 2006; Usher et al., 1999). LC neurons exhibit a phasic increase in activity shortly following task-relevant or otherwise motivationally significant stimuli (Aston-Jones et al., 2000). The resulting transient release of NE in cortical areas temporarily increases the responsivity of these areas to their input, which selectively facilitates the processing of the eliciting stimulus (Berridge and Waterhouse, 2003; Servan-Schreiber et al., 1990). Phasic increases in LC activity are followed by a brief refractory period during which LC-NE-mediated facilitation of information processing is temporarily unavailable (e.g., Aghajanian et al., 1977). These temporal dynamics of the LC-NE system suggest that the attentional blink may be mediated by the LC-NE system (Nieuwenhuis et al., 2005a; Warren et al., 2009). Consistent with this idea, β -adrenergic blockade impaired detection of the second target in an attentional-blink task (De Martino et al., 2007).

On each trial of this task, participants viewed a rapid serial visual presentation (RSVP) stream consisting of 2 target stimuli (T1 and T2; digits) and multiple distractor stimuli (letters), presented for about 100 ms each. The temporal distance between T1 and T2 was 1, 2, 3 or 7 items. Following each stream, participants were asked to report T1 and T2.

Visual-search task

This task examined attentional selection in the spatial domain. The spatially-nonspecific pattern of LC projections to the cortex suggests that the LC-NE system does not mediate spatial attention (Cohen et al., 2004; Greenwood et al., 2005; Nieuwenhuis et al., 2007). This task was

included to assess whether possible performance abnormalities of the D β H-deficient patients were restricted to NE-mediated tasks. On each trial of this task, participants searched for a target stimulus (a red vertical bar) among a variable number of distractor stimuli (green vertical bars and red horizontal bars) in a visual-search array, and indicated as quickly as possible whether the target stimulus was present or absent.

Oddball tasks combined with EEG measurement

We examined the P3, a prominent component of the scalp-recorded event-related brain potential. The P3 component is a broad, positive, large-amplitude potential which peaks between 300 and 400 ms following presentation of stimuli in any sensory modality (Sutton et al., 1965), and is largest over central-parietal midline electrodes. The amplitude of the P3 is strongly affected by the subjective probability and motivational significance of the eliciting stimulus: P3 amplitude increases with decreasing probability and with increasing motivational significance of the eliciting stimulus. In contrast, with the exception of tone intensity (Roth et al., 1984), P3 amplitude is relatively insensitive to physical stimulus properties. Several lines of evidence suggest that the P3 reflects the phasic response of the LC-NE system to the outcome of stimulus evaluation and decision making, and the consequent effects of the noradrenergic potentiation of information processing (reviewed in Nieuwenhuis et al., 2005b; see also Liu et al., 2009; Pineda et al., 1989).

The most common paradigm for studying the P3 is the oddball task, in which infrequent target stimuli are embedded in a series of frequently presented non-target stimuli (standards), and participants have to respond to each target stimulus but not to the standard stimuli. We measured participants' EEG while they performed visual and auditory versions of the oddball task, and assessed the P3 elicited by target stimuli.

Pitch-discrimination task combined with pupillometry

We examined participants' pupil diameter during performance of a pitch-discrimination task. Although the luminance level is the most important determinant of pupil diameter, there are also small but reliable changes in pupil diameter related to cognitive processing (Beatty and Wagoner, 1978; Kahneman, 1973). A large number of studies have shown that task processing is accompanied by a rapid increase in pupil diameter, and that the size of this pupil dilation reflects the information-processing load (e.g., Hess and Polt, 1964).

Several studies have reported that D β H-deficient patients have small pupils, but a normal pupillary light reflex and accommodation response (Biaggioni et al., 1990; Man in 't Veld et al., 1987a; Robertson et al., 1986). In addition, one study reported a prolonged redilation time following the light reflex in a sibling pair with D β H deficiency (Smith and Smith, 1999). The light reflex and accommodation response both produce pupil constrictions, which are subserved by the iris sphincter muscles. These muscles are innervated by cholinergic input from the parasympathetic nervous system. In contrast, pupil dilation is controlled by the iris dilator muscles which are activated

primarily via noradrenergic innervation of α -1 adrenoceptors (Hoffman and Taylor, 2001). This suggests that task-evoked pupil dilations in D β H-deficient patients might be abnormal.

On each trial of this task, a sequence of two tones was presented, and participants had to indicate whether the second tone was higher or lower in pitch than the first. We analyzed participants' baseline pupil diameter and their pupil dilation in response to the second tone.

Results

The control participants' behavioral, EEG and pupil data were analyzed by means of repeated-measures ANOVAs, with session (session 1 vs. session 2) and the independent task variables as within-subject factors. We tested whether the critical measures/effects in each patient OFF medication deviated from those in the control group using a modified *t*-test developed specifically to compare individual patients with a small control group (Crawford and Howell, 1998). In addition, we examined the effects of medication on the patients' scores, using the regression-based method developed by Crawford and Garthwaite (2006; see Appendix I for details of these analyses).

We focus our description of the results on the critical measures/effects of each task. The full factorial analyses of the data, the PANAS (i.e., subjective state) data, and results of the individual participants are reported in the Appendix II.

Catecholamine concentrations

Table 2 shows the average plasma and urine NE and DA concentrations in the patient group ON and OFF medication, and the plasma concentrations in the control group (see Supplementary Table 2 for the data from the individual patients). When OFF medication, two of the patients (patients 3 and 4) had plasma NE concentrations that were significantly lower than that in the control group [*ps* (1-tailed) < 0.03; Crawford and Howell's (1998) modified *t*-test] and the other patients had undetectable plasma NE concentrations. The apparent extremely low residual plasma NE concentration in patients 3 and 4 were likely due to technical artifacts, since plasma concentrations of the NE metabolite dihydroxyphenylglycol (DHPG) were extremely low in these patients when they were OFF medication. DHPG concentrations in patients 3 and 4 OFF medication were lower than 0.03 nmol/l, which is less than 1% of normal. As expected, all patients' plasma and urine NE concentrations were higher when ON compared to OFF medication, and this effect was especially pronounced for the urine concentrations. For the ON-medication session, the plasma NE concentrations of patient 1 and 5 did not differ significantly from the control group [*p*(1-tailed) = 0.09 and 0.08, respectively], but the plasma NE concentrations of patient 3 and 4 were still lower than that in the control group [*p*(1-tailed) = 0.048 and 0.049, respectively].

When OFF medication, all patients had higher plasma DA concentrations than the control group (all *ps* < 0.001). Although most patients' plasma DA concentrations were lower when ON compared to OFF medication, the ON medication concentration was still larger than that in the

control group for all but one patient. The medication effects on the urine DA concentrations were less consistent; patients 1 and 2 had higher urine DA concentrations when ON medication, whereas patients 3, 4 and 5 showed the opposite effect.

Table 2. Plasma and urine catecholamine concentrations in the control group and the patient group OFF and ON medication (means \pm standard deviations).

	Healthy controls ⁺	Patients OFF	Patients ON
Plasma NE	1.46 \pm 0.45	0.10 \pm 0.12	0.57 \pm 0.13
Urine NE	-	5.50 \pm 5.40	9682 \pm 4839
Plasma DA	0.06 \pm 0.02	1.28 \pm 1.43	0.40 \pm 0.40
Urine DA	-	1271 \pm 903	793 \pm 379

Notes: ⁺ plasma concentrations were determined for 6 control participants; OFF = off medication; ON = on DOPS medication; all concentrations are in nmol/l; see Supplementary Table 2 for the catecholamine concentrations of the individual patients and missing data.

Emotional working-memory performance

The critical measure in this task was the interfering effect of emotional relative to neutral distractors on reaction time (RT). As expected, the control participants responded more slowly on trials with emotional compared to neutral distractors [$F(1, 7) = 14.7, p = 0.006$]. In addition, consistent with previous studies (Oei et al., 2009, 2010), distractor type interacted with target presence [$F(1, 7) = 16.3, p = 0.005$], indicating that the emotional-interference effect on RT was significant on target-present trials [$F(1, 7) = 43.9, p < 0.001$; effect range = 80 - 299 ms] but not on target-absent trials [$F(1, 7) = 0.75, p = 0.42$].




Figure 1. Average emotional-interference effect (i.e., RT on trials with emotional relative to neutral distractors) for the control group and the patient group OFF and ON medication, as a function of target presence (error bars are standard errors of the means). Because session did not interact with distractor type or target presence in the control group, the results from the control group are averaged across the two sessions.

Figure 1 shows the average increase in correct RT on trials with emotional relative to neutral distractors as a function of target presence, in the control group and in the patient group OFF and ON medication. When OFF medication, all patients showed an emotion-related slowing of responses on target-present trials that did not differ from the effect in the control group (effect range = 72 - 226 ms; all $t_s(7) < 0.8$; $p_s > 0.24$; Table 4; see Supplementary Figure 2 for the individual effects). In addition, all patients showed a smaller emotional interference effect when they were ON compared to OFF medication, but this medication effect did not differ significantly from the control group's practice effect in any of the patients (all $p_s > 0.08$; Table 4). The normal emotional-interference effect in the patients OFF medication, and the finding that this interference effect was less pronounced when the patients were ON medication are both remarkable given the evidence that emotional-interference effects are normally mediated by NE.

The full factorial analysis of the effects of target presence, working-memory load, distractor type and session on correct RT and accuracy in the control group is reported in Appendix II and in Supplementary Figure 1.

Attentional-blink performance

Figure 2 shows the average T1 accuracy (upper panels) and T2 accuracy (lower panels; contingent on correct T1 identification) in the control group and the patient group, as a function of lag (1, 2, 3 or 7) and session. The T2 accuracy curves show a pattern that is characteristic of attentional blink data: lag-1 sparing, followed by a drop in performance for lags 2 and 3 (i.e., the attentional blink), and a recovery of performance at lag 7. This pattern was expressed in a significant effect of lag in the control group [$F(3, 27) = 12.1, p = 0.001$].

The critical measure in this task is the size of the attentional blink, which we defined as the decrease in T2 identification accuracy at lags 2 and 3, relative to lag 7 (Maclean and Arnell, 2010). When OFF medication, the patient group showed a larger attentional blink than the control group (average = 33.5% vs. 16.7%), but the difference from the control group only approached significance in patient 1 (Table 4; see Supplementary Figure 3 for the individual T2 accuracy curves). In addition, the patients showed a smaller attentional blink when they were ON compared to OFF medication: for three of the four patients tested on this task, the effect of medication on attentional-blink size was significantly larger than the practice effect in the control group ($p_s < 0.05$; Table 4). The fourth patient also showed a marked increase in T2 accuracy when ON compared to OFF medication, but this did not result in a significant effect on attentional-blink size because the enhancing effect of medication was present at lags 2, 3 and 7. Together, these findings suggest that T2 identification accuracy during the attentional blink was impaired in the patients OFF medication, and that this impairment was restored by the DOPS medication.




Figure 2. Average T1 and T2 identification accuracy in the attentional-blink task for the control group and the patient group, as a function of lag and session (error bars are standard errors of the means). Trials on which T1 and T2 were accurately identified but in the wrong order were treated as correct. As is usual, T2 accuracy is reported contingent on accurate identification of T1.

Visual-search performance

The critical measure in this task was the effect of set size (i.e., the total number of items in the search display) on RT. As expected, RT in the control group showed an increasing trend with set size [$F(2, 18) = 29.7, p < 0.001$], and set-size effects were larger for target-absent than target-present trials [$F(2, 18) = 7.8, p = 0.004$]. The variation in set size allowed us to derive the function relating RT to set size. The slope of this function measures the cost for adding additional items to the display and is often interpreted as “search efficiency,” with steeper slopes indicating slower, less efficient search.




Figure 3. Average visual-search slopes for the control group and the patient group, as a function of target presence and session (error bars are standard errors of the means).

Figure 3 shows the average slopes for the control group and the patient group, as a function of target presence and session. The average slopes in the patient group were very similar to those in the control group, both ON and OFF medication. In the OFF-medication session, none of the patients' slopes deviated significantly from the control group (all $t_s(9) < 1.2$; $p_s > 0.13$; Table 4; see Supplementary Figure 5 for the individual slopes). In addition, the effects of medication did not differ significantly from the control group's practice effect in any of the patients (all $p_s > 0.11$; Table 4). These results indicate that the patients had normal visual search efficiency, both ON and OFF medication.

The full factorial analysis of the effects of target presence, set size and session in the control group is reported in Appendix II and in Supplementary Figure 4.

The P3 component of the electroencephalogram

P3 amplitudes were maximal at electrode Pz in both the control group and the patient group; hence we focused our analyses on this electrode position. Figure 4 shows the grand average waveforms for standard and target stimuli in the visual and auditory oddball task, for the control group and the patient group ON and OFF medication. As expected, P3s were much larger for target stimuli than for standard stimuli. Figure 5 shows the P3 amplitudes of the individual participants.

When OFF medication, patient 5 showed a significantly smaller P3 amplitude than the control group in both the auditory and the visual oddball task, and patient 4 showed a significantly smaller P3 amplitude than the control group in the visual oddball task only (Table 4). For the other patients, P3 amplitude did not differ significantly from the control group. The effect of medication on P3 amplitude did not differ significantly from the control group's test-retest effect in any of the patients (all $p_s > 0.19$; Table 4). These findings suggest that some but not all patients showed a P3 that was smaller than the P3 in the normal population, independently of whether they were ON or OFF medication.

The analyses of target-detection performance (RT and accuracy) are reported in Appendix II and in Supplementary Figure 6.




Figure 4. Grand-average waveforms for electrode Pz for the control group and the patient group, time-locked to the onset of the target and standard stimuli, in the auditory and visual oddball tasks. Because P3 amplitude in the control group did not differ across sessions [$F(1, 9) = 0.1, p = 0.72$], the data for the control participants are averaged across the two sessions.




Figure 5. P3 amplitudes for the control participants and the patients in the auditory and visual oddball task. The bold lines indicate the grand average amplitudes, and the thinner lines and points indicate the amplitudes of each individual participant. Because there was no effect of session in the control group, the data for the control participants are averaged across the two sessions.

Pupil diameter during the pitch-discrimination task

The average baseline pupil diameter in the control group was 3.86 mm (SD = 0.56), and did not differ across the two sessions [$t(7) = 0.31, p = 0.77$]. When OFF medication, patient 2 had significantly smaller pupils than the control group. Patient 5 had significantly larger pupils than the control group which was due to a genetic deficit unrelated to her D β H deficiency (Erez et al., 2010). The other patients' baseline pupil diameter did not differ significantly from the control group (Table 4; see Supplementary Table 1 for each patient's baseline pupil diameter). Remarkably, patient 4 had significantly smaller pupils when he was ON compared to OFF medication. For the other patients, there was no significant effect of medication on baseline pupil diameter (Table 4).

We next assessed the magnitude of the task-evoked pupil dilations. As expected, all control participants showed a substantial pupil dilation following the comparison tone (average pupil dilation = 0.16 mm; SD = 0.04). Pupil dilation in the control group was not significantly affected by session [$F(1, 7) = 2.3, p = 0.17$] or tone-discrimination difficulty [$F(3, 21) = 2.4, p = 0.09$]. Figure 6 shows the time course of the grand-average pupil dilation following the comparison tone, for the control group and the patient group ON and OFF medication. When OFF medication, all but one patient showed significantly smaller task-evoked pupil dilations than the control group (see Supplementary Table 1 for each patient's average pupil dilation). Remarkably, patient 4 showed a significantly smaller pupil dilation when ON compared to OFF medication. The pupil dilation of patient 3 was also significantly affected by medication, but this result must be interpreted with caution because this patient's pupil dilations were negative in both sessions. For the other patients, there was no significant effect of medication on the task-evoked pupil dilation (Table 4).

The analyses of tone-discrimination performance (RT and accuracy) are reported in Appendix II and in Supplementary Figure 7.




Figure 6. Time course of the grand-average pupil dilations in response to the comparison tone, for the control group and the patient group ON and OFF medication

Brain structure

Table 3 shows the average total brain volumes and the percentages of grey matter, white matter and cerebrospinal fluid (CSF) in the patient group and the control group, separately for the male and female participants. Four of the five patients had a smaller total brain volume than the control group. However, the proportions of grey matter, white matter and CSF did not differ from the control group in any of the patients (Table 4; see Supplementary Table 1 for the data of the individual patients).

The voxel-based morphometry analysis (Appendix I) revealed no significant topographic differences in grey matter volume between the patient group and the control group. The TCFE-corrected p -values for both the controls > patients contrast and the patients > controls contrast were larger than 0.34 in all voxels, suggesting that there were no trends for a group difference in grey matter distribution in any brain region. Together, these results suggest that most of the patients had an overall smaller brain than the control group, but that this difference was not confined to a specific tissue type or brain region.

Table 3. Whole-brain volume and percentage of grey matter, white matter and cerebrospinal fluid for the control group and the patient group, separately for the male and female participants (means \pm standard deviations).

	Control group		Patient group	
	Men (N = 4)	Women (N = 5)	Men (N = 2)	Women (N = 3)
Brain volume (dm ³)	1.72 \pm 0.05	1.46 \pm 0.05	1.51 \pm 0.06	1.30 \pm 0.11
% grey matter	47.2 \pm 0.7	44.1 \pm 1.8	47.0 \pm 1.5	44.7 \pm 3.2
% white matter	38.7 \pm 1.2	39.3 \pm 2.0	38.5 \pm 0.5	38.9 \pm 2.2
% CSF	14.1 \pm 1.7	16.5 \pm 2.4	14.5 \pm 1.0	16.4 \pm 1.4

Notes: we did not collect MRI data from one female control participant; CSF = cerebrospinal fluid

Table 4. For each critical effect/measure, the *p* value reflecting the significance of the difference between each patient's OFF medication score and the average score of the control group (Crawford and Howell, 1998), and the *p* value indicating the significance of the deviation of each patient's medication effect from the control group's practice effect (Crawford and Garthwaite, 2006). *p* values < 0.05, which indicate that the estimated percentage of the normal population that would show a more extreme effect is smaller than 5%, are bold-faced.

Patient OFF medication vs. control group	patient				
	1	2	3	4	5
Emotional-interference effect on RT in target-present trials	-	0.44	0.25	0.29	0.41
Attentional-blink size	0.051	-	0.19	0.10	0.38
Visual search efficiency in target-present trials	0.36	0.15	0.14	0.15	0.41
Visual search efficiency in target-absent trials	0.41	0.23	0.36	0.24	0.50
P3 amplitude auditory oddball task	0.10	0.34	0.16	0.09	0.04
P3 amplitude visual oddball task	0.27	0.32	0.13	0.048	0.01
Baseline pupil diameter	-	0.03	0.45	0.22	0.002⁺
Pupil dilation response	-	0.03	0.003	0.21	0.001
Brain volume (dm ³)	0.29	0.006	0.03	0.01	0.02
% grey matter	0.46	0.16	0.16	0.10	0.053
% white matter	0.33	0.39	0.33	0.46	0.12
% cerebrospinal fluid	0.39	0.30	0.45	0.30	0.35
Patient's medication effect vs. control group's practice effect	1	2	3	4	5
Emotional-interference effect on RT in target-present trials	-	0.19	0.09	0.18	0.19
Attentional-blink size	0.045	-	0.003	0.049	0.24
Visual search efficiency in target-present trials	0.45	0.28	0.28	0.28	0.39
Visual search efficiency in target-absent trials	0.22	0.11	0.18	0.50	0.16
P3 amplitude auditory oddball task	0.19	0.44	0.41	0.38	0.19
P3 amplitude visual oddball task	0.21	0.27	0.29	0.36	0.35
Baseline pupil diameter	-	0.25	0.20	0.003	0.21
Pupil dilation response	-	0.34	0.03	0.04	0.08

- = no data were collected; ⁺ this patient had significantly larger pupils than the control group, which was due to a genetic defect unrelated to DBH deficiency: a deletion on the short arm of chromosome 11 (Erez et al., 2010)

Discussion

The present study was the first systematic investigation of neurocognitive function in D β H deficiency. We tested five D β H-deficient patients and a matched healthy control group on a comprehensive cognitive task battery. In addition, we examined whether the patients differed from the control group with regard to the P3 component of the electroencephalogram, pupil dynamics and brain structure.

The patients' performance on most cognitive tasks did not differ substantially from the healthy control group, irrespective of whether they were ON or OFF DOPS medication. More specifically, the patients showed normal visual-search efficiency, tone-discrimination performance and target-detection performance, and a normal emotional-interference effect. In addition, we found an intact P3 component in most patients. Since DOPS medication effectively ameliorates D β H-deficient patients' orthostatic hypotension, medication-related changes in blood pressure and consequent effects on fatigue and affective state are important factors to take into account when comparing the patients' performance ON versus OFF medication. However, it is unlikely that these factors were responsible for the lack of medication effects on cognitive performance, for the following reasons. First, potential effects of fatigue or other physical symptoms on task performance would predict impaired performance when patients were OFF relative to ON medication, which was not found in most tasks. Second, the patients reported no substantial differences in affective state between the two sessions (Supplementary Table 3). Third, the critical measures in our cognitive tasks were difference scores (i.e. differences between task conditions), hence general medication-related effects on performance would cancel out in these difference scores.

The only cognitive function that was affected in the patients OFF medication was attentional selection in the temporal domain, as reflected by an increased attentional blink (i.e., impairment in processing the second of two target stimuli that are presented in close temporal succession). The attentional blink has not only been associated with NE (De Martino et al., 2007; Nieuwenhuis et al., 2005a; Warren et al., 2009), but also with dopamine (DA; Colzato et al., 2008); Colzato et al. have provided indirect evidence that higher DA levels are associated with a smaller attentional blink. Because D β H-deficient patients do not convert DA to NE, they are not only characterized by a lack of NE but also by increased DA levels (Man in 't Veld, 1987a), and DOPS medication both increases NE levels and reduces the excessive DA levels (Man in 't Veld et al., 1987b; Thomas et al., 1998). Thus, based on the patients' DA levels, it would be predicted that the patients OFF medication would show a smaller attentional blink than the healthy control group, and that the patients would show a smaller attentional blink OFF medication than ON medication. Since the opposite effects were found, this strongly suggests that the increased attentional blink in the patients OFF medication was due to the absence of NE rather than the excess of DA.

The largely spared neurocognitive function in the D β H-deficient patients is remarkable given the large body of evidence suggesting that the LC-NE system plays an important role in many

aspects of neurocognitive function (for recent reviews see Robbins and Arnsten, 2009; Sara, 2009). For example, individual differences in noradrenergic genotype in the normal population are predictive of performance on cognitive tasks measuring attention (Greene et al., 2009) and working memory (Parasuraman et al., 2005), and have been related to vulnerability to several psychiatric disorders (e.g., Cubells and Zabetian, 2004; Roman et al., 2002). In addition, D β H-knockout mice that lack NE due to a targeted disruption of the D β H gene show several behavioral deficits, including impairments in active-avoidance learning (Thomas and Palmiter, 1997a), memory retrieval (Murchison et al., 2004), and maternal and social behavior (Marino et al., 2005; Thomas and Palmiter, 1997b). Finally, pharmacological, neurophysiological, and lesion studies in animals suggest that the LC-NE system plays a crucial role in regulating the optimization of behavioral performance (e.g., Aston-Jones and Cohen, 2005; Bouret and Sara, 2005). It must be noted, however, that our task battery did not address all aspects of cognitive function. For example, we did not assess higher-level cognitive functions such as executive control and exploratory behavior. Therefore, our results leave open the possibility that the patients have subtle cognitive deficits that were not revealed by our task battery. In addition, although our data clearly indicate that there were no substantial abnormalities in the patients' performance on our test battery, it cannot be excluded that there were some subtle differences which failed to reach significance due to a lack of power of our experimental design.

Although the patients' relatively normal performance on our cognitive task battery is striking, it is consistent with informal clinical observations that D β H-deficient patients do not have obvious cognitive impairments or psychiatric disorders. Indeed, the absence of mental problems in most D β H-deficient patients that have been encountered so far has intrigued investigators in the areas of depression and schizophrenia (Cubells and Zabetian, 2004). It is especially remarkable that the patients OFF medication did not show impaired performance on cognitive tasks that are normally mediated by the LC-NE system (e.g., the emotional working-memory task), and showed a relatively intact P3 component, which is thought to reflect the noradrenergic potentiation of information processing (Liu et al., 2009; Nieuwenhuis et al., 2005b; Pineda et al., 1989). These findings suggest that alternative neural mechanisms and/or neuromodulatory systems compensate for the absence of NE in D β H-deficient patients. Previous findings that D β H-deficient patients have a relatively normal sleep pattern (Tulen et al., 1990; 1991), although the sleep-wake cycle is normally mediated by the LC-NE system (Hobson et al., 1986; Jouvet, 1969), are consistent with this idea.

Since D β H is responsible for the conversion of DA to NE, it is thought that DA rather than NE is stored and released by noradrenergic neurons in D β H-deficient patients. Indeed, plasma DA levels in DBH-deficient patients respond to various physiological and pharmacological manipulations that normally affect plasma NE levels (Man in 't Veld, 1987a; Robertson et al., 1986), although it remains to be determined whether this also applies to DA levels in the central nervous system. Thus, a possible explanation for the spared neurocognitive function in D β H deficiency is that DA has, to some extent, taken over the function of NE in the brains of D β H-

deficient patients. Obviously, a functional replacement of NE by DA would require the presence of postsynaptic receptors with DA affinity in noradrenergic synapses. Studies in mice suggest that some α 2-adrenergic receptor subtypes have a comparable affinity for DA and NE (Zhang et al., 1999), whereas α 1- and β -adrenergic receptors have a much lower affinity for DA than for NE (Zhang et al., 2004). However, since the congenital absence of NE may have altered the affinity of adrenergic receptors, it is unknown whether the same receptor characteristics apply to D β H-deficient patients. Another possible explanation for a functional replacement of NE by DA is that D β H-deficient patients have an increased density of postsynaptic DA receptors on noradrenergic synapses. A recent positron emission tomography (PET) study in mice suggests that D β H knockout mice have a normal density of D2 dopamine receptors in the high-affinity state (Skinbjerg et al., 2010), which does not support this hypothesis. However, since results from D β H-knockout mice might not be generalizable to human D β H-deficient patients, the assessment of DA receptor densities in human D β H-deficient patients, for example using PET scanning, remains an important objective for future studies.

It is interesting to note that the first study that used gene targeting to produce D β H-deficient mice found that the majority of D β H-deficient embryos died in mid-gestation and only 5% reached adulthood (Thomas et al., 1995). To prevent embryonic lethality, subsequent studies using D β H-knockout mice have supplied the embryos with adrenergic agonists (isoproterenol and phenylephrine) and DOPS via the maternal drinking water, such that NE is present in the D β H-knockout mice until birth. Thomas et al.'s (1995) results suggest that the human D β H-deficient patients may represent the minority of D β H-deficiency cases that have survived this condition. If this is true, an interesting speculation is that these patients were able to survive because they happened to have optimal dopaminergic or noradrenergic genotypes to compensate for the absence of NE. Future studies might assess this possibility by examining whether the frequency of occurrence of specific alleles of dopaminergic and noradrenergic genes (e.g., the COMT, DAT, and the dopamine and noradrenergic receptor genes) in D β H-deficient patients deviates from those in the normal population.

In contrast to the generally normal neurocognitive function in the D β H-deficient patients, we did find clear abnormalities in their task-evoked pupil dilation response. The task-evoked pupil dilation was very small or absent in most of the patients, which might be due to a decreased noradrenergic innervation of the iris dilator muscle. However, it is also possible that the abnormal pupil dynamics in some of the patients resulted from ocular abnormalities unrelated to their DBH deficiency; this might explain why the pupil-dilation response was not restored by DOPS medication. Importantly, the patients' small or absent task-evoked pupil dilations did not reflect a decreased processing of the task-related stimuli, since their performance on the tone-discrimination task during which their pupils were measured was not impaired.

The patient group also differed from the control group with regard to total brain volume: all but one patient had a significantly smaller brain volume than the control group, but the relative proportions of grey matter, white matter and cerebrospinal fluid, and the distribution of grey matter

volume across the brain did not deviate from those in the control group. The smaller brain volume in most DBH-deficient patients is in line with recent findings suggesting that NE has a neurotrophic effect on cortical neurons (e.g., Counts and Mufson, 2010; Kalinin et al., 2007; Madrigal et al., 2007, 2009). Apparently, the patients' decreased brain volume did not result in cognitive impairments; this suggests that although the patients have a smaller number of neurons, their neurons are intact and make proper connections.

To conclude, our findings suggest that neurocognitive function in human D β H-deficient patients is largely spared, even when they are OFF medication, but that their total brain volume is smaller than that of the normal population. The normal neurocognitive function in D β H-deficient patients is striking given the important role of NE in normal cognition, but corroborates informal clinical observations that most patients do not have obvious cognitive impairments. Our findings suggest that D β H-deficient patients have developed alternative mechanisms to compensate for the absence of NE in the brain, possibly through a functional replacement of NE by DA; the nature of these compensatory mechanisms remains to be explored by future studies.

Appendix I: Supplementary Methods

Emotional working-memory task

We used a modified Sternberg item-recognition task (Oei et al., 2009; Oei et al., 2010; Sternberg, 1966). Each trial started with a fixation cross presented for 1 s in the center of the screen. Following the fixation cross, either one or four capital letters (the target set, $1.5^\circ \times 1.2^\circ$ per letter) appeared on the screen for 1 s. The target set had to be held in memory during the following 1.5-s delay period. During this delay period, a picture was presented. Pictures were selected from the International Affective Picture System (IAPS; Lang, Bradley, and Cuthbert, 2005). Half of the pictures were negatively arousing ($M \pm SE$: valence 2.4 ± 0.8 , arousal 6.6 ± 0.4), the other half were emotionally neutral ($M \pm SE$: valence 5.1 ± 0.6 , arousal 3.3 ± 0.7), as rated on a 1-9 point scale (Lang, Bradley, & Cuthbert, 2005). Following the delay period, four capital letters were presented (the probe set, $1.5^\circ \times 1.2^\circ$ per letter). On half of the trials the probe set contained one letter from the target set, and on the other half of the trials the probe set did not contain a letter from the target set. The probe set was followed by an intertrial interval of 2 s.

The participants' task was to indicate whether one of the probe letters had been part of the last target set or not, by pressing the 'z' or the 'm' key. The key assignment was balanced across participants. Participants were instructed to respond as fast and as accurately as possible. The probe set stayed on the screen until the participant made a response. If the participant did not respond within 3 s, the trial ended automatically and a 'TOO SLOW' message appeared on the screen. Prior to the start of the experimental session, participants viewed on-screen instructions and were given 8 practice trials. The experimental session consisted of 15 repetitions of the factorial combination of working-memory load (1 or 4 target letters), distractor type (neutral or negative picture) and target presence (target present or target absent). The task lasted approximately 18 minutes.

Attentional-blink task

Stimuli. Stimuli were presented in black against a light grey background. Each trial started with a fixation cross measuring $0.5 \times 0.5^\circ$, presented for 1s in the center of the screen. Subsequently, the fixation cross was replaced by a rapid serial visual presentation (RSVP) stream of 19 uppercase letters and 2 digits, each measuring approximately $0.9 \times 0.9^\circ$. Each letter was randomly drawn (without replacement) from the alphabet and presented for 74 ms, followed by a 24-ms blank interval. "I," "O," "Q," and "S" were left out as they resemble digits too much. The two digits (T1 and T2) were randomly drawn without replacement from the set 2 to 9. T1 was presented 10 to 13 temporal positions from the beginning of the stream. The temporal distance between T1 and T2 was either one, two, three or seven items, corresponding to lags of 98, 196, 294, and 686 ms.

Procedure. The participant's task was to identify both T1 and T2 by typing the digits in order on a standard keyboard after the end of the RSVP stream. Participants were instructed to guess whenever they failed to identify a digit. The two keyboard entries were followed by the presentation of a feedback stimulus for 150 ms (e.g., '+, -' to indicate that T1 was correct and T2

was incorrect). After a 1-s blank screen, the next trial started. Each participant started with 12 practice trials, three for each lag, randomly intermixed. This was followed by six blocks of 40 trials each with each block containing ten repetitions of each lag.

Visual-search task

Stimuli. Each trial started with a white fixation cross measuring $0.9 \times 0.9^\circ$ against a dark background, presented for 500 ms in the center of the screen. Subsequently, the fixation cross was replaced by a search display, which consisted of 4, 8, or 16 items that were randomly plotted in the cells of an imaginary 6×6 matrix (8.7° horizontally \times 9.6° vertically) with some random jitter within the cells. On half of the trials, the target, a vertical red bar, was present in the array. On the other half of the trials, the target was absent. The distractors were vertical green bars and horizontal red bars. Thus, the target was defined by a specific conjunction of features (color and orientation).

Procedure. On each trial, the participant's task was to report whether or not the target ($0.7 \times 1.3^\circ$) was present by giving a response with their left or right index finger using the 'z' and 'm' keys on the computer keyboard. The keyboard entry was immediately followed by a 1,000-ms blank screen after which the next trial started. Participants performed two blocks of 96 trials each, with each block containing 16 repetitions of the factorial combination of set size (4, 8, or 16) and trial type (target present or absent) presented in random order. Prior to the start of the experimental session, participants viewed on-screen instructions and were given 12 practice trials. The task instructions encouraged participants to respond as quickly as possible while minimizing the number of errors. Performance feedback was provided at the end of each block.

Oddball tasks and EEG measurement

Oddball tasks. Participants performed a visual and an auditory oddball task. In the visual oddball task, a series of black crosses and circles ($1.7 \times 1.7^\circ$) was presented on a light grey background. Each stimulus was presented for 250 ms and the interval between two successive stimuli was 2500 ms. In the auditory oddball task, a series of 1000-Hz and 2000-Hz tones (75 dB) was presented. Each tone lasted 150 ms and the interval between two successive tones was 2100 ms. Participants were instructed to make speeded key-press responses with the dominant hand to target stimuli (circles/2000-Hz tones, 20% of the trials) but not to non-target stimuli (crosses/1000-Hz tones, 80% of the trials). Each task consisted of 30 target trials and 120 non-target trials.

EEG recording. For the Dutch patients and all control participants, EEG activity was recorded from 24 Ag/AgCl scalp electrodes (Fp1, AFz, Fz, F3, F7, FCz, Cz, C3, T7, CPz, Pz, P3, P7, POz, O1, Oz, O2, P8, P4, C4, T8, F8, F4, Fp2). In addition, two electrodes were placed at the left and right mastoid. We measured the horizontal and vertical electro-oculogram (EOG) using bipolar recordings from electrodes placed approximately 1 cm lateral of the outer canthi of the two eyes and from electrodes placed approximately 1 cm above and below the participant's left eye. For the American patients, EEG, EOG and mastoid activity was recorded from a high-density array of

128 Ag/AgCl electrodes embedded in soft sponges (Geodesic Sensor Net, EGI, Inc., Eugene, OR, USA),

Signal processing and data analyses. For the Dutch patients and the control participants, the signal was DC amplified and digitized with a BioSemi ActiveTwo system (BioSemi B.V., Amsterdam, The Netherlands) at a sampling rate of 256 Hz. For the American patients, the signal was DC amplified and digitized with a Net Amps 200 amplifier at a sampling rate of 250 Hz, using Net Station 4.3 software (EGI, Inc., Eugene, OR, USA). Each active electrode was referenced offline to the average of the left and right mastoids. EEG and EOG were high-pass filtered at 0.1 Hz. We extracted single-trial epochs for a period from 200 ms before until 800 ms after stimulus onset. Ocular and eyeblink artifacts were corrected using the method of Gratton, Coles, and Donchin (1983) as implemented in Brain Vision Analyzer. Epochs with other artifacts (spike artifacts [50 μ V/2 ms] and slow drifts [200 μ V/200 ms]) were also discarded. Then, for each participant, task and stimulus type (target/standard), averaged waveforms aligned to a 200-ms prestimulus baseline were generated. The P3 amplitude was defined as the most positive peak in the 200–600-ms time window after the stimulus. We focused our analyses on the electrode position at which the P3 amplitude in response to target stimuli was largest.

Pitch-discrimination task and pupillometry

Stimuli and procedure. Participants performed an auditory pitch-discrimination task (Gilzenrat et al., 2010; Kahneman and Beatty, 1967) while their pupils were continuously measured. They were seated in front of a computer monitor displaying a blank medium gray field, and were instructed to hold gaze within a central fixation square delineated by a thin black border subtending 10° of visual angle. Participants were presented sequences of two sinusoidal tones (72 dB, 250 ms), and were instructed to indicate whether the second of the two tones was higher or lower in pitch than the first. Each trial began with an 850-Hz reference tone. This tone was followed 3 s later by the comparison tone, which ranged from 820 Hz to 880 Hz in steps of 10 Hz. Participants were instructed to respond as quickly and accurately as possible upon hearing the comparison tone. All participants pressed a left key if the second tone was lower and a right key if the second tone was higher than the first tone. Four seconds after the comparison tone, participants received a 250-ms feedback sound that informed them of their accuracy. The feedback was followed by a variable intertrial interval, chosen randomly between 4 and 8 s. Prior to the start of the experimental session, participants viewed on-screen instructions and were given a short block of practice trials at easiest discriminability to familiarize them with the task.

Participants performed two blocks of 36 trials, in counterbalanced order, with each block lasting approximately 10 minutes. In total, participants received 18 trials in which the reference tone and comparison tone were of equal pitch (i.e., impossible-discrimination trials), and they always received negative feedback on these trials. On the other trials, the comparison tone was selected randomly without replacement from the set [820, 830, 840, 860, 870, and 880 Hz], such that participants encountered all of these comparison tones nine times.

The experiment was conducted at a slightly dimmed illumination level. For the Dutch patients and the control participants, the left and right pupil diameters were recorded at 60 Hz using a Tobii T120 eye tracker, which is integrated into a 17-inch TFT monitor (Tobii Technology, Stockholm, Sweden). For the American patients, the left pupil diameter was recorded at 120 Hz using an Applied Systems Laboratory EYE-TRAC 6000 system (ASL, Bedford, MA, USA). These patients used a chinrest and headrest that positioned them 38.5 cm from a Sony Trinitron Multiscan E540 computer monitor.

Pupil analysis. Pupil data were processed and analyzed using Brain Vision Analyzer (Brain products, Gilching, Germany). Artifacts and blinks were removed using a linear-interpolation algorithm. We assessed the baseline pupil diameter prior to trial onset, and the pupil dilation following the comparison tone. To determine baseline pupil diameter, we averaged the pupil data during the two seconds immediately preceding the reference tone. The pupil dilation evoked by the comparison tone was measured as the average deviation from the baseline in the 3 s following onset of the comparison tone.

Positive Affect Negative Affect Scale (PANAS)

At the beginning and the end of each test session, participants completed the PANAS (Watson, Clark, & Tellegen, 1988; translated in Dutch by Peeters, Ponds, & Boon-Vermeeren, 1999), which consists of 10 negative and 10 positive mood terms. For each term, participants indicated to what extent they currently felt that way, using a 5-point response scale with values from 1-very slightly or not at all, to 5-extremely.

MRI

Acquisition. All MRI scans were obtained on a 3-T Philips Achieva MRI scanner, using a three-dimensional T1-weighted gradient echo sequence (TR = 9.8 ms; TE = 4.6 ms; flip angle = 8°, 140 slices). The voxel size was 0.88 x 0.88 x 1.2 mm³.

Analysis. The structural images were brain extracted (Smith, 2002), and the resulting images were segmented into grey matter, white matter and cerebrospinal fluid (CSF; Zhang 2001). We determined each participant's total brain volume, as well as the proportions of grey matter, white matter and CSF, and assessed whether these measures in each patient differed from those in the control group, using the modified independent-samples t-test developed by Crawford and Howell (1998). Because there are sex differences in total brain volume and in the proportion of grey matter (e.g., Cosgrove et al., 2007), we compared the female patients to the female control participants and the male patients to the male control participants.

To assess the presence of regionally specific differences in grey matter density between the patient group and the control group, we performed a voxel-based morphometry-style analysis (Ashburner and Friston, 2000, Good et al., 2001) using FSL tools (FMRIB's Software Library; Smith et al., 2004). The grey-matter partial volume images were aligned to MNI152 standard space using affine registration (Jenkinson and Smith 2001, Jenkinson et al., 2002), and the resulting

images of the five patients and five matched healthy controls were averaged to create a study-specific grey-matter template. The native grey matter images were then non-linearly re-registered to this grey-matter template, modulated, and smoothed with an isotropic Gaussian kernel with a sigma of 2 mm. We used permutation-based non-parametric inference within the framework of the general-linear model, to assess whether there were brain regions with a significantly lower grey matter density in the patient group than in the control group, and vice versa (5000 permutations). We used threshold-free cluster enhancement (TFCE), a new method for finding significant clusters in MRI data without having to define an initial cluster-forming threshold (Smith and Nichols, 2009). Statistical maps were thresholded at $p < 0.05$, corrected for multiple comparisons across space.

Catecholamine measurement

Plasma catecholamine concentrations were measured using high performance liquid pressure (HPLC) analysis. For the Dutch participants, fluorometric detection was used (Willemsen et al., 1995). Within- and between-run coefficients of variation for plasma norepinephrine were 4.1% and 6.1% at a level of 1.76 nmol/l, respectively, and the analytical detection limit for norepinephrine was 0.002 nmol/l. For the American patients, electrochemical detection was used (Holmes et al., 1994). The coefficient of variation for plasma norepinephrine was 4.5% at a level of 1.51 nmol/l, and the analytical detection limit for norepinephrine was 0.024 nmol/l. Catecholamines were collected in ice-chilled 10 ml Vacutainer tubes (Becton-Dickenson Co., Franklin Lakes, NJ) containing 0.2 ml of a solution of EGTA (0.25 mol/L) and glutathione (0.20 mol/l).

Statistical analyses

We compared the results of the patients OFF medication to those of the control participants using a modified independent-samples *t*-test developed specifically to compare an individual patient with a small control group (Crawford and Howell, 1998). This method maintains the Type I error rate (false positives) at the specified (5%) level regardless of the size of the control sample (Crawford and Garthwaite, 2005). The *p* value obtained by this method indicates whether the patient's score is significantly different from the control group, and also provides an unbiased point estimate of the abnormality of the patient's score; that is, it reflects the estimated proportion of the control population that would obtain a more extreme score (Crawford and Garthwaite, 2006a). We used this method to test whether the critical measures/effects in each patient OFF medication deviated from those in the control group, using a statistical threshold of $p < 0.05$ (1-tailed). To control for potential practice effects, we compared the results of the patients that were tested OFF medication on the first study day with the control group's results on the first study day, and the results of the patients that were tested OFF medication on the second study day with the control group's results on the second study day.

We next examined the effects of medication on the patients' scores, using a regression-based method developed by Crawford and Garthwaite (2006b). The control participants' data were used to

generate regression equations that predicted their scores in the second session from their scores in the first session, and vice versa (i.e., the practice effect). These regression equations were then used to predict each patient's ON-medication score from their OFF-medication score, and it was tested whether there was a significant difference between the predicted and observed ON-medication scores. Like Crawford and Howell's (1998) modified *t*-test, this method controls the Type 1 error rate even when the size of the control sample is small. The *p* value obtained by this method provides an estimate of the abnormality of the difference between each patient's predicted and observed ON-medication scores, which reflects the estimated proportion of the control population that would show a larger difference. We used this method to test for each patient's critical measures/effects whether the effect of medication was significantly larger than the practice effect in the control group, using a statistical threshold of $p < 0.05$ (1-tailed). For the patients that were tested ON medication on the first study day, the predicted ON-medication scores were based on the regression equation in which the control participants' scores on the second study day predicted their scores on the first study day. For the patients that were tested ON medication on the second study day, the opposite regression equation was used.

Appendix II: Supplementary Results

Supplementary Table 1. demographic and clinical characteristics of each patient

patient	1	2	3	4	5
Age (years)	25	41	15	22	19
Sex	female	female	male	male	female
Nationality	Dutch	Dutch	American	American	Canadian
Scaled WAIS-III vocabulary score	8	8	13	16	12
Raven's SPM score	43	45	41	53	46
Estimated IQ (based on SPM score; Peck, 1970)	98	112	100	119	107
Abnormalities unrelated to D β H deficiency				diabetic	irregular and large pupils ⁺
Order of ON and OFF medication test days	ON-OFF	ON-OFF	OFF-ON	OFF-ON	OFF-ON
Period on medication before study participation	1.5 years	6 years	-*	-*	4 years
Period on medication before ON-medication test day	1.5 years	6 years	2.5 days	3.5 days	2 days [^]
Period off medication before OFF-medication test day	13 days	7 days	whole life	whole life	4 days
Interval between test sessions (days)	13	7	6	6	6
Missing data (tasks)	EWM, pupil [#]	AB			
Baseline pupil diameter OFF medication (mm)	-	2.35	3.79	4.34	6.29
Baseline pupil diameter ON medication (mm)	-	2.39	3.46	2.83	6.70
Task-evoked pupil dilation OFF medication (mm)	-	0.064	-0.048	0.129	-0.088
Task-evoked pupil dilation ON medication (mm)	-	0.099	-0.053	0.060	-0.017
Brain volume (dm ³)	1.42	1.20	1.55	1.46	1.28
% grey matter	43.9	41.9	48.1	45.9	48.2
% white matter	40.4	40.0	38.1	38.9	36.4
% cerebrospinal fluid	15.7	18.0	13.8	15.2	15.4

Notes: patients 3 and 4 are brothers, the other patients are unrelated; AB = attentional-blink task; EWM = emotional-working memory task; + due to a genetic defect unrelated to D β H deficiency: a deletion on the short arm of chromosome 11; * these two patients had never been on DOPS medication before the study; ^ this patient had consumed 7 doses of 300 mg before she was tested ON medication, and was feeling normal at that time; # emotional-working memory and pupillometry data were not collected for the two matched control participants of patient 1 either.

Supplementary Table 2. Each patient’s plasma and urine catecholamine concentrations for the ON and OFF medication sessions

patient	Plasma NE		Urine NE		Plasma DA		Urine DA	
	ON	OFF	ON	OFF	ON	OFF	ON	OFF
1	0.72	ND	4682	ND	1.08	2.95	1163	669
2	-	-	4212	ND	0.39	2.73	187	21
3	0.46	0.22	14390	11	0.20	0.20	914	1670
4	0.47	0.17	12536	11	0.27	0.25	695	2242
5	0.64	ND	12588	6	0.06	0.25	1005	1757

Notes: all values are in nmol/l; ON = on DOPS medication, OFF = off medication; ND = not detectable; patient 2’s plasma NE concentrations were unmeasurable due to interfering peaks in the chromatogram.

Subjective state

Supplementary Table 3 shows the average positive and negative affect (PANAS) scores in the control group and the patient group at the beginning and end of each test session. The control participants’ positive affect score was higher in the first than in the second session [$F(1, 9) = 9.6, p = 0.01$], and higher at the beginning than at the end of the test sessions [$F(1, 9) = 6.6, p = 0.03$]. The control participants’ negative affect scores were very low, and were not significantly affected by session or point in time ($ps > 0.18$). The patient group reported an overall slightly lower positive affect than the control group. The patient group’s negative affect score was identical to that of the control group, except for a somewhat higher negative affect at the beginning of the OFF-medication session. This suggests that medication status did not have substantial effects on the patients’ affective state, which is surprising given previous findings that social anxiety and mood symptoms were diminished by L-DOPS treatment in two $D\beta H$ -deficient siblings (Critchley et al., 2000).

Supplementary Table 3. Positive and negative affect scores at the beginning and end of each test session in the control group and the patient group (means \pm standard deviations)


		Control group (N = 10)		Patient group (N = 5)	
		Session 1	Session 2	ON	OFF
Positive affect	beginning	3.0 \pm 0.3	2.7 \pm 0.4	2.7 \pm 0.6	2.5 \pm 0.3
	end	2.8 \pm 0.5	2.6 \pm 0.5	2.4 \pm 0.6	2.5 \pm 0.5
Negative affect	beginning	1.2 \pm 0.3	1.1 \pm 0.2	1.2 \pm 0.3	1.4 \pm 0.6
	end	1.2 \pm 0.2	1.1 \pm 0.2	1.2 \pm 0.2	1.1 \pm 0.2

Note: range of both scales = 1-5; ON = on DOPS medication, OFF = off medication

Emotional working-memory performance


A repeated-measures ANOVA on RT in the control group yielded significant main effects of working-memory load [$F(1, 7) = 65.7, p < 0.001$], target presence [$F(1, 7) = 8.9, p = 0.02$] and distractor type [$F(1, 7) = 14.7, p = 0.006$]. There was a main effect of session as well [$F(1, 7) =$

21.5, $p = 0.002$]. In addition, distractor type interacted with target presence [$F(1, 7) = 16.3, p = 0.005$], indicating that the interfering effect of emotional distractors on RT was larger on target-present than on target-absent trials. Finally, there was an interaction between session and working-memory load [$F(1, 7) = 11.9, p = 0.01$], indicating that the effect of working memory load on RT was larger in session 1 than in session 2.



Supplementary Figure 1. Average correct RT and accuracy for the control group and the patient group in the emotional working-memory task, as a function of target presence, working-memory load, distractor type and session (error bars are standard errors of the mean).

As expected, accuracy in the control group was significantly affected by working-memory load [$F(1, 7) = 30.8, p = 0.001$] and target presence [$F(1, 7) = 26.9, p = 0.001$]. In addition, there was an interaction between working-memory load and target presence [$F(1, 7) = 11.6, p = 0.01$]. There were no main effects of session ($p = 0.93$) or distractor type ($p = 0.22$) on accuracy.




Supplementary Figure 2. Each individual participant’s emotional-interference effect on RT as a function of target presence. Because session did not interact with distractor type or target presence in the control group, the data from the control participants are averaged across sessions.

Attentional-blink performance

T1-identification accuracy in the control group showed an increasing trend with lag [$F(3, 27) = 4.1, p = 0.01$], but did not differ between the two sessions [$F(1, 9) = 2.2, p = 0.17$]. Although most patients showed a numerically lower T1 identification accuracy when OFF compared to ON medication, the medication effect on T1 accuracy was not significantly different from the control group’s test-retest effect in any of the patients (all p s > 0.09 , according to Crawford & Garthwaite’s regression-based method).

T2-identification accuracy in the control group was better in session 2 than in session 1 [$F(1, 9) = 7.7, p = 0.02$]. In addition, T2-identification accuracy was affected by lag [$F(3, 27) = 12.1, p = 0.001$], reflecting the characteristic shape of the attentional-blink curve. There was no significant interaction between session and lag [$F(3, 27) = 2.1, p = 0.13$].




Supplementary Figure 3. Each individual participant's T2 accuracy as a function of lag and session.


Visual-search performance

As expected, RT in the control group was significantly affected by set size [$F(2, 18) = 29.7$, $p < 0.001$] and target presence [$F(1, 9) = 6.6$, $p = 0.03$]. In addition, there was a significant interaction of these two variables [$F(2, 18) = 7.8$, $p = 0.004$], indicating that set-size effects were larger for target-absent trials. There was no main effect of session ($p = 0.64$), but session interacted with target presence [$F(1, 9) = 10.6$, $p = 0.01$], indicating that the effect of target presence was larger in session 2. There was also a three-way interaction between session, target presence and set size [$F(2, 18) = 3.9$, $p = 0.04$], indicating that the interaction between target presence and set size was more pronounced in session 2.

Error rates were rather low (average 4.6% in the control group, 2.9% in the patient group ON medication and 2.2% in the patient group OFF medication), and were not affected by set size in the control group ($p = 0.41$), indicating that the increasing RT with set size was not due to a speed-accuracy trade-off.




Supplementary Figure 4. Average correct RT for the control group and the patient group in the visual-search task as a function of target presence, set size and session (error bars are standard errors of the mean).



Supplementary Figure 5. Slopes in the visual-search task for each individual participants, as a function of target presence and session.

Target-detection performance in the oddball tasks



Supplementary Figure 6. Correct RT for the control participants and the patients in the auditory and visual oddball task. The bold lines indicate the average values across participants, and the thinner lines and points indicate the RTs of each individual participant. Because there was no effect of session in the control group, the data for the control participants are averaged across the two sessions.


Target-detection RT in the control group did not differ across sessions [$F(1, 9) = 0.25, p = 0.63$], but was shorter in the visual task than in the auditory task [$F(1, 9) = 10.6, p = 0.01$].

Target-detection accuracy was very high (mean accuracy > 98% in both groups and sessions), and did not differ across sessions [$F(1, 9) = 1.6, p = 0.24$] or task [$F(1, 9) = 2.0, p = 0.19$] in the control group.

Pitch-discrimination performance in the tone-discrimination task

The control group's RT increased as a function of tone-discrimination difficulty [$F(2, 14) = 37.4, p < 0.001$]. There was no significant main effect of session on RT ($p = 0.08$), and no interaction between session and difficulty ($p = 0.9$).

The control group's accuracy decreased with increasing difficulty [$F(2, 14) = 14.1, p < 0.001$]. There was no significant main effect of session on accuracy ($p = 0.06$), and no interaction between session and difficulty ($p = 0.08$).



Supplementary Figure 7. Average correct RT and percentage of correct responses in the tone-discrimination task for the control group and the patient group, as a function of the pitch difference between the two tones (i.e., pitch-discrimination difficulty) and session (error bars indicate standard errors of the mean)

Chapter 5

Neural mechanisms underlying the induction and relief of perceptual curiosity

This chapter is based on: Jepma, M., Verdonschot, R.G., Van Steenbergen, H., Rombouts, S.A., & Nieuwenhuis, S. (under review). Neural mechanisms underlying the induction and relief of perceptual curiosity.

Abstract

Curiosity is one of the most basic biological drives in both animals and humans, and has been identified as a key motive for learning and discovery. Despite the importance of curiosity and related behaviors, the topic has been largely neglected in human neuroscience; hence little is known about the neurobiological mechanisms underlying curiosity. We used functional magnetic resonance imaging (fMRI) to investigate what happens in our brain during the induction and subsequent relief of perceptual curiosity. Our core findings were that (i) the induction of perceptual curiosity, through the presentation of ambiguous visual input, activated the anterior insula and anterior cingulate cortex, brain regions sensitive to conflict and arousal; (ii) the relief of perceptual curiosity, through visual disambiguation, activated regions of the striatum that have been related to reward processing; and (iii) the relief of perceptual curiosity was associated with hippocampal activation and enhanced incidental memory. These findings provide the first demonstration of the neural basis of human perceptual curiosity. Our results provide compelling neurobiological support for a classic psychological theory of curiosity, which holds that curiosity is an aversive condition of increased arousal whose termination is rewarding and facilitates memory (Berlyne, 1954).

Introduction

Curiosity is one of the most basic biological drives in both animals and humans, and has been identified as a key motive for learning and discovery. In the 1950's and 60's, curiosity and related behaviors were topics of intense investigation among experimental psychologists, resulting in an extensive theoretical framework for understanding curiosity and related behaviors (e.g., Berlyne, 1954; Berlyne, 1960; Berlyne, 1966; Loewenstein, 1994). According to a classic psychological theory of curiosity, developed by Berlyne (1954), curiosity evoked by ambiguous, complex, or conflicting stimuli is an aversive condition associated with increased levels of arousal. The theory further holds that termination of this condition, through access to relevant information, is rewarding and promotes learning. Despite the importance of curiosity in many aspects of behavior, the topic has been largely neglected in cognitive neuroscience; hence little is known about the neurobiological mechanisms underlying curiosity. To address this issue, we investigated the neural underpinnings of human curiosity using functional magnetic resonance imaging (fMRI).

Because of its many different facets, curiosity has proven difficult to define. To acknowledge the multifaceted nature of curiosity, a two-dimensional categorization of different types of curiosity has been proposed. The first dimension distinguishes between *perceptual* and *epistemic* curiosity. Perceptual curiosity is aroused by novel, strange or ambiguous stimuli, whereas epistemic curiosity refers to the desire for knowledge or intellectual information which applies mainly to humans (Berlyne, 1954). A second, orthogonal, distinction was made between *specific* and *diversive* curiosity, referring to the desire for a particular piece of information versus the more general stimulation-seeking motive that is closely related to boredom (Berlyne, 1960; see Litman, 2008, for a related distinction].

Berlyne proposed that specific curiosity results from *subjective uncertainty*, which is regarded as a form of conflict due to competing hypotheses regarding the object of uncertainty. The concept of subjective uncertainty is analogous to the information-theoretic notion of *entropy* H , which refers to the objective uncertainty of an outcome (Shannon, 1948). Entropy is defined as

$$H(X) = -\sum_i^N P(x_i) \log_2 P(x_i)$$

, where N is the total number of possible outcomes and $P(x_i)$ is the probability of outcome x_i .

Entropy thus increases with the number of possible outcomes and with the nearness in likelihood of the different possible outcomes. Similarly, Berlyne proposed that someone's subjective uncertainty about a specific stimulus or event (e.g., the identity of an object or the solution to a problem) depends on the number of alternative hypotheses he or she has, and the relative confidence placed in each hypothesis.

In the present study, we focused on specific perceptual curiosity, one of the most basic types of curiosity that is found in animals as well as humans. One way to induce specific perceptual curiosity is to present participants with blurred pictures. An early study using this method showed

that blurred pictures evoked longer EEG desynchronization (alpha-wave blocking) than clear pictures, but only when the identity of the blurred pictures was unknown, which provides preliminary evidence that perceptual curiosity causes an increase in arousal (Berlyne & Borsa, 1968). Another experiment showed that the subjective uncertainty induced by a picture, derived from the number of guesses participants made regarding the picture's identity and the relative confidence they placed in each guess, was highest for pictures with an intermediate degree of blur (Nicki, 1970). This study also showed that participants actively preferred to view the clear version of a preceding blurred picture over viewing an unrelated clear picture, and that the preference for uncertainty reduction was strongest for pictures with an intermediate degree of blur (i.e. pictures associated with maximal subjective uncertainty). Importantly, the preference for uncertainty reduction disappeared when participants knew the identity of the blurred picture. These findings are consistent with the idea that the reduction of perceptual curiosity is rewarding.

We used a modified version of the blurred-pictures paradigm to investigate the neural underpinnings of both the induction and the subsequent relief of human perceptual curiosity. More specifically, we examined whether we could find support at the neural level for the main assumptions of Berlyne's classic curiosity theory. First, the assumption that curiosity is an aversive condition of increased arousal predicts that the induction of curiosity will produce activation in brain areas sensitive to autonomic arousal, conflict and other aversive states. Second, the assumption that the reduction of curiosity is rewarding predicts that this will produce activation in brain regions involved in reward processing, such as the striatum. Third, the assumption that the reduction of curiosity promotes learning and memory predicts that uncertainty-reducing stimuli will be associated with enhanced memory performance and related increased hippocampal activation.

We scanned 19 healthy participants while they viewed sequences of two pictures, in a passive-viewing task. To manipulate the induction and reduction of perceptual uncertainty, we used the following four combinations of clear and blurred pictures (Figure 1):

- (1) A blurred picture followed by its corresponding clear picture (B-C_{corresponding})
- (2) A blurred picture followed by an unrelated clear picture (B-C_{unrelated})
- (3) A clear picture followed by an identical picture (C-C)
- (4) A clear picture followed by its corresponding blurred picture (C-B)

This design resulted in the induction of perceptual uncertainty by the first picture on half of the trials (the B-C_{corresponding} and B-C_{unrelated} conditions), which was resolved by the second picture on half of these trials (the B-C_{corresponding} condition). We used an intermediate degree of blur for all blurred pictures, because this caused maximal subjective uncertainty (Nicki, 1970). Participants' ratings after the scanning session indicated that they had indeed been curious about the blurred pictures (Supplementary Table 1).

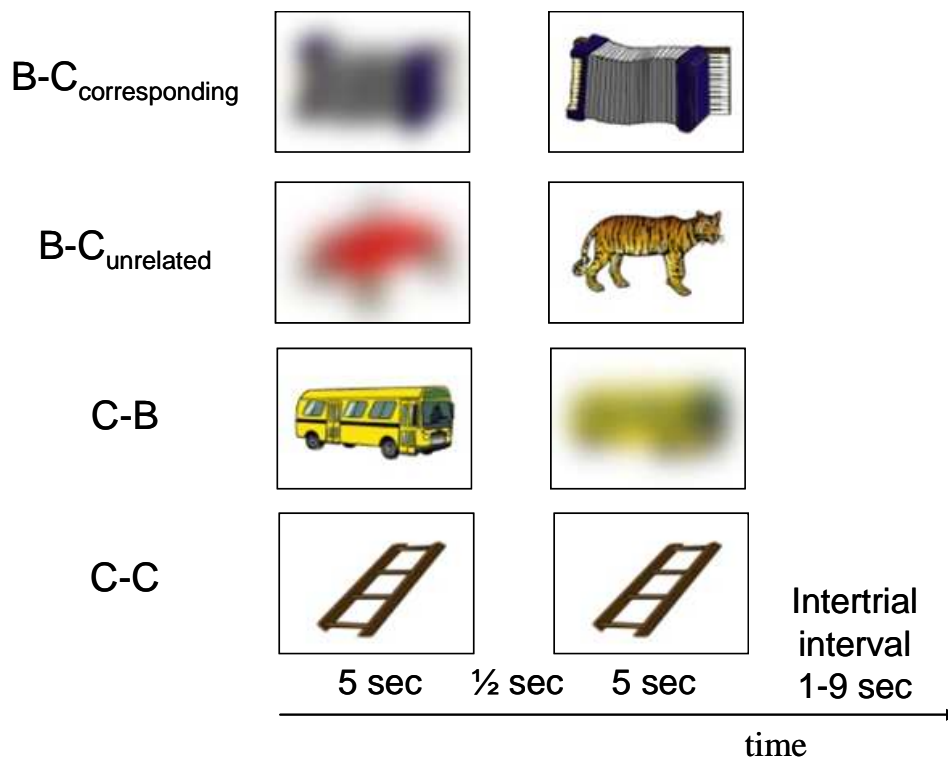


Figure 1. Examples of pictures presented in each of the four conditions. The experiment consisted of 35 trials from each condition, presented in pseudorandom order. Participants were not aware of the aim of the study; we told them that we were interested in the brain activation associated with the perceptual processing of clear and blurred pictures. After scanning, participants were surprised with an unexpected memory test in which they were asked to recall as many objects as possible from the pictures they had seen in the scanner. They also rated several statements regarding their curiosity about the blurred pictures (Supplementary Table 1), and completed the perceptual-curiosity scale [Collins, Litman, & Spielberg, 2004].

Results

Free-recall performance

The number of pictures that participants recalled in an unexpected free-recall test after the scan session was significantly affected by the condition in which the pictures had been presented [$F(3,54) = 11.5, p < 0.001$]. Participants recalled more pictures from the B-C_{corresponding} condition (mean = 10.4) than pictures from the B-C_{unrelated}, C-C and C-B conditions (mean = 6.1, 7.3 and 8.0, respectively; all $ps < 0.03$). Thus, incidental memory for uncertainty-reducing stimuli was enhanced.

Brain activation associated with perceptual uncertainty

To examine the brain activation associated with perceptual uncertainty we focused on the neural response to the first picture in each trial, and identified brain regions where activation was larger when the picture was blurred compared to clear. These regions included the bilateral anterior insular cortex (AIC) and anterior cingulate cortex (ACC; Figure 2, upper panel; see Supplementary Table 2 for all activations). Functional-connectivity studies have suggested that the AIC and AAC

are part of a putative “salience network” (Seeley et al., 2007), which has been associated with autonomic arousal (Critchley, 2005) and various aversive emotional experiences (e.g., Craig, Reiman, Evans, & Bushnell, 1996; Eisenberger, Lieberman, & Williams, 2003; Ploghaus, et al., 1999). The idea that the AIC and AAC activations are part of the same functional network was supported in our data by strong across-subject correlations between the strength of the activations in these areas ($ps < 0.001$).




Figure 2. Brain activation associated with perceptual uncertainty.

Upper panel: The colored regions were more active when the first picture in a trial was blurred (i.e. the first pictures in the B-C_{corresponding} and B-C_{unrelated} conditions) than when it was clear (i.e. the first pictures in the C-C and C-B conditions). R = right; L = left; ACC = anterior cingulate cortex; AIC = anterior insular cortex. See also Supplementary Table 2.

Lower panel: The colored regions were deactivated when the first picture in a trial was blurred compared to when it was clear. All displayed activations are whole-brain uncorrected Z statistic maps ($p < 0.001$), which were overlaid onto the standard MNI brain.

The opposite contrast, which identified brain regions that were more activated by clear pictures than by blurred pictures, revealed activation in a set of brain regions that have been associated with the ‘default-mode network’ (Figure 2, lower panel). The default-mode network, which includes regions of the precuneus, posterior lateral parieto-occipital cortex and medial prefrontal cortex, is typically stronger activated during rest than during cognitive effort (e.g., Raichle, et al., 2001). The relative deactivation of this network in response to blurred compared to clear pictures suggests that participants actively processed the blurred pictures. Consistent with this

interpretation, participants indicated that they had been curious about the blurred pictures, had tried to guess the identities of the objects depicted in them, and had been rather disappointed when a blurred picture was not followed by its corresponding clear version (Supplementary Table 1).

Several findings suggest that the AIC activation reflected a neural substrate of a negative arousal state associated with perceptual curiosity. First, the activated regions of the AIC closely overlap with areas that are typically activated in response to errors, negative feedback and other aversive events (Ullsperger, Harsay, Wessel, & Ridderinkhof, 2010). Second, the strength of participants' AIC activation was positively correlated with their trait curiosity as indexed by the perceptual-curiosity questionnaire ($r = 0.52, p = 0.02$ and $r = 0.46, p = 0.049$ for the right and left AIC, respectively; Figure 3). Third, the participants who reported more disappointment when the identity of a blurred picture was not revealed showed stronger left AIC activation than the participants who reported less disappointment [$t(17) = 2.0, p(1\text{-tailed}) = 0.03$; see Materials and Methods].

Interestingly, the strength of participants' AIC and ACC activation associated with perceptual uncertainty was predictive of the number of pictures they later recalled from the B- $C_{\text{corresponding}}$ condition ($r = 0.46, p = 0.048$ and $r = 0.59, p = 0.008$ for the right AIC and ACC, respectively; the correlation with the left AIC was marginally significant), but not of the number of pictures they recalled from the other conditions (all $ps > 0.3$). This suggests that the uncertainty-related activation of the AIC and ACC contributed to the enhanced memory for stimuli that reduced this uncertainty.




Figure 3. The individual participants' peak activation (z value) for the perceptual-uncertainty contrast in the right and left AIC plotted against their perceptual-curiosity score.

Brain activation associated with the relief of perceptual uncertainty

To examine the brain activation associated with the relief of perceptual uncertainty, we created a contrast that identified brain regions where activation was larger in response to the second picture in the B- $C_{\text{corresponding}}$ condition than in response to the second picture in the B- $C_{\text{unrelated}}$

condition. Significant activation was found in regions of the dorsal striatum (the left caudate and right putamen), and in the left lateral orbitofrontal cortex extending into the ventral striatum (ventral putamen; Figure 4; see Supplementary Table 3 for all activations). Striatal activation has been associated with reward processing, the coding of ‘reward-prediction errors’ (i.e., the difference between observed and expected reward) and reinforcement learning (Daw & Doya, 2006; Haruno & Kawato, 2006; O’Doherty, 2004). Since the uncertainty induced by a blurred picture was relieved by the following picture on only half of the trials, the reduction of perceptual uncertainty by the second picture possibly caused a (partial) reward-prediction error. Accordingly, the striatal activation could reflect the reward value and/or the reward-prediction error associated with the relief of perceptual uncertainty.




Figure 4. Brain activation associated with the reduction of perceptual uncertainty. The colored regions were more active when the second picture in a trial reduced the perceptual uncertainty induced by the preceding picture (i.e., the second picture in the B-C_{corresponding} condition) than when the second picture did not reduce the perceptual uncertainty induced by the preceding picture (i.e., the second picture in the B-C_{unrelated} condition). R = right; L = left. The displayed activations are whole-brain uncorrected Z statistic maps ($p < 0.001$), which were overlaid onto the standard MNI brain. See also Supplementary Table 3.

Confirming predictions, a region-of-interest (ROI) analysis of the hippocampus revealed that regions of the bilateral hippocampus showed stronger activation in response to the second picture in the B-C_{corresponding} than in the B-C_{unrelated} condition (Figure 5). A contrast that identified brain regions where activation was larger in response to the second picture in the B-C_{corresponding} condition than in the C-C condition also revealed significant activation in the left hippocampus (360 mm³ at $p < 0.001$, uncorrected). The event-related time courses of the BOLD signal in response to the second picture in each of the four conditions illustrate the specific increase in hippocampal activation for the B-C_{corresponding} condition (Figure 5). The increased hippocampal activation in response to uncertainty-reducing stimuli likely underlied the enhanced later recall of these stimuli. Interestingly, the strength of participants’ hippocampal activation in response to the reduction of uncertainty was positively correlated with the strength of their right-AIC activation in response to

the induction of uncertainty ($r = 0.57, p = 0.01$ and $r = 0.47, p = 0.04$ for the left and the right hippocampal ROI, respectively). In addition, there was a positive correlation between the strength of the hippocampal activation and the right putamen activation ($r = 0.48, p = 0.04$ for both hippocampal ROIs). This is consistent with the recent hypothesis that interactions between the hippocampus and midbrain dopamine neurons and their striatal targets promote memory for rewarding or otherwise motivationally significant events (Shohamy & Adcock, 2010).

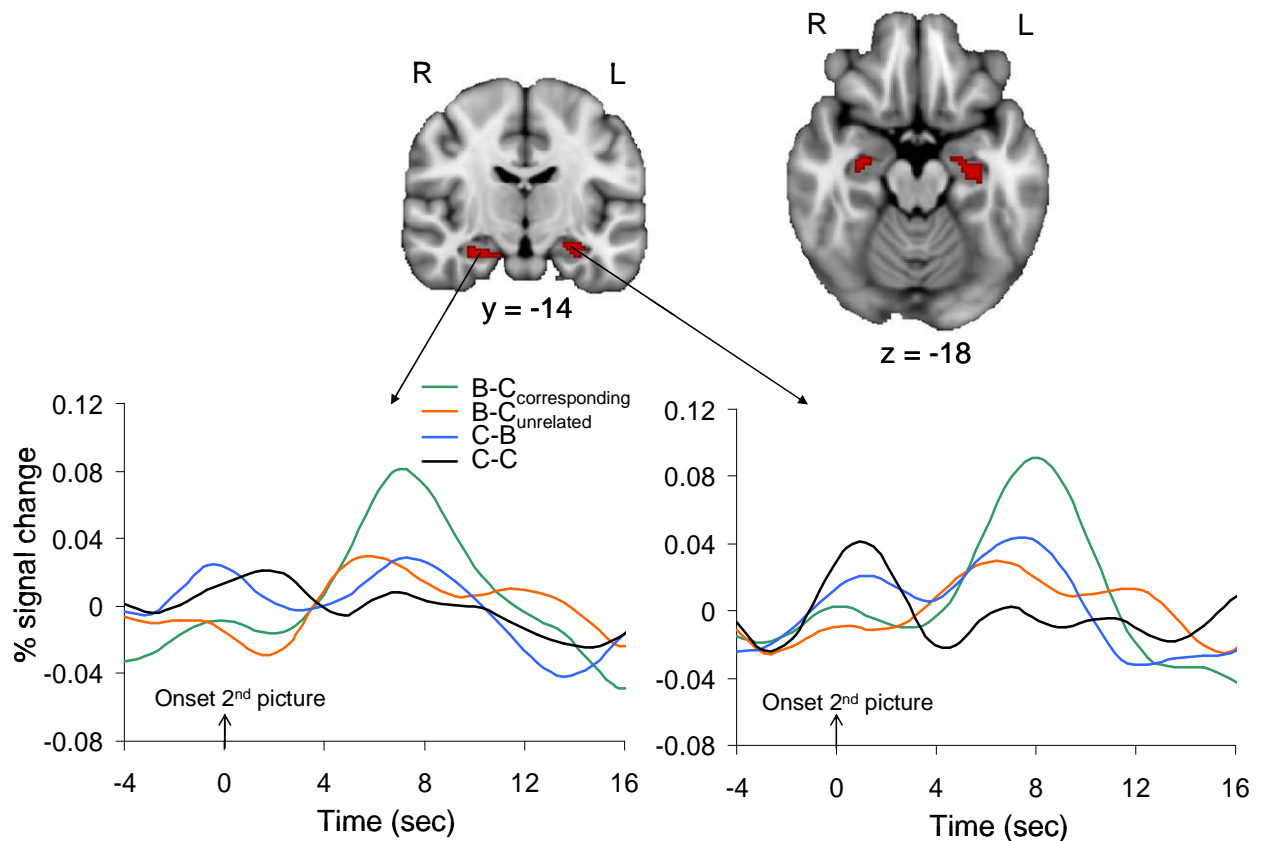


Figure 5. Hippocampal activation associated with the reduction of perceptual uncertainty.

Upper panel: The colored regions were more active when the second picture in a trial reduced the perceptual uncertainty induced by the preceding picture (i.e., the second picture in the B-C_{corresponding} condition) than when the second picture did not reduce the perceptual uncertainty induced by the preceding picture (i.e., the second picture in the B-C_{unrelated} condition). The displayed activations are the uncorrected Z statistic maps in the hippocampal ROIs ($p < 0.01$) overlaid onto the standard MNI brain.

Lower panel: Time course of hemodynamic activity in response to the second picture in each of the four conditions. Time courses were extracted from the hippocampal activation clusters shown in the upper panel.

Discussion

The present study is the first demonstration of the neurobiological basis of human perceptual curiosity. By elucidating the neural underpinnings of the induction and relief of perceptual curiosity, our study extends existing behavioral accounts of curiosity. In particular, our results provide compelling neurobiological evidence for Berlyne's classic psychological theory of curiosity

(Berlyne, 1954, 1960, 1966). First, our finding that perceptual uncertainty activated brain regions sensitive to arousal and conflict supports the assumption that curiosity evoked by ambiguous stimuli is an aversive condition, and induces an increase in arousal. Second, our finding that the reduction of perceptual uncertainty activated brain regions involved in reward processing supports the assumption that the termination of this condition, through access to relevant information, is rewarding. Third, our findings that the reduction of perceptual uncertainty was associated with increased hippocampal activation and enhanced incidental memory support the assumption that uncertainty reduction facilitates memory and learning.

Our findings are also consistent with Loewenstein's information-gap account of curiosity which proposes that curiosity is a negative feeling of deprivation that is caused by an inconsistency, or gap, between one's actual and aspired level of knowledge (Loewenstein, 1994). Since people differ in their aspired level of knowledge, the same actual level of knowledge will evoke curiosity in some people but not in others. In line with this idea, we found that inter-individual variation in trait perceptual curiosity correlated with the strength of AIC activation in response to perceptual uncertainty, suggesting that people with a higher level of aspired perceptual knowledge experience stronger negative feelings when confronted with ambiguous perceptual input.

We found that perceptual curiosity was associated with activation in the AIC and the ACC, regions of a putative salience network that is sensitive to conflict and arousal. These activations may have been modulated by the neuromodulatory locus coeruleus-norepinephrine (LC-NE) system. The LC exhibits strong activity at times of elevated arousal and exhibits a phasic increase in activity in response to motivationally significant stimuli (Aston-Jones & Cohen, 2005). The ensuing release of NE leads to an increased responsiveness of neurons in LC projection areas, including the ACC and AIC (Berridge & Waterhouse, 2003). Thus, the activation of these brain regions in response to curiosity-inducing stimuli was possibly driven by an increased noradrenergic innervation. Consistent with this idea, pharmacological studies in rats have shown that behavioral exploration of novel or unexpected objects is mediated by the LC-NE system (Devauges & Sara, 1990; Mansour, et al., 2003; Sara, Dyon-Laurent, & Hervé, 1995).

The relief of perceptual curiosity was associated with activation in regions of the striatum that are involved in reward processing, suggesting that curiosity reduction is rewarding. This idea is consistent with previous behavioral findings that people actively prefer to view the clear version of a preceding blurred picture over viewing an unrelated clear picture (Nicki, 1970). Other work has shown that people have a similar preference for exploring perceptually novel over familiar stimuli, a tendency that is also associated with striatal activation (Wittmann, Daw, Seymour, & Dolan, 2008). In the reinforcement-learning literature, this bias towards the exploration of uncertain or novel options is captured by the concept of an "exploration bonus" that is assigned to uncertain or novel stimuli to promote their exploration (Kakade & Dayan, 2002).

The relief of perceptual curiosity was also associated with enhanced incidental memory, and with increased hippocampal activation, a plausible neural substrate underlying the behavioral memory effect. Furthermore, there was a positive across-subject correlation between the strength of

the hippocampal and putamen activations in response to uncertainty reduction, suggesting that interactions between these areas contributed to the enhanced memory for curiosity-reducing stimuli. The finding that curiosity reduction leads to enhanced memory suggests that the induction of curiosity before the presentation of teaching material (e.g., by asking people to guess the meaning of foreign words before showing them the translations) can facilitate learning.

Although no prior studies have investigated the neural mechanisms underlying perceptual curiosity, one recent study has investigated the neural substrates of specific epistemic curiosity evoked by the presentation of trivia questions (Kang, et al., 2009). That study found that questions that were rated as more puzzling were associated with stronger activation in regions of the caudate. However, since the questions were always followed by their correct answers, it was unclear whether this activation reflected epistemic curiosity, feedback anticipation, or a combination of the two. In our study, the curiosity induced by blurred pictures was often not relieved, which allowed examination of the neural correlates of pure curiosity. In addition, by comparing conditions in which the second picture did versus did not reduce perceptual uncertainty, we could separately examine the neural correlates of the relief of curiosity.

We did not ask participants to rate their curiosity on each trial since we were concerned that this would confound the brain activation reflecting their natural curiosity. Therefore, a limitation of our study is that we could not take into account trial-to-trial variation in experienced curiosity. In addition, it is likely that curiosity reduction through passive exposure to uncertainty-reducing stimuli, as examined in the present study, differs from curiosity reduction that is achieved through active exploration. A recent study showed that hippocampus activation was stronger when people had volitional control over the visual exploration of pictures in a visual-learning task than when they received exactly the same visual information in a passive condition (Voss, Gonsalves, Federmeier, Tranel, & Cohen, 2011). This suggests that the hippocampus activation associated with uncertainty reduction that we found in the present study would have been even stronger if participants would have had the opportunity to actively control the exploration of uncertainty-inducing stimuli.

To conclude, our results provide evidence at the neural level that perceptual curiosity evokes an aversive state of increased arousal, whose termination is rewarding and promotes incidental memory. Because curiosity plays a key role in many aspects of human behavior, a better understanding of the psychological and neurobiological basis of curiosity may have considerable practical implications for various societal objectives. Together with previous behavioral findings (Berlyne & Normore, 1972), our results suggest that inventing ways to arouse people's curiosity could contribute to the optimization of educational systems and advertising strategies, and may promote scientific discovery.

Materials and methods

Participants

Nineteen healthy volunteers participated (14 women and 5 men; aged 19–29 years; mean age = 22.8 years; SD = 2.4), in return for €25,-. Participants gave written consent before participation, and the study was approved by the medical ethics committee of the Leiden University Medical Center. All participants had normal or corrected-to-normal vision and reported to be right-handed.

Stimuli and procedure

The stimuli used in the task were pictures of common objects selected from Rossion and Pourtois' colored picture databank (Rossion & Pourtois, 2004). This databank is a set of 260 colored line drawings of objects, provided with norms for name agreement, image agreement, familiarity, and complexity ratings. We selected 140 pictures with perfect name agreement from this databank. All pictures had a resolution of 71 dpi, and were centered on a white rectangle of 197 x 281 pixels. We created a blurred version of each picture by means of Gaussian smoothing with a radius of 20 to 22 pixels (Adobe Photoshop 5.0; all pictures can be found at www.sandernieuwenhuis.nl/SOM). By reducing the picture's high-frequency components, Gaussian smoothing acts as a low-pass filter. Results from a behavioral pilot experiment with 49 participants indicated that the objects displayed in the blurred pictures could not be identified by the majority of the participants.

On each trial, a sequence of two pictures was presented. The pictures were projected onto a screen and viewed through a mirror attached to the head coil of the scanner. Each picture was presented for 5 s in the middle of the screen on a white background, and was surrounded by a black frame (18.5 x 13.8°). The two pictures in a trial were separated by a 500-ms interval during which only the frame was presented. The intertrial interval varied between 1 and 9 s (uniform distribution).

The experiment consisted of 35 trials from each of four conditions illustrated in Figure 1, presented in pseudorandom order. For the blurred pictures in the B-C_{unrelated} condition we used blurred versions of 35 additional pictures from Rossion and Pourtois' databank (i.e. pictures of which the clear version was not used). The 140 clear pictures were presented in the same order for all participants. To exclude the possibility that differences between the conditions were caused by picture-specific effects, we divided the 140 clear pictures into four subsets of 35 pictures with comparable familiarity, complexity and imagery ratings (all $ps > 0.86$) and alternated the coupling of the four picture subsets to the four conditions across participants according to a balanced Latin-square design.

The experiment was divided into five runs of 28 trials between which we stopped the scanner and talked with the participant to verify that he or she was still attending to the pictures. Each run contained seven trials from each condition and lasted approximately eight minutes. We

told the participants that the experiment was designed to investigate the brain activation associated with the perceptual processing of clear and blurred pictures, and informed them of the four possible ways in which clear and blurred pictures could be combined.

After completing the experiment, participants were given an unexpected free-recall test outside the scanner; they were asked to type in the names of as many objects as they could recall from the pictures they had seen in the scanner. Subsequently, participants were asked to indicate, on a 5-point scale (1 = not at all; 5 = very much), the degree to which they had (1) been curious about the blurred pictures; (2) tried to guess the identity of the objects depicted on the blurred pictures; (3) been disappointed when a blurred picture was not followed by the corresponding clear version; (4) recognized the objects depicted on the blurred pictures; and (5) tried to remember the pictures. Finally, participants completed the perceptual curiosity scale (Collins, Litman, & Spielberg, 2004).

Image acquisition

Scanning was performed with a standard whole-head coil on a 3-T Philips Achieva MRI system (Best, The Netherlands). In each of the five functional runs, 210 T2*-weighted whole-brain EPIs were acquired (TR = 2.2 sec; TE = 30 ms, flip angle = 80°, 38 axial slices, 2.75 x 2.75 x 2.75 mm + 10% interslice gap). In addition, a high-resolution EPI scan and a T1-weighted anatomical scan were obtained for registration purposes (EPI scan: TR = 2.2 ms; TE = 30 ms, flip angle = 80°, 84 axial slices, 1.96 x 1.96 x 2 mm; 3D T1-weighted scan: TR = 9.7 ms; TE = 4.6 ms, flip angle = 8°, 140 axial slices, 0.88 x 0.88 x 1.2 mm).

Image analysis

MRI data analysis was carried out using FEAT (fMRI Expert Analysis Tool) version 5.98, which is part of FSL (FMRIB's Software Library; Smith et al., 2004). Image pre-processing consisted of motion correction (Jenkinson, Bannister, Brady, & Smith, 2002), non-brain removal (Smith, 2002), spatial smoothing using a Gaussian kernel of FWHM 8.0 mm, grand-mean intensity normalization of the entire 4D data set by a single multiplicative factor, and high-pass temporal filtering to remove low-frequency artifacts (Gaussian-weighted least-squares straight line fitting, with $\sigma = 100$ s). Functional scans were registered to high-resolution EPI images, which were registered to T1 images, which were registered to standard MNI space (Jenkinson, et al., 2002; Jenkinson & Smith, 2001).

The fMRI time series were analyzed using an event-related approach in the context of a general linear model with local autocorrelation correction (Woolrich, Ripley, Brady, & Smith, 2001). We constructed six explanatory variables of interest: two for the first picture in a trial (Blurred or Clear), and four for the second picture in a trial (Clear-corresponding, Clear-unrelated, Clear-double, or Blurred). Each explanatory variable was time-locked to the picture onset and had a duration of 5 s (i.e., the entire duration of the picture presentation). The hemodynamic response to each event was estimated by convolving each explanatory variable with a canonical hemodynamic

response function and its temporal derivative. The model was high-pass-filtered (Gaussian-weighted least-squares straight line fitting, $\sigma = 100$ s).

For each run, in each participant, we assessed several contrasts. The contrasts were combined across the five runs on a subject-by-subject basis using fixed-effects analyses (Beckmann, Jenkinson, & Smith, 2003; Woolrich, Behrens, Beckmann, Jenkinson, & Smith, 2004). These second-level contrast images were submitted to third-level mixed-effects whole-brain group analyses (Beckmann et al., 2003; Woolrich et al., 2004). The statistical parametric images were thresholded at $p < .001$ (uncorrected), in combination with a minimum cluster size of 200 mm^3 .

Region-of-interest analyses. In addition to the whole-brain analyses, we conducted region-of-interests (ROI) analyses to test the predicted hippocampus activation in response to the relief of perceptual uncertainty. We used anatomical ROIs of the left and right hippocampus specified by the Harvard–Oxford subcortical structural atlas, as implemented in FSLView version 3.1.2. Only the voxels that were part of the hippocampus with a probability of at least 50% were included in the ROIs, resulting in left and right hippocampus ROIs of 4016 and 4248 mm^3 , respectively. We tested for activation within these ROIs that exceeded an uncorrected threshold of $p < 0.01$. To further examine the hippocampal activation, we extracted the average time course of the hemodynamic response function in response to the second picture in each of the four conditions using PEATE (perl event-related average time course extraction), a companion tool to FSL (<http://www.jonaskaplan.com/peate/peate-tk.html>). Time courses were extracted from the hippocampal activation clusters of the curiosity-relief contrast (i.e., the regions with stronger activation in response to the second picture in the $B-C_{\text{corresponding}}$ condition than in response to the second picture in the $B-C_{\text{unrelated}}$ condition).

To examine whether individual differences in trait-perceptual curiosity and free-recall performance were predictive of individual differences in brain activation, we extracted each participant's peak z value from the activation clusters of interest (the AIC and ACC activations for the perceptual-uncertainty contrast, and the caudate, putamen and hippocampus clusters for the uncertainty-reduction contrast). We correlated these peak z values with participants' scores on the perceptual-curiosity questionnaire and with their free-recall performance. In addition, we computed the across-subject correlations between the peak z values of the different activation clusters of interest.

Disappointment median-split analysis. To examine whether participants' rated disappointment when the identity of a blurred object was not revealed predicted the strength of their AIC activation in response to perceptual uncertainty, we divided all participants into two groups based on their disappointment ratings: nine participants reported a strong disappointment (ratings of 4 or 5 on a five-point scale) and the other ten participants reported less disappointment (ratings of 2 or 3). We used a t -test to determine whether the high-disappointment group showed stronger AIC activation in response to perceptual uncertainty than the low-disappointment group.

Appendix

Supplementary Table 1. Participants' ratings of the degree to which they had been curious about the blurred pictures, recognized the blurred pictures, and had tried to remember the pictures (means \pm standard deviations). All ratings were on a scale from 1 (not at all) to 5 (very much).

I was curious about the blurred pictures	4.11 \pm 0.88
I tried to guess the identity of the objects depicted in the blurred pictures	4.53 \pm 0.70
I was disappointed when a blurred picture was not followed by its clear version	3.16 \pm 1.02
I recognized the objects depicted in the blurred pictures	2.79 \pm 0.92
I tried to remember the pictures	1.74 \pm 0.73

Supplementary Table 2. Regions showing stronger activation when the first picture in a trial was blurred compared to clear. Data are thresholded at $p < 0.001$, uncorrected, and only clusters exceeding 200 mm³ are reported.

Region	Left/Right	Cluster size (mm ³)	Z _{MAX}	MNI peak coordinates (mm)		
				x	y	z
Anterior insular cortex	R	3192	4.32	36	24	-4
Anterior insular cortex	L	1152	4.11	-28	22	-4
Anterior cingulate cortex	R	1464	4.13	10	24	48
Anterior cingulate cortex	L	488	3.45	-6	12	44
Inferior frontal gyrus	R	3240	4.00	50	16	26
Frontal pole	R	424	3.51	32	48	8
Lingual gyrus	R/L	5520	4.09	8	-80	-8
Occipital pole	L	912	4.21	-12	-94	10
Posterior cingulate gyrus	R/L	616	3.91	2	-30	24

Supplementary Table 3. Regions showing stronger activation in response to the second picture in the B-C_{corresponding} condition compared to the B-C_{unrelated} condition. Data are thresholded at $p < 0.001$, uncorrected, and only clusters exceeding 200 mm³ are reported.

Region	Left/Right	Cluster size (mm ³)	Z _{MAX}	MNI peak coordinates (mm)		
				x	y	z
Caudate (dorsal striatum)	L	224	3.26	-12	6	10
Putamen (dorsal striatum)	R	600	3.55	30	-20	10
Orbitofrontal cortex (extending into ventral putamen)	L	624	3.54	-28	6	-12
Lateral occipital cortex	R/L	1808	3.46	40	-74	0
Posterior insula	R	456	3.57	42	-4	6

Chapter 6

The effects of accessory stimuli on information processing: Evidence from electrophysiology and a diffusion-model analysis

This chapter is published as: Jepma, M., Wagenmakers, E.-J., Band, G.P.H., & Nieuwenhuis, S. (2009). The effects of accessory stimuli on information processing: Evidence from electrophysiology and a diffusion-model analysis. *Journal of Cognitive Neuroscience*, 21, 847-864.

Abstract

People typically respond faster to a stimulus when it is accompanied by a task-irrelevant accessory stimulus presented in another perceptual modality. However, the mechanisms responsible for this accessory-stimulus effect are still poorly understood. We examined the effects of auditory accessory stimulation on the processing of visual stimuli using scalp electrophysiology (Experiment 1) and a diffusion-model analysis (Experiment 2). In accordance with previous studies, lateralized readiness potentials indicated that accessory stimuli do not speed motor execution. Surface Laplacians over the motor cortex, however, revealed a bihemispheric increase in motor activation--an effect predicted by nonspecific arousal models. The diffusion-model analysis suggested that accessory stimuli do not affect parameters of the decision process, but expedite only the nondecision component of information processing. Consequently, we conclude that accessory stimuli facilitate stimulus encoding. The visual P1 and N1 amplitudes on accessory-stimulus trials were modulated in a way that is consistent with multisensory energy integration, a possible mechanism for this facilitation.

Introduction

During most everyday activities, people receive information from multiple sensory modalities. When you ride your bicycle through a city centre, for example, you see the road and the traffic around you, hear cars approaching from behind, and feel the pedals and steering wheel of your bicycle. The signals from the different modalities are not processed independently, but are integrated into coherent representational states. Cognitive psychologists have demonstrated multisensory integration in several psychological phenomena. In ventriloquism, for example, the source of an auditory signal is wrongfully perceived at the location of a visual cue (Howard & Templeton, 1966). Multisensory processing can also lead to a change in the perceived signal itself. This was illustrated in a classic experiment in which a face articulating “gaga” was presented visually, while “baba” was presented aurally. It was found that people usually combined the signals from the two sources and perceived “dada” (McGurk & McDonald, 1976). The present paper addresses another striking instance of crossmodal interaction: the phenomenon that task-irrelevant stimulation (i.e., noise) in one perceptual modality can speed up responses to stimuli concurrently presented in another perceptual modality.

It has repeatedly been found that responses in reaction time (RT) tasks are shorter when a salient but task-irrelevant *accessory stimulus* presented in another perceptual modality accompanies the imperative stimulus, compared to when the imperative stimulus is presented alone (e.g., Bernstein, Clark, & Edelstein, 1969a, 1969b). This speed-up of RTs—often without a concomitant increase in errors—has been referred to as the accessory stimulus (AS) effect. The AS effect has been found in both simple and choice RT tasks (e.g., Bernstein et al., 1969a, 1969b; Morrell, 1968), is largest for auditory stimuli accompanying visual imperative stimuli (Bernstein, 1970; Davis & Green, 1969), and increases in size with the intensity of the AS (Stahl & Rammsayer, 2005). Because the AS is typically presented simultaneously with, or in close temporal proximity to, the imperative stimulus, it has no value for the participant as a cue to start voluntary preparation. Indeed, AS effects have been found even when the AS lags the imperative stimulus (e.g., Bernstein et al., 1969a, 1969b; Stahl & Rammsayer, 2005). In addition, in most experiments the predictive value of the AS is limited by the inclusion of trials on which no AS is presented (no-AS trials), as well as trials on which the AS is not followed by an imperative stimulus (catch trials).

The various explanations of the AS effect that have been proposed so far can be divided into four types of accounts, depending on the components of information processing that are assumed to be affected. One account of the AS effect is that accessory stimuli facilitate stimulus encoding. In particular, it has been proposed that stimulus energy is combined across different modalities in such a way that adding an auditory AS is comparable to increasing the intensity of the visual imperative stimulus (Bernstein, Rose, & Ashe, 1970). According to the energy-integration hypothesis, the increased strength of the joint event speeds up the stimulus-encoding process, resulting in shorter RTs. The critical assumption of the energy-integration hypothesis is supported by the finding that auditory stimuli can increase the perceived intensity of simultaneously presented visual stimuli

(Stein, London, Wilkinson, & Price, 1996). At the neural level, an AS effect on stimulus encoding might be explained in terms of the effects of multisensory neurons—neurons that respond to stimuli from more than one modality. Such neurons exist not only in higher-order association areas, but also in low-level, modality-specific sensory areas (Ghanzafar & Schroeder, 2006), supporting the notion that multisensory interactions can influence early sensory processing.

According to the second and third account, accessory stimuli affect a critical parameter of the decision process that is based on the sensory evidence obtained during stimulus encoding. The mechanism underlying two-choice decisions is well-described by the accumulation of noisy information from a stimulus over time (Gold & Shadlen, 2007; Smith & Ratcliff, 2004). Information accumulates toward one or the other of two decision thresholds until one of the thresholds is reached; then the response associated with that threshold is initiated. One possibility is that accessory stimuli speed up the rate with which evidence is accumulated in the decision process (Hackley & Valle-Inclán, 1999), for example through an AS-triggered, rapid and transient increase in attention to the imperative stimulus. Another possibility is that accessory stimuli do not change the rate of information build-up but instead cause a lowering of the decision threshold (Posner, 1978). According to this view, decisions are made on the basis of less evidence, resulting in shorter RTs and, possibly, more errors. Such a speed-accuracy trade-off has indeed been found in some AS studies (e.g., Posner, 1978).

The fourth account of the AS effect holds that accessory stimuli speed up motor-execution processes. Apparent support for a motoric locus of the AS effect has come from studies that have found an increased response force (Miller, Franz, & Ulrich, 1999; Stahl & Rammsayer, 2005) or a speeding of reflexes (Low, Larson, Burke, & Hackley, 1996; Stafford & Jacobs, 1990) to stimuli accompanied by an acoustic AS. Other evidence that has been presented as support for the motor account is the interaction effect on RT of AS presence and some factors known to affect motor processes, such as tonic muscle tension (Sanders, 1980; Schmidt, Gielen, & Van den Heuvel, 1984). Sanders (1980, 1983) has argued, using additive factors logic (Sternberg, 1969), that such interactions indicate a motoric locus for the AS-triggered speeding of RTs. A discussion of the problems with this argument will be deferred until the General Discussion.

Despite a substantial empirical database, there is no general agreement among researchers on which of these four accounts explains most of the data. One possible source of confusion in the debate may be that the various accounts are not mutually exclusive, and hence different portions of the database may be explained by different accounts. Another reason for the lack of agreement may be that it is hard to distinguish the various accounts on the basis of behavioral performance measures alone.

Probably the most conclusive evidence to date has been reported by Hackley and Valle-Inclán (1998, 1999). These investigators recorded the electroencephalogram (EEG) from participants performing an AS task and computed the lateralized readiness potential (LRP) to investigate the timing of the AS effect. The LRP is an EEG index of hand-specific response preparation. It is computed as the difference in EEG activity over the motor cortices contralateral

and ipsilateral to the responding hand, and averages zero until the accumulated evidence at the level of the motor cortex for one of the response options is outweighing that for the other response option. Thus, the onset of the LRP reflects the point during the decision process during which, on average, stimulus-specific accumulators have gathered evidence favoring one of the two response options, and this evidence has been transmitted, first to brain areas representing the relevant stimulus-response mappings, and then to the motor cortex where it is expressed in asymmetrical activity of the two hemispheres (Gold & Shadlen, 2007; Spencer & Coles, 1999). Hackley and Valle-Inclán found that accessory stimuli shorten the interval between stimulus onset and LRP onset but not the interval between LRP onset and the overt response. This is strong evidence against the notion that accessory stimuli speed up motor-execution processes, and in support of the view that the AS effect develops during stimulus encoding and/or an early phase of the decision-making process (i.e., before the motor cortex begins to reveal the outcome of the decision).

Despite the knowledge gained by these LRP studies, several important questions remain unanswered. For example, is it possible to reconcile the conclusion that accessory stimuli do not speed up motor-execution processes with findings of an AS effect on voluntary response force and the amplitude of somatic reflexes? Can we distinguish the possibilities, suggested by LRP studies, that the AS effect develops during stimulus encoding or during an early phase of the decision-making process? And can we find indications that the AS effect is a result of energy convergence in low-level sensory brain areas? We addressed these and other questions in the two experiments reported below.

Experiment 1: Electrophysiology

The aims of this experiment were threefold. First, we tried to replicate the finding, reported by Hackley and Valle-Inclán (1998, 1999), that auditory accessory stimuli speed up visual information processing before LRP onset but not after LRP onset. As noted, this type of information provides important clues about the processing components influenced by accessory stimuli.

Second, we wanted to investigate whether accessory stimuli have an effect on central motor processes that is not revealed by the LRP methodology used in previous research. Specifically, the LRP is a relative measure, which shows the difference in activity between the contralateral and ipsilateral motor cortices, but not the respective activities of each individual motor cortex. Therefore, the LRP does not reveal potential AS-induced *nonspecific* increases in motor activity—increases in activity that are equal for the contralateral and ipsilateral motor cortices, and that are not expressed in a RT benefit. The possibility that accessory stimuli increase bilateral motor activity without speeding the actual response execution is consistent with proposals that energy-related stimulus properties (e.g., the intensity of the AS) have nonspecific arousal effects that are dissociable from the effects of translating the stimulus into the appropriate response (Sanders, 1983). It is possible to assess the activity of each individual motor cortex by estimating the surface Laplacians over the primary motor areas by means of the source-derivation method (Hjorth, 1975).

The Laplacian acts as a high-pass spatial filter by reducing the common activities between neighboring electrodes. It removes the blurring effect of current diffusion through the highly resistive skull, and is considered to give a good approximation of the corticogram (Gevins, 1989).

The third aim of Experiment 1 was to evaluate a prediction of the energy-integration hypothesis, by examining the effect of accessory stimuli on the P1 and N1, two early visual evoked ERP components recorded over the lateral occipital cortex. Previous research has shown that these components increase in amplitude with increasing stimulus brightness (i.e., energy; Blenner & Yingling, 1993). Therefore, if auditory accessory stimuli increase energy in brain areas specialized in visual processing, this energy increase (i.e., the converged energy from the visual and auditory stimuli) should manifest in increased amplitudes of the P1 and N1 associated with the visual imperative stimuli. A failure to find such amplitude enhancements would provide evidence against the energy-integration hypothesis. It is important to note that the observation of such enhancements, though consistent with the energy-integration hypothesis, would not present definitive evidence for this hypothesis; although the use of surface Laplacians improves estimates of the orientation and location (i.e., biased towards superficial sources) of intracerebral generators, this method does not solve the inverse problem. That is, it cannot exclude the possibility that the amplitude increases reflect the summation at the scalp of electrical activity from two or more different cell populations, rather than the summed activity from one source in visual areas. Nevertheless, the current results will be valuable as a basis for future studies designed to distinguish these possibilities. We also compared the latencies of the P1 and N1 components on AS trials and no-AS trials to determine to what extent the AS effect was already present at the corresponding stages of information processing.

Method

Participants. Thirteen volunteers participated (10 women; 12 right handed; aged 18-30 years; mean age = 21.5). All participants reported normal hearing and normal or corrected-to-normal vision. All participants gave informed consent and received either 15 euros or course credits for participation.

Stimuli and procedure. The task we used was a slightly modified version of the task used by Hackley and Valle-Inclán (1998). On most trials the single letter 'S' or 'T' was presented for 250 ms, in the center of the screen. The letter subtended either 1.0° or 0.8° in visual angle, on 80% and 20% of the trials, respectively. When a 1.0° letter was presented, participants were to indicate whether it was an S or a T by pressing a left or a right key (go trials). The key assignment was balanced across participants. When a 0.8° letter was presented, the response was to be withheld (nogo trials). On a randomly chosen 50% of the trials, an AS (800 Hz, 80 dB, 150 ms long tone) was presented 30 ms prior to the letter onset. The tones were presented binaurally through Epymotic air-pulse ear phones. Intertrial intervals were 2, 3 or 4 s. Unlike in Hackley and Valle-Inclán's task, we also included trials on which the AS was presented alone (catch trials). These catch trials were included to discourage premature responses to the AS, and to be able to compare ERPs to auditory-

only, visual-only, and combined visual-and-auditory stimulation. Keypress responses were made with the left and the right thumb, and participants were instructed to respond as fast as possible. An ERROR message of 1 s was displayed following incorrect go trial responses and responses on nogo trials.

Participants completed one practice block, followed by 15 experimental blocks. Ten of the experimental blocks contained 40 go trials, 10 nogo trials and 6 catch trials each. In order to obtain enough catch trials, the remaining five blocks contained 16 go trials, 4 nogo trials, and 28 catch trials each. These blocks were presented as the 3rd, 6th, 9th, 12th and 15th block of the experiment. After each block the mean RT appeared on the screen, and participants could take a short break if needed. A total of 800 trials was presented throughout the experiment, which lasted about one hour.

Instrumentation and recording. Visual stimuli were presented on a 19" computer monitor, located at a distance of about 60 cm from the participant. Presentation of the visual and auditory stimuli was controlled by a personal computer using E-prime 1.1. EEG was recorded from 64 Ag/AgCl scalp electrodes mounted in an elastic cap, and from the left and right mastoids, using a 64-channel active electrode recording system (sampling rate 512 Hz). Two additional electrodes (CMS-Common Mode Sense and DRL-Driven Right Leg) were used as reference and ground (see <http://www.biosemi.com/faq/cms&drl.htm> for details). The signal was referenced offline to the average mastoid signal. The horizontal and vertical electro-oculogram (EOG) were measured using bipolar recordings from electrodes placed approximately 1 cm lateral of the outer canthi of the two eyes and from electrodes placed approximately 1 cm above and below the participant's left eye. EEG and EOG were high-pass filtered at 0.1 Hz and low-pass filtered at 30 Hz. Electromyographic (EMG) activity of the flexor pollicis brevis was recorded with paired electrodes fixed about 2 cm apart on the skin of the Thenar eminence of each hand, bandpass-filtered (10-256 Hz), and full-wave rectified.

Signal processing and data analyses. Single-trial epochs were extracted offline for a period from 500 ms before until 800 ms after the critical event. Ocular and eyeblink artifacts were corrected using the method of Gratton, Coles, and Donchin (1983). Epochs with other artifacts (spike artifacts [50 μ V/2 ms] and slow drifts [200 μ V/200 ms]) were also discarded. Then, for each participant and each condition of interest, the EEG epochs were averaged with respect to letter onset (imaginary letter onset on catch trials) and EMG onset to create stimulus-locked and EMG-locked averages. A baseline, computed as the average signal activity across the 200 ms prior to the AS, was subtracted for each ERP. The EMG traces were visually inspected and the EMG onsets were hand-scored by an experimenter. We used this method because visual inspection is more accurate than automated algorithms (Hodges & Bui, 1996; Van Boxtel, Geraats, Van den Berg-Lenssen, & Brunia, 1993). To prevent subjective influence on the onset scoring, the experimenter who scored the onsets was unaware of the trial types to which the EMG traces corresponded.

Trials were excluded from the data analyses if the RT was shorter than 100 ms or longer than 1000 ms, or when the response was incorrect. This resulted in the exclusion of 1.4% of the trials. The EMG onset was used to divide the total RT in premotor time (interval between stimulus

onset and EMG onset) and motor time (interval between EMG onset and overt response). For the LRP analysis, we used the same procedure as Hackley and Valle-Inclán (1998): Stimulus- and EMG-locked LRPs were computed from monopolar recordings over C3 and C4, using the standard double subtraction method. LRP latency was assessed at 30%, 50% and 70% of the peak amplitude, using jackknife tests (Miller, Patterson & Ulrich, 1998). For the surface Laplacian estimation, we used the spherical spline interpolation algorithm of Perrin, Pernier, Bertrand, and Echallier (1989), as implemented in Brain Vision Analyzer. This method is based on the entire electrode array and consists of two steps: first the values recorded at each electrode are interpolated, and then the spatial second derivative of this function is computed. We used 4 as the degree of spline and 10° as the maximum Legendre polynomial. The P1 amplitude was defined as the peak amplitude of the average surface Laplacian over electrodes PO7 and PO8 in the 60-140 ms time window. The N1 amplitude was defined as the peak amplitude of the average surface Laplacian over electrodes P7 and P8 in the 100-200 ms time window.

Results

Behavioral results. In agreement with the findings of Hackley and Valle-Inclán (1998), RT on go trials was shorter on AS trials (mean = 501 ms, SD = 77 ms) than on no-AS trials (mean = 519 ms, SD = 80 ms; $t(12) = 5.0, p < 0.001$). (We verified that this AS effect was of a similar magnitude in the blocks with a high probability of catch trials [21 ms] and the blocks with a low probability of catch trials [18 ms], $F < 1$.) Accuracy on go trials did not differ between AS trials and no-AS trials (97.2% vs. 97.5%; $t(12) = 0.6, p = 0.53$). The percentage of nogo errors (false alarms) was higher on AS trials than on no-AS trials (9.6% vs. 6.8%; $t(12) = 2.2, p = 0.047$).

Responses on catch trials were very rare: one of the participants responded to a catch trial twice, whereas the other participants never responded to a catch trial. This indicates that accessory stimuli did not induce fast-guess responses.

Motor and premotor time. The premotor time was shorter on AS trials than on no-AS trials (364 ms vs. 379 ms; $t(12) = 6.9, p < 0.001$). The motor time did not differ between AS trials and no-AS trials (122 ms vs. 124 ms; $t(12) = 1.4, p = 0.18$).

Electrophysiological data. Figure 1 shows the stimulus- and EMG-locked LRPs for the AS trials and no-AS trials. Consistent with Hackley and Valle-Inclán's (1998) results, we found an AS effect on the stimulus-locked but not on the EMG-locked LRP latency. The difference on the stimulus-locked LRP latency was 16 ms for the 30% amplitude point ($t[12] = 0.62, p = 0.27$), 23 ms for the 50% amplitude point ($t[12] = 1.65, p = 0.06$), and also 23 ms for the 70% amplitude point ($t[12] = 2.04, p = 0.03$). It is interesting that these effect sizes roughly correspond to the AS effect on RT. In contrast, the EMG-locked LRPs for the AS trials and no-AS trials almost overlapped, and no significant AS effect was found for any of the three time points (all $ts < 0.2$). Taken together, this pattern of results indicates that, like RT, the LRP onset occurred earlier and was somewhat less variable in latency on AS trials than on no-AS trials. Importantly, accessory stimuli did not speed processes that followed LRP onset.

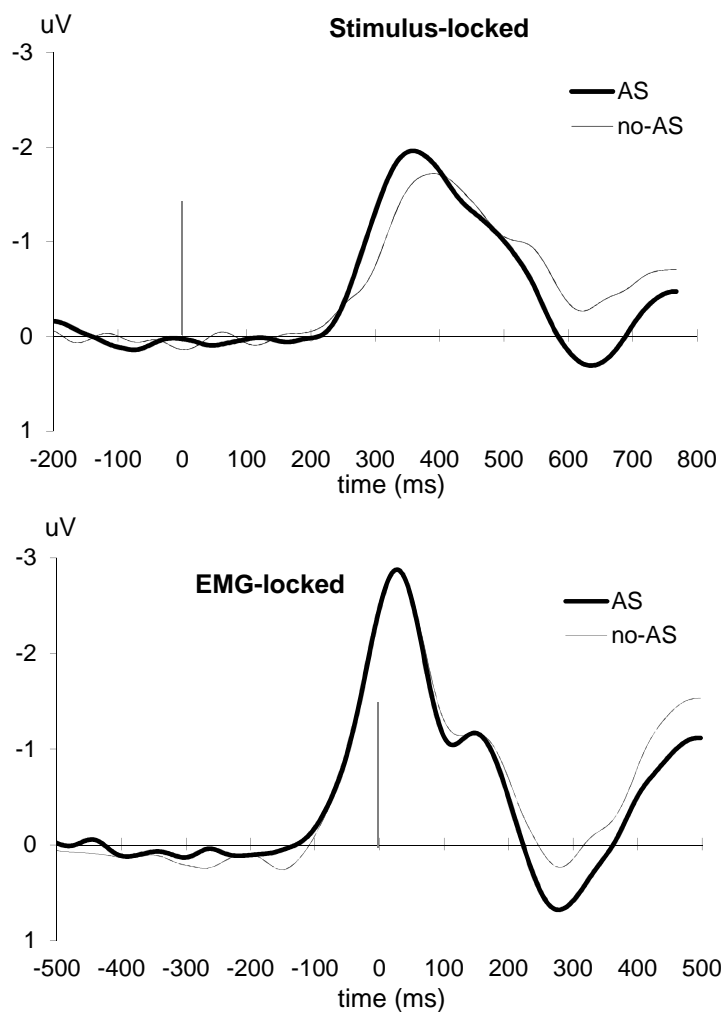


Figure 1. LRPs as a function of AS presence, time-locked to the onset of the visual imperative stimulus (upper panel) and to EMG onset (lower panel). Accessory stimuli were presented 30 ms before the imperative stimulus.

The Laplacian waveforms over the motor cortex contra- and ipsilateral to the involved hand are shown in Figure 2, separately for the AS trials and no-AS trials. On AS trials, two early peaks were observed that were absent on no-AS trials. These peaks reflect tone-related activation in the Sylvian fissure, volume-conducted to the vertex (e.g., Giard et al., 1994). Preceding EMG onset, a negative wave developed over the contralateral motor cortex and a positive wave over the ipsilateral motor cortex. This pattern has also been reported in previous studies, and is thought to reflect the activation of the involved motor cortex and the suppression of activation in the non-involved motor cortex (Burle, Vidal, Tandonnet, & Hasbroucq, 2004 ; Miller, 2007; Tandonnet, Burle, Vidal, & Hasbroucq, 2003; Vidal, Grapperon, Bonnet, & Hasbroucq, 2003). Importantly, both the ipsi- and the contralateral waves were more negative in amplitude on AS trials, suggesting that accessory stimuli induced a nonspecific (i.e., bilateral) increase in activation of the motor cortex. At the time of EMG onset, the AS effect on the Laplacian amplitude was $4.2 \mu\text{V}/\text{cm}^2$ for the contralateral

(involved) motor cortex and $2.6 \mu\text{V} / \text{cm}^2$ for the ipsilateral (non-involved) motor cortex⁴. A repeated-measures ANOVA with laterality (ipsi/contra) and AS presence as within-subject factors yielded main effects of laterality ($F[1,11] = 60.0, p < 0.001$) and AS presence [$F(1,11) = 9.5, p = 0.01$], but no significant interaction ($F[1,11] = 0.6, p = 0.47$). Follow-up contrasts indicated that the contralateral negativity, reflecting the activation of the involved motor cortex, was larger on AS trials than on no-AS trials ($17.5 \mu\text{V} / \text{cm}^2$ vs. $13.3 \mu\text{V} / \text{cm}^2$; $t[11] = 2.2, p = 0.046$). Likewise, the ipsilateral positive wave, reflecting the inhibition of the non-involved motor cortex, was smaller in amplitude on AS-trials than on no-AS trials ($14.5 \mu\text{V} / \text{cm}^2$ vs. $17.1 \mu\text{V} / \text{cm}^2$; $t[11] = 2.4, p = 0.03$). These results confirm the notion that accessory stimuli caused a nonspecific increase in motor cortex activation.

To test the prediction suggested by the energy-integration hypothesis, we tested whether accessory stimuli increased the amplitudes of early visual ERP components. More specifically, we assessed the AS effect on the stimulus-locked Laplacian components corresponding to the P1 (electrodes PO7/8) and the N1 (P7/8; see Figure 3). Consistent with the energy-integration hypothesis, the P1 amplitude was larger on AS trials than on no-AS trials ($t[12] = 4.4, p < 0.001$). The N1 amplitude was also larger on AS trials, but this effect just missed significance ($t[12] = 1.6, p = 0.065$). Interestingly, as illustrated in Figure 3 (upper panel), the P1/N1 amplitude differences between AS trials and no-AS trials were similar to the amplitudes of the P1 and N1 components elicited by the accessory stimuli on catch (i.e., auditory-only) trials. To further illustrate this, Figure 3 (lower panel) shows the waveforms on AS trials (combined visual-and-auditory), as well as the sum waveform created by adding the waveforms associated with catch trials (auditory-only) and no-AS trials (visual-only). Although they do not entirely overlap, the similarity of these waveforms is remarkable, and consistent with the energy-integration hypothesis⁵.

To assess whether the AS latency effect observed for the RTs and LRPs is already present at the time of the P1 and N1 components, we determined the AS effect on the peak latencies of these components. There was no AS effect on the P1 latency ($t[12] = 0.1, p = 0.46$). The N1 peaked 6 ms earlier on AS trials than on no-AS trials, a small but consistent difference ($t[12] = 1.9, p = 0.04$).

⁴ The analyses reported here controlled for the difference in pre-EMG baseline between AS trials and no-AS trials. This baseline difference reflects the tone-elicited negative component (see Figure 2, upper panel), smeared out in the EMG-locked averages. Thus we subtracted the baseline, defined as the amplitude of the peak immediately preceding EMG onset, from the Laplacian amplitudes at the time of EMG onset. One participant was excluded from these analyses because he did not show a clear baseline peak.

⁵ Most ERP studies on multisensory processing focus on superadditive enhancements (i.e., situations in which the multisensory response exceeds the sum of the unisensory responses) to demonstrate multisensory interactions. Cell-recording studies, however, have revealed that superadditivity is merely one facet of multisensory integration, and one that is produced under very specific circumstances, namely when the unisensory component-stimuli are weakly effective. Across the broader range of stimulus intensities, the majority of the multisensory interactions approximate linear summation, i.e., additive enhancements (reviewed in Stanford, & Stein, 2007).

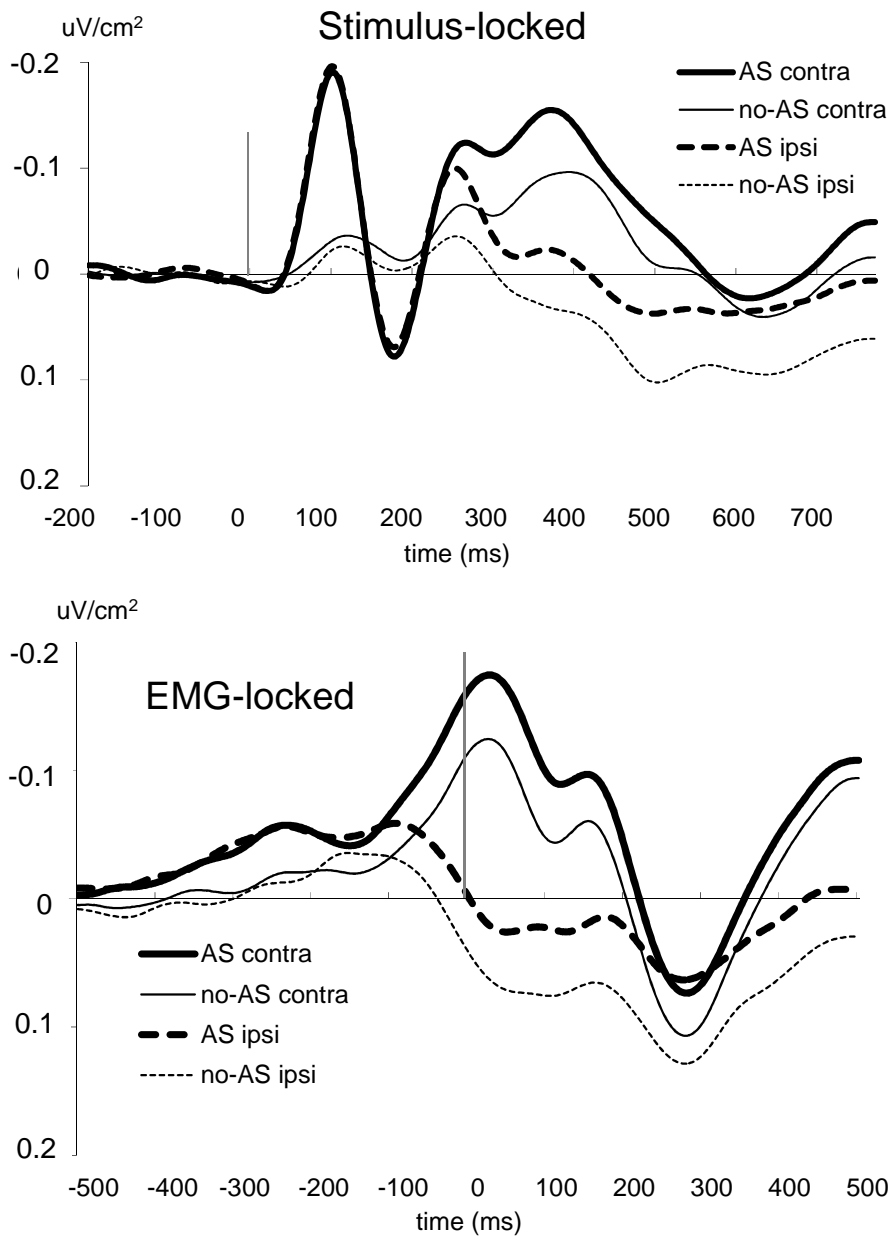


Figure 2. Surface Laplacians over the motor cortex as a function of AS presence, time-locked to the onset of the visual imperative stimulus (upper panel) and to EMG onset (lower panel). Accessory stimuli were presented 30 ms before the imperative stimulus.




Figure 3. Upper panel: Surface Laplacians over electrodes PO7/8 and P7/8 for AS trials, no-AS trials and catch trials, time-locked to stimulus onset. Lower panel: The sum wave created by adding the no-AS signal to the catch signal is similar to the waveform for AS trials.

Discussion

The principal findings of Experiment 1 may be summarized as follows. In accordance with previous studies (Hackley & Valle-Inclán, 1998, 1999), we found that the AS effect was entirely confined to the time period prior to LRP onset (~100 ms prior to EMG onset). Consistent with this finding, the AS effect was reflected in premotor times but not in motor times. A small portion (about one third) of the effect was already apparent 160 ms after stimulus onset, at the time of the N1 peak. Accordingly, most of the effect must have developed between the N1, a component associated with stimulus encoding, and LRP onset, the moment at which the motor cortex begins to reveal the outcome of the decision-making process. These findings confirm that accessory stimuli do not expedite response execution; they indicate that the AS effect reflects a speed-up of stimulus encoding or an early phase of the decision-making process (presumably in association cortices; Gold & Shadlen, 2007). Given that auditory signals can modulate cortical visual processing as early as 40 ms following their onset (Giard & Peronnet, 1999), this temporal “locus” of the AS effect seems consistent with the observation, under some circumstances, of a residual AS effect when the

auditory AS lags the imperative stimulus (up to 100 ms; e.g., Bernstein et al., 1969a, 1969b; Stahl & Rammsayer, 2005).

Interestingly, AS trials were associated with increased amplitudes of the P1 and N1 components, in a way that is consistent with the energy-integration hypothesis. Specifically, the P1/N1 amplitudes on AS trials (combined visual-and-auditory) were of a similar magnitude as the summed amplitudes observed on no-AS trials (visual-only) and catch trials (auditory-only). Thus, it is possible that the speed-up of RTs on AS trials reflects the effects of energy integration in visual processing areas, a possibility that is consistent with anatomical and physiological findings (Ghazanfar & Schroeder, 2006). However, the data do not rule out an alternative interpretation, namely that the increased P1/N1 amplitudes reflect the summation at the scalp of signals originating from visual and auditory processing areas. Other methods are necessary to distinguish between these possibilities.

Previous work has found that accessory stimuli increase response force and reflex magnitude, and that, in general, these response-amplitude measures correlate poorly with RT (Low et al., 1996; Miller et al., 1999; Stahl & Rammsayer, 2005). These findings have been viewed as support for the proposal by Sanders (1983) that accessory stimuli trigger a phasic burst of arousal that leads to nonspecific priming of low-level motor pathways, and that this effect occurs independently from the stimulus-response translation processes contributing to RT. Sanders' proposal dovetails nicely with another principal result of the current study—the finding that accessory stimuli evoked a nonspecific (i.e., bilateral) increase in motor-cortex activity that, as noted above, was not expressed in a RT benefit. This finding seems to provide direct evidence for an AS-induced nonspecific increase in motor activation, and, furthermore, suggests a possible explanation of why this nonspecific effect is expressed in higher response force (as determined in previous studies; a similar explanation may apply to reflex magnitude) but not in shorter RTs. According to this explanation, response force is determined by the activation of the relevant (i.e., contralateral) motor cortex, which is higher on AS trials. This assumption is consistent with neuroimaging studies and neurophysiological recordings (Cramer et al., 2002; Maier, Bennett, Hepp-Reymond, & Lemon, 1993). In contrast, choice RT is dependent on (or at least scales with) the *difference* between the activity in the relevant and irrelevant motor cortex, which is not affected, due to the nonspecificity of the AS effect. This assumption is consistent with previous results indicating that the LRP amplitude at the time of EMG onset is constant across spontaneous variations in RT (Gratton, Coles, Sirevaag, Eriksen, & Donchin, 1988; Mordkoff & Grosjean, 2001), and with the present finding that the EMG onsets on AS trials and no-AS trials were associated with the same LRP amplitude. In any case, the assumption is in accordance with an influential class of decision-making models (e.g., Laming, 1968), which assumes that a response is initiated when the difference between the evidence for each of the two possible responses reaches a certain criterion value. One of these models is the diffusion model (Ratcliff, 1978), which will be used in the next study.

Experiment 2: Diffusion-Model Analysis

In this experiment we aimed to further clarify which components of information processing are affected by accessory stimuli on the basis of a diffusion-model analysis of AS effects on RT and accuracy. The diffusion model is a model of two-choice decision making that defines the decision process as the continuous accumulation of noisy stimulus information over time, from a starting point towards one of two decision criteria or thresholds (Ratcliff, 1978; see Figure 4). When one of the two thresholds is reached, the corresponding response is initiated. There are several reasons to assume that the diffusion model gives an accurate reflection of how the decision process is implemented in the brain. First, the diffusion process is the optimal decision process: it provides the fastest responses for a fixed level of accuracy, or the highest accuracy for a fixed response time (Wald, 1947). Second, the diffusion model explains the dynamics of neuronal activity during decision-making behavior (Gold & Shadlen, 2007; Smith & Ratcliff, 2004). And third, the diffusion model successfully accounts for RT distributions and error rates in a variety of two-alternative forced-choice tasks (e.g., Ratcliff, 2002; Ratcliff, Van Zandt, & McKoon, 1999).

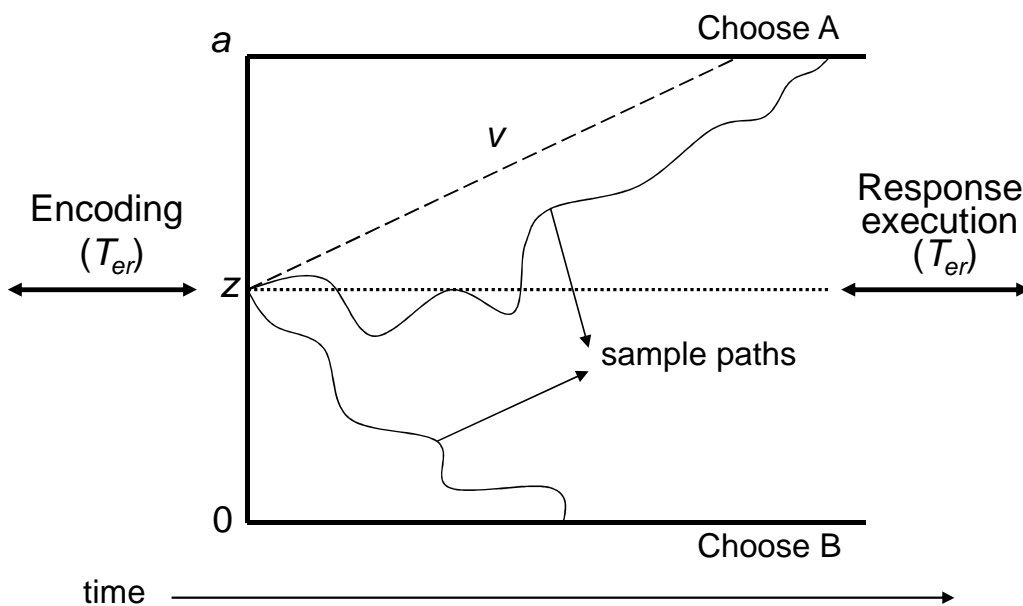


Figure 4. An illustration of the diffusion model. The parameters are: a = boundary separation, z = starting point, v = drift rate, T_{er} = mean nondesision time. The sample paths represent moment-by-moment fluctuations in the evidence favoring the two possible responses, which is due to noise in the decision process. The decision process starts at z and terminates when one of the two boundaries is reached. The duration of T_{er} determines the additional time needed for stimulus encoding and response execution.

The diffusion model can be helpful in evaluating the various accounts of the AS effect because some of the main model parameters correspond closely to the different processing components emphasized by these accounts. The three most important parameters of the model in this respect are the drift rate, the boundary separation, and the nondesision component. The drift rate (v) is the mean rate of evidence accumulation in the decision process, which depends on the

quality of the stimulus and the perceptual system. The higher the absolute value of the drift rate, the faster a decision threshold is reached. If accessory stimuli increase the drift rate of the diffusion model, this would support the idea that accessory stimuli induce a faster build-up of information. The boundary separation (a) is the distance between the two decision criteria. This parameter determines on how much evidence a decision is based, and can be controlled strategically by the decision maker. If accessory stimuli lower the boundary separation, this would provide support for the notion that the AS effect reflects a lowering of the decision threshold (Posner, 1978). As noted above, a speed-accuracy trade-off in the empirical data also provides an important diagnostic criterion for a change in decision threshold. Besides the decision process, there are other components of processing involved in a two-choice RT task, namely stimulus encoding and response execution which, respectively, precede and follow the decision process. In the diffusion model, these nondecision processes are combined into one nondecision component, T_{er} . A shortening of the nondecision component by accessory stimuli would indicate that stimulus encoding and/or motor execution are speeded.

We applied the diffusion model to data from a standard lexical-decision task, in which participants were asked to classify letter strings as a word or a nonword, with task instructions emphasizing reaction speed in half of the blocks and response accuracy in the other half of the blocks. The diffusion model has been shown to provide a good fit of lexical-decision data, accounting for the effects of the experimental variables on RTs for correct and error responses, shapes of the RT distributions, and accuracy values (Ratcliff, Gomez, & McKoon, 2004; Wagenmakers, Ratcliff, Gomez, & McKoon, 2008). Importantly, on half of the trials, the letter string was preceded by an auditory AS, and our major aim was to examine which model parameter(s) could best account for the corresponding differences in task performance. In particular, this approach allowed us to test between the two possible interpretations of the AS effect suggested by Experiment 1: speeding of stimulus encoding or speeding of evidence accumulation.

Method

Participants. Twenty-one students participated (18 women; 19 right-handed; aged 18-31 years; mean age = 22; all native Dutch speakers). All participants reported normal hearing and normal or corrected-to-normal vision. Each participant completed two sessions of approximately 75 minutes each, on separate days. Participants received either 15 euros or course credits for participation.

Stimuli. The stimuli were 800 Dutch words and 800 nonwords. Both the words and the nonwords consisted of 4, 5 or 6 letters (195 4-letter, 251 5-letter and 354 6-letter words as well as nonwords). The frequency of the words ranged from 0.07 to 5.48 per million (mean = 3.47, SD = 1.28; Baayen, Piepenbrock, & Gulikers, 1995). The nonwords were generated by replacing one letter of an existing word; vowels were replaced by vowels and consonants by consonants. The words that were used to generate the nonwords were not used as word stimuli.

A 200-ms long, 80 dB, 1000-Hz sine-wave tone was used as the AS. The tones were presented binaurally through headphones.

Procedure. Participants were tested individually in a dimly lit room. Stimuli were presented on a personal computer screen, with responses collected from the keyboard. On-screen instructions were provided. On most trials a letter string was presented (Courier New font; visual angle = 2.7° for 4-letter words and 4.0° for 6-letter words), and participants were instructed to decide whether or not each letter string was a Dutch word by pressing the z or the / key. The key assignment was balanced across participants. The letter string remained on the screen until a response was made, and was followed by an intertrial interval of 2, 3, or 4 s. On a randomly chosen 50% of the trials the AS was presented 100 ms prior to the onset of the letter string. Participants were informed that the tones were irrelevant to the task and could be ignored. On 11% of the trials the AS was presented alone (catch trials), to discourage premature responses to the AS.

In each of the two sessions, participants completed two practice blocks of 27 trials, followed by 20 experimental blocks of 45 trials. Each experimental block consisted of 20 trials on which a letter string was presented alone, 20 trials on which a letter string was presented together with the AS, and 5 catch trials.

Speed-accuracy instructions alternated across blocks. In speed blocks, participants were instructed to respond as quickly as possible, but without making a lot of errors, and responses slower than 750 ms were followed by a message TOO SLOW of 1 s. When a response was faster than 250 ms, the message TOO FAST was displayed for 1 s. No accuracy feedback was given in these blocks. In accuracy blocks, participants were instructed to respond as accurately as possible, but without taking more time to respond than necessary, and incorrect responses were followed by a message ERROR of 1 s. No speed feedback was given in these blocks. Each block started with an on-screen announcement of the upcoming speed-accuracy instruction, which was displayed for 2 s. At the end of each block the mean RT and the proportion of correct responses appeared on the screen, and participants could take a short break before initiating the next block.

Results

Behavioral results. Figure 5 shows the mean correct RT and mean proportions correct as a function of word type, instruction and AS presence. RTs smaller than 300 ms or larger than 2500 ms were excluded from analysis, which resulted in the exclusion of 0.5% of the trials. In accordance with previous studies, RTs were shorter on AS trials than on no-AS trials (636 ms vs. 660 ms; $F(1,20) = 75.7, p < 0.001$), yielding a reliable AS effect. Furthermore, RTs were shorter following speed instructions than following accuracy instructions (599 ms vs. 697 ms; $F(1,20) = 42.6, p < 0.001$), and shorter for words than for nonwords (627 ms vs. 669 ms; $F(1,20) = 84.0, p < 0.001$). AS presence did not interact with instruction ($p = .37$) or word type ($p = .83$). However, the latter two variables showed a significant interaction, indicating that the RT difference between the speed and accuracy instructions was larger for nonwords than for words ($F[1,20] = 6.9, p = 0.016$).

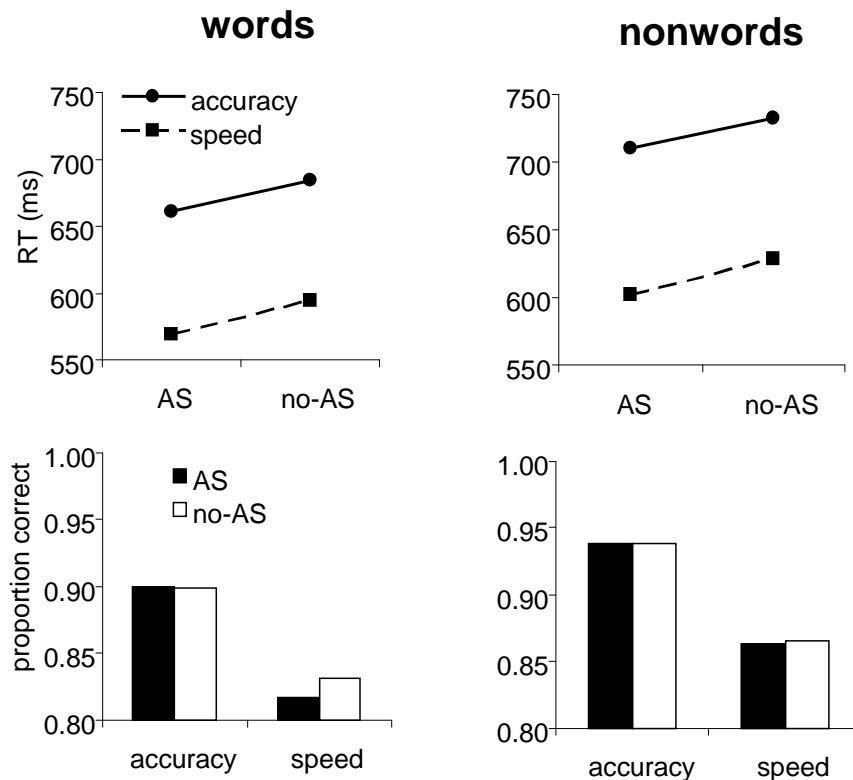


Figure 5. Mean correct RT and proportion correct as a function of word type, instruction (speed/accuracy) and AS presence.

Proportion correct showed no reliable difference between AS trials and no-AS trials (both 0.88; $F(1,20) = 1.4, p = 0.26$). As expected, proportion correct was higher when the instruction emphasized accuracy than when it emphasized speed (0.92 vs. 0.85; $F(1,20) = 44.5, p < 0.001$). In addition, proportion correct was higher for nonwords than for words (0.90 vs. 0.86; $F(1,20) = 9.8, p = 0.005$). None of the interactions between the three variables were significant (all $ps > 0.09$).

Finally, responses on catch trials were practically absent: One of the participants responded to a catch trial once, whereas the other participants never responded to a catch trial.

Diffusion-model analysis. For fitting the diffusion model to the data we used the Diffusion Model Analysis Toolbox (DMAT; Vandekerckhove & Tuerlinckx, 2007, 2008). DMAT estimates parameters by maximizing a multinomial likelihood function. The data that are used to fit the diffusion model are the RT distributions for correct and incorrect responses, and the proportion correct responses. To assess the processing components that are affected by accessory stimuli, four different models were fitted to the data. The four models differed with regard to the parameters that were free to vary as a function of AS presence. In one model (the All free model), T_{er} , a , and v were all left free to vary. In addition, there were three models in which either T_{er} , a , or v could vary, whereas the other parameters were held constant (the T_{er} model, a model, and v model, respectively).

The following parameter settings were the same for all models: 1) The intertrial variability in nondecision time (st) was held constant across all conditions. 2) The starting point of the diffusion process (z) was set at a fixed proportion of the boundary separation, such that the bias in starting point was constant across conditions. 3) Boundary separation (a) and the intertrial variability in starting point (sz) were free to vary between the speed and accuracy conditions (Ratcliff & Rouder, 1998, Experiment 1; Ratcliff, Thapar, & McKoon, 2001, Experiment 2). 4) Mean drift rate (v) and intertrial variability in drift rate (η) were free to vary between the word and nonword trials (Ratcliff, Thapar, Gomez, & McKoon, 2004).

The models were fitted to the data in two ways. First, the models were fitted to each participant's data individually. When a participant made 10 or fewer errors in a condition, the participant's error data for this condition were not included in the fitting procedure. Second, the models were fitted to the averaged data. The averaged data was obtained by calculating the accuracy and the RTs for correct and error trials associated with the .1, .3, .5, .7 and .9 quantiles for each individual participant, and then averaging these values across participants. (Note that the quantile RTs are not the mean RTs within bins [Ratcliff, 1979], but the boundary RTs of each quantile) The codes that were used to fit the models can be found at <http://users.fmg.uva.nl/ewagenmakers/papers.html>.

AS effects on the diffusion-model parameters. To assess which parameters were affected by AS presence, we analyzed the AS effect on the estimates of the T_{er} , a and v parameters in the All free model. Table 1 shows both the average parameter estimates across participants and the parameter estimates resulting from fits of the models to the averaged data. The parameter estimates obtained by the two fitting methods were very similar, which replicates findings from previous studies (e.g., Ratcliff et al., 2001, 2004). The average parameter estimates across participants and the parameter estimates resulting from fits to the averaged data were within one SD of each other for all parameters. As expected, the boundary separation was smaller when the instruction emphasized speed than when it emphasized accuracy ($F(1,20) = 48.1, p < 0.001$). In addition, drift rates were higher for words than for nonwords ($F(1,20) = 16.6, p = 0.001$). Importantly, neither boundary separation nor drift rate was affected by AS presence (both $F(1,20) < 1$). In contrast, the nondecision component, T_{er} , was significantly smaller on AS trials than on no-AS trials ($t(20) = 5.7, p < 0.001$). These results suggest that accessory stimuli shorten one or more nondecision processes, but do not affect the decision process itself.

Table 1. Parameter estimates for the fit of the All free model (SD in parentheses)

	parameter	AS	No AS
average values across participants	T_{er}	.471 (.027)	.488 (.027)
	a (speed)	.097 (.018)	.099 (.018)
	a (accuracy)	.146 (.037)	.148 (.041)
	v (words)	.404 (.169)	.391 (.128)
	v (nonwords)	-.331 (.101)	-.313 (.064)
fits to averaged data	T_{er}	.475	.494
	a (speed)	.089	.091
	a (accuracy)	.130	.133
	v (words)	.318	.327
	v (nonwords)	-.286	-.287

Model selection. To further assess the AS effect on the different model parameters, we tested which model had the best fit to the data. To compare the adequacy of the four models (i.e., the All free model, T_{er} model, a model, and v model) in explaining the observed data we used the Bayesian Information Criterion (BIC; Raftery, 1996), a statistical criterion for model selection. The BIC is an increasing function of the residual sum of squares from the estimated model, and an increasing function of the number of free parameters to be estimated. Thus, the best model is the model with the lowest BIC value. In addition, the raw BIC values were transformed to a probability scale, enabling a more intuitive comparison of the probabilities of each model being the best model (Wagenmakers & Farrell, 2004). The transformation of BIC values to probability values consists of three steps. First, for each model i , the difference in BIC with respect to the model with the lowest BIC value is computed (i.e., $\Delta_i(\text{BIC})$). Second, the relative likelihood L of each model i is estimated by means of the following transformation: $L(M_i | \text{data}) \propto \exp[-0.5 \Delta_i(\text{BIC})]$, where \propto stands for “is proportional to”. Last, the model probabilities are computed by normalizing the relative model likelihoods, which is done by dividing each model likelihood by the sum of the likelihoods of all models. Table 2 summarizes the BIC values and probabilities of each of the four models. Again, both the average values across participants and the values resulting from fits of the model to the averaged data are displayed. The T_{er} model had by far the best fit, both for the individually fitted data and for the averaged data. In the individual analyses, the T_{er} model yielded the best fit for 18 of the 21 participants. For the sake of completeness we also examined the models in which combinations of two parameters (T_{er} and a ; T_{er} and v ; a and v) were free to vary as a function of AS presence. The BIC values of these three models were all worse than that of the T_{er} model.

Table 2. BIC values for each model. Note: p = BIC model probability

		<i>df</i>	BIC	<i>p</i>
average values across participants	All free model	20	6,725	< 0.0001
	T_{er} model	12	6,680	> 0.9998
	<i>a</i> model	15	6,703	< 0.0001
	<i>v</i> model	15	6,706	< 0.0001
fits to averaged data	All free model	20	139,653	< 0.0001
	T_{er} model	12	139,583	> 0.9998
	<i>a</i> model	15	139,714	< 0.0001
	<i>v</i> model	15	139,878	< 0.0001

Model fits. To examine the RT distributions, the .1, .3, .5, .7 and .9 quantile RTs of each participant were averaged across participants. Figure 6 shows the mean correct quantile RTs as well as the mean proportions correct in each condition. The predicted quantile RTs and proportions correct from the best fitting model (the T_{er} model) are indicated as well. Figure 6 shows that all five quantile RTs of the correct responses were shorter on AS trials than on no-AS trials. However, the absolute AS effect was small relative to the differences between the quantile RTs, which makes visual inspection difficult. To examine the AS effect in more detail, we calculated the RT difference between AS trials and no-AS trials (i.e., the AS effect) for each of the five correct RT quantiles. The resulting *delta plot* provides a way of zooming in on the AS effect at different points of the RT distribution (e.g., de Jong, Liang, & Lauber, 1994; Ridderinkhof, 2002). Figure 7 shows the delta plots for the observed data and for the data produced by the best-fitting T_{er} , *a* and *v* models. The AS effect is rather constant across the .1 - .7 quantiles, as is predicted by the T_{er} model, but is somewhat increased for the .9 quantile. The *a* and *v* models both predict that the AS effect gradually increases as RTs become longer. Most of the conditions in the observed data did not show this pattern, which explains why the T_{er} provided a better account of the data than the *a* and *v* models. In addition, an AS effect on *a* or *v* would lead to different proportions of correct responses in AS trials and no-AS trials, which was not found in the data.

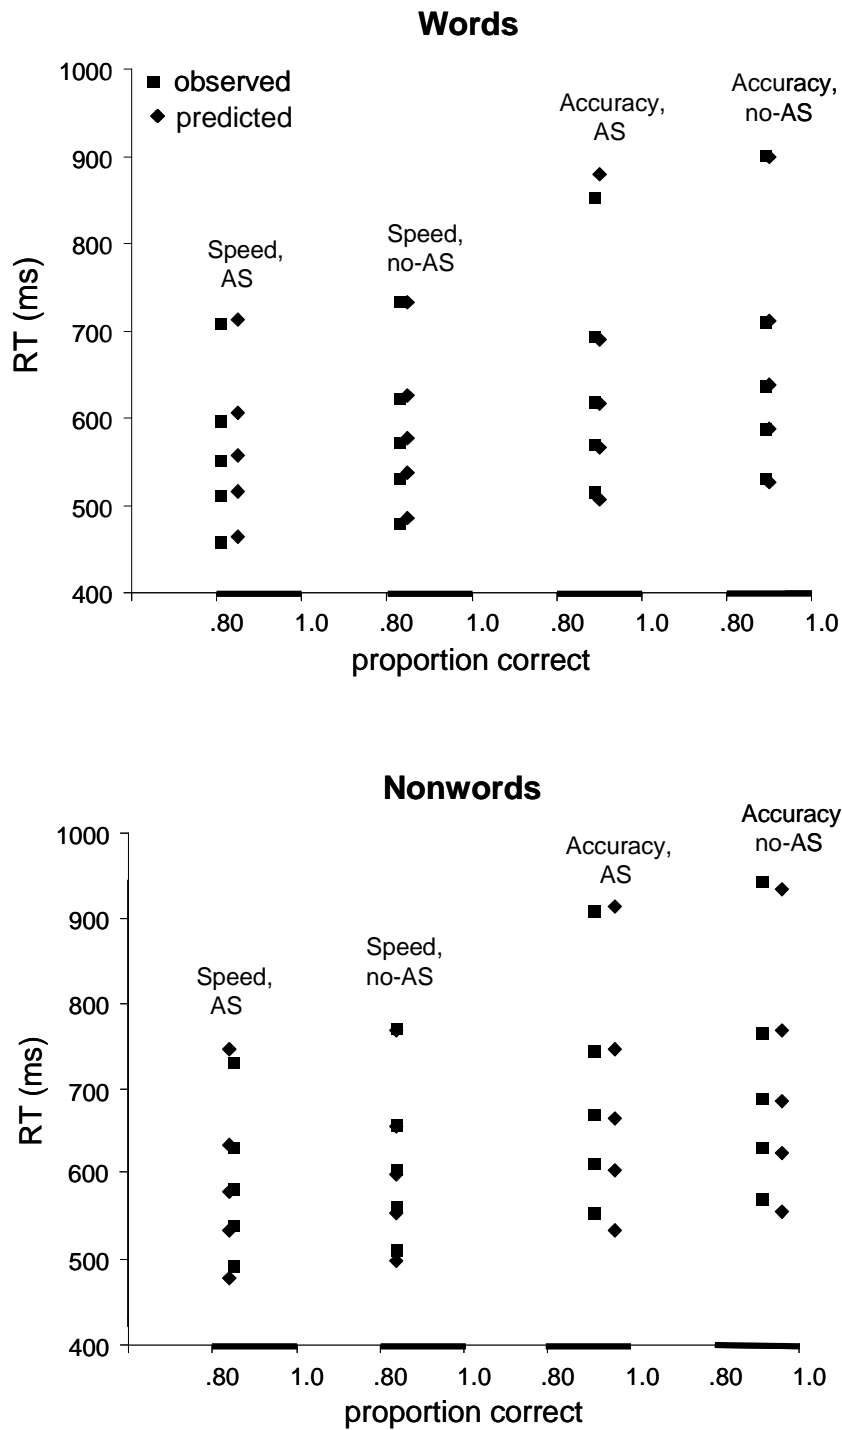


Figure 6. The observed and predicted .1, .3, .5, .7 and .9 correct quantile RTs plotted against the corresponding proportions correct, as a function of word type, instruction (speed/accuracy) and AS presence.

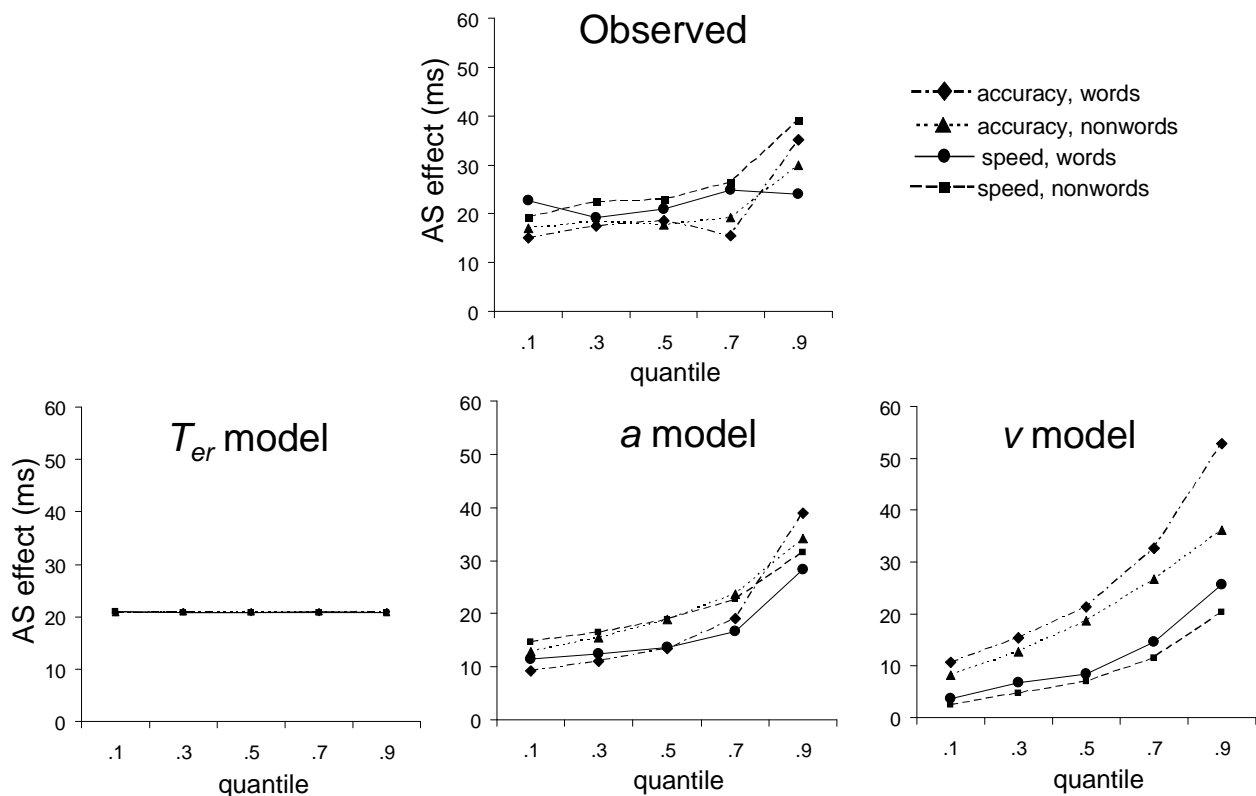


Figure 7. Observed and predicted delta plots for the correct RT distribution as a function of instruction and word type. Note that in the v model, the AS effect does vary with boundary separation a . This occurs because the changes in v have a larger impact on RT when a is large (accuracy instruction) than when a is small (speed instruction)

Discussion

We applied the diffusion model to the data from a lexical-decision experiment in which the visual imperative stimuli (letter strings) were accompanied by an auditory AS or not. The diffusion-model analysis of these data provided important evidence regarding the source of the AS effect. The fit of a model in which all critical parameters were left unconstrained showed that the AS effect was largely accounted for by a change in the nondecision component T_{er} . In contrast, the decision parameters drift rate and boundary separation, although sensitive to other experimental variables, were not affected by AS presence. In the regular behavioral analyses, we also found no indications for an AS effect on boundary separation: there was no speed-accuracy trade-off between AS trials and no-AS trials; and no interaction between the effects of AS presence and instruction (emphasis on speed or accuracy), a variable which affected boundary separation. A comparison of models in which only one parameter was allowed to vary between AS trials and no-AS trials pointed in the same direction: for almost all of the participants the T_{er} model was best able to explain the data. The T_{er} model was also significantly better than models in which combinations of two parameters or all three parameters were free to vary as a function of AS presence. Finally, the AS effect was relatively constant across the RT distribution. This implies that accessory stimuli did not alter the

shape of the RT distribution but shifted the complete distribution to the left, which is consistent with an effect on the nondecision component.

These results strongly suggest that accessory stimuli do not affect the decision process itself, but instead speed up nondecision processes. Based on the diffusion-model analysis alone, it cannot be determined whether the shortening of the nondecision component reflects a speeding of stimulus encoding or response execution, or both. However, the electrophysiological results of Experiment 1 and previous work (Hackley & Valle-Inclán, 1998, 1999) rule out a speeding of response execution. Therefore, the combined results from Experiments 1 and 2 suggest that the AS effect reflects speeding of the stimulus-encoding process.

General Discussion

We conducted two experiments to assess which components of information processing are affected by accessory stimuli. The combined results of the two experiments have led us to the following three main conclusions. First, accessory stimuli speed up encoding of the imperative stimulus. This is possibly the result of energy integration in visual-processing areas. Second, accessory stimuli cause a bilateral (nonspecific) increase in cortical motor activation, which is not expressed in a RT benefit. Third, accessory stimuli have little or no effect on the decision process. Each conclusion will be addressed below.

Accessory stimuli speed up encoding of the imperative stimulus

The EEG results and diffusion-model analyses reported here support the stimulus-encoding account of the AS effect. The EEG results indicated that some of the effect was already present at the time of the N1 peak, and that most of the effect developed in the interval between the N1 and LRP onset. The diffusion-model analyses suggested that the effect occurred before the start of the decision process, which is presumably some tens of milliseconds before LRP onset, which marks the moment when asymmetric evidence accumulation is revealed at the level of the motor cortex. The notion that accessory stimuli speed up stimulus encoding seems consistent with behavioral studies demonstrating that auditory signals, when presented concurrently with the visual imperative stimulus, can facilitate spatial visual search (van der Burg, Olivers, Bronkhorst, & Theeuwes, 2008) and target detection in rapid serial visual presentation streams (Dalton & Spence, 2007; Vroomen & de Gelder, 2000), and increase the perceived intensity of visual stimuli (Stein et al., 1996). An interesting goal for future research will be to investigate whether these seemingly similar phenomena are indeed caused by a common mechanism.

The energy-integration hypothesis has been forwarded as a specific account of how accessory stimuli might speed up stimulus encoding (Bernstein, 1970). According to this hypothesis, stimulus energy is integrated across different modalities in such a way that adding an auditory AS is comparable to increasing the intensity of the visual imperative stimulus. The notion of intermodal energy convergence—even in presumptive unimodal sensory areas—is consistent

with the existence of direct connections between auditory cortex and primary visual cortex (Falchier, Clavagnier, Barone, & Kennedy, 2002; Rockland & Ojima, 2003), and multisensory neurons in low-level sensory areas, such as auditory-sensitive neurons in visual cortex (Morrell, 1972). We found that accessory stimuli increased the amplitudes of early ERP components (P1/N1) over visual-processing areas in a way that is consistent with the energy-integration hypothesis. Previous studies have found that amplitude increases of early visual ERP components are associated with faster target-detection RTs and forward shifts in the perceived onset of visual stimuli (McDonald, Teder-Sälejärvi, Di Russo, & Hillyard, 2005; Talsma, Mulckhuyse, Slagter, & Theeuwes, 2007), suggesting that increased strength of neural activity in visual cortex speeds up downstream perceptual processing. Thus, accessory stimuli might have led to increased neural activity in visual cortex (reflected in P1/N1), which in turn might have speeded up subsequent encoding processes. This possibility is consistent with our finding that the first AS-induced increase in ERP amplitude (~ 100 ms after stimulus onset, at the time of the P1 peak) preceded the beginning of the latency effect (~160 ms after stimulus onset, at the time of the N1 peak). However, due to the inherent limitations of EEG methods (i.e., the “inverse problem”), the ERP findings cannot be taken as conclusive evidence for energy integration in visual-processing areas. They provide merely a motivation for future research designed to determine the mechanism underlying the AS effect.

The effect of accessory stimuli on stimulus encoding might be related to stochastic resonance in sensory systems. Stochastic resonance is the counterintuitive phenomenon that adding a certain level of noise to a nonlinear system enhances its response to a weak (subthreshold) input signal (Benzi, Sutera, & Vulpiani, 1981). A possible explanation for stochastic resonance in perceptual systems is that the addition of noise pushes subthreshold stimuli across their threshold, resulting in improved detection of the stimuli (Moss, Ward, & Sannita, 2004). Stochastic resonance effects on stimulus detection have also been demonstrated when the signal and the noise were of different modalities (Manjarrez, Mendez, Martinez, Flores, & Mirasso, 2007). Manjarrez and colleagues found that continuous auditory noise improved the detection of subthreshold visual stimuli, which was explained by an increased response of multisensory neurons to the converged auditory and visual input. Along similar lines, the joint presentation of imperative and accessory stimuli might cause a faster increase in neural activation in visual-processing areas than the imperative stimulus alone, thereby precipitating detection of the imperative stimulus. Whether indeed similar neural mechanisms are involved in the AS effect and stochastic resonance is an interesting question for future research.

The conclusion that accessory stimuli facilitate stimulus encoding may be important for a better understanding of other phenomena reported in the attentional literature. A prominent example is the warning effect, which is also referred to as the temporal preparation effect. In the temporal preparation paradigm, a warning stimulus announces the onset of an imperative stimulus. Unlike in the AS paradigm, the interval (or foreperiod) between warning stimulus and imperative stimulus is long enough to enable deliberate preparation (usually > 500 ms). When foreperiods are constant within blocks but vary between blocks, the typical finding is that RT increases with increasing

foreperiod length (Niemi & Näätänen, 1981). This is thought to reflect a more difficult estimation of the timing of the imperative stimulus for longer foreperiods (Klemmer, 1956). LRP studies and psychophysical measurements have yielded evidence for a pre-motoric locus of the effect (Müller-Gethmann, Ulrich, & Rinkebaumer, 2003; Rolke & Hofmann, 2007; but see Rudell & Hu, 2001). Furthermore, animal research has indicated that during the foreperiod interval there is a gradual increase in the firing rate of visual neurons (Ghose & Maunsell, 2002), suggesting that the benefit of temporal preparation is at least in part due to perceptual changes. Although the warning effect does not reflect motoric changes, the degree of temporal preparation is known to affect response force (Mattes & Ulrich, 1997) and reflex amplitude (Brunia & van Boxtel, 2000). Thus, in several regards there is a marked similarity between the effects of temporal preparation and accessory stimulation. Indeed, Bernstein, Chu, Briggs, and Schurman (1973) have suggested that enhanced preparation is one of the mechanisms underlying the AS effect. While warning stimuli cause a gradual increase in the firing rate of visual neurons, accessory stimuli might cause an immediate increase in firing rate. This would imply that the warning effect and the AS effect correspond to, respectively, endogenous and exogenous instances of the same process (cf. Hackley & Valle-Inclán, 2003).

Accessory stimuli cause a nonspecific increase in motor activation

Besides an effect on stimulus encoding, accessory stimuli induced a bilateral (nonspecific) increase in motor activation, which had no effect on RT. This finding supports the proposal by Sanders (1983) that accessory stimuli trigger a phasic burst of arousal that leads to nonspecific priming of low-level motor pathways, and that this effect occurs independently from the stimulus-response translation processes contributing to RT. It also has important implications for previous findings of AS effects on motor processes. Interactions of AS presence with manipulations that influence motor processes (e.g., instructed tonic muscle tension) have been interpreted, using additive-factors logic, as evidence that accessory stimuli affect the speed of motor processes (Sanders, 1980; Schmidt et al., 1984). One problem with this line of reasoning is that the critical assumptions underlying the additive-factors logic are highly disputed. For example, researchers have challenged the assumption that information processing consists of a sequence of discrete nonoverlapping stages (e.g., Spencer & Coles, 1999). But even setting aside the problems with these assumptions, an interaction between accessory stimulation and motor manipulations only indicates that accessory stimuli influence motor processes; the interaction does not specify the nature of this influence and whether it is associated with a change in the *duration* of motor processes. An AS-induced nonspecific increase in motor activation, even when having no direct effect on RT, may modulate the effects of other variables on the duration of motor processes (and hence RT), and therefore could have been responsible for the interactions that were found in studies using additive-factors logic.

As discussed above, the conclusion that accessory stimuli caused a bilateral increase in motor cortex activation also offers an explanation for previous findings that accessory stimuli

increase response force, independently from their effects on RT (Miller et al., 1999; Stahl & Rammsayer, 2005). According to this explanation, the AS-induced stronger activation of the relevant (contralateral) motor cortex causes an increase in response force. Conversely, there is no evidence that a bilateral increase in motor activation affects choice RT. Instead, it appears that choice RT is dependent on the *difference* between the activity in the relevant and irrelevant motor cortex (Gratton et al., 1988; Mordkoff & Grosjean, 2001), which was not substantially affected by accessory stimuli in Experiment 1. It is plausible that the AS-evoked nonspecific arousal effect also increases the excitability of other motor systems. If so, this may explain the finding of an increased photic blink reflex when the reflex-eliciting stimulus was accompanied by an acoustic AS (Low et al., 1996)

Accessory stimuli might activate the motor cortex either directly, via connections between the auditory cortex and the motor cortex (Buser & Imbert, 1961; Ermolaeva, Tolchenova, & Brukhanskaya, 1981), or indirectly. One possible indirect way in which accessory stimuli could activate the motor cortex is via the locus coeruleus, the main noradrenergic nucleus in the brainstem. Locus coeruleus neurons exhibit a rapid increase in activity following motivationally significant or salient stimuli (Aston-Jones, Rajkowski, & Cohen, 2000). This causes the release of norepinephrine in cortical and subcortical projection areas, which increases the responsiveness of efferent neurons to their input (Servan-Schreiber, Printz, & Cohen, 1990). It is plausible that the high-intensity auditory accessory stimuli that were used in the current study, by virtue of their salience, caused a phasic locus coeruleus response. The resulting release of norepinephrine may have caused the AS-induced increase in motor activation. In line with this hypothesis, it has been shown that the availability of norepinephrine is critical for an AS-induced increase of the masseteric-reflex amplitude (Stafford & Jacobs, 1990). It remains to be determined whether the noradrenergic system is also involved in AS-induced changes in voluntary motor responses.

Accessory stimuli have little or no effect on the decision process

Our diffusion-model analyses suggested that AS presence did not affect the main parameters of the decision process: the rate of evidence accumulation and the decision threshold. In addition, no AS-induced speed-accuracy trade-off was found in either of the two experiments. These findings suggest that accessory stimuli did not have a substantial effect on the decision process. However, the increased number of nogo errors (i.e., false alarms) suggests that accessory stimuli induced a lowering of the decision threshold for the go-nogo decision (Gomez, Ratcliff, & Perea, 2007). Note that go responses were much more frequent than nogo responses (80% vs. 20%), which probably resulted in a bias towards the go response. In terms of the diffusion model, this means that the starting point for the go-nogo decision was closer to the go threshold than to the nogo threshold. In contrast, the decision which hand to respond with was unlikely to be biased towards one of the decision thresholds, because left and right responses occurred equally often. The effect of a lowering of the decision threshold on the probability that the diffusion process reaches that threshold by mistake is larger as the threshold is closer to the starting point. Therefore, it is possible

that accessory stimuli caused a lowering of the decision thresholds that was too small to significantly affect the number of errors in the left-right decision, but large enough to increase the number of incorrect go responses in the go-nogo decision.

The above hypothesis predicts that accessory stimuli only induce a speed-accuracy trade-off in situations in which the decision threshold is close to the starting point of the decision process. Aside from circumstances that induce a strong response bias, this is likely to be the case in easy choice RT tasks. Previously studies provide strong support for this prediction: Significantly increased error rates on AS trials have generally been found in studies using relatively simple tasks (e.g., requiring a spatially compatible stimulus-response mapping) with very short mean RTs (< 350 ms), suggesting that the response threshold was close to the starting point (Low et al., 1996; Posner, Klein, Summers, & Buggie, 1973; Schmidt et al., 1984). In contrast, the absence of a significant AS effect on error rate has been found in more complex tasks that produced intermediate to long mean RTs (> 500 ms; e.g., Hackley & Valle-Inclán, 1999; De Jong, 1991, Experiment 1; the present two experiments). To prevent too many errors, the decision thresholds in these more complex tasks were probably at a relatively large distance from the starting point. Thus, previous findings of AS effects on error rates are consistent with the hypothesis that accessory stimuli cause a small lowering of the decision thresholds, which is only expressed in an increased error rate when the threshold is close to the starting point.

Our LRP findings showed that AS presence did not affect the response-locked LRP. According to the continuous flow theory (Eriksen & Schultz, 1979), stimulus evaluation and response activation proceed largely in parallel, and response activation is continuously influenced by the output of the stimulus-evaluation process. This suggests that the LRP is an accurate reflection of the accumulated evidence in the decision process, and corresponds to the drift rate in the diffusion model. Although systematic evidence for this view is still missing, important support has been provided by electrophysiological data (Coles, Gratton, & Donchin, 1988; Gratton et al., 1988) and computational considerations (Usher & McClelland, 2001). To the extent that the LRP indexes an evidence-accumulation process, the absence of an effect of AS presence on the response-locked LRP suggests that neither the rate of evidence accumulation nor the decision threshold was affected by accessory stimuli. This would be consistent with our diffusion-model analyses.

Summary of conclusions

Our findings suggest that accessory stimuli facilitate encoding of the imperative stimulus. A possible mechanism for this facilitation, consistent with anatomical and physiological findings, is energy integration in visual-processing areas. To further investigate this possibility, a closer link with the multisensory-integration literature and associated methods is warranted. In addition, we found that accessory stimuli induce a bilateral increase in motor activation that is independent of the RT benefit. This finding provides new and direct support for nonspecific arousal models, and offers an explanation for previously reported AS effects on response-amplitude measures. Finally, we found no evidence that accessory stimuli affect the rate of evidence accumulation in the decision

process. An AS-induced lowering of the decision threshold, if present at all, is small, and is translated in increased error rates only for decisions with a starting point that is already close to the decision threshold. We believe that these findings, obtained by a combination of electrophysiology and diffusion-model analyses, provide an important contribution to our understanding of the effects of accessory stimuli on information processing. One important aim for future research will be to combine these two methods in a single experiment, such that the various types of results can be more easily integrated.

Chapter 7

Temporal expectation and information processing: A model-based analysis

This chapter is based on: Jepma, M., Wagenmakers, E.-J., & Nieuwenhuis, S. (under revision).
Temporal expectation and information processing: A model-based analysis.

Abstract

People are able to use temporal cues to anticipate the timing of an event, enabling them to process that event more efficiently. We conducted two experiments, using the fixed-foreperiod paradigm (Experiment 1) and the temporal-cueing paradigm (Experiment 2), to assess which components of information processing are speeded when subjects use such temporal cues to predict the onset of a target stimulus. We analyzed the observed temporal expectation effects on task performance using sequential-sampling models of decision making: the Ratcliff diffusion model and the shifted-Wald model. The results from the two experiments were consistent: temporal expectation affected the duration of nondecision processes (target encoding and/or response preparation) but had little effect on the two main components of the decision process: response-threshold setting and the rate of evidence accumulation. Our findings provide novel evidence about the psychological processes underlying temporal expectation effects on simple- and choice-reaction time.

Introduction

People are able to use temporal cues to anticipate with great precision the timing of an event, enabling them to optimize the processing of that event. For example, people can use the onset of amber traffic lights to direct the temporal focus of attention towards the moment in time in which the lights will turn green (or red, depending on the region of the world they are in), allowing them to speed up their response to the green signal. Experimental psychologists have long known that response times (RTs) are faster if a target is preceded by a warning signal that is presented at a constant, or at least predictable, temporal delay (reviewed in Hackley, 2009; Niemi & Näätänen, 1981; Nobre, Correa, & Coull, 2007). This beneficial effect is also observed for choice RTs, even though the warning signal contains no information about the identity of the upcoming stimulus. The ability of people to use temporal cues is also evident in the brain: neurons in several brain areas encode the probability that a stimulus will occur at any given point in time (Ghose & Maunsell, 2002; Janssen & Shadlen, 2005; Riehle, Grün, Diesmann, & Aertsen, 1997). The goal of the current study was to increase our understanding of *which components of information processing* are speeded when people can predict the onset of a target stimulus.

The effects of temporal expectation on task performance have been studied with two different paradigms, developed in largely separate literatures. One is the foreperiod paradigm, in which the warning signal is a perfect predictor of the interval (or *foreperiod*) between the onsets of the warning signal and the target. Foreperiod duration is typically varied between blocks of trials. The typical finding in this paradigm is that RTs increase progressively as the duration of the foreperiod is increased and therefore harder to estimate (Klemmer, 1956; Niemi & Näätänen, 1981)⁶. The other paradigm is the temporal-cueing paradigm, in which a cue predicts with some certainty (e.g., 80%) the interval between the onsets of the cue and the target. The *cue-target interval* is varied within blocks of trials. The typical finding in this paradigm is that RTs are faster when the cue-target interval is validly cued (i.e., confirms the participant's expectation) than when the interval is invalidly cued (Correa, Lupiáñez, Milliken, & Tudela, 2004; Coull & Nobre, 1998). The manipulation within blocks of cue-target intervals and the dissociation of expected and actual cue-target intervals (on invalidly cued trials) make the temporal-cueing paradigm more suitable for event-related fMRI studies, which have examined the brain areas that are activated when people process the temporal cue and orient their attention (reviewed in Coull, 2004). However, it seems reasonable to assume that the key behavioral effects obtained in the two paradigms reflect similar underlying mechanisms: in both paradigms participants are required to voluntarily orient their attention to particular moments in time; and experimental manipulations (foreperiod duration or cue validity) affect the degree to which participants are prepared at the moment when the target is presented.

⁶ We only consider foreperiods ≥ 400 ms. At shorter foreperiods target processing can benefit from the phasic increase in arousal elicited by the presentation of the warning stimulus: an *accessory stimulus effect* (Hackley & Valle-Inclán, 1998; Jepma, Wagenmakers, Band, & Nieuwenhuis, 2009)

Which aspects of information processing are responsible for the decrease in RTs as temporal certainty increases? One possible account is that temporal certainty facilitates encoding of the target (cf. Jepma, Wagenmakers, Band, & Nieuwenhuis, 2009; Niemi & Näätänen, 1981). Another possibility is that temporal certainty affects a critical parameter of the decision process that is based on the sensory evidence obtained during stimulus encoding. The mechanism underlying two-choice decisions is well described by the accumulation of noisy information from a stimulus over time (Gold & Shadlen, 2007; Smith & Ratcliff, 2004). Information accumulates toward one or the other of two decision thresholds until one of the thresholds is reached; then the response associated with that threshold is initiated. It is possible that orienting attention to the moment of target onset, or a timed phasic increase in arousal, speeds up the rate with which evidence is accumulated in the decision process (cf. Grosjean, Rosenbaum, & Elsinger, 2001). Another possibility is that increased temporal certainty does not change the rate of information build-up but instead causes a lowering of the decision threshold (or, equivalently, a rise in starting point; Bogacz, Wagenmakers, Forstmann, & Nieuwenhuis, 2010). That is, participants begin to decrease the threshold in anticipation of the target. As a result, responses are faster because decisions are made on the basis of less evidence (Posner, 1978). A final account assumes that increased certainty about the timing of an upcoming target can be used to prepare the motor system, without committing to any particular response (Bertelson, 1967; Sanders, 1980). This may speed up the execution of a specific motor response to the target, much like a pre-heated engine will make a car start quicker in any direction.

Previous research has found substantial evidence regarding the locus of temporal certainty effects: To examine the response execution account, researchers have conducted choice-RT experiments that examined the effect of foreperiod on the lateralized readiness potential (LRP), a difference wave that indexes hand-specific response preparation. The onset of the LRP indicates the moment at which the motor cortex associated with the responding hand becomes more active than the ipsilateral motor cortex, an early indication of the forthcoming motor response. The general finding is that the effect of foreperiod on the interval between LRP onset and the overt response is small or absent, which has led researchers to conclude that there is very little evidence for a foreperiod effect on the duration of motor preparation and execution (e.g., Hackley, Schankin, Wohlschlaeger, & Wascher, 2007; Müller-Gethmann, Ulrich, & Rinkeauer, 2003). However, Tandonnet and colleagues have suggested that these LRP findings may be misleading. They examined the Laplacian-transformed event-related potential (ERP) waveforms to obtain separate estimates of the ipsilateral and contralateral motor cortex response. Although effect sizes were modest, Tandonnet and colleagues found that increased temporal certainty decreased the time between the onset of the contralateral negativity indexing the motor command and the electromyographic (EMG) onset (Tandonnet, Burle, Vidal, & Hasbroucq, 2003, 2006), suggesting a speedup of motor preparation. When they used the same data to compute the monopolar (i.e. standard) and the Laplacian LRPs, they found no foreperiod effect on the LRP-to-response interval. This suggests that the double-subtraction methods used to compute the LRP can obscure subtle latency effects present in the constituent ERP waveforms. Tandonnet and colleagues further found

that increased temporal certainty shortened the time between EMG onset and the actual key press (Tandonnet et al., 2003; see also Hasbroucq, Akamatsu, Mouret, & Seal, 1995). This indicates that temporal certainty can also influence the duration of motor execution.

While there are small but robust effects of temporal certainty on the duration of motor processes, these effects cannot fully account for temporal expectation effects on RT. In particular, several studies have found that increased temporal certainty reduces the interval between the stimulus and the P3/LRP onset, two established markers of the combined duration of stimulus encoding and decision making (Correa, Lupiáñez, Madrid, & Tudela, 2006; Müller-Gethmann et al., 2003). These studies suggest that temporal expectation effects on RT must also have an earlier locus.

Temporal certainty improves various aspects of perception (Bausenhart, Rolke, & Ulrich, 2008; Martens & Johnson 2005; reviewed in Nobre et al., 2007). Importantly, it also improves performance in psychophysical variants of the two paradigms discussed above, in which target stimuli are briefly presented and then masked: Increased temporal certainty enhances perceptual sensitivity (*d*-prime) in both the foreperiod paradigm (Rolke, 2008; Rolke & Hofmann, 2007) and the temporal-cueing paradigm (Correa, Lupiáñez, & Tudela, 2005). However, although highly informative, these findings cannot adjudicate between effects on encoding and the rate of evidence accumulation (cf. Rolke & Hofmann, 2007). That is, perceptual sensitivity may be enhanced because encoding lasts shorter and evidence accumulation can start earlier, or because evidence accumulation progresses at a faster rate; both scenarios result in more evidence by the time the target stimulus is masked and subjects must make a decision. Bausenhart and colleagues have tried to distinguish between these accounts by investigating the foreperiod effect on the shape of speed-accuracy tradeoff functions obtained with the response-signal method (Bausenhart, Rolke, Seibold, & Ulrich, 2010). They found that foreperiod affected the intercept but not the slope of these functions, providing evidence for changes in encoding duration but not the rate of evidence accumulation. Together, these and other behavioral findings (Seifried, Ulrich, Bausenhart, Rolke, & Osman, 2010) provide substantial evidence that temporal certainty affects the duration of stimulus encoding.

While there is substantial evidence that temporal certainty affects the duration of encoding and motor processes, the picture is less clear for the two main components of the decision process: threshold setting and rate of evidence accumulation. According to the response-threshold account, increased temporal certainty results in a well-timed lowering of the response threshold, such that decisions are made on the basis of less evidence. A straightforward prediction of this account is that the faster RTs should be accompanied by a higher proportion of errors—choice errors in choice-RT tasks and false alarms in simple-RT tasks in catch trials. Unfortunately, studies with simple-RT tasks generally do not report false-alarm proportions, or do not include catch trials in the design. Furthermore, response accuracy in choice-RT tasks is generally near ceiling, which necessarily results in negligible and non-significant foreperiod effects. In the rare two-choice RT studies in which accuracy was off ceiling, foreperiod effects on accuracy were small or absent. An exception

is an experiment reported by Posner, Klein, Summers, and Buggie (1973), who found a speed-accuracy tradeoff when comparing foreperiods of 400 ms and 800 ms. A distinct feature of the results in this study were the extremely fast responses, due to speed emphasis in the task instructions. As we have discussed elsewhere (Jepma et al., 2009), the effect of lowering the decision threshold on the probability that the evidence-accumulation process reaches that threshold by mistake (i.e., resulting in an error), is larger when the threshold is closer to the starting point, as is the case when instructions emphasize speed. Thus, the increase in error rates with higher temporal certainty (foreperiod = 400 ms) in the experiment of Posner and colleagues is consistent with the response-threshold account and may have become apparent because of a small distance between starting point and threshold. Taken together, a review of speed-accuracy tradeoff data yields little evidence for or against the response-threshold account. Furthermore, as we will discuss later, although the response-threshold account predicts a speed-accuracy trade-off, the observation of a speed-accuracy trade-off is not uniquely diagnostic of shifts in response threshold.

Finally, as noted above, there is preliminary evidence that temporal certainty in the foreperiod paradigm does not affect the rate of evidence accumulation (Bausenhardt et al., 2010). Aside from those results, there are no data informing the evidence-accumulation account, in part because standard behavioral indices predicted by the evidence-accumulation account cannot be distinguished from predictions of the encoding account (cf. Rolke & Hofmann, 2007). Therefore, other methods are needed to test whether temporal certainty affects components of the decision process.

We conducted two experiments using the two paradigms that are most commonly used in temporal-certainty research: the fixed-foreperiod paradigm (Experiment 1) and the temporal-cueing paradigm (Experiment 2). Previous work has always focused on either the foreperiod paradigm or the temporal-cueing paradigm, which explains the lack of integration of the two literatures. To enable a comparison of the temporal-certainty effects in the two paradigms, we identified the psychological process(es) underlying the observed temporal-certainty effects in both paradigms using two sequential-sampling models for distributions of response times and error rates. One goal was to confirm the hypothesis that temporal certainty affects the duration of nondecision processes, as suggested by the literature reviewed above. However, the models we used were particularly useful for testing the evidence-accumulation and response-threshold accounts, because each of these components of decision making corresponds with a unique parameter in both models. Therefore, our primary goal was to examine whether the values of these decision-making parameters changed as a function of temporal certainty.

Experiment 1

In Experiment 1 we investigated which components of information processing are affected by temporal certainty using a diffusion-model analysis of the foreperiod effect on RT and accuracy. The diffusion model is a model of two-choice decision making that defines the decision process as

the continuous accumulation of noisy stimulus information over time, from a starting point towards one of two decision criteria or thresholds (Ratcliff & Rouder, 1998; see Figure 1). When one of the two thresholds is reached, the corresponding response is initiated. There are several reasons to assume that the diffusion model gives an accurate reflection of how the decision process is implemented in the brain. First, the diffusion process is the optimal decision process: it provides the fastest responses for a fixed level of accuracy, or the highest accuracy for a fixed response time (Wald, 1947). Second, the diffusion model explains the dynamics of neuronal activity during decision-making behavior (Gold & Shadlen, 2007; Smith & Ratcliff, 2004). And third, the diffusion model successfully accounts for RT distributions and error rates in a variety of two-alternative forced-choice tasks (e.g., Ratcliff, Van Zandt, & McKoon, 1999).




Figure 1. An illustration of the diffusion model. The parameters are: a = boundary separation, z = starting point, v = drift rate, T_{er} = mean nondesideration time. The sample paths represent moment-by-moment fluctuations in the evidence favoring the two possible responses, which is due to noise in the decision process. The decision process starts at z and terminates when one of the two boundaries is reached. The duration of T_{er} determines the additional time needed for stimulus encoding and response execution.

The diffusion model can be helpful in evaluating the various accounts of the foreperiod effect because some of the main model parameters correspond closely to the different processing components emphasized by these accounts. The three most important parameters of the model in this respect are the drift rate, the boundary separation, and the nondesideration component. The drift rate (v) is the mean rate of evidence accumulation in the decision process, which depends on the quality of the stimulus and the perceptual system. The higher the absolute value of the drift rate, the faster a decision threshold is reached. If accurate predictions of target onset time increase the drift rate of the diffusion model, this would support the idea that high temporal certainty induces a faster build-up of information. The boundary separation (a) is the distance between the two decision criteria. This parameter determines on how much evidence a decision is based, and can be controlled strategically by the decision maker. If the decision maker uses temporal prediction to briefly lower the boundary separation, this would provide support for the notion that the foreperiod

effect reflects a lowering of the decision threshold (Posner, 1978). Besides the decision process, there are other components of processing involved in a two-choice RT task, namely stimulus encoding and response execution which, respectively, precede and follow the decision process. In the diffusion model, these nondecision processes are combined into one nondecision component, T_{er} . A shortening of the nondecision component by accurate prediction of target onset would indicate that stimulus encoding and/or motor execution are speeded.

We applied the diffusion model to data from a standard lexical-decision task, in which participants were asked to classify letter strings as a word or a nonword, with task instructions emphasizing response speed in half of the blocks and response accuracy in the other half of the blocks. The diffusion model has been shown to provide a good fit of lexical-decision data, accounting for the effects of the experimental variables on RTs for correct and error responses, shapes of the RT distributions, and accuracy values (Ratcliff, Gomez, & McKoon, 2004; Wagenmakers, Ratcliff, Gomez, & McKoon, 2008). Importantly, each letter string was preceded by a warning signal, and the stimulus-onset asynchrony between the two stimuli, the foreperiod, was varied between blocks (500 or 2700 ms). Our major aim was to examine which model parameter(s) could best account for the corresponding differences in task performance.

Method

Participants. Fourteen students participated (11 women; aged 18-29 years; mean age = 21.5; all native Dutch speakers). All participants reported normal hearing and normal or corrected-to-normal vision. Each participant completed two sessions of approximately 90 minutes each, on separate days. Participants received either 18 euros or course credits for participation.

Design and procedure. Participants were tested individually in a dimly lit room. Stimuli were presented in silver on a navy blue background on a personal-computer screen. Each trial started with the presentation of a 200-ms asterisk symbol (visual angle = 0.8°) in the center of the screen, which marked the onset of the foreperiod. This warning signal was followed by the remainder of the foreperiod (300 ms or 2500 ms) during which a fixation plus (0.3°) was on the screen. Then a letter string was presented (Courier New font; visual angle = 2.7° for 4-letter words and 4.0° for 6-letter words), and participants were instructed to decide whether or not the letter string was a Dutch word by pressing the z or the / key. The key assignment was balanced across participants. The letter string remained on the screen until a response was made, after which the fixation plus reappeared for an intertrial interval of $(1.1 + X)$ s, with X being a random variable that followed an exponential distribution with a mean of 1 s. This random interval was used to emphasize the importance of the warning signal as a temporal reference for preparation (cf. Rolke & Hofmann, 2007).

The word stimuli were 800 Dutch words and 800 nonwords. Both the words and the nonwords consisted of 4, 5 or 6 letters (195 4-letter, 251 5-letter and 354 6-letter words as well as nonwords). The frequency of the words ranged from 0.07 to 5.48 per million (mean = 3.47, SD = 1.28; Baayen, Piepenbrock, & Gulikers, 1995). The nonwords were generated by replacing one

letter of an existing word; vowels were replaced by vowels and consonants by consonants. The words that were used to generate the nonwords were not used as word stimuli.

In each of the two sessions, participants completed two practice blocks of 24 trials, followed by 16 experimental blocks of 50 trials: 25 with a word and 25 with a nonword. The combination of speed-accuracy instructions and foreperiod changed after every two blocks according to an ABCD DCBA order that was the same in both sessions and varied across participants. Before the start of each block, participants received an on-screen announcement of the upcoming foreperiod (long or short) and speed-accuracy instructions (focus on accuracy or speed), after which they could press the space bar to start the block. In speed blocks, participants were instructed to respond as quickly as possible, but without making a lot of errors, and responses slower than 650 ms were followed by a message TOO SLOW of 1 s. When a response was faster than 250 ms, the message TOO FAST was displayed for 1 s. No accuracy feedback was given in these blocks. In accuracy blocks, participants were instructed to respond as accurately as possible, but without taking more time to respond than necessary, and incorrect responses were followed by a message ERROR of 1 s. Responses faster than 250 ms or slower than 1200 ms were followed by a TOO FAST or TOO SLOW message. At the end of each block the mean RT and the proportion of correct responses appeared on the screen, and participants could take a short break before initiating the next block.

Diffusion-model analysis. For fitting the diffusion model to the data we used the Diffusion Model Analysis Toolbox (DMAT; Vandekerckhove & Tuerlinckx, 2008). DMAT estimates parameters by maximizing a multinomial likelihood function. The data that are used to fit the diffusion model are the RT distributions for correct and incorrect responses, and the percentage correct responses.

We fitted four different diffusion models to the data. The following parameter settings applied to all models: 1) The intertrial variability in nondecision time (st) was held constant across all conditions. 2) The starting point of the diffusion process (z) was set at a fixed proportion of the boundary separation, such that the bias in starting point was constant across conditions. 3) Boundary separation (a) and the intertrial variability in starting point (sz) were free to vary between the speed and accuracy conditions (Ratcliff & Rouder, 1998, Experiment 1; Ratcliff, Thapar, & McKoon, 2001, Experiment 2). 4) Mean drift rate (v) and intertrial variability in drift rate (η) were free to vary between the word and nonword trials (Ratcliff, Thapar, Gomez, & McKoon, 2004). The four models differed with regard to the parameters that were free to vary as a function of foreperiod duration. In one model (the All free model), T_{er} , a , and v were all left free to vary. In addition, there were three models in which either T_{er} , a , or v could vary, whereas the other parameters were held constant (the T_{er} model, a model, and v model, respectively).

The models were fitted to the data in two ways. First, the models were fitted to each participant's data individually. When a participant made 10 or fewer errors in a condition, the participant's error data for this condition were not included in the fitting procedure. Second, the models were fitted to the averaged data. The averaged data was obtained by calculating the accuracy and the RTs for correct and error trials associated with the .1, .3, .5, .7 and .9 quantiles for

each individual participant, and then averaging these values across participants. (Note that the quantile RTs are not the mean RTs within bins [Ratcliff, 1979], but the boundary RTs of each quantile.)

Results

Behavioral results. Figure 2 shows the mean correct RT and mean percentage correct as a function of foreperiod duration, instruction and word type. RTs shorter than 250 ms or longer than 2500 ms were excluded from analysis, which resulted in the exclusion of 0.6% of the trials. In accordance with previous studies, RTs were shorter on short-foreperiod trials than on long-foreperiod trials (573 ms vs. 625 ms; $F(1,13) = 53.6, p < 0.001$), yielding a reliable foreperiod effect of 52 ms. Furthermore, RTs were shorter following speed instructions than following accuracy instructions (563 ms vs. 635 ms; $F(1,13) = 23.3, p < 0.001$), and shorter for words than for nonwords (584 ms vs. 615 ms; $F(1,13) = 10.6, p = 0.006$). There were no significant interactions between the three variables.

Percentage correct was lower on short-foreperiod trials than on long-foreperiod trials (80.7 vs. 82.4%, indicating that the increased speed on short-foreperiod trials was accompanied by a small but reliable drop in accuracy ($F(1,13) = 12.6, p = 0.004$). This drop in accuracy on short-foreperiod trials was present in the accuracy condition (83.9 vs. 87.2%) but not in the speed condition (77.6 vs. 77.7%), as reflected in a significant interaction between foreperiod duration and instruction, $F(1,13) = 8.2, p = 0.013$. As expected, percentage correct was higher when the instruction emphasized accuracy than when it emphasized speed (85.5% vs. 77.7%; $F(1,13) = 32.0, p < 0.001$).




Figure 2. Mean correct RT and proportion correct in Experiment 1 as a function of word type, instruction (speed/accuracy) and foreperiod duration.

Experimental effects on the diffusion-model parameters. To assess which parameters were affected by foreperiod duration, we analyzed the foreperiod effect on the estimates of the T_{er} , a and v parameters in the All free model. Table 1 shows the average parameter estimates across participants. As expected, the boundary separation was smaller when the instruction emphasized speed than when it emphasized accuracy ($F(1,13) = 32.9, p < 0.001$). In addition, (absolute) drift rates were higher for words than for nonwords ($F(1,13) = 72.8, p < 0.001$). Importantly, neither boundary separation ($F(1,13) = 2.2, p = 0.16$) nor drift rate ($F(1,13) = 0.01, p = 0.91$) was affected by foreperiod. In contrast, the nondecision component, T_{er} , was significantly smaller on trials with a short foreperiod than on trials with a long foreperiod ($t(13) = 6.0, p < 0.001$). These results suggest that reducing temporal uncertainty shortens one or more nondecision processes, but does not substantially affect the decision process itself.

Table 1. Parameter estimates for the fit of the All free model (SD in parentheses) in Experiment 1. T_{er} = non-decision time (in seconds) comprising stimulus encoding and response execution; a = boundary separation; v = drift rate

parameter	Short FP	Long FP
T_{er}	.446 (.045)	.482 (.045)
a (speed)	.078 (.011)	.080 (.010)
a (accuracy)	.102 (.018)	.109 (.019)
v (words)	.295 (.139)	.314 (.171)
v (nonwords)	-.253 (.109)	-.269 (.108)

Model selection. To further assess the effect of foreperiod duration on the different model parameters, we tested which model had the best fit to the data. To compare the adequacy of the four models (i.e. the All free model, T_{er} model, a model, and v model) in explaining the observed data we used the Bayesian Information Criterion (BIC), a statistical criterion for model selection. The BIC is a decreasing function of the goodness-of-fit for the estimated model, and an increasing function of the number of free parameters to be estimated. Thus, the best model is the model with the lowest BIC value. In addition, the raw BIC values were transformed to a probability scale, enabling a more intuitive comparison of the probabilities of each model being the best model (Wagenmakers & Farrell, 2004). The transformation of BIC values to probability values consists of three steps. First, for each model i , the difference in BIC with respect to the model with the lowest BIC value is computed (i.e., $\Delta_i(\text{BIC})$). Second, the relative likelihood L of each model i is estimated by means of the following transformation: $L(M_i | \text{data}) \propto \exp[-0.5 \Delta_i(\text{BIC})]$, where \propto stands for “is proportional to”. Third, the model probabilities are computed by normalizing the relative model likelihoods, which is done by dividing each model likelihood by the sum of the likelihoods of all models. Table 2 summarizes the average BIC values and probabilities of each of the four models. The T_{er} model was by far the best model ($F(3,39) = 18.3, p < 0.001$). In the individual analyses, the T_{er} model was the best model for 10 of the 14 participants.

Since boundary separation (a) varied as a function of instruction, and drift rate (v) as a function of word type, the a model and v model had more free parameters (13) than the T_{er} model (12). To examine the possibility that the T_{er} model was favored because of its fewer free parameters, we fitted an additional a model to each participant's data in which the effects of instruction and foreperiod on a were additive instead of fully free. Similarly, we fitted an additional v model to each participant's data in which the effects of word type and foreperiod on v were additive. These additive a and v models had the same number of free parameters as the T_{er} model. The average BICs of the additive a and v models were somewhat larger but did not differ significantly from the fully-free versions of these models ($p = 0.21$ and $p = 0.98$ for the a and v models, respectively). Importantly, the additive a and v models had higher BICs than the T_{er} model ($ps < 0.01$), suggesting that the conclusion in favor of the T_{er} model was not due to the fewer free parameters of this model. For the sake of completeness we also examined the models in which combinations of two parameters (T_{er} and a ; T_{er} and v ; a and v) were free to vary as a function of foreperiod duration. The average BIC values of these three models were all higher than that of the T_{er} model, suggesting that the effects of temporal uncertainty could be explained best by a change in nondecision time alone.

Table 2. BIC values for each model in Experiment 1 (SD in parentheses).

	Df	BIC	$p(BIC)$
All free model	16	7,112 (492)	< 0.01
T_{er} model	12	7,102 (492)	> 0.99
a model	13	7,131 (491)	< 0.0001
v model	13	7,195 (512)	< 0.0001

Note: p = BIC model probability

Model fits. To examine the RT distributions, we averaged the .1, .3, .5, .7 and .9 quantile RTs across participants. Figure 3 shows the mean correct quantile RTs as well as the mean proportions correct in each condition. The predicted quantile RTs and proportions correct from the best model (the T_{er} model) are indicated as well. Figure 3 shows that all five quantile RTs of the correct responses were shorter on short-foreperiod trials than on long-foreperiod trials. However, the absolute foreperiod effect was small relative to the differences between the quantile RTs, which makes visual inspection difficult. To examine the foreperiod effect in more detail, we calculated the RT difference between short-foreperiod trials and long-foreperiod trials (i.e., the foreperiod effect) for each of the five correct RT quantiles. We then plotted the foreperiod effect as a function of response speed (the average of the quantile RTs in the long-foreperiod trials and short-foreperiod trials).

The resulting *delta plot* provides a way of zooming in on the foreperiod effect at different points of the RT distribution (e.g., Ridderinkhof, 2002). Figure 4 shows the delta plots for the observed data and for the data produced by the best-fitting T_{er} , a and v models. The foreperiod effect is rather constant across the .1 - .7 quantiles, as is predicted by the T_{er} model, but is somewhat

increased for the .9 quantile for the word conditions. The a and v models both predict that the foreperiod effect gradually increases as RTs become longer. Most of the conditions in the observed data did not show this pattern, which explains why the T_{er} provided a better account of the data than the a and v models.




Figure 3. The observed and predicted (by T_{er} model) .1, .3, .5, .7 and .9 correct quantile RTs in Experiment 1, plotted against the corresponding proportions correct, as a function of word type, instruction (speed/accuracy) and foreperiod duration.




Figure 4. Observed and predicted delta plots showing the foreperiod effect on RT as a function of mean correct quantile RT, instruction (speed/accuracy) and word type in Experiment 1. The upper panels show the observed delta plots and the fit of the best model, the T_{er} model. The lower panels show the fits of the v model and the a model. Note that in the v model, the foreperiod effect varies with boundary separation a . This occurs because the changes in v have a larger impact on RT when a is large (accuracy instruction) than when a is small (speed instruction).

Discussion

We applied the diffusion model to the data from a lexical-decision experiment in which the visual imperative stimuli (letter strings) were preceded by a short or long foreperiod. The diffusion-model analysis of these data provided important evidence regarding the source of the foreperiod effect. The fit of a model in which all critical parameters were left unconstrained showed that the foreperiod effect was largely accounted for by a change in the nondecision component T_{er} . A comparison of models in which only one parameter was allowed to vary between short and long-foreperiod trials pointed in the same direction: for almost all of the participants the T_{er} model was best able to explain the data. The T_{er} model was also significantly better than a model in which all three parameters were free to vary as a function of foreperiod duration. Finally, consistent with previous studies (Hohle, 1965; Leth-Steensen, 2009), the foreperiod effect was relatively constant across the RT distribution. This implies that increased temporal certainty did not alter the shape of the RT distribution but shifted the complete distribution to the left, which is consistent with an effect on the nondecision component.

In contrast, the decision parameters drift rate and boundary separation, although sensitive to other experimental variables, were not substantially affected by foreperiod duration. In the behavioral analyses, we did find a potential indication for a foreperiod effect on boundary separation: there was a speed-accuracy trade-off between short and long-foreperiod trials when instructions emphasized accuracy, but not when instructions emphasized speed. As noted above, a speed-accuracy trade-off in the empirical data can provide a diagnostic criterion for a change in decision threshold. However, because this empirical pattern was not accompanied by a reliable foreperiod effect on the threshold model parameter, we propose another explanation of the speed-accuracy tradeoff. Laming (1979) has suggested that subjects may anticipate the arrival of a stimulus by starting sampling information from the perceptual display at the moment when they think the stimulus will be presented. If subjects start sampling too early, responses will be fast but also less accurate because they start with sampling noise. We assume that subjects use this strategy in blocks when the foreperiod is short and the anticipated timing of the stimulus is relatively good, but not in blocks with a fixed long foreperiod, when the stimulus onset is much harder to anticipate. In long-foreperiod blocks, subjects always wait with sampling until the target occurs, and errors due to premature sampling do not occur. According to this account, accuracy is reduced on short-foreperiod trials not because of a reduction in boundary separation but because subjects engage in premature sampling (of noise) on a proportion of the trials.

A prediction of the premature-sampling hypothesis in terms of diffusion-model parameters is that the inter-trial variability in starting point (sz) will be larger in short-foreperiod blocks than in long-foreperiod blocks, since premature sampling will inflate estimates of starting-point variability. To test this prediction, we fitted a diffusion model to the data in which not only boundary separation, drift rate and nondecision time, but also starting-point variability was free to vary as a function of foreperiod duration. This analysis revealed that estimated starting-point variability was significantly larger when instructions emphasized speed than when instructions emphasized

accuracy ($F(1,13) = 62.2, p < 0.001$). In addition, there was a trend-level effect of foreperiod duration ($F(1,13) = 3.72, p = 0.076$), as well as a significant interaction between foreperiod duration and instruction on estimated starting-point variability ($F(1,13) = 16.8, p = 0.001$)⁷. Follow-up contrasts indicated that starting-point variability was larger in short-foreperiod blocks than in long-foreperiod blocks when instructions emphasized accuracy (0.034 vs. 0.009; $t(13) = 3.14, p = 0.008$), but not when instructions emphasized speed (0.062 vs. 0.063; $t(13) = 0.21, p = 0.84$). Importantly, these effects of instruction and foreperiod on starting-point variability parallel the effects of instruction and foreperiod on behavioral accuracy (a drop in accuracy on short-foreperiod trials in the accuracy condition but not in the speed condition). These results support the idea that the observed speed-accuracy trade-off between short and long-foreperiod trials in the accuracy condition was due to premature sampling on a proportion of the short-foreperiod trials. Interestingly, this proportion of premature-sampling trials may also be responsible for a part of the observed decrease in the nondecision component T_{er} : on average, sampling (evidence accumulation) starts earlier on short-foreperiod trials than on long-foreperiod trials (when subjects always await the onset of the stimulus), resulting in a shorter encoding phase. However, this account cannot explain why perceptual sensitivity is improved on short-foreperiod trials (Correa et al., 2005; Rolke & Hofmann, 2007), indicating that there must be an additional, effective, shortening of encoding time.

The results from Experiment 1 strongly suggest that increased temporal certainty does not affect the decision process itself, but instead speeds up nondecision processes, consistent with our literature review. However, based on the diffusion-model analysis alone, it cannot be determined whether the shortening of the nondecision component reflects a speeding of stimulus encoding or response execution, or both.

Experiment 2

Besides the foreperiod paradigm, the effects of temporal expectation on task performance have been studied extensively with the temporal-cueing paradigm. In Experiment 2, we examined whether our conclusion that temporal certainty in the fixed-foreperiod paradigm affects mainly nondecision processes can be generalized to the temporal-cueing paradigm. The temporal-cueing paradigm is comparable to the variable-foreperiod paradigm in the sense that the foreperiod varies from trial to trial, but has the additional feature that the warning signal (cue) predicts the foreperiod duration with a large degree of certainty. In Experiment 2, these temporal cues were presented in the context of a simple-RT task, requiring rapid target detection; in choice-RT tasks temporal-cueing effects are generally absent, presumably because target discrimination interferes with the processing of the cue (Correa et al., 2004). On each trial, a cue predicted with a validity of 75% whether the cue-target interval was 400 or 1400 ms. Target brightness (bright or dim) was also

⁷ For this model, there were also significant effects of speed vs. accuracy instruction on boundary separation, of word type on drift rate, and of foreperiod duration on nondecision time.

varied between trials. We expected to find a cue-validity effect on RT for the short cue-target interval, because of differences in temporal preparation at the moment the target appears. A similar validity effect is generally not observed for the long cue-target interval (Correa et al., 2004; Coull & Nobre, 1998), because subjects have time to reorient their attention to the long cue-target interval after they realize that a cue indicating the short-cue interval is invalid (Correa et al., 2004; Karlin, 1959).

We analyzed the data using the shifted-Wald model (Wald, 1947; Figure 5), a model based on the Wald distribution, which represents the density of the first passage times of a Wiener diffusion process toward a single absorbing boundary. The shifted-Wald model conceptualizes the decision process as a single-boundary diffusion process, and successfully accounts for RT distributions in paradigms in which there is only a single response boundary, such as simple-RT tasks (Luce, 1986, pp. 51–57), go/no-go tasks (Heathcote, 2004; Schwarz, 2001; see Carpenter & Williams, 1995, for a comparable, ballistic approach).

This shifted-Wald distribution can be characterized by three parameters that correspond closely to the three main parameters of the diffusion model: the drift rate of the diffusion process (γ), the separation between the starting point of the diffusion process and the absorbing barrier (i.e. the decision threshold; α), and a parameter that shifts the entire RT distribution and thus quantifies the time needed for nondecision processes (θ).




Figure 5. An illustration of the shifted-Wald model. The parameters are: a = distance between the starting point and the decision threshold, γ = drift rate, and θ = mean nondecision time.

Method

Participants. Sixteen students participated (14 women; aged 19-28 years; mean age = 21.8; all native Dutch speakers). Each participant completed one session of approximately 100 minutes in return for 13 euros or course credits.

Design and procedure. All stimuli were presented in the center of the screen on a black background. Each trial started with a white fixation point that was displayed for a quasi-random duration between 500-1500 ms (in steps of 200 ms). This was followed by the 50-ms presentation

of a grey, short (visual angle = $0.6^\circ \times 0.2^\circ$) or long ($1.4^\circ \times 0.2^\circ$) horizontal rectangular bar in the center of the screen. This cue provided information about the subsequent cue-target interval. Specifically, the short bar indicated that the target would appear early (i.e., cue-target interval = 400 ms) on 75% of the trials (valid cue) and late (cue-target interval = 1400 ms) on 25% of the trials (invalid cue). The long bar indicated that the target would appear late (cue-target interval = 1400 ms) on 75% of the trials (valid cue) and early (cue-target interval = 400 ms) on 25% of the trials (invalid cue). The cue was followed by a blank screen for the remainder of the cue-target interval (350 ms or 1300 ms). Then the target, a white (bright) or dark grey (dim) circle (visual angle = 1.0°) was presented for 100 ms, followed by a blank screen until the participant made a response. Then the next trial began. When no response was registered within 2 s of target onset, the message “You have not responded” was presented for 1 s. If a response with $RT < 100$ ms was registered, the message “Too fast! Wait with responding until the circle appears” was presented for 2 s.

Before the start of the experiment, participants were dark-adapted for 5 min in a room sealed from light. Dark adaptation increases the difference in RTs between bright and near-threshold stimuli (cf. Jaśkowski, Kurczewska, Nowik, van der Lubbe, & Verleger, 2007). The actual experiment started with 16 practice trials, followed by 16 blocks of 112 trials. Each block contained 28 trials with each combination of cue-target interval (short, long) and target brightness (bright, dim), 7 (25%) of which were invalidly cued. There was a 1-minute break between blocks and a 5-minute break halfway through the experiment. Participants were instructed to press the space bar as soon as they detected the target. They were encouraged to use the cue to optimize performance. At the end of each block the mean RT and the proportion of correct responses (= non-anticipations) appeared on the screen.

Shifted-Wald-model analysis. To assess the processing components that are affected by temporal uncertainty, the parameters γ , α , and θ were left free to vary as a function of cue validity and cue-target interval. In addition, the parameters γ and θ were free to vary as a function of target brightness, but α was not, reflecting the notion that subjects cannot instantaneously adjust the decision threshold once the (dim or bright) target is presented.

To reduce the impact of a few very short and long reaction times on the parameter estimates, we fitted to the data a mixture of the shifted-Wald distribution and a uniform distribution of response contaminants (e.g., Ratcliff & Tuerlinckx, 2002; Zeigenfuse & Lee, 2010). The uniform distribution of contaminants ranged from 100 ms to 1000 ms – the RTs below and above these boundaries were excluded from analysis.

Participant heterogeneity in the parameter estimates for the mixture-shifted-Wald model was taken into account using hierarchical Bayesian modeling (e.g., Farrell & Ludwig, 2008; Gelman & Hill, 2007; Rouder, Lu, Speckman, Sun, & Jiang, 2005; Rouder, Sun, Speckman, Lu, & Zhou, 2003; Shiffrin, Lee, Kim, & Wagenmakers, 2008). Hierarchical Bayesian methods reduce the variability in the recovered parameters and produce more accurate parameter estimates than single-level maximum likelihood estimation (Farrell & Ludwig, 2008; Rouder et al., 2005). The hierarchical Bayesian approach assumes that the parameters of individual participants are drawn

from group-level distributions that specify how the individual parameters are distributed in the population. The group-level distributions thus define the between-subjects variations of the parameters and can themselves be characterized by a set of parameters. One of the benefits of hierarchical modeling is that knowledge from the group-level distribution serves to shrink noisy estimates for individual participants to less extreme values.

In the Bayesian hierarchical model, individual parameters γ_i , α_i , θ_i –for the shifted-Wald distribution– and π_i –the mixture proportion– are assumed to come from group-level distributions with means μ_γ , μ_α , μ_θ and μ_π . These distributions were assumed to be normal, both for the shifted-Wald parameters and for the probit-transformed mixture proportion. The mean and standard deviation of the group-level distributions needed to be assigned prior distributions; these distributions were uninformative in the sense that the posterior distributions were not noticeably influenced by increasing or decreasing the width of the prior distributions⁸.

Parameter estimation for the mixture-shifted-Wald model was carried out by means of Markov chain Monte Carlo (MCMC) sampling in the WinBUGS program (Lunn, Thomas, Best, & Spiegelhalter, 2000; Lunn, Spiegelhalter, Thomas, & Best, 2009). The WinBUGS code that was used to fit the model can be found at www.ejwagenmakers.com. For reasons of speed and robustness, the likelihood function for the mixture between uniform and shifted-Wald distributions was coded separately and made available via the WinBUGS Development Interface (WBDev; e.g., Wetzels, Lee, & Wagenmakers, 2010). The MCMC sampling used three separate chains; each chain had a burn-in of 20,000 iterations, after which 20,000 further samples were drawn with a thinning factor of 10. This left 2,000 samples per chain for a total of 6,000 samples for each posterior distribution. Visual inspection and calculation of the R-hat statistic (Gelman & Rubin, 1992) confirmed that the three chains had converged to the same distribution (i.e., for all group-level parameters, R-hat = 1.00).

The results showed that the probability of a response contaminant was very low; for the group-level mean parameter, the mode of the posterior distribution was only .004. Nevertheless, inclusion of the contaminant distribution had a pronounced effect on the estimated nondecision time θ – without the contaminant distribution, θ was estimated to be implausibly low. Note that in the absence of a contaminant distribution, the entire distribution of θ has to be lower than the minimum observed RT. Thus, the inclusion of the contaminant distribution made the model more robust to misspecification due to the presence of outliers, even though the probability of observing an outlier was very low.

⁸ Because of numerical underflow errors for the likelihood, the Wald distribution does not allow one to use completely uninformative prior distributions. For this reason, we used prior distributions that were uninformative within a range that is plausible for data from a simple-RT task. See www.ejwagenmakers.com for a precise specification of the prior distribution, the model code, and the model output.

Results

Behavioral results. RTs shorter than 100 ms and longer than 1000 ms were excluded from analysis, which resulted in the exclusion of 1.8% of the trials. The proportion of trials on which participants failed to respond was 0.6% (<1.5% misses for all participants). The number of misses was larger on invalid than on valid trials, $F(1,15)=2.3$, $p=.002$. Figure 6 shows mean RT as a function of cue-target interval, cue validity, and target brightness. RTs were faster for bright targets than for dim targets (282 ms vs. 351 ms; $F(1,15) = 274.6$, $p < 0.001$); and faster for the long cue-target interval than for the short cue-target interval (307 ms vs. 326 ms; $F(1,15) = 13.2$, $p = 0.002$). Furthermore, as expected, RTs were faster on validly cued than on invalidly cued trials (308 ms vs. 324 ms; $F(1,15) = 29.5$, $p < 0.001$), indicating that participants used the cues to optimize their performance. As expected, the effect of cue validity was much larger for the short cue-target interval (28 ms) than for the long cue-target interval (6 ms; $F(1,15) = 12.2$, $p = 0.003$). This 2-way interaction effect was qualified by a significant 3-way interaction ($F(1,15) = 13.6$, $p = 0.002$), indicating that the cue-validity effect was the largest when the cue-target interval was short (i.e. when participants could not reorient their attention on invalidly cued trials) and the target was dim (i.e. when there was room for improvement in RT performance).




Figure 6. Mean RT in Experiment 2 as a function of cue-target interval (CTI), cue validity, and target brightness (bright, dim).

Experimental effects on the shifted-Wald model parameters. The results for the parameters of substantive interest are shown in Table 3. As expected, drift rate γ was higher for bright than for dim targets, $F(1,15) = 76.91$, $p < 0.001$, but was not affected by cue-target interval and cue validity ($ps > 0.8$). Decision threshold α tended to be lower for the long cue-target interval than for the short cue-target interval, $F(1,15) = 3.711$, $p = 0.07$. Importantly, cue validity only affected the nondecision component θ . The θ parameter was significantly smaller on validly cued trials than on invalidly cued trials, $F(1,15) = 7.7$, $p = 0.01$. Furthermore, there was an interaction between cue validity and cue-target interval, $F(1,15) = 24.4$, $p < 0.001$, indicating that the validity effect on θ was present on trials with a short cue-target interval (mean = 24 ms) but not on trials with a long

cue-target interval (mean = -3 ms), mimicking the validity effects on RT (28 ms and 6 ms, respectively). Finally, θ was smaller for bright targets than dim targets, $F(1,15) = 32.6$, $p < 0.001$.

Table 3. Parameter estimates for the fit of the mixture-shifted-Wald model in Experiment 2. The upper half of the table reports the averages of the individual parameter values (SD in parentheses), and the lower half of the table reports the posterior means of the group-level normal distributions from which the individual parameters were drawn. θ = non-decision time (in seconds) comprising stimulus encoding and response execution; α = decision threshold; γ = drift rate.

		Short cue-target interval		Long cue-target interval	
Parameter		Invalid cue	Valid cue	Invalid cue	Valid cue
Averages of the individual parameter values	θ (dim)	.156 (.153)	.131 (.066)	.152 (.068)	.156 (.075)
	θ (bright)	.118 (.036)	.095 (.050)	.122 (.040)	0.123 (.054)
	γ (dim)	6.26 (1.12)	6.44 (1.50)	6.40 (1.31)	6.31 (1.10)
	γ (bright)	8.01 (1.84)	7.58 (1.61)	7.79 (2.06)	8.01 (1.85)
	α	1.37 (.26)	1.34 (.26)	1.17 (.34)	1.14 (.34)
Means of the group-level parameter distributions	μ_θ (dim)	.156 (.016)	.132 (.019)	.152 (.020)	.156 (.021)
	μ_θ (bright)	.118 (.011)	.095 (.015)	.123 (.012)	.124 (.015)
	μ_γ (dim)	6.28 (.35)	6.47 (.43)	6.44 (.41)	6.33 (.33)
	μ_γ (bright)	8.01 (.55)	7.59 (.47)	7.81 (.59)	8.01 (.53)
	μ_α	1.37 (0.09)	1.34 (.09)	1.17 (0.11)	1.14 (.10)

Note: dim = dim target; bright = bright target

Model fit. Figure 7 shows the mean observed .1, .3, .5, .7 and .9 quantile RTs in each condition, as well as those predicted by the shifted-Wald model. The model provided a generally good fit to the empirical RT quantiles; the largest difference between the observed and model-predicted quantile RTs was 15 ms, and the average difference was 6 ms. For the short cue-target interval, all five quantile RTs associated with bright and dim targets were shorter on validly cued trials than on invalidly cued trials. This cue-validity effect was less pronounced or absent for the long cue-target interval. To examine in more detail the cue-validity effect at different points of the RT distribution, we plotted the observed and predicted cue-validity effect for each of the five RT quantiles as a function of response speed (the average of the quantile RTs in the validly cued trials and invalidly cued trials). The resulting delta plots are shown in Figure 8. For the bright targets, the cue-validity effect at the short cue-target interval was rather constant across the five RT quantiles, which suggests that cue validity mainly affected the nondecision time (i.e., parameter θ of the shifted-Wald model). For the dim targets, however, the cue-validity effect at the short cue-target interval increased with increasing RTs. To assess whether this increase was significant, we subjected the cue-validity effect at each quantile to a linear-regression analysis with mean quantile RT and a constant as explanatory factors, separately for each participant (Burle, van den Wildenberg, & Ridderinkhof, 2005; De Jong, Liang, & Lauber, 1994). We then tested whether the average regression coefficient of mean quantile RT (i.e., the slope of the delta plot) was

significantly different from 0, using a one-sample t -test. This test just reached significance (mean regression coefficient = 0.12; SD = 0.21; $t(15) = 2.2$, $p = 0.044$), suggesting that, for the dim targets, part of the cue-validity effect was attributable to an effect on the decision process. Because target brightness was varied on a trial-by-trial basis, processing of bright and dim targets could differ in drift rate but not in decision threshold. Therefore, the increasing validity effect with increasing RT for the dim targets was likely due to a cue-validity effect on drift rate (i.e., parameter γ of the shifted-Wald model).

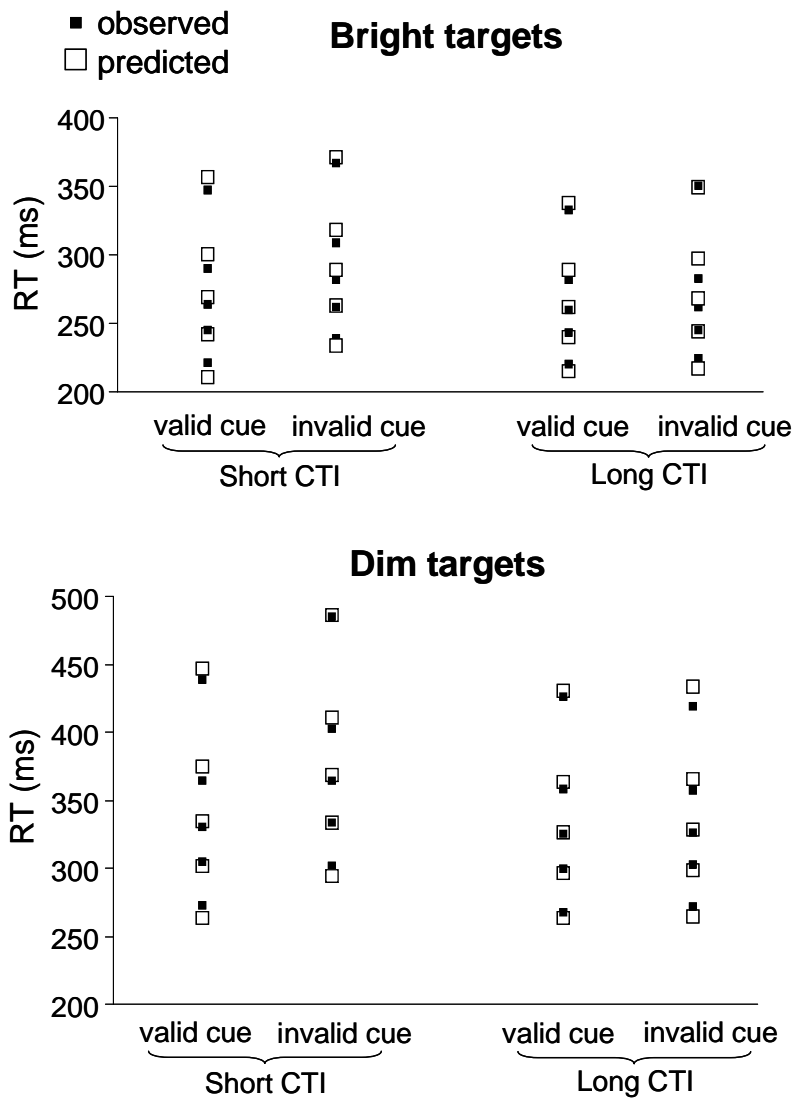


Figure 7. The observed and predicted .1, .3, .5, .7 and .9 quantile RTs in Experiment 2 as a function of target brightness (bright, dim), cue-target interval (CTI) and cue validity.




Figure 8. Observed and predicted delta plots showing the cue-validity effect on RT as a function of mean quantile RT, cue-target interval (CTI) and target brightness (bright, dim) in Experiment 2.

Discussion

We applied the shifted-Wald model to the data from a simple-RT experiment in which the targets were preceded by a cue that validly or invalidly indicated the cue-target interval: short (400 ms) or long (1400 ms). As expected we found a substantial cue-validity effect on RT for the short cue-target interval but not for the long cue-target interval, because the participants had time to reorient their attention to the long cue-target interval after they realized that a cue indicating the short-cue interval was invalid (Correa et al., 2004; Karlin, 1959). The model analysis provided useful evidence regarding the source of the cue-validity effect: Cue validity significantly affected the θ parameter, but not the parameters of the decision process, γ and α . Indeed, the effects of cue validity on the estimated duration of nondecision processes were very similar in size to the cue-validity effects on RT, both at the short and the long cue-target interval. The delta plots showed a somewhat more complicated pattern: in one condition (short cue-target interval, bright targets) the cue-validity effect was relatively constant across the RT distribution, suggesting that increased temporal certainty decreased the duration of the nondecision component. In another condition (short cue-target interval, dim targets) the cue-validity effect showed an increase across RT bins, suggesting that increased temporal certainty increased the rate of evidence accumulation. However, the absence of an effect of cue validity on the γ parameter, and the nonsignificant interactions between cue validity and the other variables on γ , suggest that this effect was relatively minor. Thus, the cue-validity effect was largely accounted for by a change in the nondecision component θ .

Another interesting finding in Experiment 2 was the trend-level effect of cue-target interval on estimated decision threshold (parameter α ; $p = .07$): the decision threshold was lower for the long cue-target interval than for the short cue-target interval. At first blush, this finding seems inconsistent with the absence of an effect of foreperiod on boundary separation in Experiment 1. But on closer thought, the two experiments are rather different in terms of the effect of warning interval. In the simple-RT task of Experiment 2, participants could substantially lower the decision threshold if the target had not appeared after the short cue-target interval: there was no more uncertainty about the timing of the target, the long cue-target interval was relatively short (1400 ms)

and hence easier to anticipate, and the identity of the response was known. Accordingly, in Experiment 2 participants responded faster when the cue-target interval was long compared to when it was short. Conversely, in the choice-RT task of Experiment 1, the long foreperiod was relatively long (2700 ms) and hence harder to anticipate, and there was always the risk of making choice errors if the boundary separation was set too small. Accordingly, in this experiment participants responded slower when the foreperiod was long compared to when it was short. Therefore, the effects of warning interval on temporal certainty went in opposite directions in the two experiments, which mirrors the opposing effects of long and short foreperiods on temporal certainty in fixed and variable-foreperiod paradigms (Bertelson, & Tisseyre, 1968; Vallesi, McIntosh, & Stuss, 2009).

General Discussion

Preparing the system to respond to an upcoming stimulus is energy-consuming and maintaining such a state of readiness, because stimulus onset time is uncertain, can be experienced as an aversive state (Gottsdanker, 1975; Näätänen, 1972). This indicates the importance of using cues to predict stimulus onset and time the system's preparation accordingly. We conducted two experiments, using the foreperiod paradigm (Experiment 1) and the temporal-cueing paradigm (Experiment 2), to assess which components of information processing are speeded when subjects use such temporal cues to reduce uncertainty. The results from these two experiments were consistent: temporal certainty affected the duration of nondecision processes but had little effect on the two critical components of the decision process—decision-threshold setting and the rate of evidence accumulation.

Our findings are consistent with two previous studies that examined correct-RT distributions (Leth-Steensen, 2009) and speed-accuracy tradeoff functions (Bausenhart, et al., 2010) in the foreperiod paradigm, and found that manipulations of foreperiod mainly shifted these distributions while having very little effect on their shapes. The novel contributions of our study are that we analyzed temporal certainty effects in both the foreperiod paradigm and the temporal-cueing paradigm, using sequential-sampling models of decision making that took into account response accuracy and RT distributions on correct and error trials. Our findings therefore provide the strongest evidence to date that temporal certainty effects on simple-RT and choice-RT reflect a change in the duration of nondecision processes, not changes in the decision process.

Although our results cannot distinguish between effects of temporal certainty on encoding and motor processes, the literature suggests that the duration of both types of processes is affected. Temporal certainty modulates many aspects of perception (Nobre et al., 2007), including perceptual sensitivity in the foreperiod paradigm and temporal-cueing paradigm (Correa et al., 2005; Rolke, 2008; Rolke & Hofmann, 2007), and the duration of perceptual processing in a clock paradigm (Seifried et al., 2010). Temporal certainty also modulates various aspects of motor preparation (Davranche et al. 2007; Miniussi, Wilding, Coull, & Nobre, 1999; Riehle et al., 1997) and decreases the duration of motor preparation and execution, although effect sizes are small (Tandonnet et al.,

2003, 2006). Näätänen's (1971) model suggests that subjects may use temporal cues to anticipate the arrival of the imperative stimulus by increasing the level of 'motor readiness', such that the time to reach the 'motor action limit' is reduced⁹. This model is supported by findings that increased temporal certainty reduces the force generated to execute the response (Mattes & Ulrich, 1997) and the activation of the corresponding primary motor cortex (Tandonnet et al., 2006).

Our conclusions stand in sharp contrast with those reached by Hackley (2009) on the basis of a review of ERP studies with the foreperiod paradigm. Hackley's main argument against an encoding account is that temporal certainty has little or no effect on the latency of the P1 and N1, two early perceptual brain potentials (reviewed in Correa et al. 2006; see Hackley et al., 2007 for a significant but very small effect), and the latency of the N2pc (Hackley et al., 2007), an electrophysiological index of the allocation of spatial attention. A possible explanation of these findings is that the latency of early ERP components is not a reliable index of the duration of task-relevant encoding processes. For example, because of the parallel organization of the visual system, the processes underlying these ERP components may not lie in the pathway that determines the RT. Hackley et al. (2007) reject this hypothesis with the argument that N2pc latency has been found to correlate highly with RT in a number of studies. However, a problem with this argument is that these N2pc-RT correlations were found in experiments with an important spatial component (i.e., requiring the N2pc process to perform the task), whereas the non-significant effect of foreperiod on N2pc latency was found in a study in which stimulus location played a negligible role (Hackley et al., 2007). Another possible explanation, suggested by Hackley (2009), is that increased temporal certainty leads to increased visual-cortex activation in response to the visual imperative stimulus, as reflected in increased P1 and N1 amplitudes (i.e., a nonchronometric change), changes that then propagate forward to produce a greater speed of subsequent encoding processes. This proposal is consistent with P1/N1 amplitude effects of temporal certainty (Correa et al., 2006) and with our recent proposal regarding the temporal locus of the accessory stimulus effect (Jepma et al., 2009).

Hackley's (2009) argument against the motor preparation account is that foreperiod duration has little or no effect on the interval between LRP onset and the overt response (e.g., Hackley et al., 2007; Müller-Gethmann, Ulrich, & Rinkenauer, 2003). However, a limitation of the LRP double-subtraction measure is that it is blind to the respective contributions of each individual motor cortex. To address this limitation, Tandonnet et al. (2003, 2006) used Laplacian-transformed ERP waveforms to obtain separate estimates of the ipsilateral and contralateral motor cortex response, and found that increased temporal certainty decreased the duration of motor preparation of the responding hand (by 25 ms and 18 ms in the 2003- and 2006-studies, respectively). Because the same data showed no reliable effect of foreperiod on the (Laplacian and monopolar) LRP-to-response interval, Tandonnet and colleagues suggested that the effect of interest (i.e., timing of preparation of the responding hand) may be masked in the LRP by preparatory effects on activation in the ipsilateral, non-involved motor cortex. It is unclear whether similar subtle latency effects

⁹ This change in the distance to the motor threshold, which is not modeled in the sequential-sampling models used here, must be distinguished from response-threshold changes in the decision process.

were also present in the constituent ERP waveforms in the LRP studies reviewed by Hackley (2009). In any case, future ERP research needs to resolve the issue of how the onset of motor preparation processes should be measured.

Our conclusion that temporal certainty has little effect on the decision threshold might appear incompatible with occasional findings (including in Experiment 1) that temporal certainty causes a speed-accuracy tradeoff, a phenomenon that has been taken as diagnostic of decision-threshold modulations. However, on the basis of results in Experiment 1, we have suggested that a speed-accuracy tradeoff can be explained by an alternative hypothesis: when temporal certainty is high and the moment of stimulus onset is relatively easy to anticipate, subjects may engage in premature sampling of stimulus information on a proportion of the trials. Such premature sampling will lead to faster but less accurate responses because subjects will start with sampling noise. Thus, behavioral speed-accuracy trade-offs may be explained not only by changes in decision threshold but also by changes in premature sampling, which is an interesting topic for future research.

Our results provide important clues about the components of information processing that are speeded when people use temporal cues to anticipate the onset of an imperative stimulus. But ultimately we also need to understand the neural mechanisms underlying this voluntary control of temporal expectation. Neuroimaging and patient studies have suggested an important role for prefrontal structures in controlling temporal expectation in the foreperiod paradigm (Hackley et al., 2009) and the temporal-cueing paradigm (Coull & Nobre, 1998; Triviño, Correa, Arnedo, & Lupiáñez, 2010). Other studies have identified norepinephrine as a key neuromodulator underlying temporal certainty effects (Coull, Nobre, & Frith, 2001; Witte & Marrocco, 1997), consistent with the finding that the firing rate of locus coeruleus neurons increases during the warning interval in the foreperiod paradigm (Yamamoto & Ozawa, 1989). It will be a challenge for future studies to determine the exact mechanisms by which prefrontal structures and/or the locus coeruleus-norepinephrine system control the duration of nondecision processes as a function of degree of temporal expectation.

Chapter 8

Summary

The studies described in this thesis address a range of topics related to arousal, exploration, temporal attention, and the locus coeruleus-norepinephrine (LC-NE) system.

Chapters 2 and 3 reported two studies that investigated the central tenets of the adaptive gain theory of LC function in human participants. According to the adaptive gain theory, the LC-NE system is critical for optimizing behavioral performance by regulating the balance between exploitative and exploratory control states. However, crucial empirical tests of this theory in humans have been lacking. Recent studies have suggested that, under controlled circumstances, pupil diameter provides an indirect index of LC activity. Inspired by this evidence, we used pupillometry to examine whether the relationship between pupil diameter, control state and task-related utility is consistent with the assumptions of the adaptive gain theory (Chapter 2). The results from this study provide indirect evidence for the idea that LC mode regulates the exploration-exploitation trade-off, and that transitions between LC modes are driven by assessments of task-related utility.

To provide a more direct test of the adaptive gain theory's assumption that the tonic LC mode promotes an exploratory control state, we investigated the effects of a pharmacological manipulation of the LC-NE system on measures of exploration and task (dis)engagement (Chapter 3). Contrary to predictions of the adaptive gain theory, we found no evidence that the elevated NE levels induced by our pharmacological manipulation were associated with increased task disengagement or exploratory behavior in our experimental tasks. In contrast to the results reported in Chapter 2, these results suggest that the LC-NE system may not be involved in the regulation of the exploration-exploitation trade-off in humans, at least not within the context of a single task. These results leave open the possibility that the LC-NE system is involved in random exploration exceeding the current task context. Possible reasons for the discrepancy between the conclusions of our pupillometry study and our pharmacological study are that (i) the difference in baseline pupil size preceding exploitative versus exploratory choices did reflect something else than a difference in tonic LC activity; (ii) our pharmacological manipulation did not produce a net increase in tonic NE levels; (iii) our pharmacological manipulation had unexpected effects on the phasic LC response or on other neuromodulatory systems. These and other possible explanations remain to be explored in future studies.

Chapter 4 reported a study on neurocognitive function in patients with dopamine- β -hydroxylase (D β H) deficiency, a rare genetic disorder characterized by the absence of NE in the peripheral and the central nervous system. NE levels in D β H-deficient patients can be restored by the administration of a synthetic precursor of NE. Informal clinical observations suggest that D β H-deficient patients do not have obvious cognitive impairments, but systematic studies on neurocognitive function in D β H deficiency have been lacking. We tested five D β H-deficient patients and a healthy control group on a comprehensive neurocognitive task battery; the patients were tested once on and once off medication. We found that the patients' neurocognitive performance was largely spared, even when they were off medication, which is striking given the

important role of NE in normal cognition. The results from this study suggest that other neuromodulators have taken over the function of NE in the brains of these patients.

Exploratory behavior in both animals and humans is often driven by curiosity. Chapter 5 reported a study on the neural correlates of human perceptual curiosity, using functional magnetic resonance imaging (fMRI). We found that the induction of perceptual curiosity, through the presentation of ambiguous visual input, activated brain regions sensitive to conflict and arousal. The relief of curiosity, through visual disambiguation, activated regions of the striatum that are involved in reward processing. In addition, the relief of curiosity was associated with increased hippocampal activation and enhanced incidental memory. Our results suggest that perceptual curiosity evoked by ambiguous visual input is an aversive condition. In addition, our results suggest that termination of this condition is rewarding and promotes learning. The results from this study provide compelling neurobiological evidence for Berlyne's (1954, 1966) classic psychological theory of curiosity.

Chapter 6 and 7 reported several experiments investigating the effects of accessory stimuli (Chapter 6) and temporal preparation (Chapter 7) on information processing, using scalp electrophysiology and sequential-sampling models of decision making. The results from these experiments suggest that both accessory stimulation and temporal preparation affect nondecision processes (stimulus encoding and/or response preparation) but have little effect on the two main components of the decision process: decision-threshold setting and the rate of evidence accumulation.

Attention and decision making are essential components of cognition that are involved in almost all aspects of behavior. In the last decades, reward-based decision making and spatial attention have been topics of extensive investigation among cognitive neuroscientists. However, other less specific dimensions of decision making and attention, such as arousal, exploration and temporal attention, have received less recognition and interest. The research described in this thesis aims to contribute to a better understanding of these constructs. Taken together, the studies reported in this thesis suggest that arousal, exploration and temporal attention are closely related, which is likely due to a shared neural basis.

References

- Abercrombie, E.D., & Jacobs, B.L. (1987). Single-unit response of noradrenergic neurons in the locus coeruleus of freely moving cats. II. Adaptation to chronically presented stressful stimuli. *Journal of Neuroscience*, 7, 2844-2848.
- Aghajanian, G.K., Cedarbaum, J.M., & Wang, R.Y. (1977). Evidence for norepinephrine-mediated collateral inhibition of locus coeruleus neurons. *Brain Res*, 136, 570-7.
- Alvarez, V.A., Chow, C.C., Van Bockstaele, E.J., & Williams, J.T. (2002). Frequency-dependent synchrony in locus ceruleus: role of electrotonic coupling. *Proc Natl Acad Sci USA*, 99, 4032-4036.
- Anderson, B.D.O., & Moore, J.B. (1979). *Optimal Filtering*. Englewood Cliffs, NJ: Prentice-Hall.
- Arango, V., Underwood, M.D., & Mann, J.J. (1996). Fewer pigmented locus coeruleus neurons in suicide victims: preliminary results. *Biol Psychiatry*, 39, 112-20.
- Ashburner, J., & Friston, K.J. (2000). Voxel-based morphometry-the methods. *Neuroimage*, 11, 805-21.
- Aston-Jones, G. (2004). Locus coeruleus, A5 and A7 noradrenergic cell groups. In G. Paxinos (Ed.), *The rat nervous system* (3rd edition, pp. 259–294). San Diego: Elsevier Academic Press.
- Aston-Jones, G., & Bloom, F.E. (1981a) Activity of norepinephrine-containing locus coeruleus neurons in behaving rats anticipates fluctuations in the sleep-waking cycle. *Journal of Neuroscience*, 1, 876-86.
- Aston-Jones, G., & Bloom, F.E. (1981b). Norepinephrine-containing locus coeruleus neurons in behaving rats exhibit pronounced responses to non-noxious environmental stimuli. *Journal of Neuroscience*, 1, 887-900.
- Aston-Jones, G., Foote, S.L., & Bloom, F.E. (1984). Anatomy and physiology of locus coeruleus neurons: functional implications. In: Ziegler, M., Lake, C.R. (eds). *Frontiers of Clinical Neuroscience: Vol. 2, Norepinephrine*. Williams and Wilkins: Baltimore. pp 92-116.
- Aston-Jones, G., Iba, M., Clayton, E., Rajkowski, J., & Cohen, J.D. (2007). Locus coeruleus and regulation of behavioral flexibility and attention: Clinical implications. In Ordway, G.A., Schwartz, M., & Frazer, A. (eds.), *Norepinephrine: Neurobiology and Therapeutics for the 21st Century*, (pp. 196-235). Cambridge University Press: Cambridge, UK.
- Aston-Jones, G., Rajkowski, J., Kubiak, P., Valentino, R.J., & Shipley, M.T. (1996). Role of the locus coeruleus in emotional activation. *Prog Brain Res*, 107, 379-402.
- Aston-Jones, G., & Cohen, J.D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annual Review of Neuroscience*, 28, 403-450.
- Aston-Jones, G., Ennis, M., Pieribone, V.A., Nickell, W.T., & Shipley, M.T. (1986). The brain nucleus locus coeruleus: restricted afferent control of a broad efferent network. *Science*, 234, 734-737.
- Aston-Jones, G., Rajkowski, J., & Cohen, J. (2000). Locus coeruleus and regulation of behavioral flexibility and attention. *Progress in Brain Research*, 126, 165-182.
- Aston-Jones, G., Rajkowski, J., & Cohen, J. (1999). Role of locus coeruleus in attention and behavioral flexibility. *Biological Psychiatry*, 46, 1309-20.

- Aston-Jones, G., Rajkowski, J., & Kubiak, P. (1997). Conditioned responses of monkey locus coeruleus neurons anticipate acquisition of discriminative behavior in a vigilance task. *Neuroscience*, 80, 697-715.
- Aston-Jones, G., Rajkowski, J., Lu, W., Zhu, Y., Cohen, J.D., & Morecraft, R.J. (2002). Prominent projections from the orbital prefrontal cortex to the locus coeruleus in monkey. *Society for Neuroscience Abstracts*, 28, 86-89.
- Aston-Jones, G., Rajkowski, J., Kubiak, P., Alexinsky, T. (1994). Locus coeruleus neurons in monkey are selectively activated by attended cues in a vigilance task. *Journal of Neuroscience*, 14, 4467-80.
- Baayen, R.H., Piepenbrock, R., & Gulikers, L. (1995). CELEX-2. CD-ROM available from the Linguistic Data Consortium, University of Pennsylvania.
- Bausenhart, K.M., Rolke, B., & Ulrich, R. (2008). Temporal preparation improves temporal resolution: evidence from constant foreperiods. *Perception and Psychophysics*, 70, 1504-1514.
- Bausenhart, K.M., Rolke, B., Seibold, V.C., & Ulrich, R. (2010). Temporal preparation influences the dynamics of information processing: Evidence for early onset of information accumulation. *Vision Research*, 50, 1025-34.
- Beatty, J., & Wagoner, B.L. (1978). Pupillometric signs of brain activation vary with level of cognitive processing. *Science*, 199, 1216-8.
- Beckmann, C.F., Jenkinson M., & Smith, S.M. (2003). General multilevel linear modeling for group analysis in fMRI. *Neuroimage* 20, 1052–63.
- Behrens TE, Woolrich MW, Walton ME, & Rushworth MF (2007). Learning the value of information in an uncertain world. *Nat Neurosci* 10: 1214-1221.
- Benzi, R., Sutera, A., & Vulpiani, A. (1981). The mechanism of stochastic resonance. *Journal of Physics A*, 14, L453–L457.
- Berlyne, D.E. (1954). A theory of human curiosity. *Br J Psychol*, 45, 180-91.
- Berlyne, D.E. (1960). *Conflict, Arousal, and Curiosity*, New York: McGraw Hill.
- Berlyne, D.E. (1966). Curiosity and exploration. *Science*, 153, 25-33.
- Berlyne, D.E., & Borsa, D.M. (1968). Uncertainty and the orientation reaction. *Percept Psychophys*, 3, 77-79.
- Berlyne, D.E., & Normore, L.F. (1972). Effects of prior uncertainty on incidental free recall. *J Exp Psychol*, 96, 43-48.
- Bernstein, I.H. (1970). Can we see and hear at the same time? Intersensory facilitation of reaction time. *Acta Psychologica*, 33, 21-35.
- Bernstein, I.H., Chu, P.K., Briggs, P., & Schurman, D.L. (1973). Stimulus intensity and foreperiod effects in intersensory facilitation. *Quarterly Journal of Experimental Psychology*, 25, 171-81.
- Bernstein, I.H., Clark, M.H., & Edelman, B.A. (1969a). Effects of an auditory signal on visual reaction time. *Journal of Experimental Psychology*, 80, 567-569.
- Bernstein, I.H., Clark, M.H., & Edelman, B.A. (1969b). Intermodal effects in choice reaction time. *Journal of Experimental Psychology*, 81, 405-407.
- Bernstein, I.H., Rose, R.G., & Ashe, V. (1970). Energy integration in intersensory facilitation. *Journal of Experimental Psychology*, 86, 196-203.
- Bernstein, P.S., Scheffers, M.K., & Coles, M.G. (1995). "Where did I go wrong?" A psychophysiological analysis of error detection. *Journal of Experimental Psychology. Human Perception and Performance*, 21, 1312-1322.

- Berridge, C.W., & Waterhouse, B.D. (2003). The locus coeruleus-noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. *Brain Research. Brain Research Reviews*, 42, 33-84.
- Bertelson, P. (1967). The time course of preparation. *Quarterly Journal of Experimental Psychology*, 19, 272-279.
- Bertelson, P., & Tisseyre, F. (1968). The time course of preparation with regular and irregular foreperiods. *Quarterly Journal of Experimental Psychology*, 20, 297-300.
- Biaggioni, I., Goldstein, D.S., Atkinson, T., & Robertson, D. (1990). Dopamine-beta-hydroxylase deficiency in humans. *Neurology*, 40, 370-3.
- Biaggioni, I., & Robertson, D. (1987). Endogenous restoration of noradrenaline by precursor therapy in dopamine beta-hydroxylase deficiency. *Lancet*, 2, 1170-2.
- Blenner, J.L., and Yingling, C.D. (1993). Modality specificity of evoked potential augmenting/reducing. *Electroencephalography and Clinical Neurophysiology*, 88, 131-142.
- Bogacz, R., Wagenmakers, E.-J., Forstmann, B.U., & Nieuwenhuis, S. (2010). The neural basis of the speed-accuracy tradeoff. *Trends in Neurosciences*, 33, 10-16.
- Bond, A., & Lader, M. (1974). The use of analogue scales in rating subjective feelings. *Br J Psychol*, 47, 211-218.
- Boorman, E.D., Behrens, T.E., Woolrich, M.W., & Rushworth, M.F. (2009). How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron*, 62, 733-743.
- Borland, R.G., & Nicholson, A.N. (1984). Visual motor co-ordination and dynamic visual acuity. *Br J Clin Pharmacol*, 18(suppl 1), 69S-72S.
- Botvinick, M.M. (2007). Conflict monitoring and decision making: reconciling two perspectives on anterior cingulate function. *Cognitive, Affective & Behavioral Neuroscience*, 7, 356-366.
- Bourdaud, N., Chavarriaga, R., Galan, F., & Millán, Jdel R. (2008). Characterizing the EEG correlates of exploratory behavior. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 16, 549-556.
- Bouret, S., & Sara, S.J. (2004). Reward expectation, orientation of attention and locus coeruleus-medial frontal cortex interplay during learning. *European Journal of Neuroscience*, 20, 791-802.
- Bouret, S., & Sara, S.J. (2005). Network reset: a simplified overarching theory of locus coeruleus noradrenaline function. *Trends Neurosci*, 28, 574-582.
- Brunia, C.H.M., & van Boxtel, G.J.M. (2000). Motor preparation. In J. T. Caccioppo, L. G. Tassinary, and G. G. Berntson (Eds.), *Handbook of Psychophysiology* (2nd ed., pp. 507-532). New York: Cambridge University Press.
- Buchner, A., Rothermund, K., Wentura, D., & Mehl, B. (2004). Valence of distractor words increases the effects of irrelevant speech on serial recall. *Mem Cognit*, 32, 722-31.
- Burle, B., van den Wildenberg, W., & Ridderinkhof, K.R. (2005). Dynamics of facilitation and interference in cue-priming and Simon tasks. *European Journal of Cognitive Psychology*, 17, 619-641.
- Burle, B., Vidal, F., Tandonnet, C., & Hasbroucq, T. (2004). Physiological evidence for response inhibition in choice reaction time tasks. *Brain and Cognition*, 56, 153-164.
- Buser, P., & Imbert, M. (1961). Sensory projections to the motor cortex in cats: a microelectrode study. In W. A. Rosenblith (Ed.), *Sensory communication* (pp. 607-626). Cambridge, Mass., Massachusetts Institute of Technology.

- Bymaster, F.P., Zhang, W., Carter, P.A., Shaw, J., Chernet, E., Phebus, L., et al (2002). Fluoxetine, but not other selective serotonin uptake inhibitors, increases norepinephrine and dopamine extracellular levels in prefrontal cortex. *Psychopharmacology (Berl)*, 160, 353-361.
- Cahill, L. & McGaugh, J.L. (1998). Mechanisms of emotional arousal and lasting declarative memory. *Trends Neurosci*, 21, 294-9.
- Carpenter, R.H.S., & Williams, M.L.L. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature*, 377, 59-62.
- Carter, M.E., Yizhar, O., Chikahisa, S., Nguyen, H., Adamantidis, A., Nishino, S., Deisseroth, K., & de Lecea, L. (2010). Tuning arousal with optogenetic modulation of locus coeruleus neurons. *Nature Neuroscience*, 13, 1526-33.
- Chamberlain SR, Müller U, Blackwell AD, Robbins TW, & Sahakian BJ. (2006). Noradrenergic modulation of working memory and emotional memory in humans. *Psychopharmacology (Berl)*, 188, 397-407.
- Chamberlain, S.R., Müller, U., Blackwell, A.D., Clark, L., Robbins, T.W., & Sahakian, B.J. (2006). Neurochemical modulation of response inhibition and probabilistic learning in humans. *Science*, 311, 861-863.
- Cheshire, W.P. Jr, Dickson, D.W., Nahm, K.F., Kaufmann, H.C., & Benarroch, E.E. (2006). Dopamine beta-hydroxylase deficiency involves the central autonomic network. *Acta Neuropathol*, 112, 227-9.
- Clayton, E.C., Rajkowski, J., Cohen, J.D., & Aston-Jones, G. (2004). Phasic activation of monkey locus coeruleus neurons by simple decisions in a forced choice task. *Journal of Neuroscience*, 24, 9914-9920.
- Cohen, A.F., Ashby, L., Crowley, D., Land, G., Peck, A.W., & Miller, A.A. (1985). Lamotrigine (BW430C), a potential anticonvulsant. Effects on the central nervous system in comparison with phenytoin and diazepam. *Br J Clin Pharmacol*, 20, 619-629.
- Cohen, J.D., Aston-Jones, G., & Gilzenrat, M.S. (2004). A systems-level theory on attention and cognitive control: Guided activation, adaptive gating, conflict monitoring, and exploitation versus exploration. In: Posner MI, editor. *Cognitive neuroscience of attention*. New York: Guilford Press, p. 71–90.
- Cohen, J.D., McClure, S.M., & Yu, A.J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 362, 933-942.
- Coles, M.G., Gratton, G., & Donchin, E. (1988). Detecting early communication: using measures of movement-related potentials to illuminate human information processing. *Biological Psychology*, 26, 69-89.
- Collins, R.P., Litman, J.A., & Spielberger, C.D. (2004). The measurement of perceptual curiosity. *Pers Individ Dif*, 36, 1127–1141.
- Colzato, L.S., Slagter, H.A., Spapé, M.M., & Hommel, B. (2008). Blinks of the eye predict blinks of the mind. *Neuropsychologia*, 46, 3179-83.
- Correa, A., Lupiáñez, J., & Tudela, P. (2005). Attentional preparation based on temporal expectancy modulates processing at the perceptual level. *Psychonomic Bulletin & Review*, 12, 328-334.
- Correa, A., Lupiáñez, J., Madrid, E., & Tudela, P. (2006). Temporal attention enhances early visual processing: a review and new evidence from event-related potentials. *Brain Research*, 1076, 116-128.

- Correa, A., Lupiáñez, J., Milliken, B., & Tudela, P. (2004). Endogenous temporal orienting of attention in detection and discrimination tasks. *Perception & Psychophysics*, 66, 264-278.
- Cosgrove, K.P., Mazure, C.M., & Staley, J.K. (2007). Evolving knowledge of sex differences in brain structure, function, and chemistry. *Biol Psychiatry*, 62, 847-55.
- Coull, J.T., Nobre, A.C., & Frith, C.D. (2001). The noradrenergic alpha2 agonist clonidine modulates behavioural and neuroanatomical correlates of human attentional orienting and alerting. *Cerebral Cortex*, 11, 73-84
- Coull, J.T. (2004). fMRI studies of temporal attention: allocating attention within, or towards, time. *Cognitive Brain Research*, 21, 216-226.
- Coull, J.T., & Nobre, A.C. (1998). Where and when to pay attention: the neural systems for directing attention to spatial locations and to time intervals as revealed by both PET and fMRI. *Journal of Neuroscience*, 18, 7426-7435.
- Counts, S.E., & Mufson, E.J. (2010). Noradrenaline activation of neurotrophic pathways protects against neuronal amyloid toxicity. *J Neurochem*, 113, 649-60.
- Craig, A.D., Reiman, E.M., Evans, A., & Bushnell, M.C. (1996). Functional imaging of an illusion of pain. *Nature*, 384, 258-60
- Cramer, S.C., Weisskoff, R.M., Schaechter, J.D., Nelles, G., Foley, M., Finklestein, S.P., & Rosen, B.R. (1992). Motor cortex activation is related to force of squeezing. *Human Brain Mapping*, 16, 197-205.
- Crawford, J.R., & Garthwaite, P.H. (2006). Comparing patients' predicted test scores from a regression equation with their obtained scores: a significance test and point estimate of abnormality with accompanying confidence limits. *Neuropsychology*, 20, 259-71.
- Crawford, J.R., & Garthwaite, P.H. (2006a). Methods of testing for a deficit in single case studies: Evaluation of statistical power by Monte Carlo simulation. *Cognitive Neuropsychology*, 23, 877-904.
- Crawford, J.R., & Garthwaite, P.H. (2005). Testing for suspected impairments and dissociations in single-case studies in neuropsychology: Evaluation of alternatives using Monte Carlo simulations and revised tests for dissociations. *Neuropsychology*, 19, 318-31.
- Crawford, J.R. & Howell, D.C. (1998). Comparing an individual's test score against norms derived from small samples. *Clin Neuropsychol*, 12, 482-86.
- Critchley, H., Mathias, C.J., & Dolan, R.J. (2000). Congenital dopamine- β -hydroxylase deficiency and schizotypy. *Schizophr Res*, 41, 254.
- Critchley, H.D. (2005). Neural mechanisms of autonomic, affective, and cognitive integration. *J Comp Neurol*, 493, 154-66.
- Critchley, H.D., Tang, J., Glaser, D., Butterworth, B., & Dolan, R.J. (2005). Anterior cingulate activity during error and autonomic response. *Neuroimage*, 27, 885-895.
- Cubells, J.F., & Zabetian, C.P. (2004). Human genetics of plasma dopamine beta-hydroxylase activity: applications to research in psychiatry and neurology. *Psychopharmacology (Berl)*, 174, 463-76.
- Daffner, K.R., Mesulam, M.M., Scinto, L.F., Cohen, L.G., Kennedy, B.P., West, W.C., & Holcomb, P.J. (1998). Regulation of attention to novel stimuli by frontal lobes: an event-related potential study. *Neuroreport*, 9, 787-91.
- Dalton, P., & Spence, C. (2007). Attentional capture in serial audiovisual search tasks. *Perception & Psychophysics*, 69, 422-438.
- Davis, R., & Green, F.A. (1969). Intersensory differences in the effect of warning signals on reaction time. *Acta Psychologica*, 30, 155-167.

- Davranche, K., Tandonnet, C., Burle, B., Meynier, C., Vidal, F., & Hasbroucq, T. (2007). The dual nature of time preparation: neural activation and suppression revealed by transcranial magnetic stimulation of the motor cortex. *European Journal of Neuroscience*, 25, 3766-3774.
- Daw, N.D. & Doya, K. (2006). The computational neurobiology of learning and reward. *Curr Opin Neurobiol*, 16, 199-204.
- Daw, N.D., O'Doherty, J.P., Dayan, P., Seymour, B., & Dolan, R.J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441, 876-479.
- Dayan P & Abbott LF (2001): *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems* (Chapter 9). London: The MIT Press,
- Dayan, P. & Yu, A.J. (2006). Phasic norepinephrine: a neural interrupt signal for unexpected events. *Network*, 17, 335-50.
- De Jong, R. (1991). Partial information or facilitation? Different interpretations of results from speed-accuracy decomposition. *Perception & Psychophysics*, 50, 333-350.
- De Jong, R., Liang, C. C., & Lauber, E. (1994). Conditional and unconditional automaticity: a dual-process model of effects of spatial stimulus response correspondence. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 731-750.
- De Martino, B., Strange, B.A., & Dolan, R.J. (2008). Noradrenergic neuromodulation of human attention for emotional and neutral stimuli. *Psychopharmacology (Berl)*, 197, 127-136.
- Devauges, V. & Sara, S.J. (1990). Activation of the noradrenergic system facilitates an attentional shift in the rat. *Behavioural Brain Research*, 39, 19-28.
- Dolcos, F., & McCarthy, G. (2006). Brain systems mediating cognitive interference by emotional distraction. *J Neurosci*, 26, 2072-9.
- Dostert, P., Benedetti, M.S., & Poggesi, I. (1997). Review of the pharmacokinetics and metabolism of reboxetine, a selective noradrenaline reuptake inhibitor. *Eur Neuropsychopharmacol*, 7 (suppl 1), S23-S35.
- Doya, K. (2002). Metalearning and neuromodulation. *Neural Networks*, 15, 495-506.
- Dreisbach, G., Müller, J., Goschke, T., Strobel, A., Schulze, K., Lesch, K.P., & Brocke, B. (2005). Dopamine and cognitive control: the influence of spontaneous eyeblink rate and dopamine gene polymorphisms on perseveration and distractibility. *Behavioral Neuroscience*, 119, 483-490.
- Duffy, E. (1957). The psychological significance of the concept of arousal or activation. *Psychological Review*, 64, 265-75
- Edwards, D.M., Pellizzoni, C., Breuel, H.P., Berardi, A., Castelli, M.G., Frigerio, E., et al (1995). Pharmacokinetics of reboxetine in healthy volunteers. Single oral doses, linearity and plasma protein binding. *Biopharm Drug Dispos*, 16, 443-460.
- Efron, B., & Tibshirani, R.J. (1993). *An Introduction to the Bootstrap*. London: Chapman & Hall.
- Einhäuser, W., Stout, J., Koch, C., & Carter, O. (2008). Pupil dilation reflects perceptual selection and predicts subsequent stability in perceptual rivalry. *Proceedings of the National Academy of Sciences of the USA*, 105, 1704-1709.
- Eisenberger, N.I., Lieberman, M.D., & Williams, K.D. (2003). Does rejection hurt? An fMRI study of social exclusion. *Science* 302, 290-2.
- Elam, M., Svensson, T.H., & Thorén, P. (1986). Locus coeruleus neurons and sympathetic nerves: activation by cutaneous sensory afferents. *Brain Research*, 366, 254-261.
- Ennaceur A., & Delacour J. (1988). A new one-trial test for neurobiological studies of memory in rats. 1: Behavioral data. *Behavioural Brain Research*, 31, 47-59.

- Erez, A., Li, J., Geraghty, M.T., Ben-Shachar, S., Cooper, M.L., Mensing, D.E., et al. (2010). Mosaic deletion 11p13 in a child with dopamine beta-hydroxylase deficiency--case report and review of the literature. *Am J Med Genet A*, 152A, 732-6.
- Eriksen, C.W., & Schultz, D.W. (1979). Information processing in visual search: a continuous flow conception and experimental results. *Perception & Psychophysics*, 25, 249-263.
- Ermolaeva, V.Y., Tolchenova, G.A., & Brukhanskaya, N.A. (1981). Cortico-cortical connections between the auditory fields and the sensorimotor region of the cortex. *Neuroscience and Behavioral Physiology*, 11, 514-520.
- Falchier, A., Clavagnier, S., Barone, P., & Kennedy, H. (2002). Anatomical evidence of multimodal integration in primate striate cortex. *Journal of Neuroscience*, 22, 5749-5759.
- Farrell, S., & Ludwig, C.J.H. (2008). Bayesian and maximum likelihood estimation of hierarchical response time models. *Psychonomic Bulletin & Review*, 15, 1209-1217.
- Fernandez-Duque, D., & Posner, M.I. (1997). Relating the mechanisms of orienting and alerting. *Neuropsychologia*, 35, 477-86.
- Flicker, C., & Geyer, M.A. (1982). Behavior during hippocampal microinfusions. I. Norepinephrine and diversive exploration. *Brain Research*, 257, 79-103.
- Frank, M.J., Doll, B.B., Oas-Terpstra, J., & Moreno, F. (2009). Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nature Neuroscience*, 12, 1062-1068.
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel/hierarchical models*. Cambridge: Cambridge University Press.
- Gelman, A., & Rubin, D.B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 7, 457-472.
- Gevens, A. (1989). Dynamic functional topography of cognitive tasks. *Brain Topography*, 2, 37-56.
- Ghazanfar, A.A., and Schroeder, C.E. (2006). Is neocortex essentially multisensory? *Trends in Cognitive Sciences*, 10, 278-285.
- Ghose, G.M., & Maunsell, J.H. (2002). Attentional modulation in visual cortex depends on task timing. *Nature*, 419, 616-620.
- Giard, M.H., Perrin, F., Echallier, J. F., Thévenet, M., Froment, J.C., & Pernier, J. (1994). Dissociation of temporal and frontal components in the human auditory N1 wave: a scalp current density and dipole model analysis. *Electroencephalography and Clinical Neurophysiology*, 92, 238-252.
- Giard, M.H., & Peronnet, F. (1999). Auditory-visual integration during multimodal object recognition in humans: a behavioral and electrophysiological study. *Journal of Cognitive Neuroscience*, 11, 473-490.
- Gilzenrat, M.S., Nieuwenhuis, S., Jepma, M., & Cohen, J.D. (2010). Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function. *Cognitive, Affective & Behavioral Neuroscience*, 10, 252-269.
- Gold, J.I., & Shadlen, M.N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30, 535-574.
- Goldstein, D.S. (2006). L-Dihydroxyphenylserine (L-DOPS): a norepinephrine prodrug. *Cardiovasc Drug Rev*, 24, 189-203.
- Gomez, P., Ratcliff, R., & Perea, M. (2007). A model of the go/no-go task. *Journal of Experimental Psychology. General*, 136, 389-413.

- Gompf, H.S., Mathai, C., Fuller, P.M., Wood, D.A., Pedersen, N.P., Saper, C.B., & Lu, J. (2010). Locus ceruleus and anterior cingulate cortex sustain wakefulness in a novel environment. *Journal of Neuroscience*, 30, 14543-51.
- Good, C.D., Johnsrude, I.S., Ashburner, J., Henson, R.N., Friston, K.J., & Frackowiak, R.S. (2001). A voxel-based morphometric study of ageing in 465 normal adult human brains. *Neuroimage*, 14, 21-36.
- Gottsdanker, R. (1975). The attaining and maintaining of preparation. In: P. Rabbitt and S. Dornic (Eds.) *Attention and Performance V*. London: Academic Press, pp. 33-49.
- Grant, S.J., Aston-Jones, G., & Redmond, D.E. Jr. (1988). Responses of primate locus coeruleus neurons to simple and complex sensory stimuli. *Brain Research Bulletin*, 21, 401-10.
- Gratton, G., Coles, M.G., & Donchin, E. (1983). A new method for off-line removal of ocular artifact. *Electroencephalography and Clinical Neurophysiology*, 55, 468-484.
- Gratton, G., Coles, M.G., Sirevaag, E.J., Eriksen, C.W., & Donchin, E. (1988). Pre- and poststimulus activation of response channels: a psychophysiological analysis. *Journal of Experimental Psychology: Human Perception and Performance*, 14, 331-344.
- Greene, C.M., Bellgrove, M.A., Gill, M., & Robertson, I.H. (2009). Noradrenergic genotype predicts lapses in sustained attention. *Neuropsychologia*, 47, 591-4.
- Greenwood, P.M., Fossella, J.A., & Parasuraman, R. (2005). Specificity of the effect of a nicotinic receptor polymorphism on individual differences in visuospatial attention. *J Cogn Neurosci*, 17, 1611-20.
- Grosjean, M., Rosenbaum, D.A., & Elsinger, C. (2001). Timing and reaction time. *Journal of Experimental Psychology: General*, 130, 256-272.
- Hackley, S.A., & Valle-Inclán, F. (1998). Automatic alerting does not speed late motoric processes in a reaction-time task. *Nature*, 391, 786-788.
- Hackley, S.A., & Valle-Inclán, F. (1999). Accessory stimulus effects on response selection: Does arousal speed decision making? *Journal of Cognitive Neuroscience*, 11, 321-329.
- Hackley, S.A., & Valle-Inclán, F. (2003). Which stages of processing are speeded by a warning signal? *Biological Psychology*, 64, 27-45.
- Hackley, S.A. (2009). The speeding of voluntary reaction by a warning signal. *Psychophysiology*, 46, 225-233.
- Hackley, S.A., Langner, R., Rolke, B., Erb, M., Grodd, W., & Ulrich, R. (2009). Separation of phasic arousal and expectancy effects in a speeded reaction time task via fMRI. *Psychophysiology*, 46, 163-171.
- Hackley, S.A., Schankin, A., Wohlschlaeger, A., & Wascher, E. (2007). Localization of temporal preparation effects via trisected reaction time. *Psychophysiology*, 44, 334-338.
- Haruno, M. & Kawato, M. (2006). Different neural correlates of reward expectation and reward expectation error in the putamen and caudate nucleus during stimulus-action-reward association learning. *J Neurophysiol* 95, 948-59
- Hasbroucq, T., Akamatsu, M., Mouret, I., & Seal, J. (1995). Finger pairings in choice reaction time tasks: does the between-hands advantage reflect response preparation? *Journal of Motor Behavior*, 27, 251-262.
- Heathcote, A. (2004). Fitting Wald and ex-Wald distributions to response time data: An example using functions for the S-PLUS package. *Behavior Research Methods, Instruments, & Computers*, 36, 678-694.
- Hess, E.H., & Polt, J.M. (1964). Pupil Size in Relation to Mental Activity during Simple Problem-Solving. *Science*, 143, 1190-2.

- Hjorth, B. (1975). An on-line transformation of EEG scalp potentials into orthogonal source derivations. *Electroencephalography and Clinical Neurophysiology*, 39, 526-30.
- Hobson, J.A., Lydic, R., & Bahdoyan, H.A. (1986). Evolving concepts of sleep cycle generation: from brain centers to neuronal population. *Behav Brain Sci*, 9, 371-448.
- Hobson, J.A., McCarley, R.W., & Wyzinski, P.W. (1975). Sleep cycle oscillation: reciprocal discharge by two brainstem neuronal groups. *Science*, 189, 55-8.
- Hodges, P.W., & Bui, B.H. (1996). A comparison of computer-based methods for the determination of onset of muscle contraction using electromyography. *Electroencephalography and Clinical Neurophysiology*, 101, 511-9.
- Hoffman, B.B., & Taylor, P. (2001). Neurotransmission: the autonomic and somatic motor nervous system. In: Hardman, J.G., Limbird, L.E., Molinoff, P.B., Gilman, A.G. Goodman & Gilman's the Pharmacological Basis of Therapeutics. New York: McGraw-Hill. p. 115–53.
- Hohle, R.H. (1965). Inferred components of reaction time as functions of foreperiod duration. *Journal of Experimental Psychology*, 69, 382–386.
- Holmes, C., Eisenhofer, G., & Goldstein, D.S. (1994). Improved assay for plasma dihydroxyphenylacetic acid and other catechols using high-performance liquid chromatography with electrochemical detection. *J Chromatogr B Biomed Appl*, 653, 131-8.
- Hommel, B., & Akyürek, E.G. (2005). Lag-1 sparing in the attentional blink: benefits and costs of integrating two events into a single episode. *Q J Exp Psychol A*, 58, 1415-33.
- Howard, I.P., & Templeton, W.B. (1966). *Human Spatial Orientation*. Wiley, New York.
- Hughes RN. (2007). Neotic preferences in laboratory rodents: issues, assessment and substrates. *Neuroscience and Biobehavioral Reviews*, 31, 441-64.
- Hyttel, J. (1994). Pharmacological characterization of selective serotonin reuptake inhibitors (SSRIs). *Int Clin Psychopharmacol*, 9 (suppl 1), 19-26.
- Invernizzi, R.W., & Garattini, S. (2004). Role of presynaptic alpha2-adrenoceptors in antidepressant action: recent findings from microdialysis studies. *Prog Neuropsychopharmacol Biol Psychiatry*, 28, 819-827.
- Ishikawa, Y., Kato, Y., Murakami, Y., Inoue, T., Koshiyama, H., & Imura, H. (1987). Effect of L-threo-3,4-dihydroxyphenylserine (L-DOPS) on catecholamine levels in plasma and cerebrospinal fluid (CSF) in anesthetized rats. *Proc Soc Exp Biol Med*, 184, 197-200.
- Itil, T.M., Menon, G.N., Bozak, M.M., & Itil, K.Z. (1984). CNS effects of citalopram, a new serotonin inhibitor antidepressant (a quantitative pharmacoelectroencephalography study). *Prog Neuropsychopharmacol Biol Psychiatry*, 8, 397-409.
- Janisse, M.P. (1977). *Pupillometry: The Psychology of the Pupillary Response*. Washington, D.C.: Hemisphere Publishing Co.
- Janssen, P., & Shadlen, M.N. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. *Nature Neuroscience*, 8, 234–241.
- Jaśkowski, P., Kurczewska, M., Nowik, A., van der Lubbe, R.H., & Verleger, R. (2007). Locus of the intensity effect in simple reaction time tasks. *Perception and Psychophysics*, 69, 1334-1343.
- Jenkinson, M., & Smith, S.M. (2001). A global optimisation method for robust affine registration of brain images. *Med Image Anal*, 5, 143-56.
- Jenkinson, M., Bannister, P., Brady, M., & Smith, S.M. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*, 17, 825-41.

- Jepma, M., & Nieuwenhuis, S. (in press). Pupil diameter predicts changes in the exploration-exploitation trade-off: Evidence for the adaptive gain theory. *J Cogn Neurosci*
- Jepma, M., Wagenmakers, E.-J., Band, G.P.H., & Nieuwenhuis, S. (2009). The effects of accessory stimuli on information processing: Evidence from electrophysiology and a diffusion-model analysis. *Journal of Cognitive Neuroscience*, 21, 847-864.
- Jouvet, M. (1969). Biogenic amines and the states of sleep. *Science*, 163, 32-41.
- Kahneman, D. & Beatty, J. (1967). Pupillary responses in a pitch-discrimination task. *Percept Psychophys*, 2, 101-5.
- Kahneman, D. (1973). *Attention and Effort*. Englewood Cliffs, NJ: Prentice-Hall.
- Kakade, S., & Dayan, P. (2002). Dopamine: generalization and bonuses. *Neural Netw*, 15, 549-59.
- Kalinin, S., Gavriljuk, V., Polak, P.E., Vasser, R., Zhao, J., Heneka, M.T., et al. (2007). Noradrenaline deficiency in brain increases beta-amyloid plaque burden in an animal model of Alzheimer's disease. *Neurobiol Aging*, 28, 1206-14.
- Kang, M.J., Hsum, M., Krajcich, I.M., Loewenstein, G., McClure, S.M., Wang, J.T., & Camerer, C.F. (2009). The wick in the candle of learning: epistemic curiosity activates reward circuitry and enhances memory. *Psychol Sci*, 20, 963-73.
- Karlin, L. (1959). Reaction time as a function of foreperiod duration and variability. *Journal of Experimental Psychology*, 58, 185-191.
- Kato, T., Karai, N., Katsuyama, M., Nakamura, M., & Katsube, J. (1987a). Studies on the activity of L-threo-3,4-dihydroxyphenylserine (L-DOPS) as a catecholamine precursor in the brain. Comparison with that of L-dopa. *Biochem Pharmacol*, 36, 3051-7.
- Kato, T., Katsuyama, M., Karai, N., Nakamura, M., & Katsube, J. (1987b). Studies on the central action of L-threo-3,4-dihydroxyphenyl-serine (L-threo-DOPS) in FLA-63-treated mice. *Pharmacol Biochem Behav*, 26, 407-11.
- Klemmer, E.T. (1956). Time uncertainty in simple reaction time. *Journal of Experimental Psychology*, 51, 179-84.
- Kubiak, P., Rajkowski, J., & Aston-Jones, G. (1992). Behavioral performance and sensory responsiveness of LC neurons in a vigilance task varies with tonic LC discharge rate. *Soc Neurosci Abstr*, 18, 538.
- Lagarias, J.C., Reeds, J.A., Wright, M.H., & Wright, P.E. (1998). Convergence Properties of the Nelder-Mead Simplex Method in Low Dimensions. *SIAM J. Optim*, 9, 112-147.
- Laming, D. (1979). Choice reaction performance following an error. *Acta Psychologica*, 43, 199-224.
- Laming, D.R.J. (1968). *Information theory of choice reaction time*. New York: Wiley.
- Lang, P.J., Bradley, M.M., & Cuthbert, B.N. (2005). *International Affective Picture System (IAPS): Digitized photographs, instruction manual, and affective ratings (Tech. Rep. No. A-6)*. Gainesville: University of Florida, Center for Research in Psychophysiology.
- Lapiz, M.D., & Morilak, D.A. (2006). Noradrenergic modulation of cognitive function in rat medial prefrontal cortex as measured by attentional set shifting capability. *Neuroscience*, 137, 1039-1049.
- Lapiz, M.D., Bondi, C.O., & Morilak, D.A. (2007). Chronic treatment with desipramine improves cognitive performance of rats in an attentional set-shifting test. *Neuropsychopharmacology (Berl, Ger)*, 32, 1000-1010.
- Leth-Steensen, C. (2009). Lengthening fixed preparatory foreperiod durations within a digit magnitude classification task serves mainly to shift distributions of response times upwards. *Acta Psychologica*, 130, 72-80.

- Litman, J.A. (2008). Interest and deprivation factors of epistemic curiosity, *Pers Individ Dif*, 44, 1585–95.
- Liu, J., Kiehl, K.A., Pearlson, G., Perrone-Bizzozero, N.I., Eichele, T., & Calhoun, V.D. (2009). Genetic determinants of target and novelty-related event-related potentials in the auditory oddball response. *Neuroimage*, 46, 809-16.
- Loewenfeld, I. (1993). *The Pupil: Anatomy, Physiology, and Clinical Applications*. Detroit: Wayne State University Press.
- Loewenstein, G. (1994). The Psychology of curiosity: A review and reinterpretation. *Psychol Bull*, 116, 75-98.
- Lorch, R.F., & Myers, J.L. (1990). Regression-analyses of repeated measures data in cognitive research. *Journal of Experimental Psychology: Learning Memory and Cognition*, 16, 149-157.
- Low, K.A., Larson, S.L., Burke, J., & Hackley, S.A. (1996). Alerting effects on choice reaction time and the photic eyeblink reflex. *Electroencephalography and Clinical Neurophysiology*, 98, 385-393.
- Luce, R.D. (1986). *Response times: Their role in inferring elementary mental organization*. New York: Oxford University Press.
- Luksys, G., Gerstner, W., & Sandi, C. (2009). Stress, genotype and norepinephrine in the prediction of mouse behavior using reinforcement learning. *Nat Neurosci*, 12, 1180-1186.
- Lunn, D.J., Spiegelhalter, D., Thomas, A., & Best, N. (2009). The BUGS project: Evolution, critique and future directions. *Statistics in Medicine*, 28, 3049-3067.
- Lunn, D.J., Thomas, A., Best, N., & Spiegelhalter, D. (2000). WinBUGS—A Bayesian modelling framework: Concepts, structure, and extensibility. *Statistics & Computing*, 10, 325-337.
- Lynn, R. (1966). *Attention, arousal, and the orientation reaction*. Oxford: Pergamon Press.
- Maclean, M.H., & Arnell, K.M. (2010). Personality predicts temporal attention costs in the attentional blink paradigm. *Psychon Bull Rev*, 17, 556-62.
- Madrigal, J.L., Kalinin, S., Richardson, J.C., & Feinstein, D.L. (2007). Neuroprotective actions of noradrenaline: effects on glutathione synthesis and activation of peroxisome proliferator activated receptor delta. *J Neurochem*, 103, 2092–2101.
- Madrigal, J.L., Leza, J.C., Polak, P., Kalinin, S., & Feinstein, D.L. (2009). Astrocyte-derived MCP-1 mediates neuroprotective effects of noradrenaline. *J Neurosci*, 29, 263-267.
- Maier, M.A., Bennett, K.M., Hepp-Reymond, M.C., & Lemon, R.N. (1993). Contribution of the monkey corticomotoneuronal system to the control of force in precision grip. *Journal of Neurophysiology*, 69, 772-785.
- Man in 't Veld, A.J., Boomsma, F., Moleman, P., & Schalekamp, M.A. (1987a). Congenital dopamine-beta hydroxylase deficiency. A novel orthostatic syndrome. *Lancet*, 1, 183-8.
- Man in 't Veld, A.J., Boomsma, F., van den Meiracker, A.H., & Schalekamp, M.A. (1987b). Effect of unnatural noradrenaline precursor on sympathetic control and orthostatic hypotension in dopamine-beta-hydroxylase deficiency. *Lancet*, 2, 1172-5.
- Manjarrez, E., Mendez, I., Martinez, L., Flores, A., & Mirasso, C.R. (2007). Effects of auditory noise on the psychophysical detection of visual signals: cross-modal stochastic resonance. *Neuroscience Letters*, 415, 231-236.
- Mann, T.A. & Walker, P. (2003). Autism and a deficit in broadening the spread of visual attention. *Journal of Child Psychology and Psychiatry*, 44, 274-84.
- Mansour, A.A., Babstock, D.M., Penney, J.H., Martin, G.M., McLean, J.H., & Harley, C.W. (2003). Novel objects in a holeboard probe the role of the locus coeruleus in curiosity: support for two modes of attention in the rat. *Behav Neurosci*, 117, 621-31.

- Marino, M.D., Bourdélát-Parks, B.N., Cameron Liles, L., & Weinshenker, D. (2005). Genetic reduction of noradrenergic function alters social memory and reduces aggression in mice. *Behav Brain Res*, 161, 197-203.
- Martens, S., & Wyble, B. (2010). The attentional blink: past, present, and future of a blind spot in perceptual awareness. *Neurosci Biobehav Rev*, 34, 947-57.
- Martens, S., & Johnson, A. (2005). Timing attention: cuing target onset interval attenuates the attentional blink. *Memory and Cognition*, 33, 234-240.
- Mathias, C.J., Bannister, R.B., Cortelli, P., Heslop, K., Polak, J.M., Raimbach, S., et al. (1990). Clinical, autonomic and therapeutic observations in two siblings with postural hypotension and sympathetic failure due to an inability to synthesize noradrenaline from dopamine because of a deficiency of dopamine beta hydroxylase. *Q J Med*, 75, 617-33.
- Mattes, S., & Ulrich, R. (1997). Response force is sensitive to the temporal uncertainty of response stimuli. *Perception and Psychophysics*, 59, 1089-1097.
- McClure, S., Gilzenrat, M., & Cohen, J. (2005). An exploration-exploitation model based on norepinephrine and dopamine activity. In Y. Weiss, B. Schölkopf, & J. Platt (Eds.), *Advances in neural information processing systems*, 18, pp 867-874. Cambridge, MA: MIT Press.
- McDonald, J.J., Teder-Sälejärvi, W.A., Di Russo, F., & Hillyard, S.A. (2005). Neural basis of auditory-induced shifts in visual time-order perception. *Nature Neuroscience*, 8, 1197-1202.
- McGaughy, J., Ross, R.S., & Eichenbaum, H. (2008). Noradrenergic, but not cholinergic, deafferentation of prefrontal cortex impairs attentional set-shifting. *Neuroscience*, 153, 63-71.
- McGurk, H., & MacDonald, J. (1976). Hearing lips and seeing voices. *Nature*, 264, 746-748.
- Miller, J. (2007). Contralateral and ipsilateral motor activation in visual simple reaction time: a test of the hemispheric coactivation model. *Experimental Brain Research*, 176, 539-558.
- Miller, J., Franz, V., & Ulrich, R. (1999). Effects of auditory stimulus intensity on response force in simple, go-no-go, and choice tasks. *Perception & Psychophysics*, 61, 107-119.
- Miller, J., Patterson, T., & Ulrich, R. (1998). Jackknife-based method for measuring LRP onset latency differences. *Psychophysiology*, 35, 99-115.
- Miniussi, C., Wilding, E.L., Coull, J.T., & Nobre, A.C. (1999). Orienting attention in time: modulation of brain potentials. *Brain*, 122, 1507-1518.
- Miskowiak, K., Papadatou-Pastou, M., Cowen, P.J., Goodwin, G.M., Norbury, R., & Harmer, C.J. (2007). Single dose antidepressant administration modulates the neural processing of self-referent personality trait words. *Neuroimage*, 37, 904-911.
- Mordkoff, J.T., & Grosjean, M. (2001). The lateralized readiness potential and response kinetics in response-time tasks. *Psychophysiology*, 38, 777-786.
- Morrell, F. (1972). Visual system's view of acoustic space. *Nature*, 238, 44-46.
- Morrell, L.K. (1968). Temporal characteristics of sensory interaction in choice reaction times. *Journal of Experimental Psychology*, 77, 14-18.
- Moss, F., Ward, L.M., & Sannita, W.G. (2004). Stochastic resonance and sensory information processing: a tutorial and review of application. *Clinical Neurophysiology*, 115, 267-281.
- Müller-Gethmann, H., Ulrich, R., & Rinkenauer, G. (2003). Locus of the effect of temporal preparation: evidence from the lateralized readiness potential. *Psychophysiology*, 40, 597-611.
- Murchison, C.F., Zhang, X.Y., Zhang, W.P., Ouyang, M., Lee, A., & Thomas, S.A. (2004). A distinct role for norepinephrine in memory retrieval. *Cell*, 117, 131-43.

- Näätänen, R. (1971). Non-aging fore-periods and simple reaction time. *Acta Psychologica*, 35, 316–327.
- Näätänen, R. (1972). Time uncertainty and occurrence uncertainty of the stimulus in a simple reaction-time task. *Acta Psychologica*, 36, 492-503.
- Neiss, R. (1988). Reconceptualizing arousal: Psychological states in motor performance. *Psychological Bulletin*, 103, 345–366.
- Newman, L.A., Darling, J., & McGaughy, J. (2008). Atomoxetine reverses attentional deficits produced by noradrenergic deafferentation of medial prefrontal cortex. *Psychopharmacology (Berl)*, 200, 39-50.
- Nicki, R.M. (1970). The reinforcing effect of uncertainty reduction on a human operant. *Can J Psychol*, 24, 389-399.
- Niemi, P., & Näätänen, R. (1981). Foreperiod and simple reaction time. *Psychological Bulletin*, 89, 133–162.
- Nieuwenhuis, S., Aston-Jones, G., & Cohen, J.D. (2005b). Decision making, the P3, and the locus coeruleus norepinephrine system. *Psychol Bull*, 131: 510-32.
- Nieuwenhuis, S., Gilzenrat, M.S., Holmes, B.D., & Cohen, J.D. (2005a). The role of the locus coeruleus in mediating the attentional blink: a neurocomputational theory. *J Exp Psychol Gen*, 134, 291-307.
- Nieuwenhuis, S., de Geus, E.J., & Aston-Jones, G. (in press). The anatomical and functional relationship between the P3 and autonomic components of the orienting response. *Psychophysiology*.
- Nieuwenhuis, S., van Nieuwpoort, I.C., Veltman, D.J., & Drent, M.L. (2007). Effects of the noradrenergic agonist clonidine on temporal and spatial attention. *Psychopharmacology (Berl)*, 193, 261-9.
- Noble, S., & Benfield, P. (1997). Citalopram: A review of its pharmacology, clinical efficacy and tolerability in the treatment of depression. *CNS Drugs*, 8, 410–431.
- Nobre, A., Correa, A., & Coull, J. (2007). The hazards of time. *Current Opinion in Neurobiology*, 17, 465-470.
- O'Doherty, J.P. (2004). Reward representations and reward-related learning in the human brain: insights from neuroimaging. *Curr Opin Neurobiol*, 14, 769-76.
- Oei, N.Y., Tollenaar, M.S., Elzinga, B.M., & Spinhoven, P. (2010). Propranolol reduces emotional distraction in working memory: a partial mediating role of propranolol-induced cortisol increases? *Neurobiol Learn Mem*, 93, 388-95.
- Oei, N.Y., Tollenaar, M.S., Spinhoven, P., & Elzinga, B.M. (2009). Hydrocortisone reduces emotional distracter interference in working memory. *Psychoneuroendocrinology*, 34, 1284-93.
- Ordway, G.A., Schenk, J., Stockmeier, C.A., May, W., & Klimek, V. (2003). Elevated agonist binding to alpha2-adrenoceptors in the locus coeruleus in major depression. *Biological Psychiatry*, 53, 315-23.
- Page, M.E., & Lucki, I. (2002). Effects of acute and chronic reboxetine treatment on stress-induced monoamine efflux in the rat frontal cortex. *Neuropsychopharmacology*, 27, 237-247
- Parasuraman, R., Greenwood, P.M., Kumar, R., & Fossella, J. (2005). Beyond heritability: neurotransmitter genes differentially modulate visuospatial attention and working memory. *Psychol Sci*, 16, 200-7.
- Parmentier, F.B., Elford, G., Escera, C., Andrés, P., & San Miguel, I. (2008). The cognitive locus of distraction by acoustic novelty in the cross-modal oddball task. *Cognition*, 106, 408-32.

- Pavlov, I.P. (1927). *Conditioned reflexes*. Oxford: Clarendon Press.
- Peck, O.F. (1970). The conversation of Progressive Matrices and Mill Hill Vocabulary raw scores into deviation IQs. *Journal of Clinical Psychology*, 26, 67-70.
- Peeters, F.P.M.L., Ponds, R.W.H.M., & Vermeeren, M.T.G. (1996). Affectivity and self-report of depression and anxiety (Dutch). *Tijdschr Psych*, 38, 240-50.
- Perrin, F., Pernier, J., Bertrand, O., & Echallier, J.F. (1989). Spherical splines for scalp potential and current density mapping. *Electroencephalography and Clinical Neurophysiology*, 72, 184-187.
- Phillips, M.A., Bitsios, P., Szabadi, E., & Bradshaw, C.M. (2000). Comparison of the antidepressants reboxetine, fluvoxamine and amitriptyline upon spontaneous pupillary fluctuations in healthy human volunteers. *Psychopharmacology (Berl)*, 149, 72-76.
- Pineda, J.A., Foote, S.L., & Neville, H.J. (1989). Effects of locus coeruleus lesions on auditory, long-latency, event-related potentials in monkey. *J Neurosci*, 9, 81-93.
- Ploghaus, A., Tracey, I., Gati, J.S., Clare, S., Menon, R.S., Matthews, P.M., & Rawlins, J.N. (1999). Dissociating pain from its anticipation in the human brain. *Science*, 284, 1979-81.
- Posner, M.I. (1978). *Chronometric explorations of mind*. Hillsdale, NJ: Erlbaum.
- Posner, M.I., Klein, R., Summers, J. & Buggie, S. (1973). On the selection of signals. *Memory and Cognition*, 1, 2-12.
- Pribram, K.H., & McGuinness, D. (1975). Arousal, activation, and effort in the control of attention. *Psychological Review*, 82, 116-149.
- Raftery, A.E. (1996). Approximate Bayes factors and accounting for model uncertainty in generalized linear models. *Biometrika*, 83, 251-266.
- Raichle, M.E., MacLeod, A.M., Snyder, A.Z., Powers, W.J., Gusnard, D.A., & Shulman, G.L. (2001). A default mode of brain function. *Proc Natl Acad Sci USA*, 98, 676-82.
- Rajkowski, J., Kubiak, P., & Aston-Jones, G. (1993). Correlations between locus coeruleus (LC) neural activity, pupil diameter and behavior in monkey support a role of LC in attention. *Society for Neuroscience Abstracts*, 19, 974.
- Rajkowski, J., Kubiak, P., & Aston-Jones, G. (1994). Locus coeruleus activity in monkey: phasic and tonic changes are associated with altered vigilance. *Brain Research Bulletin*, 35, 607-616.
- Rajkowski, J., Lu, W., Zhu, Y., Cohen, J.D., & Aston-Jones, G. (2000). Prominent projections from the anterior cingulate cortex to the locus coeruleus (LC) in rhesus monkey. *Society for Neuroscience Abstracts*, 26, 2230.
- Rajkowski, J., Majczynski, H., Clayton, E., & Aston-Jones, G. (2004). Activation of monkey locus coeruleus neurons varies with difficulty and performance in a target detection task. *Journal of Neurophysiology*, 92, 361-71.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59-108.
- Ratcliff, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin*, 86, 446-461.
- Ratcliff, R. (2002). A diffusion model account of response time and accuracy in a brightness discrimination task: fitting real data and failing to fit fake but plausible data. *Psychonomic Bulletin & Review*, 9, 278-291.
- Ratcliff, R., & Rouder, J.N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, 9, 347-356.
- Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, 9, 438-481.

- Ratcliff, R., Gomez, P., & McKoon, G. (2004). A diffusion model account of the lexical decision task. *Psychological Review*, 111, 159-82
- Ratcliff, R., Thapar, A., & McKoon, G. (2001). The effects of aging on reaction time in a signal detection task. *Psychology and Aging*, 16, 323-341.
- Ratcliff, R., Thapar, A., Gomez, P., & McKoon, G. (2004). A diffusion model analysis of the effects of aging in the lexical-decision task. *Psychology and Aging*, 19, 278-289.
- Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, 106, 261-300.
- Raven, J.C., Court, J.H., & Raven, J. Manual for Raven's progressive matrices and vocabulary scales: Section 3 Standard progressive matrices. London: Lewis; 1988.
- Raymond, J.E., Shapiro, K.L., & Arnell, K.M. (1992). Temporary suppression of visual processing in an RSVP task: an attentional blink? *J Exp Psychol Hum Percept Perform*, 18, 849-60.
- Reiner, P.B. (1986). Correlational analysis of central noradrenergic neuronal activity and sympathetic tone in behaving cats. *Brain Research*, 378, 86-96.
- Ressler, K.J., & Nemeroff, C.B. (2001). Role of norepinephrine in the pathophysiology of neuropsychiatric disorders. *CNS Spectr*, 6, 663-6.
- Ridderinkhof, K.R. (2002). Activation and suppression in conflict tasks: Empirical clarification through distributional analyses. In W. Prinz & B. Hommel (Eds.), *Common mechanisms in perception and action. Volume 19: Attention & performance* (pp. 494–519). Oxford: Oxford University Press.
- Ridderinkhof, K.R., Ullsperger, M., Crone, E.A., & Nieuwenhuis, S. (2004). The role of the medial frontal cortex in cognitive control. *Science*, 306, 443-447.
- Riehle, A., Grün, S., Diesmann, M., & Aertsen, A. (1997). Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*, 278, 1950-1953.
- Robbins, T.W., & Arnsten, A.F. (2009). The neuropsychopharmacology of fronto-executive function: monoaminergic modulation. *Annu Rev Neurosci*, 32, 267-87.
- Robbins, T.W. (1997). Arousal systems and attentional processes. *Biological Psychology*, 45, 57-71.
- Robertson, D., Goldberg, M.R., Onrot, J., Hollister, A.S., Wiley, R., Thompson, J.G. Jr, et al. (1986). Isolated failure of autonomic noradrenergic neurotransmission. Evidence for impaired beta-hydroxylation of dopamine. *N Engl J Med*, 314, 1494-7.
- Rockland, K.S., and Ojima, H. (2003). Multisensory convergence in calcarine visual areas in macaque monkey. *International Journal of Psychophysiology*, 50, 19-26.
- Rolke, B. (2008). Temporal preparation facilitates perceptual identification of letters. *Perception & Psychophysics*, 70, 1305-1313.
- Rolke, B., & Hofmann, P. (2007). Temporal uncertainty degrades perceptual processing. *Psychonomic Bulletin and Review*, 14, 522-526.
- Roman, T., Schmitz, M., Polanczyk, G.V., Eizirik, M., Rohde, L.A., & Hutz, M.H. (2002). Further evidence for the association between attention-deficit/hyperactivity disorder and the dopamine-beta-hydroxylase gene. *Am J Med Genet*, 114, 154-8.
- Rossion, B., & Pourtois, G. (2004). Revisiting Snodgrass and Vanderwart's object pictorial set: the role of surface detail in basic-level object recognition. *Perception*, 33, 217-36.
- Roth, W.T., Dorato, K.H., & Kopell, B.S. (1984). Intensity and task effects on evoked physiological responses to noise bursts. *Psychophysiology*, 21, 466-81.
- Rouder, J.N., Lu, J., Speckman, P., Sun, D., & Jiang, Y. (2005). A hierarchical model for estimating response time distributions. *Psychonomic Bulletin & Review*, 12, 195-223.

- Rouder, J.N., Sun, D., Speckman, P.L., Lu, J., & Zhou, D. (2003). A hierarchical Bayesian statistical framework for response time distributions. *Psychometrika*, 68, 589-606.
- Rudell, A.P., and Hu, B. (2001). Does a warning signal accelerate the processing of sensory information? Evidence from recognition potential responses to high and low frequency words. *International Journal of Psychophysiology*, 41, 31-42.
- Saletu, B., Anderer, P., Saletu-Zyhlarz, G.M., Arnold, O., & Pascual-Marqui, R.D. (2002). Classification and evaluation of the pharmacodynamics of psychotropic drugs by single-lead pharmaco-EEG, EEG mapping and tomography (LORETA). *Methods Find Exp Clin Pharmacol*, 24(Suppl. C), 97S-120S.
- Sanders, A.F. (1980). Stage analysis of reaction process. In G.E. Stelmach, and J. Requin (Eds.), *Tutorials in Motor Behavior* (pp. 331-354). Amsterdam: North Holland.
- Sanders, A.F. (1983). Towards a model of stress and human performance. *Acta Psychologica*, 53, 61-97.
- Sara, S.J. (2009). The locus coeruleus and noradrenergic modulation of cognition. *Nature Reviews Neuroscience*, 10, 211-223.
- Sara, S.J., Dyon-Laurent, C., & Hervé, A. (1995). Novelty seeking behavior in the rat is dependent upon the integrity of the noradrenergic system. *Brain Res Cogn Brain Res* 2, 181-7.
- Satterthwaite, T.D., Green, L., Myerson, J., Parker, J., Ramaratnam, M., Buckner, R.L. (2007). Dissociable but inter-related systems of cognitive control and reward during decision making: Evidence from pupillometry and event-related fMRI. *Neuroimage*, 37, 1017-1031.
- Schmidt, R.A, Gielen, S.C., and van den Heuvel, P.J. (1984). The locus of intersensory facilitation of reaction time. *Acta Psychologica*, 57, 145-164.
- Schmitt, J.A., Riedel, W.J., Vuurman, E.F., Kruizinga, M., & Ramaekers, J.G. (2002). Modulation of the critical flicker fusion effects of serotonin reuptake inhibitors by concomitant pupillary changes. *Psychopharmacology (Berl)*, 160, 381-386.
- Schröger, E. & Wolff, C. (1998). Behavioral and electrophysiological effects of task-irrelevant sound change: a new distraction paradigm. *Brain Research. Cognitive Brain Research*, 7, 71-87.
- Schwarz, W. (2001). The ex-Wald distribution as a descriptive model of response times. *Behavior Research Methods, Instruments, & Computers*, 33, 457-469.
- Seeley, W.W., Menon, V., Schatzberg, A.F., Keller, J., Glover, G.H., Kenna, H., Reiss, A.L., & Greicius, M.D. (2007). Dissociable intrinsic connectivity networks for salience processing and executive control. *J Neurosci*, 27, 2349-56.
- Seifried, T., Ulrich, R., Bausenhardt, K.M., Rolke, B., & Osman, A. (2010). Temporal preparation decreases perceptual latency: Evidence from a clock paradigm. *Quarterly Journal of Experimental Psychology*, 63, 2432-51.
- Semba, J., & Takahashi, R. (1985). The effects of L-threo-dihydroxyphenylserine on norepinephrine metabolism in rat brain. *Psychiatry Res*, 15, 319-26.
- Servan-Schreiber, D., Printz, H., & Cohen, J.D. (1990). A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior. *Science*, 249, 892-895.
- Seu, E., Lang, A., Rivera, R.J., & Jentsch, J.D. (2008). Inhibition of the norepinephrine transporter improves behavioral flexibility in rats and monkeys. *Psychopharmacology (Berl, Ger)*, 202, 505-519.
- Shapiro, K.L., Caldwell, J., & Sorensen, R.E. (1997). Personal names and the attentional blink: a visual "cocktail party" effect. *J Exp Psychol Hum Percept Perform*, 23, 504-14.

- Shiffrin, R.M., Lee, M.D., Kim, W., & Wagenmakers, E.-J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical Bayesian methods. *Cognitive Science*, 32, 1248-1284.
- Shohamy, D., & Adcock, R.A. (2010). Dopamine and adaptive memory. *Trends Cogn Sci*, 14, 464-72.
- Siegle, G.J., Steinhauer, S.R., Stenger, V.A., Konecky, R., & Carter, C.S. (2003). Use of concurrent pupil dilation assessment to inform interpretation and analysis of fMRI data. *Neuroimage*, 20, 114-124.
- Siever, L.J., & Davis, K.L. (1985). Overview: toward a dysregulation hypothesis of depression. *Am J Psychiatry*, 142, 1017-31.
- Skinbjerg, M., Seneca, N., Liow, J.S., Hong, J., Weinshenker, D., Pike, V.W., et al. (2010). Dopamine beta-hydroxylase-deficient mice have normal densities of D(2) dopamine receptors in the high-affinity state based on in vivo PET imaging and in vitro radioligand binding. *Synapse*, 64, 699-703.
- Smith, P.L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, 27, 161-168.
- Smith, S.M. (2002). Fast robust automated brain extraction. *Hum Brain Mapp*, 17, 143-55.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E., Johansen-Berg, H., et al. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*, 23 Suppl 1, S208-19.
- Smith, S.M., & Nichols, T.E. (2009). Threshold-free cluster enhancement: addressing problems of smoothing, threshold dependence and localisation in cluster inference. *Neuroimage*, 44, 83-98.
- Smith, S.A. & Smith, S.E. (1999). Bilateral Horner's syndrome: detection and occurrence. *J Neurol Neurosurg Psychiatry*, 66, 48-51.
- Smith, S.M., Zhang, Y., Jenkinson, M., Chen, J., Matthews, P.M., Federico, A., et al. (2002). Accurate, robust, and automated longitudinal and cross-sectional brain change analysis. *Neuroimage*, 17, 479-89.
- Sokolov, E.N. (1963). *Perception and the Conditioned Reflex*. Oxford: Pergamon Press.
- Spencer, K.M. & Coles, M.G. (1999). The lateralized readiness potential: relationship between human data and response activation in a connectionist model. *Psychophysiology*, 36, 364-370.
- Stafford, I.L. & Jacobs, B.L. (1990). Noradrenergic modulation of the masseteric reflex in behaving cats. *Journal of Neuroscience*, 10, 91-107.
- Stahl, J. & Rammsayer, T.H. (2005). Accessory stimulation in the time course of visuomotor information processing: Stimulus intensity effects on reaction time and response force. *Acta Psychologica*, 120, 1-18.
- Stanford, T.R., & Stein, B.E. (2007). Superadditivity in multisensory integration: putting the computation in context. *Neuroreport*, 18, 787-792.
- Steere, J.C. & Arnsten, A.F. (1997). The alpha-2A noradrenergic receptor agonist guanfacine improves visual object discrimination reversal performance in aged rhesus monkeys. *Behavioral Neuroscience*, 111, 883-891.
- Stein, B.E., London, N., Wilkinson, L.K., & Price, D.D. (1996). Enhancement of perceived visual intensity by auditory stimuli: A psychophysical analysis. *Journal of Cognitive Neuroscience*, 8, 497-506.
- Sternberg S. (1966). High-speed scanning in human memory. *Science*, 153, 652-4.
- Sternberg, S. (1969). Memory-scanning: mental processes revealed by reaction-time experiments.

- Strange, B.A. & Dolan, R.J. (2004). Beta-adrenergic modulation of emotional memory-evoked human amygdala and hippocampal responses. *Proc Natl Acad Sci USA*, 101, 11454-8.
- Strange, B.A., Hurlmann, R., & Dolan, R.J. (2003). An emotion-induced retrograde amnesia in humans is amygdala- and beta-adrenergic-dependent. *Proc Natl Acad Sci USA*, 100, 13626-31.
- Sutton, R.S., & Barto, A.G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Sutton, S., Braren, M., Zubin, J., & John, E.R. (1965). Evoked-potential correlates of stimulus uncertainty. *Science*, 150, 1187-8.
- Sutton, R.S. (1990). Integrated architectures for learning, planning and reacting based on approximating dynamic programming. In *Proceedings of the Seventh International Conference on Machine Learning*, p. 216–224, San Mateo, CA: Morgan Kaufmann.
- Szabo, S.T., & Blier, P. (2001). Effect of the selective noradrenergic reuptake inhibitor reboxetine on the firing activity of noradrenaline and serotonin neurons. *Eur J Neurosci*, 13, 2077-2087.
- Tait, D.S., Brown, V.J., Farovik, A., Theobald, D.E., Dalley, J.W., & Robbins, T.W. (2007). Lesions of the dorsal noradrenergic bundle impair attentional set-shifting in the rat. *Eur J Neurosci*, 25, 3719-3724.
- Talsma, D., Mulckhuyse, M., Slagter, H.A., & Theeuwes, J. (2007), Faster, more intense! The relation between electrophysiological reflections of attentional orienting, sensory gain control, and speed of responding. *Brain Research*, 1178, 92-105.
- Tandonnet, C., Burle, B., Vidal, F., & Hasbroucq, T. (2003). The influence of time preparation on motor processes assessed by surface Laplacian estimation. *Clinical Neurophysiology*, 114, 2376-2384.
- Tandonnet, C., Burle, B., Vidal, F., & Hasbroucq, T. (2006). Knowing when to respond and the efficiency of the cortical motor command: a Laplacian ERP study. *Brain Research*, 1109, 158-163.
- Thomas, S.A., Marck, B.T., Palmiter, R.D., & Matsumoto, A.M. (1998). Restoration of norepinephrine and reversal of phenotypes in mice lacking dopamine beta-hydroxylase. *J Neurochem*, 70, 2468-76.
- Thomas, S.A., Matsumoto, A.M., & Palmiter, R.D. (1995). Noradrenaline is essential for mouse fetal development. *Nature*, 374, 643-6.
- Thomas, S.A., & Palmiter, R.D. (1997a). Disruption of the dopamine beta-hydroxylase gene in mice suggests roles for norepinephrine in motor function, learning, and memory. *Behav Neurosci*, 111, 579-89.
- Thomas, S.A., & Palmiter, R.D. (1997b). Impaired maternal behavior in mice lacking norepinephrine and epinephrine. *Cell*, 91, 583-92.
- Thompson, J.M., O'Callaghan, C.J., Kingwell, B.A., Lambert, G.W., Jennings, G.L., & Esler, M.D. (1995). Total norepinephrine spillover, muscle sympathetic nerve activity and heart-rate spectral analysis in a patient with dopamine beta-hydroxylase deficiency. *J Auton Nerv Syst*, 55, 198-206.
- Timmers, H.J., Deinum, J., Wevers, R.A., & Lenders, J.W. (2004). Congenital dopamine-beta-hydroxylase deficiency in humans. *Ann N Y Acad Sci*, 1018, 520-3.
- Triviño, M., Correa, A., Arnedo, M., & Lupiáñez, J. (2010). Temporal orienting deficit after prefrontal damage. *Brain*, 133, 1173-1185.
- Tse, W.S., & Bond, A.J. (2002). Difference in serotonergic and noradrenergic regulation of human social behaviours. *Psychopharmacology (Berl, Ger)*, 159, 216-221.

- Tulen, J.H., Man in 't Veld, A.J., Dzoljic, M.R., Mechelse, K., & Moleman, P. (1991). Sleeping with and without norepinephrine: effects of metoclopramide and D,L-threo-3,4-dihydroxyphenylserine on sleep in dopamine beta-hydroxylase deficiency. *Sleep*, 14, 32-8.
- Tulen, J.H., Man in't Veld, A.J., Mechelse, K., & Boomsma, F. (1990). Sleep patterns in congenital dopamine beta-hydroxylase deficiency. *J Neurol*, 237, 98-102.
- Twa, M.D., Bailey, M.D., Hayes, J., & Bullimore, M. (2004). Estimation of pupil size by digital photography. *J Cataract Refract Surg*, 30, 381-389.
- Ullsperger, M., Harsay, H.A., Wessel, J.R., & Ridderinkhof, K.R. (2010). Conscious perception of errors and its relation to the anterior insula. *Brain Struct Funct*, 214, 629-43.
- Usher, M., Cohen, J.D., Servan-Schreiber, D., Rajkowski, J., & Aston-Jones, G. (1999). The role of locus coeruleus in the regulation of cognitive performance. *Science*, 283, 549-554.
- Usher, M., & McClelland, J.L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, 108, 550-592.
- Usher, M., Cohen, J.D., Servan-Schreiber, D., Rajkowski, J., & Aston-Jones, G. (1999). The role of locus coeruleus in the regulation of cognitive performance. *Science*, 283, 549-554.
- Vallesi, A., McIntosh, A.R., & Stuss, D.T. (2009). Temporal preparation in aging: a functional MRI study. *Neuropsychologia*, 47, 2876-81.
- Valls-Solé, J., Solé, A., Valldeoriola, F., Muñoz, E., Gonzalez, L.E., & Tolosa, E.S. (1995). Reaction time and acoustic startle in normal human subjects. *Neuroscience Letters*, 195, 97-100.
- Van Boxtel, G.J., Geraats, L.H., Van den Berg-Lenssen, M.M., and Brunia, C.H. (1993). Detection of EMG onset in ERP research. *Psychophysiology*, 30, 405-412.
- Van der Burg, E., Olivers, C.N., Bronkhorst, A.W., & Theeuwes, J. (2008). Pip and pop: Non-spatial auditory signals improve spatial visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 34, 1053-65.
- Van Steveninck, A.L., Cohen, A.F., & Ward, T. (1989). A microcomputer based system for recording and analysis of smooth pursuit and saccadic eye movements. *Br J Clin Pharmacol*, 27, 712-713.
- Van Steveninck, A.L., Gieschke, R., Schoemaker, H.C., Pieters, M.S., Kroon, J.M., Breimer, D.D., et al (1993). Pharmacodynamic interactions of diazepam and intravenous alcohol at pseudo steady state. *Psychopharmacology (Berl)*, 110, 471-478.
- Van Steveninck, A.L., Schoemaker, H.C., Pieters, M.S., Kroon, R., Breimer, D.D., & Cohen, A.F. (1991). A comparison of the sensitivities of adaptive tracking, eye movement analysis and visual analog lines to the effects of incremental doses of temazepam in healthy volunteers. *Clin Pharmacol Ther*, 50, 172-180.
- Van Steveninck, A.L., Van Berckel, B.N., Schoemaker, R.C., Breimer, D.D., Van Gerven, J.M., & Cohen, A.F. (1999). The sensitivity of pharmacodynamic tests for the central nervous system effects of drugs on the effects of sleep deprivation. *J Psychopharmacol*, 13, 10-17.
- Vandekerckhove, J., & Tuerlinckx, F. (2008). Diffusion model analysis with MATLAB: a DMAT primer. *Behavior Research Methods*, 40, 61-72
- Vandekerckhove, J., & Tuerlinckx, F. (2007). Fitting the Ratcliff diffusion model to experimental data. *Psychonomic Bulletin & Review*, 16, 1011-1026.
- Vidal, F., Grapperon, J., Bonnet, M., & Hasbroucq, T. (2003). The nature of unilateral motor commands in between-hand choice tasks as revealed by surface Laplacian estimation. *Psychophysiology*, 40, 796-805.

- Voss, J.L., Gonsalves, B.D., Federmeier, K.D., Tranel, D., & Cohen, N.J. (2011). Hippocampal brain-network coordination during volitional exploratory behavior enhances learning. *Nat Neurosci*, 14, 115-20.
- Vroomen, J., & de Gelder, B. (2000). Sound enhances visual perception: cross-modal effects of auditory organization on vision. *Journal of Experimental Psychology: Human Perception and Performance*, 26, 1583-1590.
- Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11, 192-196.
- Wagenmakers, E. -J., Ratcliff, R., Gomez, P., & McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *Journal of Memory and Language*, 58, 140-159.
- Wald, A. (1947). *Sequential analysis*. New York: Wiley.
- Wang, L.E., Fink, G.R., Dafotakis, M., & Grefkes, C. (2009). Noradrenergic stimulation and motor performance: differential effects of reboxetine on movement kinematics and visuomotor abilities in healthy human subjects. *Neuropsychologia*, 47, 1302-1312
- Warren, C.M., Breuer, A.T., Kantner, J., Fiset, D., Blais, C., & Masson, M.E. (2009). Target-distractor interference in the attentional blink implicates the locus coeruleus-norepinephrine system. *Psychon Bull Rev*, 16, 1106-11.
- Watson, D., Clark, L.A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: the PANAS scales. *J Pers Soc Psychol*, 54, 1063-70.
- Wechsler, D. (1997). *Wechsler Adult Intelligence Scale-III*. San Antonio, TX: The Psychological Corporation.
- Wetzels, R., Lee, M.D., & Wagenmakers, E.-J. (2010). Bayesian inference using WBDDev: A tutorial for social scientists. *Behavior Research Methods*, 42, 884-897.
- Willemsen, J.J., Ross, H.A., Jacobs, M.C., Lenders, J.W., Thien, T., Swinkels, L.M., & Benraad, T.J. (1995). Highly sensitive and specific HPLC with fluorometric detection for determination of plasma epinephrine and norepinephrine applied to kinetic studies in humans. *Clin Chem*, 41, 1455-60.
- Witte, E.A., & Marrocco, R.T. (1997). Alteration of brain noradrenergic activity in rhesus monkeys affects the alerting component of covert orienting. *Psychopharmacology*, 132, 315-323.
- Wittmann, B.C., Daw, N.D., Seymour, B., & Dolan, R.J. (2008). Striatal activity underlies novelty-based choice in humans. *Neuron*, 58, 967-73.
- Woolrich, M.W., Behrens, T.E., Beckmann, C.F., Jenkinson, M., & Smith, S.M. (2004). Multilevel linear modelling for FMRI group analysis using Bayesian inference. *Neuroimage* 21, 1732-47.
- Woolrich, M.W., Ripley, B.D., Brady, M., & Smith, S.M. (2001). Temporal autocorrelation in univariate linear modeling of FMRI data. *Neuroimage*, 14, 1370-86.
- Yamamoto, K., & Ozawa, N. (1989). Increased firing of locus coeruleus neurons associated with preparatory set in rats. *Neuroscience Letters*, 106, 112-118.
- Yerkes, R.M., & Dodson, J.D. (1908). The relation of strength of stimulus to rapidity of habit formation. *Journal of Comparative Neurology & Psychology*, 18, 459-482.
- Yu, A.J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron* 46, 681-692.
- Zeigenfuse, M.D., & Lee, M.D. (2010) A general latent assignment approach for modeling psychological contaminants. *Journal of Mathematical Psychology*, 54, 352-362.
- Zhang, W., Klimek, V., Farley, J.T., Zhu, M.Y., & Ordway, G.A. (1999). Alpha2C adrenoceptors inhibit adenylyl cyclase in mouse striatum: potential activation by dopamine. *J Pharmacol Exp Ther*, 289, 1286-92.

- Zhang, W.P., Ouyang, M., & Thomas, S.A. (2004). Potency of catecholamines and other L-tyrosine derivatives at the cloned mouse adrenergic receptors. *Neuropharmacology*, 47, 438-49.
- Zhang, Y., Brady, M., & Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans Med Imaging*, 20, 45-57.
- Zhu, Y., Iba, M., Rajkowski, J., & Aston-Jones, G. (2004). Projection from the orbitofrontal cortex to the locus coeruleus in monkeys revealed by anterograde tracing. *Society for Neuroscience Abstracts*, 30, 211-213.

Nederlandse samenvatting


Het onderzoek dat is beschreven in dit proefschrift bestrijkt een breed scala aan onderwerpen, variërend van exploratiegedrag en nieuwsgierigheid tot de effecten van temporele zekerheid op informatieverwerking. Toch is er één begrip dat als een rode lijn door het proefschrift loopt, en dat is *arousal*. Arousal kan worden gedefinieerd als de activatietoestand van het centrale en autonome zenuwstelsel, en is gerelateerd aan de mate van alertheid van een persoon of dier. Er wordt vaak een globaal onderscheid gemaakt tussen lage, gemiddelde en hoge arousalniveaus: een laag arousalniveau is geassocieerd met slaperigheid en onoplettendheid, een gemiddeld arousalniveau met een ontspannen waaktoestand, en een hoog arousalniveau met toestanden zoals stress, angst en opwinding. Bepaalde veranderingen in arousalniveau, zoals de veranderingen in arousal die optreden tijdens de slaap-waak cyclus, ontwikkelen zich langzaam en gaan gepaard met geleidelijke veranderingen in hersenactiviteit. Daarnaast zijn er ook veranderingen in arousal die heel snel optreden, bijvoorbeeld als je schrikt van een onverwacht hard geluid; deze gaan gepaard met acute en vaak kortdurende veranderingen in hersenactiviteit.

Arousal en het noradrenalinesysteem

Er zijn meerdere lichamelijke en hersensystemen betrokken bij de regulatie van arousal. Eén systeem, dat vaak voorkomt in dit proefschrift, is het centrale noradrenalinesysteem. De belangrijkste component van dit hersensysteem is de *locus coeruleus (LC)*, een cluster van neuronen (hersencellen) in de hersenstam dat verbindingen maakt met bijna alle hersengebieden. Wanneer LC-neuronen actief zijn komt de chemische stof *noradrenaline* vrij in al deze hersengebieden, waardoor de neuronen in deze gebieden sterker gaan reageren op hun input. Het noradrenalinesysteem wordt al sinds lange tijd geassocieerd met basale processen zoals de regulatie van de slaap-waak cyclus, maar in de afgelopen twintig jaar zijn wetenschappers zich gaan realiseren dat dit systeem ook een meer specifieke rol speelt in cognitieve processen zoals aandacht en geheugen, en in de regulatie van gedrag.

Bijna alles wat we weten over de LC komt voort uit onderzoek met proefdieren, waarin de dieren simpele cognitieve taken uitvoerden terwijl de activiteit van hun LC werd gemeten via ingebrachte elektroden. Zulke onderzoeken in apen hebben laten zien dat er langzame fluctuaties in het activiteitsniveau van de LC optreden gedurende het verloop van een experiment (veranderingen die optreden over tientallen seconden of minuten) en dat het niveau van LC-activiteit gerelateerd is aan het gedrag van de aap. Tijdens periodes met een laag niveau van LC-activiteit is de aap slaperig en onoplettend en presteert hij slecht op de cognitieve taak. Periodes met een gemiddeld niveau van LC-activiteit zijn geassocieerd met een alerte toestand waarin de aap gefocust is op de cognitieve taak en goed presteert op deze taak. Periodes met hoog niveau van LC-activiteit zijn geassocieerd met afgeleid en onrustig gedrag en slechte taakprestatie. Er is dus sprake van een niet-lineaire

relatie tussen het activiteitsniveau van de LC en taakprestatie: taakprestatie is optimaal bij een gemiddelde activiteit en neemt af bij zowel verlaagde als verhoogde niveaus van activiteit (Figuur 1). Zoals je misschien is opgevallen lijkt deze relatie tussen LC-activiteit en gedrag/taakprestatie veel op de eerder beschreven relatie tussen arousalniveau en gedrag.



Figuur 1. Relatie tussen LC-activiteit en taakprestatie. Een gemiddeld niveau van LC-activiteit is geassocieerd met optimale taakprestatie en een grote phasic LC response na taakrelevante stimuli (phasic LC mode). Een hoog niveau van LC-activiteit is geassocieerd met slechte taakprestatie en de afwezigheid van de phasic LC response (tonic LC mode). De pijltjes geven het moment aan waarop een taakrelevante stimulus wordt aangeboden. Volgens Aston-Jones en Cohen (2005) leiden de phasic en de tonic LC mode tot, respectievelijk, exploitatie en exploratie. Figuur is gebaseerd op Aston-Jones en Cohen (2005).

Naast deze langzaam optredende veranderingen in het niveau van LC-activiteit, hebben de bovengenoemde apenstudies aangetoond dat er ook hele snelle en kortdurende toenames in LC-activiteit optreden (in de orde van een paar honderd milliseconden) zodra de aap een taakrelevante stimulus waarneemt waarop gereageerd moet worden. Zo'n snelle toename in LC-activiteit wordt de *phasic LC response* genoemd. De phasic LC response resulteert in een tijdelijke toename van noradrenaline in de hersenen, wat zorgt voor een snellere en efficiëntere verwerking van de stimulus. Omdat de phasic LC response alleen optreedt op momenten dat een taakrelevante stimulus wordt waargenomen werkt deze als een temporeel filter dat selectief de verwerking van taakrelevante stimuli faciliteert. Een interessante bevinding van de apenstudies was dat er een relatie bleek te zijn tussen het niveau van LC-activiteit wanneer er geen stimulus werd waargenomen (de basisactiviteit) en de phasic LC response na een taakrelevante stimulus (Figuur 1). In periodes met een gemiddelde LC-basisactiviteit (waarin de aap goed presteert op de cognitieve taak) treedt er een grote phasic LC response op na taakrelevante stimuli, maar niet op

andere momenten. In periodes met een verhoogde LC-basisactiviteit daarentegen (waarin de aap snel afgeleid is en slecht presteert op de cognitieve taak) is de stimulusgerelateerde phasic LC response veel kleiner of afwezig. Deze twee toestanden van LC-activiteit worden aangeduid als de *phasic LC mode* (gemiddelde basisactiviteit/grote phasic LC response) en de *tonic LC mode* (verhoogde basisactiviteit/ kleine phasic LC response).

Het noradrenalesysteem en de balans tussen exploitatie en exploratie

Volgens een recente theorie die gebaseerd is op de hierboven beschreven bevindingen hebben de phasic en de tonic LC mode tegengestelde maar aanvullende functies in de regulatie van gedrag (Aston-Jones & Cohen, 2005; beschreven in Hoofdstuk 1 van dit proefschrift). Deze theorie stelt dat de verschillende patronen van LC-activiteit (en hiermee geassocieerde noradrenalineniveaus in de hersenen) gerelateerd zijn aan de balans tussen twee gedragsstrategieën: exploitatie en exploratie. Exploitatie is het uitbaten van bekende situaties/stimuli/acties waarvan je op basis van eerdere ervaringen weet dat dit gunstig zal uitpakken, en exploratie is het uitproberen van nieuwe of minder bekende situaties/stimuli/acties. Een goede balans tussen exploitatie (of stabiliteit) en exploratie (of flexibiliteit) is erg belangrijk om je goed aan te kunnen passen in een complexe en veranderende omgeving. Het dilemma tussen het exploiteren van bekende opties en het exploreren van nieuwe mogelijkheden speelt ook een rol in veel alledaagse beslissingen. Een voorbeeld hiervan is het beslissen van wat je vanavond gaat eten. Je kunt hierbij kiezen voor een bekend gerecht waarvan je zeker weet dat het lekker is (exploitatie). Je kunt echter ook beslissen om een nieuw gerecht uit te proberen (exploratie). In dit laatste geval loop je het risico dat de maaltijd tegenvalt, maar het is ook mogelijk dat het nieuwe gerecht je zo goed bevalt dat het je nieuwe favoriete gerecht wordt—iets dat je niet ontdekt zou hebben als je voor het bekende gerecht had gekozen.

Volgens de theorie van Aston-Jones en Cohen (2005) verbetert de phasic LC mode de exploitatie van de taak die op dat moment wordt uitgevoerd (door de selectieve verhoging van LC/noradrenaline-activiteit na taakrelevante stimuli) en het uitfilteren van irrelevante informatie (door de lage LC-basisactiviteit). Dit leidt tot een sterke aandachtsfocus op de huidige taak en optimale prestatie op deze taak. De verhoogde LC-basisactiviteit in de tonic LC mode, daarentegen, verhoogd de kans op willekeurige reacties op alle stimuli, ongeacht hun relevantie voor de huidige taak. Dit vergroot de afleidbaarheid van het cognitieve systeem, waardoor de prestatie op de huidige taak verslechtert maar de exploratie van nieuwe stimuli en acties toeneemt. De theorie stelt verder dat de LC informatie over de kosten en baten van de huidige gedragsstrategie ontvangt via verbindingen met frontale hersengebieden (waar informatie over beloningen en kosten verwerkt wordt) en dat deze informatie wordt gebruikt om veranderingen in LC-activiteit aan te sturen. Wanneer het huidige gedrag gepaard gaat met hoge baten en lage kosten zal dit leiden tot een sterke phasic LC mode en exploitatie van de huidige taak. Maar wanneer de baten van het huidige gedrag sterk afnemen en/of de kosten hiervan toenemen zal dit leiden tot een overgang naar de tonic LC

mode en toenemende exploratie. Dit laatste is erg nuttig in situaties waarin de huidige gedragsstrategie niet veel meer oplevert, want door minder te focussen op dit gedrag en meer te exploreren is de kans groot dat betere alternatieven voor het huidige gedrag worden ontdekt.

Hoewel bovengenoemde theorie veel invloed heeft gehad op wetenschappelijke ideeën over de rol van het noradrenalesysteem in de regulatie van gedrag, is deze theorie bijna volledig gebaseerd op onderzoek met proefdieren. Tests van deze theorie in mensen, die nodig zijn om de theorie te toetsen en verder te ontwikkelen, zijn tot kortgeleden uitgebleven. Om hier verandering in te brengen hebben wij de belangrijkste voorspellingen van de theorie getest bij menselijke proefpersonen (Hoofdstuk 2 en 3). Omdat het om ethische redenen niet mogelijk is om LC/noradrenaline-activiteit bij menselijke proefpersonen rechtstreeks te meten hebben we niet-invasieve methoden gebruikt om deze activiteit op een indirecte manier te onderzoeken.

Recent onderzoek heeft uitgewezen dat pupilgrootte een betrouwbare voorspeller van LC-activiteit lijkt te zijn: hoe groter de pupil, hoe hoger de LC-activiteit. Hierdoor geïnspireerd hebben we in het onderzoek dat is beschreven in Hoofdstuk 2 onderzocht of de relatie tussen pupilgrootte en exploitatie- versus exploratiegedrag consistent is met de hierboven beschreven theorie. De pupilgrootte van de proefpersonen in dit experiment werd gemeten terwijl ze een goktaak deden waarin de balans tussen exploitatie en exploratie een belangrijke rol speelt. Bij deze taak kregen de proefpersonen vier gokmachines te zien waarvan ze er steeds één moesten kiezen; de gekozen machine leverde dan een bepaalde hoeveelheid punten op. De proefpersonen maakten in totaal 150 keer een keuze, en hun doel was om zo veel mogelijk punten te verdienen. Dit was niet gemakkelijk want de opbrengsten van de gokmachines veranderden gedurende het experiment op een onvoorspelbare manier. De proefpersonen moesten daarom een goede balans zien te vinden tussen het kiezen voor de machine(s) waarvan ze wisten dat die op dat moment veel opleverden (exploitatie) en het uitproberen van de andere machines om erachter te komen of er inmiddels geen betere optie was ontstaan (exploratie). Met behulp van een computermodel konden we elke keuze classificeren als een exploitatieve of exploratieve keuze.

De resultaten van dit experiment toonden aan dat proefpersonen periodes van exploitatie afwisselden met periodes van exploratie, en dat de overgang naar een exploratieperiode meestal werd voorafgegaan door een afname in de opbrengst van de tot dan toe geëxploiteerde gokmachine en een toename in keuzemoeilijkheid. De pupildata toonden aan dat proefpersonen grotere pupillen hadden als ze exploreerden dan als ze exploiteerden. Dit resultaat ondersteunt het idee dat exploratie wordt gestimuleerd door hoge LC-basisactiviteit (de tonic LC mode; zoals gereflecteerd in een grotere pupil), terwijl exploitatie wordt gestimuleerd door lagere LC-basisactiviteit (de phasic LC mode; zoals gereflecteerd in een kleinere pupil). Een andere bevinding was dat veranderingen in de kosten/baten van de huidige keuzestrategie die optraden rondom de overgang tussen exploitatie- en exploratieperiodes gepaard gingen met geleidelijke veranderingen in pupildiameter. Dit resultaat ondersteunt het idee dat veranderingen in LC-activiteit worden veroorzaakt door veranderingen in taakgerelateerde kosten en baten. De resultaten van deze studie

ondersteunden dus de eerder beschreven theorie over de rol van het noradrenalinesysteem in de regulatie van de exploratie-exploitatie balans.

Het bewijs voor de theorie dat is geleverd door onze pupilstudie is echter heel indirect (we maten geen LC-activiteit maar pupilgrootte, een veronderstelde index hiervan). Bovendien toonde deze studie correlatieve maar geen causale verbanden aan. Om deze redenen hebben we in een volgende studie onderzocht of we de balans tussen exploratie en exploitatie konden beïnvloeden door middel van een farmacologische manipulatie van het noradrenalinesysteem (Hoofdstuk 3 van dit proefschrift). In dit onderzoek gaven we proefpersonen een drug die de hoeveelheid noradrenaline in het brein verhoogt; we creëerden dus als het ware een tonic LC mode in deze proefpersonen. Onze verwachting was dat dit tot een toenemende mate van exploratie zou leiden. Rondom het tijdstip dat de effecten van deze drug maximaal waren, voerden de proefpersonen een aantal cognitieve taken uit die ontworpen waren om exploratie/exploitatiedrag te onderzoeken, waaronder de goktaak die ook in de pupilstudie werd gebruikt. We vergeleken de resultaten van deze proefpersonen met de resultaten van andere proefpersonen waarbij het noradrenalinesysteem niet was beïnvloed (de controlegroep).

Tegen onze verwachtingen in vonden we in deze studie geen aanwijzingen dat de proefpersonen met verhoogde noradrenalineniveaus meer exploreerden dan de controlegroep. Deze bevinding kan worden gezien als problematisch voor de theorie, en suggereert dat er misschien andere mechanismen ten grondslag liggen aan de relatie tussen LC/noradrenaline-activiteit en gedrag/taakprestatie die werd gevonden in de apenstudies. Een andere mogelijkheid is dat het noradrenalinesysteem wel een rol speelt in de exploratie van stimuli/acties buiten de huidige taakcontext (wat niet werd onderzocht in onze studie), maar niet in de exploratie van verschillende opties binnen één taak. Ook is het mogelijk dat het effect van onze farmacologische manipulatie erg verschilde van persoon tot persoon, bijvoorbeeld vanwege individuele verschillen in noradrenalineniveaus voordat de drug werd ingenomen. Deze en andere mogelijkheden kunnen worden onderzocht in toekomstig onderzoek.

Nieuwsgierigheid

De hierboven beschreven theorie stelt dat veranderingen in LC/noradrenaline-activiteit worden veroorzaakt door veranderingen in de kosten en baten van de huidige gedragsstrategie. De verhoogde LC-activiteit in de tonic LC mode, en de hieruit voortkomende toename in exploratiedrag, worden volgens de theorie veroorzaakt door een afname in de opbrengsten en/of toename in de kosten van het huidige gedrag. Veranderingen in taakgerelateerde kosten en baten zijn echter niet de enige factoren die aanleiding kunnen geven tot exploratiedrag. Een andere belangrijke factor is onzekerheid over de identiteit van een stimulus of de gevolgen van een actie. Deze onzekerheid veroorzaakt vaak nieuwsgierigheid, wat kan worden gedefinieerd als de emoties en cognities die geassocieerd zijn met het verlangen om iets nieuws of onbekends te weten te

komen. Hoewel nieuwsgierigheid een belangrijke rol speelt in veel aspecten van gedrag is er nog maar weinig bekend over de onderliggende hersenmechanismen.

Volgens een klassieke psychologische theorie over nieuwsgierigheid (Berlyne, 1954) ontstaat nieuwsgierigheid naar een bepaalde stimulus/gebeurtenis wanneer men onzeker is over de identiteit of de uitkomst van deze stimulus/gebeurtenis. De nieuwsgierigheid is groter naarmate iemand meer tegenstrijdige hypothesen over de onzekere stimulus/gebeurtenis heeft, en naarmate de waarschijnlijkheden van de verschillende hypothesen dichter bij elkaar liggen. Volgens Berlyne is nieuwsgierigheid een aversieve toestand die gepaard gaat met een verhoogd arousalniveau. Een andere assumptie van Berlyne's theorie is dat de opheffing van nieuwsgierigheid, door het verkrijgen van relevante informatie, een belonend effect heeft. In het onderzoek dat is beschreven in Hoofdstuk 5 van dit proefschrift onderzochten we de hersenactivatie die optreedt tijdens nieuwsgierigheid, en de hersenactivatie die optreedt wanneer deze nieuwsgierigheid wordt opgeheven. Om perceptuele nieuwsgierigheid op te wekken lieten we proefpersonen onduidelijke plaatjes bekijken. In sommige gevallen werd de nieuwsgierigheid hierna opgeheven door het bijbehorende duidelijke plaatje te vertonen. Terwijl de proefpersonen de plaatjes bekeken maten we hun hersenactivatie met behulp van functional magnetic resonance imaging (fMRI).

De fMRI resultaten toonden aan dat nieuwsgierigheid gepaard ging met activatie van hersengebieden die gevoelig zijn voor arousal, conflict en andere negatieve ervaringen. De activatie van deze hersengebieden was het sterkst bij proefpersonen met een nieuwsgierige persoonlijkheid (gemeten met een nieuwsgierigheidsvragenlijst). De opheffing van nieuwsgierigheid leidde tot activatie in hersengebieden die betrokken zijn bij de verwerking van beloningen. Daarnaast vonden we dat de opheffing van onzekerheid de hippocampus activeerde, een hersengebied dat een belangrijke rol speelt in geheugenvorming. Ook werden de plaatjes die in eerste instantie geassocieerd waren met onzekerheid later beter onthouden. Deze resultaten leveren bewijs op een neurobiologisch niveau voor het idee dat nieuwsgierigheid een aversieve conditie is die leidt tot een verhoogd arousalniveau, en dat de opheffing van deze conditie belonend werkt. Ook hebben de resultaten een aantal praktische implicaties; de bevinding dat plaatjes die in eerste instantie met onzekerheid geassocieerd waren beter werden onthouden suggereert bijvoorbeeld dat mensen lesmateriaal beter zullen onthouden als ze van te voren nieuwsgierig zijn gemaakt over dit materiaal.

Acute arousal effecten, temporele zekerheid en informatieverwerking

De veranderingen in LC/noradrenaline-activiteit die samenhangen met veranderingen in de exploitatie-exploratie balans ontwikkelen zich over periodes van meerdere seconden of minuten; dit worden wel tonische effecten genoemd. Zoals eerder beschreven treden er naast deze relatief langzame veranderingen ook hele snelle toenames in LC-activiteit op wanneer er een taakrelevante stimulus wordt waargenomen (de phasic LC reponse). De phasic LC response treedt niet alleen op na taakrelevante stimuli, maar ook na stimuli die erg intens, onverwacht of emotioneel geladen zijn;

dus eigenlijk na alle stimuli waarop een snelle reactie van belang zou kunnen zijn voor overleving. In het onderzoek dat is beschreven in Hoofdstuk 6 van dit proefschrift onderzochten we de effecten van één categorie van deze stimuli, namelijk onverwachte harde tonen.

Experimenten waarbij proefpersonen zo snel mogelijk een beslissing moeten maken over een visuele stimulus (bijvoorbeeld de simpele beslissing of een reeks letters wel of geen bestaand woord vormt) hebben aangetoond dat deze beslissingen sneller worden gemaakt als er tegelijkertijd met de visuele stimulus een opvallende stimulus in een andere modaliteit, zoals een harde toon, wordt aangeboden. Hoewel deze extra stimulus compleet irrelevant is voor de beslissing die genomen moet worden over de visuele stimulus, versnelt deze toch de reactie op de visuele stimulus; dit wordt het *accessory stimulus effect* genoemd. Het accessory stimulus effect wordt toegeschreven aan een tijdelijke verhoging van arousal veroorzaakt door de extra stimulus, maar het is niet duidelijk welk gedeelte van het informatieverwerkingsproces hierdoor versneld wordt. In de experimenten die zijn beschreven in Hoofdstuk 6 van dit proefschrift onderzochten we deze vraag. Eerst onderzochten we met behulp van elektro-encefalografie (EEG) hoe de hersenactiviteit tijdens het maken van simpele beslissingen over visuele stimuli werd beïnvloed door irrelevante harde tonen. Met EEG is het mogelijk om het tijdsverloop van de hersenactiviteit die gepaard gaat met de verwerking van de visuele stimulus heel precies te bekijken. De resultaten lieten zien dat bepaalde componenten van hersenactiviteit die redelijk snel na de stimulus optreden, en te maken hebben met de perceptuele verwerking van de visuele stimulus, sneller plaatsvonden als er ook een toon werd aangeboden. De timing van latere componenten, die betrokken zijn bij het maken van een beslissing en het reageren op de stimulus, werd echter niet beïnvloed door de tonen. Dit duidt erop dat de irrelevante tonen vroege perceptuele processen versnelden, maar geen effect hadden op de snelheid van latere stadia van de informatieverwerking. Deze conclusie werd ook ondersteund door de resultaten van een tweede experiment, waarin we een model gebruikten om de effecten van de harde tonen op verschillende onderdelen van de informatieverwerking te onderscheiden.

Omdat de tonen in de experimenten uit Hoofdstuk 6 tegelijkertijd werden aangeboden met de visuele stimuli waarop gereageerd moest worden, konden de tonen niet gebruikt worden om het moment waarop de visuele stimulus zou verschijnen te voorspellen. De versnelde reactie op visuele stimuli die gepaard gingen met een harde toon (het accessory stimulus effect) was daarom geen gevolg van een bewuste strategie van de proefpersoon, maar trad automatisch op. In Hoofdstuk 7 van dit proefschrift onderzochten we een ander effect dat hier mogelijk aan gerelateerd is: het fenomeen dat mensen sneller op een stimulus reageren als ze het moment waarop deze stimulus verschijnt goed kunnen voorspellen. Een alledaags voorbeeld hiervan zijn wachttijdindicatoren bij stoplichten die aangeven hoe lang het nog duurt voor het stoplicht op groen springt; hierdoor kunnen mensen anticiperen op het moment dat het stoplicht groen wordt en heel snel reageren wanneer dit gebeurt. Een aantal studies heeft aangetoond dat temporele anticipatie samengaat met veranderingen in alertheid, en dat het noradrenalesysteem hier een belangrijke rol in speelt. De experimenten in Hoofdstuk 7 toonden aan dat temporele zekerheid (voorspelbaarheid van het moment waarop de stimulus verschijnt) de vroege perceptuele verwerking van een stimulus en/of de

late motorische reactie op een stimulus versnelt. Temporele zekerheid bleek echter geen invloed te hebben op tussenliggende stadia van de informatieverwerking die betrokken zijn bij het maken van een beslissing over de stimulus. De overeenkomsten tussen de effecten van temporele zekerheid (Hoofdstuk 7) en de effecten van een automatische toename in arousal veroorzaakt door accessory stimuli (Hoofdstuk 6) suggereren dat beide effecten door dezelfde hersenmechanismen worden gemoduleerd.

D β H-deficiëntie en cognitie

Alle tot nu toe beschreven studies onderzochten arousal-gerelateerde cognitieve processen en hersenfunctie in gezonde personen. Gezien de belangrijke rol van het noradrenalesysteem in de regulatie van cognitie en gedrag is het niet verrassend dat disfuncties van dit systeem kenmerkend zijn voor verschillende psychiatrische aandoeningen, zoals depressie en ADHD. Een interessante disfunctie van het noradrenalesysteem in dit opzicht is *dopamine- β -hydroxylase (D β H) deficiëntie*. D β H-deficiëntie is een zeldzame stoornis waarbij het enzym dat verantwoordelijk is voor de omzetting van dopamine in noradrenaline niet werkt. Patiënten met deze stoornis hebben hierdoor totaal geen noradrenaline in hun lichaam of hersenen, maar dit gebrek aan noradrenaline kan worden hersteld door middel van medicatie. Patiënten met D β H-deficiëntie hebben zonder medicatie veel last van lichamelijke klachten, maar ze hebben opmerkelijkwijs geen opvallende cognitieve afwijkingen. Het cognitief functioneren van patiënten met D β H-deficiëntie was tot kortgeleden echter nooit systematisch onderzocht, waardoor het niet duidelijk was of deze patiënten misschien toch subtiele afwijkingen op neurocognitief gebied hebben.

Het onderzoek dat is beschreven in Hoofdstuk 4 van dit proefschrift was het eerste onderzoek naar het neurocognitief functioneren van patiënten met D β H-deficiëntie. We namen een testbatterij van cognitieve taken af bij vijf patiënten met D β H-deficiëntie, één keer met en één keer zonder medicatie. Ook onderzochten we de hersenactiviteit van de patiënten met behulp van EEG, en maakten we MRI-scans van hun hersenen om eventuele afwijkingen in hersenstructuur te onderzoeken. Alle resultaten van de patiënten werden vergeleken met die van een gezonde controlegroep. De MRI-resultaten toonden aan dat de patiënten een kleiner hersenvolume hadden dan de controlegroep, wat klopt met het idee dat noradrenaline de afbraak van neuronen tegengaat. De prestatie van de patiënten op de meeste taken van de cognitieve testbatterij week echter niet af van die van de controlegroep, en werd niet beïnvloed door medicatie. De patiënten presteerden zelfs normaal op cognitieve taken waarbij noradrenaline normaalgesproken een belangrijke rol speelt. Dit suggereert dat de hersenen van patiënten met D β H-deficiëntie compensatiemechanismen hebben ontwikkeld voor de afwezigheid van noradrenaline; wat deze compensatiemechanismen precies zijn is een interessante vraag voor toekomstig onderzoek.

Dankwoord

Veel mensen hebben bijgedragen aan de totstandkoming van dit proefschrift. Dank aan:

Sander voor je grote betrokkenheid, ideeën en feedback tijdens alle stadia van het onderzoek, en voor het meedenken over de toekomst.

Iedereen met wie ik de afgelopen jaren heb samengewerkt, in het bijzonder Eric-Jan, Guido, Erik en Joop van het CHDR, Serge, Dilène en Wery, Chris en alle artsen die aan het patiëntenonderzoek hebben meegewerkt.

Rinus en Jun voor de leuke tijd die we hebben gehad in 2A45 en daarbuiten. Rinus voor (woord)grappen en het beantwoorden van Matlab en E-prime vragen. Jun, thanks for always being a nice and positive roommate, and for conspiring with me against Rinus.

Henk, Anne, Kalinka, Stef, Stephen, Jesse, Mikael, Pascal, Soghra, Mischa, Saskia, André, Michiel, Merel, Nelleke en andere (ex)collega's voor alle gezellige activiteiten binnen en buiten FSW.

Henk, ook bedankt dat je altijd klaar stond om te helpen met technische zaken; jouw ingenieuze scriptjes hebben veel van mijn analyses een stuk aangenamer gemaakt.

Albertien voor je hulp en daadkrachtige optreden als er iets geregeld moest worden.

Familie en vrienden voor het meeleven, afleiding en gezelligheid buiten het werk.

Bedankt allemaal!

Curriculum Vitae

Marieke Jepma was born in Utrecht on April 30th, 1983. In 2001 she completed secondary education at the Openbare Scholengemeenschap de Driemark in Winterswijk. She graduated in Cognitive Psychology at the University of Groningen in 2005, after which she completed a Master in Cognitive Neuroscience at the University of York, UK. In October 2006 she started to work as a PhD student at Leiden University where she studied the neurocognitive mechanisms involved in arousal, exploration and temporal attention, which resulted in this thesis. In June 2011 she will start working as a postdoc at the Cognitive and Affective Neuroscience Lab in Boulder, Colorado, in collaboration with the Laboratory of Neural Computation and Cognition at Brown University, to investigate the neurocognitive and computational mechanisms involved in pain-avoidance learning.