University of Groningen

Less is more

Wierda, Stefan Mark

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2014

Citation for published version (APA):
Wierda, S. M. (2014). *Less is more*. [Thesis fully internal (DIV), University of Groningen]. s.n.

Less is more

Strategic restrictions in temporal attention

Stefan M. Wierda

rijksuniversiteit
groningen

Less is more
Strategic restrictions in temporal attention

Proefschrift

ter verkrijging van de graad van doctor aan de
Rijksuniversiteit Groningen
op gezag van de
rector magnificus, dr. E. Sterken,
en volgens besluit van het College voor Promoties.

De openbare verdediging zal plaatsvinden op

woensdag 5 februari 2014 om 11.00 uur

door

Stefan Mark Wierda

geboren op 4 februari 1984
te Steenwijk

voor pa, voor moeders

Contents

In which I introduce the attentional blink and discuss the outline of this thesis

The environment that we cope with every day is complex and ever changing. With the current pace of development in technology, the stream of information entering our minds from tactile, auditory, and visual modalities rapidly increases. Relevant and irrelevant information enters our cognitive system, but only a fraction of that information is stored in memory or is used to determine actions. It is remarkable how the human mind can filter a vast amount of information and selectively process relevant information in a split second. How do we recognize and identify the useful pieces of information? An important mechanism in this process is that of attention. Temporal attention, particularly its underlying mechanisms and limitations, will be the focus of this thesis.

In 1987, Broadbent and Broadbent set out to investigate how much relevant information we can process in a short amount of time (Broadbent & Broadbent, 1987). They conducted an experiment in which one or two relevant words were presented in a rapid visual presentation (RSVP) stream of irrelevant words. Each word was presented for only 100 milliseconds, and within the list of irrelevant words (i.e., the distractors), either one or two relevant words were embedded (i.e., the targets). They found that whenever the second target (T2) was presented within approximately half a second after the first target (T1), performance in identifying the T2 dropped drastically. Since the effect critically depends on the presence of the T1, they argued that processing the T1 interfered with that of T2 for at least half a second.

The attentional blink

Half a decade later, Raymond, Shapiro, and Arnell (1992) performed several experiments to investigate the dip in T2-performance found by Broadbent and Broadbent in more detail. Raymond et al. replicated the effect, but also showed that the phenomenon was absent when they instructed participants to ignore T1. This subtle change in task instructions showed that the deficit could not be due to a sensory limitation as the visual input was identical in both condition, and thus the deficit was thought to be attention-related. Therefore, they dubbed the phenomenon the attentional blink (AB, see Box 1.1 for a typical AB experiment). Note that when T2 directly followed T1, without any intervening stimuli, the deficit did not occur. Because T2 is apparently spared from the AB when it is presented directly after T1 (also called "lag 1", in which lag codes for the serial position of T2 relative to T1), this effect became known as lag-1 sparing (Potter, Chun, Banks, & Muckenhoupt, 1998; Potter, Staub, & O Connor, 2002). To explain the AB, Raymond et al. proposed an attentional-gate theory. Whenever an item of interest (e.g., a target) is encountered in a stream of stimuli, an attentional gate assumedly opens in order to enhance attention and allow identification of the item. However, the gate closes during the identification process, effectively inhibiting all incoming stimuli. This process was thought to take approximately 500 ms, matching the duration of the AB. Whenever a T2 directly follows T1, it slips through the gate together with T1, thus explaining the lag-1 sparing.

Following the experiments of Raymond et al., a wide range of researchers performed hundreds of AB experiments, and many competing theories explaining the AB were proposed (Bowman & Wyble, 2007; Chun & Potter, 1995; Dehaene, Sergent, & Changeux, 2003; Dux & Harris, 2007; Hommel et al., 2006; Jolicœur & Dell'Acqua, 1998; Olivers & Meeter, 2008; Potter et al., 2002; Shih, 2008; Taatgen, Juvina, Schipper, Borst, & Martens, 2009; Wyble, Bowman, & Nieuwenstein, 2009). To summarize the theories, I will roughly divide them in two

types of theories: Limited resources theories and Attentional control theories. To give a brief overview of both types of theories, I will discuss a selection of theories that are representative for the type of theory in more detail, including their challenges. For an extensive review of data and theory, I recommend the reviews by Dux and Marois (2009) and the more recent review of Martens and Wyble (2010).

Box 1.1. The attentional blink paradigm
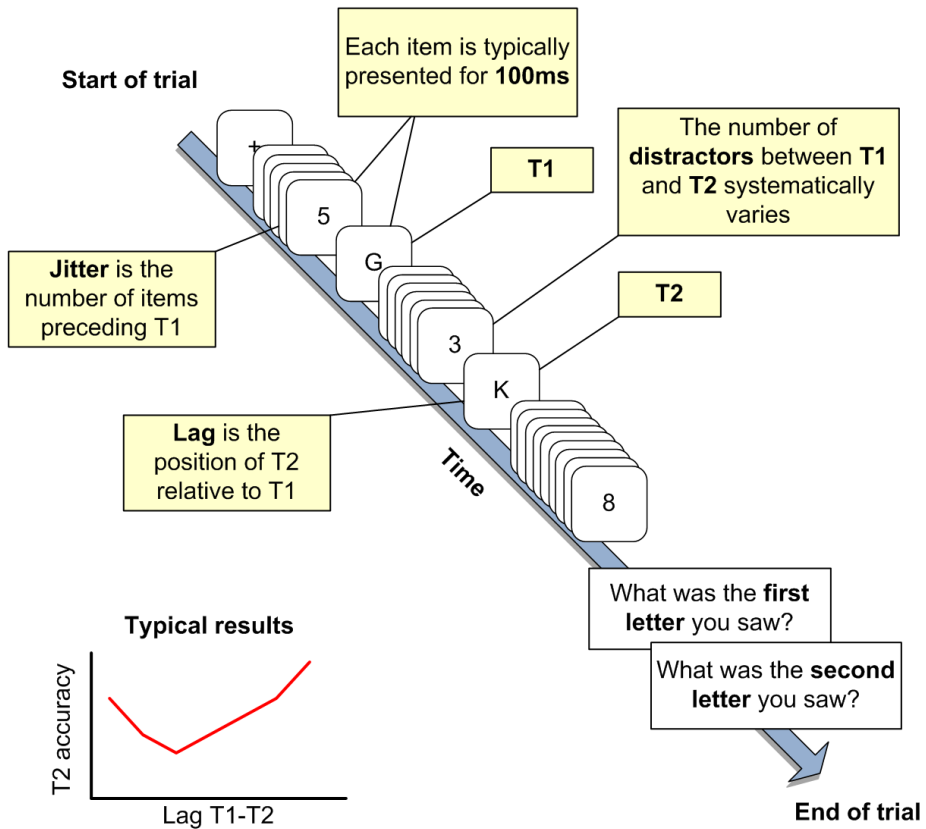
In an AB task, a RSVP stream of distractors (e.g., digits) and targets (e.g., letters) is presented in the middle of a computer screen. Typically, first a fixation cross is presented in the middle of the stream, and the participant starts the task by pressing a key. Then each item appears in the middle of the screen for about 100 ms. After a certain amount of items (jitter) T1 appears. Then, a variable number of distractors are presented followed by T2. Following T2, usually a number of additional distractors are presented. After presentation of the stream, participants have to report which targets they saw. The position of T2 relative to T1 is referred to as lag. When the lag is short (< 500 ms), performance deteriorates. However, when T2 is presented directly after T1, T2 is typically spared from the AB. This effect is called Lag-1 sparing. The combination of the sparing effect at the one hand, and the drop in performance at the other hand, typically results in a U-shaped performance curve.

Limited resources theories: A fundamental AB

In limited-resource theories, the AB is caused by competition for resources amongst targets and, in some theories, distractors. For example, in the Interference theory by Shapiro, Raymond, and Arnell (1994), targets and non-targets compete for resources in working-memory. In their theory, targets that match a template enter working memory together with the trailing non-target. Due to the limited resources in working memory, items that quickly succeed one another receive less activation, as do items that are similar to each other. Therefore, during the AB critical period, the T2 does not receive as much activation as T1, and directly competes with T1+1 and T2+1 because of the similar amount of activation they receive.

Most other limited-resources theories of the AB (e.g., Dux & Harris, 2007; Jolicœur & Dell'Acqua, 1998; Potter et al., 2002) are based on the two-stage processing theory as proposed by Chun and Potter (Chun & Potter, 1995). In this theory, Stage 1 concerns the detection of targets, which is a fast process that is executed in parallel. However, the resulting representations are fragile and prone to decay. Stage 2 is responsible for forming a reportable perception, and thus can be considered as the consolidation of the target into working memory. The stage of consolidation does not begin until the first stage of identification is finished. In addition, the second stage has a limited-capacity in the sense that targets are consolidated serially. All items that are available at the start of Stage 2 can be processed together. That is, the target and the item immediately succeeding the target will often both be processed by this stage. This explains how the presentation of two successive targets (i.e., lag-1) results in accurate performance for both targets (sparing). On the other hand, if an item is presented while an earlier Stage 2 processing cycle is not yet finished, that item is more likely to decay, resulting in the AB.

Challenges to limited resource theories

A common theme in limited resource theories is that the AB is a fundamental bottleneck—a temporal lack of resources, selective attention, or working memory capacity—in our cognitive system. This view seems to be supported by earlier findings that the AB is not affected by extensive practice (Maki & Padmanabhan, 1994; Taatgen et al., 2009). However, a recent study showed that if you train individuals using a saliency-training paradigm, it is possible to eliminate the AB (Choi, Chang, Shibata, Sasaki, & Watanabe, 2012) under certain conditions (Tang, Badcock, & Visser, In Press). In addition, studies that focus on individual differences reveal a more nuanced picture of the AB, one that might indicate that it is not necessarily a fundamental bottleneck. For example, several studies have shown that some people do not show the deficit at all (Feinstein, Stein, Castillo, & Paulus, 2004; Martens, Munneke, Smid, & Johnson, 2006; Martens, Johnson, Bolle, & Borst, 2009; Martens, Korucuoglu, Smid, & Nieuwenstein, 2010; Willems, Wierda, van Viegen, & Martens, 2013). Other challenges come from studies in which a change in task instructions drastically improves performance on the AB task. For example, a study showed that the AB was attenuated when participants had to combine both targets and give a single response (e.g., the targets were digits and the participants had to report the sum). Furthermore, when participants are distracted or an additional task is added consequently to the AB task, performance surprisingly improves (Arend, Johnston, & Shapiro, 2006; Olivers & Nieuwenhuis, 2005; Olivers & Nieuwenhuis, 2006; Taatgen et al.,

2009; Wierda, van Rijn, Taatgen, & Martens, 2010). The aforementioned studies show that there is more to the AB than merely a cognitive bottleneck, and that there is thus more to the AB than what is proposed by limited resources theories.

Attentional control theories: A strategic AB

Due to the findings that pose serious challenges for bottleneck accounts of the AB, theories were proposed that explain the AB in terms of attentional control. In these theories, the AB is more seen as a functional mechanism that is strategically deployed by the cognitive system. For example, in the Boost and Bounce theory by Oliver and Meeters (2008), the AB is a result of an attentional filter that seeks to enhance relevant information and suppresses irrelevant information. Whenever a relevant stimulus (i.e., a target) is encountered, the activation of that stimulus is enhanced (i.e., the boost). However, upon encountering an irrelevant stimulus (i.e., a distractor), incoming stimuli are inhibited to filter out the irrelevant information (i.e., the bounce).

Another example of an attentional control theory with a function explanation of the AB is the episodic simultaneous type, serial token (eSTST) model of Wyble and colleagues (Wyble et al., 2009). In their model, the AB is due to a mechanism that facilitates the encoding of episodically distinct visual representations. When a target is encountered in an RSVP stream, activity for that target is accumulated through a gating mechanism. When the accumulated activity exceeds a certain threshold, the target is recognized and consolidated by the model. During the encoding process and before the threshold is reached, the blaster—a mechanism that can amplify the activation of incoming stimuli—facilitates processing of the target by enhancing the input. During this stage, incoming stimuli can still be processed, but at a cost of episodic distinctiveness (i.e., you cannot tell which target was presented first). However, further in the encoding stage, the blaster starts to inhibit incoming stimuli to create distinct episodes between the target and trailing items. As a result, a second target that occurs during the inhibition of the blasting will be missed, resulting in an AB.

A final example is the threaded cognition model of the AB, which is based on the ACT-R cognitive architecture (Taatgen et al., 2009). When the model encounters a target, it consolidates the target into working memory. When an incoming distractor item is recognized during the consolidation of a target item, the processing of incoming items is strategically blocked to protect the consolidation process of T1. When a target item is presented during the period that processing is blocked, an AB occurs. If a second target immediately follows a first target, they are processed serially and thus lag-1 sparing occurs. The key assumption is that the assertion of control to protect consolidation of T1 is in fact unnecessary, and that this overexertion of control—that happens in most people for some unknown reason—leads to the AB.

Challenges to attentional control theories

Although the evidence in favor of attentional control theories is strong, there are concerns that need to be addressed regarding this type of theories. For example, under conditions in which T1 strongly captures attention, an AB is observed even when targets are presented

without intervening distractors (Dux, Asplund, & Marois, 2008; Dux, Asplund, & Marois, 2009). In addition, performance severely deteriorates when multiple targets are presented in a short period, compared to when targets are presented more distant in time, indicating some kind of resource limitation (Dell'Acqua, Dux, Wyble, & Jolicœur, 2012). Thus, even though the AB might be due to an attention control strategy, one still has to take into account certain structural limitations imposed by the cognitive system that can influence processing of targets both during and outside the period of the AB.

Overview of remaining chapters

The current chapter served as an introduction to the AB. In the remaining part of the thesis, I will discuss behavioral and electrophysiological findings that point towards a control strategy theory of the AB. In Chapter 2, an ERP study is presented in which a concurrent secondary task was added to the AB task. Evidence is presented that favors an attentional-control explanation of the AB over a limited-capacity explanation of the AB. Next, the difference between non-blinkers and blinkers is explored by comparing their temporal profile in Chapter 3. While non-blinkers showed no sign of suppression, blinkers showed only a modest amount of suppression. Furthermore, it seems that the temporal selection mechanisms of non-blinkers are more precise and faster than that of blinkers. Chapter 4 puts the threaded-cognition model to the test, and shows that with some minor modification, the model can explain counter-intuitive results that pose a problem for limited-capacity theories. Finally, I describe a novel method called pupil dilation deconvolution in Chapter 5. The method can isolate the slow pupillary response per cognitive event, and can thus reveal changes in attention and workload at a high temporal resolution. The results demonstrate that the method can reveal even subtle changes in attention in an AB task, along with expectancy effects. To conclude, the implications and impact of the four empirical chapters are discussed in Chapter 6.

Chapter 2: Distracting the mind improves performance

An ERP study

In which we show that less effort increases performance

Chapter was previous published in:

PLoS One, 2010

Stefan M. Wierda

Hedderik van Rijn

Niels Taatgen

Sander Martens

Abstract

Background

When a second target (T2) is presented in close succession of a first target (T1), people often fail to identify T2, a phenomenon known as the attentional blink (AB). However, the AB can be reduced substantially when participants are distracted during the task, for instance by a concurrent task, without a cost for T1 performance. The goal of the current study was to investigate the electrophysiological correlates of this paradoxical effect.

Methodology/Principal Findings

Participants successively performed three tasks, while EEG was recorded. The first task (standard AB) consisted of identifying two target letters in a sequential stream of distractor digits. The second task (grey dots task) was similar to the first task with the addition of an irrelevant grey dot moving in the periphery, concurrent with the central stimulus stream. The third task (red dot task) was similar to the second task, except that detection of an occasional brief color change in the moving grey dot was required. AB magnitude in the latter task was significantly smaller, whereas behavioral performance in the standard and grey dots tasks did not differ. Using mixed effects models, electrophysiological activity was compared during trials in the grey dots and red dot tasks that differed in task instruction but not in perceptual input. In the red dot task, both target-related parietal brain activity associated with working memory updating (P3) as well as distractor-related occipital activity was significantly reduced.

Conclusions/Significance

The results support the idea that the AB might (at least partly) arise from an overinvestment of attentional resources or an overexertion of attentional control, which is reduced when a distracting secondary task is carried out. The present findings bring us a step closer in understanding why and how an AB occurs, and how these temporal restrictions in selective attention can be overcome.

Introduction

Although the human mind is quite capable of performing multiple tasks at the same time, the multitasking brain does not always react accurate or fast enough in complex situations. An obvious example is that the likelihood of traffic accidents increases when driving is combined with the concurrent use of a cellular phone, especially when the level of complexity increases and additional attentional control is required (Strayer & Johnston, 2001). However, under some circumstances described below, an increase in cognitive load can reduce temporal restrictions in attention, such that one's multitasking performance improves rather than deteriorates.

Within the lab, restrictions in temporal attention are for instance revealed when two targets (e.g., letters) are presented in close temporal proximity within a sequential stream of distractor stimuli (e.g., digits). When the second target (T2) is presented within ~200 to 500 ms after the onset of the first target (T1), participants often fail to report the second target, reflecting the occurrence of an attentional blink (AB; Raymond et al., 1992). It has been shown, however, that identification performance can increase when a second task is added to this so-called AB task (Arend et al., 2006; Olivers & Nieuwenhuis, 2005; Olivers & Nieuwenhuis, 2006; Taatgen et al., 2009). The aim of the current study was to investigate this paradoxical phenomenon in more detail.

For two decades, the AB has been a major topic in attention research because it is informative about the rate at which stimuli can be encoded into consciously accessible representations. As the AB can be obtained using a variety of stimuli and task conditions, it is thought to reflect a very general property of perceptual awareness (for a review, see Martens & Wyble, 2010).

However, recent findings suggest that the AB is unlikely to reflect a hard-wired bottleneck. For instance, the presence of perceptual, spatial, or temporal cues has been found to reduce the magnitude of the AB (defined as the percentage of decrement in T2 performance within the AB period relative to T1 performance), presumably by redistribution or accelerated allocation of attention (e.g., Martens & Johnson, 2005; Nieuwenstein, Chun, van der Lubbe, & Hooge, 2005; Olivers & Meeter, 2008). As mentioned above, it has also been found that adding a source of distraction during the AB task can attenuate the AB. In an experiment by Olivers and Nieuwenhuis (2005), participants had to do an AB task in four different conditions. In one condition, participants did a standard AB task. In the other three conditions participants were financially rewarded for their performance, do free association (e.g., think about their most recent holiday) or listen to music during the AB task. In the latter two conditions, the AB was substantially reduced, while rewarding the participants did not make any difference. Although the initial finding could not be fully replicated (Olivers & Nieuwenhuis, 2006), others have shown that task-irrelevant visual motion or flicker (Arend et al., 2006), a change in task instruction (Ferlazzo, Lucido, Di Nocera, Fagioli, & Sdoia, 2007), or even a concurrent secondary task (Taatgen et al., 2009) can attenuate the AB as well.

In the study of Taatgen et al. (2009), participants had to perform an AB task together with a concurrent secondary task that required participants to detect the occurrence of a red dot. The AB task consisted of identifying target letters amongst distractor digits. The detection task contained a grey dot presented in peripheral vision that moved in a circular direction concurrent to the stimulus stream of the AB task. In 25% of the trials, the grey dot turned red for a brief moment. At the end of a trial, participants were instructed to report whether or

not a red dot occurred during the trial. Trials in which the red dot actually turned red were excluded from analyses, as the red dot itself induced an AB. AB magnitude was found to be smaller in the concurrent task condition than in a control condition without dots.

Supported by computer simulations, Taatgen et al. (2009) argued that an AB might be caused by an overexertion of cognitive control, which is suspended when a secondary task is concurrently performed during the AB task. In addition, their computational model provided a first explicit account of target selection processes in individuals who do not show an AB, referred to as 'non-blinkers' (Martens et al., 2006).

Whereas individual differences in AB research are usually ignored, Martens, Munneke, Smid, and Johnson (2006) demonstrated the existence of large individual differences in the magnitude of the AB, with some individuals showing little or no AB whatsoever ('non-blinkers'). Measuring event-related potentials (ERP), they found significant electrophysiological differences between non-blinkers and individuals who do show a strong AB ('blinkers'). More specifically, non-blinkers showed earlier target-related parietal activity (reflected in the P3, a component associated with the updating of working memory). In addition, they exhibited more target-related frontal activity (reflected in the FSP, a component associated with early target selection processes), as well as reduced distractor-related frontal activity.

Martens, Elmallah, London, and Johnson (2006) correlated the magnitude of the AB with the size of the P3 amplitude. In their experiment, the frequency of the first target in a standard AB paradigm was manipulated. Whenever an infrequent first target was presented, the P3 amplitude evoked by T1 and the AB magnitude increased. Martens and colleagues suggested that the amplitude of the P3 of T1 reflects the amount of attention or resources allocated to consolidate T1 (Kok, 2001; Martens et al., 2006; McArthur, Budd, & Michie, 1999; Shapiro, Schmitz, Martens, Hommel, & Schnitzler, 2006) and that the more resources allocated to T1, the smaller the chance for T2 to receive sufficient attention, reflected in the occurrence of an AB.

According to Taatgen and colleagues (2009), most individuals tend to protect memory consolidation of a first target by temporarily blocking the detection of subsequent targets, resulting in an AB. They suggested that non-blinkers refrain from exerting this protective control, thus avoiding the AB. This notion could explain the paradoxical finding of improved performance during dual-tasking: When blinkers are given a secondary task during the AB, the brain might be taxed to such an extent that the exertion of protective control is relaxed, resulting in an attenuated AB (also see Olivers & Nieuwenhuis, 2006).

An intriguing question is whether these 'induced non-blinkers' (i.e. individuals showing a reduced AB magnitude due to distraction manipulations) adopt a similar processing strategy as the natural non-blinkers reported by Martens and colleagues (2006). If so, induced non-blinkers should show similar patterns of frontal and parietal brain activity as previously observed in natural non-blinkers. That is, a shift in target-related P3 latencies would be expected, along with a larger target-related FSP component, as well as reduced distractor-related activity (Martens et al., 2006). In addition, because a lower P3 amplitude of the first target is associated with a smaller AB magnitude, one would expect the P3 amplitude to be lower in the concurrent tasks condition than in the standard AB task condition.

However, different patterns of brain activity might be expected if natural and induced non-blinkers are qualitatively different in the way they perform the AB task. The aim of the current study was to test this and to reveal what influence the addition of a secondary task to a standard AB task has on the various components in the ERP. To that end, we adapted the

task used by Taatgen et al. (2009) to include event-related potential (ERP) recordings, and investigated the FSP, P3, and overall distractor-related activity. Contrary to Taatgen et al., we compared performance on trials in the experimental condition with that of trials in an additional control condition that featured identical input and only differed in task instructions, allowing a direct comparison of behavioral as well as ERP data.

Methods

Participants

Thirty psychology students from the University of Groningen were recruited via an online sign-up program. They received course credits for their participation in the experiment. The Neuroimaging Center Institutional Review Board and the Ethics Committee of Psychology at the university of Groningen approved the experimental protocol. Written informed consent was obtained prior to the experiment. Five participants were excluded from the analyses due to bad performance (i.e. they had an accuracy <50% on the first target). Another participant was excluded from the analyses due to artifacts in more than 50% of the EEG data segments. The remaining twenty-four participants were aged 18 to 33 (mean = 21.25) with normal or correct-to-normal visual acuity.

Stimuli and Apparatus

The software package E-Prime was used to generate stimuli and to collect responses, running under Windows XP on a PC with a 17-inch monitor. Stimuli consisted of consonants (excluding "Q", "V", and "Y") and digits (excluding "0" and "1") and were presented in black on a white background in a bold 12-point Courier New font subtending 0.3° by 0.4° of visual angle at a viewing distance of approximately 50 cm. In the AB task with grey dots distractors and in the AB task with red dot detection (described below), grey (40.2 cd/m2) dots with a diameter of 10 pixels were used. The red (28.8 cd/m2) dot used in the red dot detection task had the same size.

Procedure

The experiment consisted of four practice blocks and three experimental blocks containing one of three different tasks: a standard AB task, an AB task with a concurrently moving irrelevant peripheral grey dot (grey dots task), or an AB task in which a short color change in the peripheral dot had to be detected (red dot task), as described below. The order of these blocks was counterbalanced between participants.

The first practice block always contained the standard AB task, and consisted of 108 trials. In each subsequent testing block, participants performed one of the three different tasks, preceded by a practice block of that specific task. One testing block (containing the red dot task) consisted of 288 trials whereas the other two blocks consisted of 216 trials. Each

of the preceding practice blocks consisted of 12 trials. Between each block, the participants could take a short break. After half of the trials in each testing block, participants could take another short break.

The standard AB task required the identification of two letter targets amongst a rapid stream of digit distractors. Prior to each trial, a fixation cross was presented in the middle of the screen. Participants were instructed to fixate on the cross and to press the space bar to start the trial. After pressing the space bar, the fixation cross remained on the screen for a duration of 750 ms, followed by a blank screen. After 100 ms, a rapid serial visual presentation (RSVP) stream appeared in the middle of the screen, consisting of 22 sequentially presented stimuli. Each stimulus appeared for 90 ms without inter stimulus interval. In two-thirds of the trials, two targets were presented. In one-third of the trials, no targets were presented. In all target trials, the first target (T1) was presented as the fifth item in the stream, with the second target (T2) appearing as either the third (lag 3) or eighth (lag 8) item following T1. Known from the literature, an AB is likely to occur when T2 is presented at lag 3 (i.e. 270 ms after T1 onset), and unlikely to occur when presented at lag 8 (i.e. 720 ms after T1 onset).

The targets were randomly chosen from the set of target letters, with the only constraint that T1 and T2 were always different letters. The distractors were randomly chosen from the set of digits with the constraint that two successive digits were never identical. Participants were instructed to report the two targets at the end of the stream by pressing the corresponding letters on the keyboard. Participants were instructed to take sufficient time in making their responses to ensure that typing errors were not made. Whenever participants did not see a target, they were instructed to press the spacebar instead. Participants were encouraged to type in the letters in the order in which they had been presented, but responses were accepted and counted correct in either order. No feedback was given during the experiment.

The grey dots task was identical to the task described above, except that an irrelevant grey dot circled around the RSVP stream in synchrony with the presentation of each stimulus (i.e. the dot moved every ~90 ms). The grey dot started randomly at one of 39 possible positions in a radius of 11.3° from the middle of the screen. The dot skipped two positions in clockwise rotation each time it moved. Participants were instructed to ignore the grey dot.

The red dot task was identical to the AB task with grey dots, except for the following changes. Participants were instructed to identify the targets in the AB task, while concurrently paying attention to the circling grey dot. In 25% of the trials, a red dot instead of a grey dot was presented for 90 ms, appearing once at a random time interval throughout the trial. After having responded to the target letters, participants were required to report whether or not they had seen the occurrence of a red dot ("press 'j' in case of a red dot, press 'n' in case of no red dot"). Importantly, only the trials during which no red dot appeared (75% of the trials) were considered in the analyses. Note that the perceptual input during these trials is identical to that in the AB task with irrelevant grey dots, with only the instructions being different. As mentioned above, the block of the red dot task contained more trials to keep the number of trials without a red dot occurrence equal to the number of trials in the blocks of the standard and grey dots tasks.

Although the order of the trials within each block was randomized, each condition (lag 3 or 8, targets or no targets, red dot present or absent) was balanced within a testing block. To reduce eye blink artifacts in the EEG recordings, participants were instructed to avoid making eye blinks during a trial until they were prompted to give their responses.

EEG Recording

The EEG signal was recorded using a 64-channel electro-cap with tin electrodes. The electro-cap was organized according to the international 10/20 system and connected to an REFA 8–64 average reference amplifier. Impedance was reduced to less than 5kΩ for all electrodes. Data was sampled with a frequency of 2kHz and digitally reduced to 500Hz. Two electrodes connected to the mastoids served as an offline reference. The horizontal electrooculogram (EOG) was recorded from tin electrodes attached approximately 1 to 2 cm to the left and right of the outside corner of each eye. The vertical EOG was measured from two tin electrodes placed approximately 3 cm below the left eye and 1 cm above the brow of the left eye, respectively. Brain Vision Recorder (Brian Products GmbH, Munich, Germany) was used to control the data acquisition.

Data Analysis

The preprocessing of the EEG data was done using Brain Vision Analyzer (Brain Products). The ERPs were time locked to the onset of the RSVP stream, had a duration of 2000 ms, and were calculated relative to a 200-ms prestream baseline, yielding a total length of 2200 ms. The ERP-segments were 40-Hz low-pass and 0.2-Hz high-pass filtered, and corrected for eye-movements. Segments with value differences greater than 100μV (i.e., containing artifacts) were excluded in the analysis. Artifact rejection excluded 5.7% of all trials (ranging from 0.19% to 25.6% per participant, SD = 6.4).

Both the performance and mean activity of the EEG data were analyzed using linear mixed effects models. The peak latencies were determined by peak detection over the averages per participant, therefore repeated measures ANOVAs were used to analyze peak latencies. The use of mixed effect models in the field of AB research is new, but the method is widely used in other fields (e.g., psycholinguistics, Baayen, Davidson, & Bates, 2008; eye movement data, Barr, 2008; or memory research, Kliegl, Masson, & Richter, 2009). Given that there are large individual differences in the AB and that our hypothesis predicted a different number of observations per cell, it was inappropriate to analyze results using a method that assumes an equal number of observations per cell (Baayen et al., 2008). Therefore, mixed-effects models were used to analyze most datasets, but repeated measures ANOVAs were used to analyze the averaged latency data. Another advantage of mixed effect models is that these models are suited to analyze data from non-normal distributions, such as data based on accuracy scores(-Quené & van den Bergh, 2008). For the ERP-based analyses, the default, normal distributions were used. For the behavioral data, a binominal model was used to fit the accuracy scores. The p-values reported for the non-binominal models of the EEG data were calculated by performing 10000 Markov Chain Monte Carlo (MCMC) samplings. Analyses were performed using the lmer and pvals.fnc functions in the lme4 and languageR packages (Baayen et al., 2008; Bates & Sarkar, 2008) for the statistical software R (2009).

For the analysis of the behavioral data, two mixed-effects models were fitted. In the first model, both the grey dots and red dot tasks were compared to the standard task. Performance on lag 8 trials of the standard task was taken as reference in the first model. In order to investigate differences between the grey dots task and the red dot task directly, a separate mixed-effects model was fitted on the accuracy data with exclusion of the standard task. The

grey dots task at lag 8 was taken as reference in this model. The standard task was excluded from all EEG-data analyses because of the differences in perceptual input induced by the standard task compared to the grey dots and red dot task. As the perceptual input of the grey dots and red dot task is identical, any differences found between the two tasks must be attributed to the manipulation of task instructions.

Results and Discussion

Behavioral Results

Trials in the red dot task during which the red dot actually appeared were excluded from the analysis. On such trials, the red dot itself could possibly induce an AB on the first or second target. Detection accuracy for presence of the red dot was 65.1% (SD = 29.3). Figure 2.1 shows the identification performance on the AB tasks as a function of lag for T1 and T2 given that T1 was correctly reported (T2|T1).

T1 Accuracy

Two binominal mixed effects model were fitted on the accuracy of T1 with Task entered as fixed factor. Subject was entered as a random factor. The first model was fitted on all three tasks, using performance on the standard task at lag 8 as reference. In the second model, data of the standard task were left out in order to directly compare the grey dots task versus the red dot task. Table 2.1 lists the estimates and z-statistics of the first model. The estimates of the coefficients are reported in log odds. The coefficients of Task for the grey dots task and red dot task were not significant in the model, suggesting that both tasks did not affect overall performance. In the second model, the grey dots task was directly compared to the red dot task. In this model, also no significant effect of Task was found.

Table 2.1. The estimates and z-values of the mixed-effects model for T1 accuracy.

	Mixed-effects model T1			
	Estimate	Standard Error	z-value	p-value
Standard (Intercept)	2.111	0.167	12.643	0.000
Grey dots	0.056	0.075	0.749	0.454
Red dot	0.120	0.076	1.587	0.113

T2 Accuracy

In order to examine the effect of Task on the AB, a binominal mixed effect model was fitted on accuracy scores of T2 given correct report of T1. As in the analyses of T1, two models were fitted with Lag and Task entered as fixed factors, and Subject as random factor. As in the analysis of T1 accuracy, the data of the Standard task were left out of the analyses in the second model. The estimates and z-statistics of the first model are listed in Table 2.2. A main
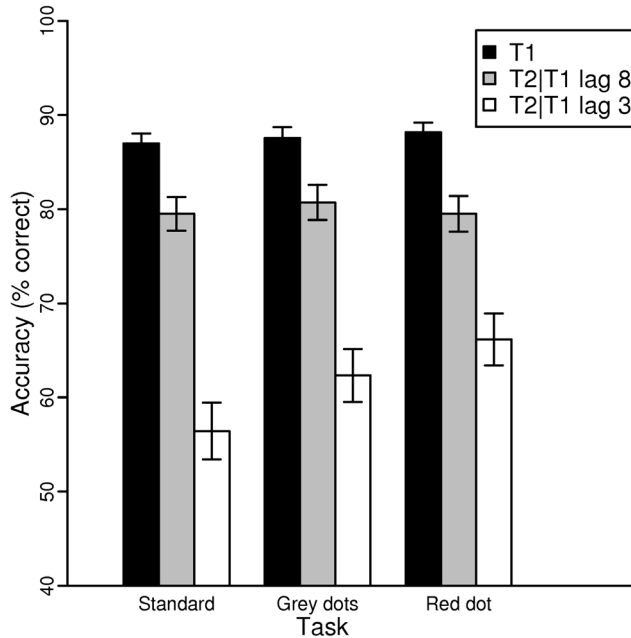
Figure 2.1. Accuracy scores of the standard, grey dots and red dot task for T2 given T1 is correct on lag 8 (dashed line) and on lag 3 (solid line).

effect of Lag was found, such that performance was lower at lag 3 than at lag 8, reflecting the occurrence of an AB. No main effect of Task was found. However, significant Lag×Task interaction effects were found with a positive direction for the red dot task relative to performance on the standard task, reflecting that at lag 3 target performance was higher in the red dot task than in the standard task. No significant differences were found between the standard task and the grey dots task. A second model comparing the grey dots task and the red dot task

Table 2.2. The estimates and z-values of the mixed-effects model for T2|T1 accuracy.

	Mixed-effects model T2|T1			
	Estimate	Standard Error	z-value	p-value
Standard, Lag 8 (Intercept)	1.477	0.145	10.159	0.000
Grey dots	0.086	0.095	0.914	0.361
Red dot	0.017	0.094	0.186	0.852
Lag 3	-1.192	0.086	-13.791	0.000
Grey dots, Lag 3	0.184	0.123	1.496	0.135
Red dot, Lag 3	0.432	0.123	3.523	0.000

directly reveals similar differences as found in the first model. Again, a main effect of Lag is found (z = −11.524, p = 0.000), indicating the presence of the AB effect. In addition, a significant Lag×Task interaction was found, (z = 2.047, p = 0.041), such that lag 3 performance was better in the red dot task than in the grey dots task. No main effect of task was found (p = 0.442). Taken together, a smaller AB effect was observed in the red dot condition compared to the standard AB task and the grey dots task. Although performance seemed better in the grey dots task than in the standard task (see Figure 2.1), no significant difference was found. The grey dots task thus seemed to be a suitable control task for comparison with the red dot task in the ERP analyses.

Electrophysiological Results

As mentioned above, ERP amplitudes were analyzed using mixed effects models. As behavioral performance was similar in the standard AB and the grey dots tasks, we included only the grey dots task and the red dot task in the ERP analyses, as both task included the same perceptual input (as mentioned before, red-dot present trials were excluded from all analyses). Because the P3 was maximal at the Pz electrode, this electrode was used to analyze P3 activity. Based on visual inspection of the grand averages, a time window of 320–640 ms after target onset was used to determine P3 amplitude (in terms of mean activity). For each individual and condition, the latencies of the P3 peaks were obtained by searching for the maximum peak in the before-mentioned time window per condition-based average. To determine the FSP amplitudes, target- and distractor-related activity was analyzed in a window of 180–340 ms after the onset of T1. Target-related activity was obtained by taking activity within that window on trials during which T1 was successfully reported, whereas distractor-related activity was obtained by taking the activity of the non-target trials within the same time window.

The P3 On Lag 8 Trials

Figure 2.2 shows the grand averages of activity at Pz during lag 8 trials, given that T1 and T2 were both correctly identified. The P3 activity of both T1 and T2 was analyzed using a mixed effect model. Task (Grey dots or Red dot) and Target (T1 or T2) were entered in the model as fixed factors. As in the analyses of the behavioral data, Subject was entered as a random

Table 2.3. The estimates and t-values of the mixed-effects model for P3 amplitude (measured in mean activity) at lag 8 trials.

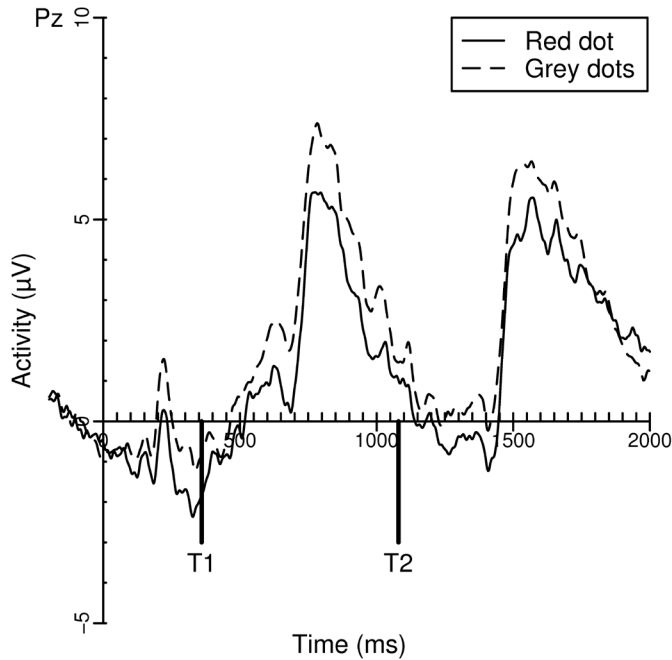|	Mixed-effects model			
	Estimate	Standard Error	t-value	p-value
Grey dots, T1 (Intercept)	5.037	0.653	7.709	0.000
Red dot	-1.455	0.353	-4.120	0.000
T2	-0.284	0.346	-0.821	0.410
Red dot, T2	0.081	0.499	0.162	0.870

Figure 2.2. Grand averages of Pz on lag 8 correct trials for the grey dots (dashed line) and red dot (solid line) task.

factor. The estimates of the model are listed in Table 2.3. In the model, there was a significant effect of the red dot task. Whenever participants were engaged in the red dot task, P3 mean activity was lower than in the grey dots task. To determine whether the decrement in activity was specific for the P3 component, a second analysis was performed with a baseline from −200 to the onset of the P3 component (i.e. 680 ms). Again, there was a significant difference in P3 amplitude between the grey dots task and the red dot task (t = −3.386, p = 0.000), suggesting that there was indeed a target-specific decrement in activity on top of a possible overall reduction in activity. A separate repeated measures ANOVA on the P3 peak latencies of the first target with Task as within-subjects factor did not reveal any significant difference

Table 2.4. The estimates and t-values of the mixed-effects model for P3 amplitude (measured in mean activity) at lag 3 trials.

	Mixed-effects model			
	Estimate	Standard Error	t-value	p-value
Grey dots, T1 (Intercept)	5.336	0.760	7.028	0.000
Red dot	-0.933	0.375	-2.492	0.013
T2	-0.475	0.376	-1.264	0.205
Red dot, T2	0.4286	0.529	0.810	0.416

between the tasks (F<1).

The P3 On Lag 3 Trials

The grand averages of Pz for lag 3 trials can be found in Figure 2.3. To investigate the effect of task on the amplitude P3 of T1 and T2, all lag 3 trials in which T1 was correctly reported were analyzed using a mixed-effects model. Target (T1 or T2) and Task were entered as fixed factors. Subject was entered as a random factor. The estimates and t-values of the coefficients of both models can be found in Table 2.4. A significant effect of Task was found, such that the P3 was smaller in the red dot task than in the grey dots task. Neither a significant difference was found for Target nor for the interaction between Task and Target. Again, a second analysis was performed with a baseline from -200 to 680 ms to see whether an analysis corrected for a possible overall negativity would yield similar results. A significant effect of the red dot task on P3 amplitude was obtained ($t = -2.205$, $p = 0.027$), suggesting that the target-specific decrement in activity found on lag 8 trials was also present on lag 3 trials. A repeated measures ANOVA of the P3 latencies with Task as within-factor did not reveal any significant main effects or interactions (F<1).
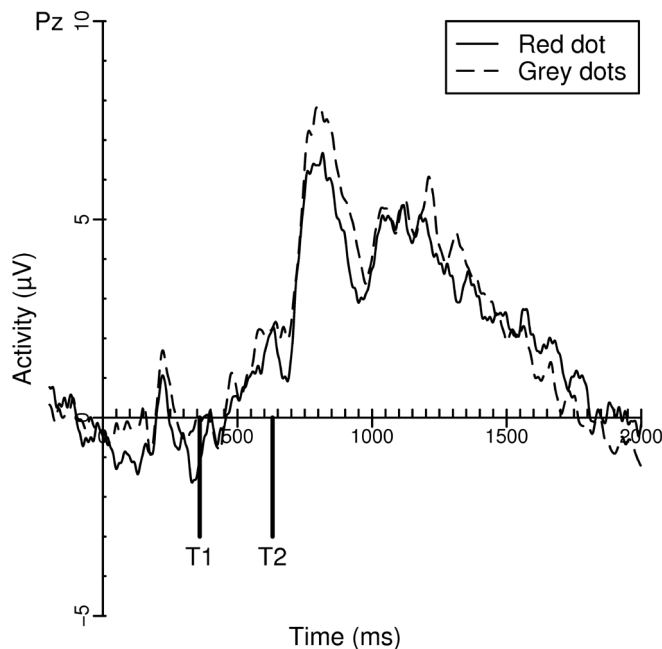
Distractor-related Activity.

Figure 2.3. Grand averages of Pz on lag 3 correct trials for the grey dots (dashed line) and red dot (solid line) task.

Table 2.5. The estimates and z-values of the mixed-effects model for Oz activity when no targets are presented.

	Mixed-effects model A			
	Estimate	Standard Error	t-value	p-value
Grey dots (Intercept)	0.217	0.514	0.421	0.692
Red dot	-0.724	0.216	-3.354	0.001
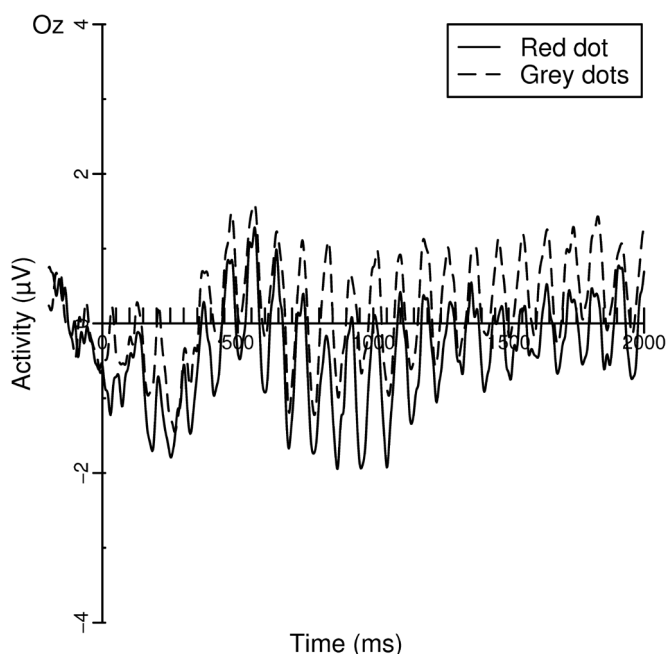
Figure 2.4. Grand averages of Pz on lag 3 correct trials for the grey dots (dashed line) and red dot (solid line) task.

To investigate distractor-related activity, mean activity on no-target trials was analyzed using mixed effect models. Distractor-related activity differed the most at the Oz electrode. Figure 2.4 shows the mean activity at Oz on the no-target trials for all three tasks. A mixed-effects model with Task entered as fixed factor and Subject entered as Random factor revealed a significant effect of the red dot task. As shown in Table 2.5, mean activity was significantly lower in the red dot task when compared to the grey dots task.

The FSP

Activity on F7 and F8 electrodes was analyzed using a mixed effect model. Task, Hemisphere and Targets (where in this case Targets coded for either the presence or absence of a target) were entered in the model as fixed factors. Subject was entered as random factor. A significant

effect was found for Targets (t = 3.328, p = 0.002), indicating the presence of a significant FSP in all three tasks. No other main effects or interactions were found to be significant (ps>0.17). To determine peak latency, the latency of the maximum positive peak within the time window of 180–340 ms after target onset was taken. A repeated measures ANOVA with Task as within-subjects factor did not reveal a significant difference in peak latency (p>0.32)

General Discussion

Previous research has shown that distracting participants during an AB task improves performance in identifying T2 during the blink period (Arend et al., 2006; Olivers & Nieuwenhuis, 2005; Olivers & Nieuwenhuis, 2006; Taatgen et al., 2009). In the current study, behavioral as well as electrophysiological effects of a red dot detection task on a concurrent AB task were investigated. Performance in this dual task setting (referred to as the red dot task) was compared to that in a single AB task (standard task). As a second control condition, a task with similar perceptual input as in the red dot task was also included in the experiment (grey dots task). In contrast to the red dot detection task, the dots that appeared in the latter task were simply to be ignored.

Compared to both control tasks, participants showed an attenuated AB in the red dot task. This result is in line with previous findings reported by Taatgen and colleagues (2009) , despite various small methodological differences between the experiments (i.e., in counterbalancing the order of conditions, the presence of feedback, and the number of lags, distractors, and targets). Most importantly, we found evidence that the P3 amplitudes induced by successfully identified targets were reduced during both lag 3 and lag 8 trials of the red dot task. The finding that P3 amplitude decreases when engaged in a second task is in line with these results (Kok, 2001).

The finding of reduced P3 amplitudes corresponds nicely with previous reports of a relationship between P3 amplitude and AB magnitude (Martens et al., 2006; McArthur et al., 1999). Whereas Martens et al. manipulated P3 amplitude using cuing and stimulus probability, McArthur et al. changed the response set size of T1. Both research groups found that a larger P3 coincided with a larger AB magnitude. More evidence is provided by a MEG study, in which a positive correlation was found between an individual's P3 (or M300) amplitude and AB magnitude (Shapiro et al., 2006).

Martens et al. (2006) specifically focused on individual differences in AB magnitude by studying a group of so-called 'non-blinkers', who show little or no AB, and a group of strong 'blinkers', who show a large AB magnitude. Compared to blinkers, non-blinkers showed more target-related frontal activity (reflected in the FSP component), less distractor-related frontal activity (on target-absent trials), and earlier peak latencies of the target-induced P3s. Martens and colleagues argued that non-blinkers consolidate relevant information quicker (evidenced by the earlier P3 peak latencies), presumably because they are capable of making an early selection, discriminating targets from distractors at an early stage of processing (reflected in increased target-related and decreased distractor-related frontal activity).

An important question that we wanted to address in the current study was whether distraction during the AB task turns blinkers into non-blinkers. The answer to that question seems to be a tentative "no". Firstly, although there was a significant reduction in AB magnitude in the red dot task, participants in the current study still showed a much larger AB than the

non-blinkers did in previous studies (Martens & Johnson, 2005; Martens et al., 2006; Martens & Valchev, 2009). Although, like the non-blinkers, participants in the red dot detection task showed a significant reduction in distractor-related activity, this reduction was most prominent over occipital rather than frontal brain areas. In addition, neither a change in P3 peak latency nor in FSP amplitude was observed in the current study. These findings may thus be taken to suggest that 'natural' non-blinkers and the 'induced non-blinkers' from the current study differ fundamentally in the way they perform the AB task.

As suggested in a recent study by Martens, Korucuoglu, Smid, and Nieuwenstein (2010), natural non-blinkers might be more efficient than blinkers to take advantage of over-learned category-level features to select targets prior to full identification, allowing natural non-blinkers to ignore rather than suppress distractors, thereby avoiding an AB (also see Taatgen et al., 2009).

The picture that tentatively emerges is that the amplitude of the P3 induced by T1 plays an important role in determining whether an AB occurs. Polich (2007) argued that P3 amplitude is related to inhibition processes occurring in the brain. The function of this inhibition process is speculated to protect the target stimulus from interference by competing stimuli. Given the existence of non-blinkers and supported by computer simulations, Taatgen et al. (Taatgen et al., 2009) hypothesized that this inhibition/protection is in fact unnecessary. Any reduction in P3 amplitude (or shift in latency) may imply that less inhibition is exerted on the items following T1, including T2, thereby increasing the chance of successful T2 report. Not only is this in line with what we observed in the red dot task, this view is also consistent with many theories that attribute the AB to an inhibition process that is induced while the first target is being consolidated (e.g., Loach & Marí-Beffa, 2003; Olivers & Watson, 2006; Olivers & Meeter, 2008; Taatgen et al., 2009; Wyble et al., 2009).

In addition to a reduction of target-related parietal activity, we also found a reduction of distractor-related occipital activity. Research on sensory-evoked responses in early visuo-cortical processing has shown that increased perceptual load results in smaller P1 amplitudes (a component associated with early attention processes, see Mangun & Hillyard, 1991) at the occipital sites, accompanied by a reduction of distractor interference on target processing (Handy, Soltani, & Mangun, 2001; Lavie & Tsal, 1994).

To summarize, we found that the AB is attenuated by the addition of a secondary task. In addition, a target-specific reduction in P3 amplitude was found on top of a general decrement in overall distractor-related activity. The presence of a secondary task in the present study presumably increased the perceptual/cognitive load, thus leading to reduced distractor processing (reflected in the reduced distractor-related occipital activity). We conclude that this reduction in distractor processing in turn leads to a reduced 'need' for inhibitory processes later in the processing pathway (reflected in reduced P3s), and therefore a reduced AB. The present findings bring us a step closer in understanding why and how an AB occurs, and how these temporal restrictions in selective attention can be overcome.

Author Contributions

Conceived and designed the experiments: SMW SM HvR NAT. Performed the experiments: SMW. Analyzed the data: SMW HvR SM. Contributed reagents/materials/analysis tools: SMW SM HvR NAT. Wrote the paper: SMW SM HvR NAT.

In which we show that there are differences in selection errors between blinkers and non-blinkers

Chapter was previous published in:

PLoS One, 2013

Charlotte Willems

Stefan M. Wierda

Eva van Viegen

Sander Martens

Abstract

Background

When two targets are presented in close temporal succession, the majority of people frequently fail to report the second target. This phenomenon, known as the 'attentional blink' (AB), has been a major topic in attention research for the past twenty years because it is informative about the rate at which stimuli can be encoded into consciously accessible representations. An aspect of the AB that has long been ignored, however, is individual differences.

Methodology/Principal Findings

Here we compare a group of blinkers (who show an AB) and non-blinkers (who show little or no AB), and investigate the boundary conditions of the non-blinkers' remarkable ability. Second, we directly test the properties of temporal selection by analyzing response errors, allowing us to uncover individual differences in suppression, delay, and diffusion of selective attention across time. Thirdly, we test the hypothesis that information concerning temporal order is compromised when an AB is somehow avoided. Surprisingly, compared to earlier studies, only a modest amount of suppression was found for blinkers. Non-blinkers showed no suppression, were more precise in selecting the second target, and made less order reversals than blinkers did. In contrast, non-blinkers made relatively more intrusions and showed a selection delay when the second target immediately followed the first target (at lag 1).

Conclusion/Significance

The findings shed new light on the mechanisms that may underlie individual differences in selective attention. The notable ability of non-blinkers to accurately perceive targets presented in close temporal succession might be due to a relatively faster and more precise target selection process compared to large blinkers.

Introduction

Restrictions to concurrent attention and awareness are revealed by the interference that commonly results when two sensory inputs must be identified closely in time. For instance, the majority of people typically fail to report the second of two targets when presented in close temporal succession (200–500 ms) amongst a sequential stream of distractors, a phenomenon known as the attentional blink (AB; Martens & Wyble, 2010; Raymond et al., 1992).

In the past two decades, the AB has been a major topic in attention research because it is informative about the rate at which stimuli can be encoded into consciously accessible representations. Although the effect is robust and can be obtained under a variety of task conditions (Martens & Wyble, 2010), large individual differences exist in the magnitude of the effect (Arnell, Howe, Joanisse, & Klein, 2006; Martens et al., 2006; McLaughlin, Shore, & Klein, 2001). Such differences have long been considered as irrelevant noise, until we demonstrated that for some individuals (referred to as 'non-blinkers') the AB can be completely absent (Martens et al., 2006). Given that there is currently much debate about the cause of the AB (see Dux & Marois, 2009; Martens & Wyble, 2010 for recent reviews), several subsequent studies have focused on individual differences in AB magnitude in an attempt to shed new light on the underlying mechanism of the AB (Arnell & Stubitz, 2010; Arnell, Stokes, Maclean, & Gicante, 2010; Colzato, Hommel, & Shapiro, 2010; Colzato, Spapè, Pannebakker, & Hommel, 2007; Colzato et al., 2008; Dale & Arnell, 2010; Dux & Marois, 2008; Green & Bavelier, 2003; Maclean & Arnell, 2010; Martens et al., 2006; Martens & Johnson, 2009; Martens et al., 2009; Martens & Wyble, 2010; Martens, Dun, Wyble, & Potter, 2010; Martens, Kandula, & Duncan, 2010; Martens et al., 2010; Shapiro et al., 2006; Slagter et al., 2007; Slagter, Johnstone, Beets, & Davidson, 2010; Taatgen et al., 2009).

Representing the extreme end on a continuum of individual AB magnitudes, non-blinkers continue to show little or no AB when identification of targets is made more difficult by either increasing the overall rate of stimulus presentation (Martens et al., 2006) or specifically reducing the duration of the targets (Martens et al., 2009; Martens et al., 2010; Martens et al., 2010). In comparison to regular 'blinkers' (individuals who do show an AB), it has been found that non-blinkers neither seem to differ in short-term memory capacity, working memory capacity, nor in general intelligence level (Martens & Johnson, 2009) (but see Arnell & Stubitz, 2010; Colzato et al., 2007, which do report a relation between WM capacity and AB magnitude).

In contrast, however, EEG measurements have revealed differences in frontal and parietal brain activity, reflecting differences in target processing (Martens et al., 2006). In particular, more target-related activity was found over the ventrolateral prefrontal cortex (assumed to play a role in a wide range of cognitive processes, including the selection of non-spatial information), whereas blinkers showed more distractor-related prefrontal activity. Regardless of the time interval between the targets, non-blinkers were also found to be quicker in consolidating the identity of targets than blinkers, showing earlier peak latencies of the P3 ERP components—associated with the updating of working memory (WM)—induced by successfully identified targets (Martens et al., 2006). In line with this result, evidence was recently found that the magnitude of the AB is related to striatal dopamine functioning, which is associated with regulating the threshold for WM updating (Slagter et al., 2012). Taken together, these findings suggest that non-blinkers are more efficient in distinguishing targets from distractors at a relatively early processing stage. Indeed, behavioral studies have provided converging

evidence showing that non-blinkers are better in ignoring distractors than blinkers are (Dux & Marois, 2008; Martens et al., 2009; Martens & Valchev, 2009).

It must be noted though that this early selection seems to be specific for alphanumeric, visual targets. AB magnitude was found to be similar for blinkers and non-blinkers when using pictures rather than alphanumeric stimuli (Martens et al., 2010). Also when using auditory alphanumeric stimuli, non-blinkers showed a substantial AB effect, although overall performance was still better than that of blinkers (Martens et al., 2009).

It was therefore suggested that in an alphanumeric AB task non-blinkers might take advantage of overlearned category-level features to select targets prior to full identification, allowing them to mostly ignore distractors and to avoid an AB. Indeed, an ERP study subsequently showed that when alphanumeric category information was unavailable (only letters were presented) and target selection could only be based on information that is processed relatively late (rotation), non-blinkers again showed a substantial AB effect (Martens et al., 2010). Delayed target-related occipito-parietal activity as well as increased distractor-related prefrontal brain activity was observed. Also, when alphanumeric category information was not available, the difference in P3 peak latency between the two groups disappeared. However, non-blinkers continued to outperform blinkers across all conditions by showing a smaller AB, suggesting that early selection processes based on category information alone cannot fully explain the observed differences between the two groups.

Nevertheless, it has been suggested that a major source of individual variability in AB magnitude must lie in processes of selective attention that are involved in determining which objects are selected for further processing and memory consolidation (Martens et al., 2006; Martens & Johnson, 2009; Martens et al., 2009; Martens & Valchev, 2009; Martens et al., 2010; Martens et al., 2010). In this regard, the insights derived from studies examining individual differences in the AB converge with recent ideas regarding the source of the AB. Whereas the earliest studies claimed that the AB is the result of capacity limitations (Chun & Potter, 1995; Shapiro et al., 1994), alternatively, the AB is lately often regarded as a problem to time or control attention (Di Lollo, Kawahara, Shahab Ghorashi, & Enns, 2005; Dux & Marois, 2009; Martens & Wyble, 2010; Taatgen et al., 2009; Wyble et al., 2009). This shift in the theoretical landscape was motivated by a number of key findings. For instance, it was found that people are capable of reporting an undisrupted stream of letters, but typically fail when required to report only a subset of this stream, as reflected in the AB task (Nieuwenstein & Potter, 2006; Olivers, van der Stigchel, & Hulleman, 2007). Furthermore, it has been found that the AB is attenuated when participants perform a second task concurrently with the primary AB task (Olivers & Nieuwenhuis, 2005; Olivers & Nieuwenhuis, 2006; Taatgen et al., 2009; Wierda et al., 2010). Together, these studies provide evidence against theories assuming resource depletion, since according to these limited-capacity theories an additional task load should increase rather than decrease the magnitude of the AB. Given these findings, the temporal selection mechanism seems important for explaining the AB, although it must be noted that recent findings also suggest a role for capacity limitations (Dell'Acqua et al., 2012; Dell'acqua, Jolicoeur, Luria, & Pluchino, 2009; Dux et al., 2008; Dux et al., 2009).

The aim of the present study was to further investigate this temporal selection mechanism by contrasting the performance of blinkers and non-blinkers. In the abovementioned studies, non-blinkers showed an AB when visual target selection was based on a target-defining feature that was processed relatively late, such as rotation (Martens et al., 2010) or semantic category (Martens et al., 2010). To test the generality of this finding, an AB experiment

was set up that featured only letter stimuli with targets defined by color, a stimulus feature that is available relatively early (Rotte, Heinze, & Smid, 1997; Smid & Heinze, 1997; Wijers, 1989). This way, early target selection should be possible, and non-blinkers should still be able to avoid an AB on the majority of trials. However, if their temporal selection ability specifically relies on the presence of alphanumeric category information—which is unavailable—the occurrence of an AB is to be expected.

To study the temporal dynamics of attention in more detail, another important goal of the current study was to investigate the temporal profile of non-blinkers and blinkers using three measures of temporal selection, namely 'suppression', 'delay', and 'diffusion', originally proposed by Vul, Nieuwenstein, & Kanwisher (2008) and Chun (1997). Since each stimulus letter was presented only once within each stream, the serial position of any reported letter was known, thus allowing us to highlight and contrast these three dimensions of target selection in blinkers and non-blinkers (Vul et al., 2008). Following Vul and colleagues, if a response consists of a letter that does not correspond with any of the letters presented within a certain temporal window around a target, we assume that the relevant information was likely to be suppressed ('suppression'). If a response corresponds with a letter that was presented after a target, it can be inferred that temporal target selection was delayed ('delay'). Finally, if distractors strongly interfere with the processing of targets, selection will be less precise, reflected in selection errors that are temporally more distant from the target ('diffusion'). Vul et al. (2008) found that the temporal selection process was suppressed, delayed, and diffused during the AB.

Both the concepts of suppression and delay have previously been associated with the AB. Regarding suppression, many studies emphasized its important role during the AB (Botella, Privado, de Liano, & Suero, 2011; Chun, 1997; Dux, Coltheart, & Harris, 2006; Dux & Harris, 2007; Dux & Marois, 2008; Harris, Benito, & Dux, 2010; Kihara, Yagi, Takeda, & Kawahara, 2011; Loach & Marí-Beffa, 2003; Martens & Valchev, 2009; Olivers, 2007; Olivers & Watson, 2008; Popple & Levi, 2007; Vogel, Luck, & Shapiro, 1998; Vul et al., 2008). In EEG studies, suppression is reflected in the P3 component that is absent or strongly attenuated during the AB (Vogel et al., 1998), and also the n2pc (associated with the allocation of attention) is known to be affected (Dell'Acqua, Sessa, Jolicoeur, & Robitaille, 2006; Jolicoeur, Sessa, Dell'Acqua, & Robitaille, 2006; Jolicœur, Sessa, Dell'Acqua, & Robitaille, 2006). Similarly there is quite some evidence supporting the idea that attentional selection is delayed during the AB, provided by behavioral studies (Botella et al., 2011; Chun, 1997; Chun & Potter, 1995; Nieuwenstein et al., 2005; Nieuwenstein, 2006; Vul et al., 2008; Vul, Hanus, & Kanwisher, 2008) and EEG studies (Martens et al., 2006; Vogel et al., 1998; Vogel & Luck, 2002), where the latter have revealed that when the second target was reported correctly at short time intervals, the P3 component was delayed in comparison to longer intervals.

Combined with our previous findings on individual differences in the AB, we predicted that non-blinkers would continue to outperform the blinkers, and would show less suppression, delay, and diffusion. Interestingly however, although many papers suggested that information processing is suppressed during the attentional blink (e.g., Kihara et al., 2011; Olivers, 2007; Vogel et al., 1998; Vul et al., 2008), a number of papers have claimed that the AB is due to a failure to suppress distractor stimuli (Dux et al., 2006; Dux & Harris, 2007; Dux & Marois, 2008; Harris et al., 2010), which implies that we should find the opposite effect; individuals with little or no AB should show relatively strong suppression, whereas individuals with a large AB should show relatively little suppression.

A final prediction concerning non-blinker performance comes from a simulation study suggesting the AB to reflect a cognitive strategy of enforcing an episodic distinction between successive stimuli (Wyble et al., 2009). When the occurrence of an AB is somehow avoided, information concerning temporal order and the correct binding of features into targets might be compromised (Martens & Wyble, 2010). In other words, non-blinkers might lack the episodic distinction between successive stimuli, and subsequently make more order reversals (i.e., reporting the second target before the first target) than blinkers do. If however, non-blinkers are generally quicker to select and consolidate targets (see e.g., Martens et al., 2006), one would expect to find fewer order reversals in non-blinkers than in blinkers. A final aim was thus to test these latter predictions.

In summary, we tested whether non-blinkers can avoid an AB when targets are to be selected on the basis of color rather than alphanumeric category information. Second, we tested whether non-blinkers show less suppression, delay, and diffusion than blinkers do. And third, we investigated whether avoiding an AB comes at a cost, reflected in non-blinkers making relatively more order reversals.

Methods

Experiment 1a consisted of an AB task with alphanumeric stimuli, requiring detection and identification of two target letters presented in a rapid serial visual presentation (RSVP) stream of 16 distractor digits. Participants were tested for the presence or absence of a sizeable AB, with the purpose of forming separate groups of consistent blinkers and non-blinkers for inclusion in Experiment 1b. Experiment 1b contained only letter stimuli, targets were defined by color, and its goal was to test the temporal profile of blinkers and non-blinkers in terms of suppression, delay, and diffusion. The purpose of Experiment 2 was to replicate the findings in a larger sample of participants.

Experiment 1a

In Experiment 1a, participants performed an AB task requiring the identification of two letter targets amongst a sequential stream of digit distractors. The purpose of this experiment was to test selected participants for the presence or absence of a sizeable AB in a classical alphanumeric AB task. In addition, we aimed to systematically study possible differences between blinkers and non-blinkers in terms of order reversals.

Participants

Twenty-nine volunteers (16 women; aged 20–31, mean = 25.0) recruited from the University of Groningen community participated in the experiment, had normal or corrected-to-normal visual acuity, normal hearing, and no history of neurological problems. One participant was excluded due to RSI problems. Thirteen participants were included because they had shown little or no AB in previous studies in our laboratory, and were therefore regarded as potential non-blinkers. The other 15 participants had previously shown a regular to large AB, and

were therefore regarded as potential blinkers. The Neuroimaging Center Institutional Review Board approved the experimental protocol and each participant signed a written consent prior to the experiment. All volunteers participated in both Experiment 1a and 1b in a single session, and received payment of € 7 in total.

Stimuli and apparatus

The generation of stimuli and the collection of responses were controlled by using E-prime 1.2 software running under Windows XP. Target stimuli consisted of uppercase consonant letters excluding 'Q', 'V', and 'Y'. Distractor stimuli consisted of digits (2 to 9). All stimuli were centrally presented in black (2 cd/m^2) on a white background (88 cd/m^2) in uppercase 14-point Monaco font on a 19-inch CRT monitor with a 100-Hz refresh rate. Viewing distance was approximately 50 cm.

Procedure

Each trial began with a message at the bottom of the screen, prompting participants to press the space bar to initiate the trial. When the space bar was pressed, the message disappeared immediately and a central fixation cross appeared. It remained on the screen for 100 ms, followed by the RSVP stream consisting of 18 items (i.e., 2 targets and 16 distractors).

All stimuli were presented for 80 ms without inter stimulus interval. The first target (T1) was always presented as the sixth item in the stream. The second target (T2) was the first, second, third, or eighth item following T1, and was thus presented at lag 1, 2, 3, or 8, respectively. In other words, the stimulus onset asynchrony (SOA) between the targets randomly varied from 80, 160, 240, to 640 ms. Each lag was presented equally often. Target letters were pseudo-randomly selected with the constraint that T1 and T2 were always different letters. Digit distractors were pseudo-randomly selected with the constraint that no single digit was presented twice in succession.

After the presentation of the stimulus stream, participants were prompted by a message at the bottom of the screen to indicate the letters they had seen by using the corresponding keys on the computer keyboard. Participants were instructed to take sufficient time in making their responses to ensure that typing errors were avoided. Participants were encouraged to type in their responses in the order in which the letters had been presented, but responses were accepted and counted correct in either order. Participants were instructed to guess if they had not seen the targets.

The experiment contained one practice block of 24 trials and two testing blocks of 144 trials each, and took approximately 30 minutes to complete. After the first testing block, participants were allowed to take a short break. At the end of the experiment, participants took another short break before continuing with Experiment 1b.

Experiment 1b

The purpose of Experiment 1b was twofold. First, we wanted to test whether non-blinkers

continue to show little or no AB when targets are defined by color rather than alphanumeric category. To that end, all stimuli consisted of letters, with targets presented in red, and distractors in black. Second, following Vul et al. (2008), we directly tested the properties of temporal selection by analyzing the distribution of reported letters, allowing us to study the suppression, delay, and diffusion of selective attention across time in blinkers and non-blinkers.

Participants

All participants of Experiment 1a volunteered to participate in Experiment 1b. Participants were assigned to the same groups of blinkers and non-blinkers as in Experiment 1a. Note that the individuals who consistently show no AB in an alphanumeric AB task as demonstrated in Experiment 1a (i.e., non-blinkers) might show an AB under the experimental conditions of Experiment 1b. To consistently refer to these individuals in Experiments 1a and 1b, we will continue to label them as 'non-blinkers', keeping in line with the literature on non-blinkers (Martens et al., 2009; Martens et al., 2010).

Stimuli and apparatus

The same stimuli and apparatus were used as in Experiment 1a, except that all stimuli consisted of consonant letters. Again 'V', 'Q', 'Y' were excluded. Targets were presented in red, whereas distractors were presented in black.

Procedure

The procedure was the same as in Experiment 1a, except that all stimuli were presented for 120 ms, such that a similar level of difficulty was obtained as in Experiment 1a. Furthermore, the RSVP consisted of 16 stimuli, and T1 was always presented as the fifth item in the stream. Experiment 1b took approximately 35 minutes to complete.

Experiment 2

The aim of Experiment 2 was to strengthen the results found in Experiment 1b by replicating the results in a larger sample of participants, enabling us to study a wider range of individual differences.

Participants

A total of 132 volunteers (98 women) recruited from the University of Groningen participated in the experiment in return for course credits. Unfortunately, due to technical problems, the age related information of the participants was lost for this experiment. However, because participants were selected from a similar pool of participants as in Experiment 1, it can be as-

sumed that the average age of the participants in both experiments was equivalent. They had normal or corrected-to-normal visual acuity, normal hearing, and no history of neurological problems. The Neuroimaging Center Institutional Review Board approved the experimental protocol and each participant signed a written consent prior to the experiment.

Stimuli and apparatus.

The stimuli and apparatus were the same as in Experiment 1b.

Procedure

The procedure was similar to that in Experiment 1b. The experiment consisted of one practice block of 14 trials and three testing blocks of 96 trials each. Participants were allowed to take a short break between blocks. They completed the experiment in approximately 45 minutes.

Results and Discussion

When appropriate, Greenhouse-Geisser-corrected p-values are reported ($\varepsilon<0.75$). In addition, a Bonferroni-correction was applied when independent t-tests were performed serving as post-hoc test.

Experiment 1a

To assure that participants were assigned to the appropriate group, AB magnitude was first computed for each individual by calculating the percentage decline in T2 accuracy at lags 2 and 3 relative to T1 accuracy across lags. Following previous non-blinker studies (Chua, 2005; Martens & Johnson, 2009; McLaughlin et al., 2001), the AB magnitude was calculated as a function of T1 accuracy by using the following formula:

$$\text{AB magnitude} = \left(\frac{\overline{T1} - T2 \mid T1_{lag2}}{\overline{T1}} + \frac{\overline{T1} - T2 \mid T1_{lag3}}{\overline{T1}} \right) / 2 * 100\% ,$$

where $\overline{T1}$ is the mean accuracy of T1, and $T2|T1_{lag}$ is the mean accuracy of T2 at a specific lag given that T1 was correctly reported. We used this particular method to assure that individuals with a high T1 accuracy, but overall low T2 accuracy were not erroneously classified as non-blinkers. However, alternative ways to calculate AB magnitude, for instance by relating T2 accuracy at lags 2 and 3 to T2 accuracy at lag 8 produced comparable results. Mean AB magnitude was 8.7% for the non-blinkers, ranging from 2.5% to 15.3%, suggesting that each individual within this group indeed showed little or no AB. For the blinkers, mean AB magnitude was 32.6%, ranging from 17.0% to 50.6%, suggesting that they showed a moderate to large AB.
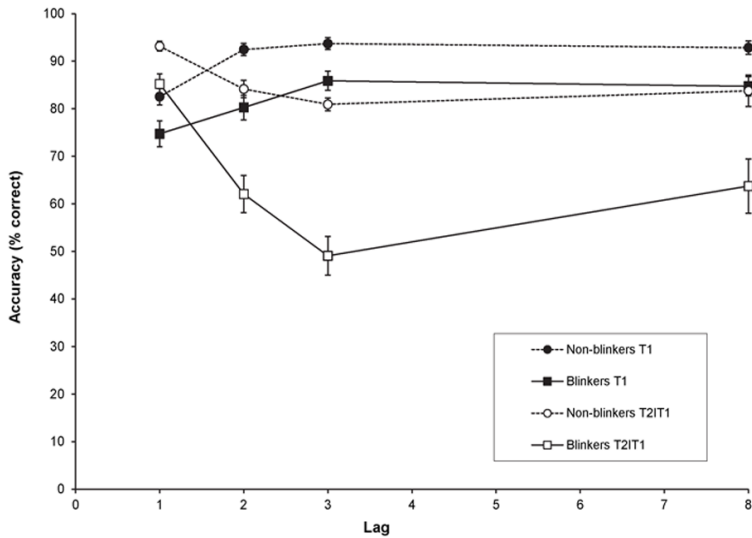
Figure 3.1. Target accuracy in Experiment 1a. Mean percentage correct report of T1 (black symbols) and T2 given correct report of T1 (white symbols) as a function of lag, for non-blinkers (circles) and blinkers (squares). Error bars reflect standard error of the mean.

Figure 3.1 shows target accuracy as a function of the interval between the two targets (lag), for non-blinkers (circle symbols) and blinkers (square symbols). A repeated measures analysis of variance (RM-ANOVA) of T1 accuracy with group (non-blinkers and blinkers) as a between-subjects factor and lag (1, 2, 3, and 8) as a within-subjects factor revealed a significant effect of group, $F(1, 26) = 13.49$, MSE = 166.64, $p = 1$, $\eta^2p = .34$, reflecting mean accuracy to be higher for non-blinkers (90.4%) than for blinkers (81.4%). In addition, a main effect of lag was found, $F(2.17, 56.49) = 33.27$, MSE = 29.43, $p < .001$, $\eta^2p = .56$, such that performance at lag 1 was relatively low. The Group × Lag interaction was not significant ($p = .23$).

A RM-ANOVA of T2 performance given correct report of T1 (T2|T1) with group as a between-subjects factor and lag as a within-subjects factor revealed a significant effect of group, $F(1, 26) = 28.75$, MSE = 406.53, $p < .001$, $\eta^2p = .53$; lag, $F(3, 78) = 33.63$, MSE = 84.49, $p < .001$, $\eta^2p = .56$; and a significant Group × Lag interaction, $F(3, 78) = 7.94$, MSE = 84.49, $p = .001$, $\eta^2p = .23$. Separate analyses in which lag 1 was excluded revealed that non-blinkers did not show a significant AB ($p = .38$), whereas blinkers did, $F(2, 28) = 8.11$, MSE = 119.22, $p = .002$, $\eta^2p = .37$.

Order reversals

We calculated the relative percentage of order reversals over the trials where T1 and T2 were both correctly reported, providing a measure of order reversals that is irrespective of individual differences in identification accuracy. Interestingly, there was a significant effect of

group, $F(1, 26) = 5.96$, MSE = 152.26, p = .022, $\eta^2 p$ = .19, such that non-blinkers showed relatively fewer order reversals than blinkers did (11.1% vs. 16.8%, respectively). In addition, we found an effect of lag, $F(3, 78) = 83.99$, MSE= 51.77, p<.001, $\eta^2 p$ = .76, as the number of order reversals decreased as a function of lag (30.8%, 14.8%, 10.0%, and .9% at lags 1, 2, 3, and 8, respectively). Also a marginally significant Group × Lag interaction was found, $F(3, 78)$ = 2.71, MSE = 51.77, p = .051, $\eta^2 p$ = .09, such that particularly at lags 2 and 3, non-blinkers seemed to show fewer order reversals than blinkers did.

Experiment 1b

Figure 3.2 shows target accuracy as a function of lag, for non-blinkers and blinkers. Mean T1 accuracy was 90.0% for the blinkers and 91.9% for the non-blinkers. A RM-ANOVA of T1 performance revealed no significant effects (ps>.10).

A RM-ANOVA of T2|T1 revealed a significant effect of group, $F(1, 26) = 8.98$, MSE = 296.76, p = .006, $\eta^2 p$ = .26; lag, $F(3, 78) = 73.40$, MSE = 99.27, p<.001, $\eta^2 p$ = .74; and a significant Group × Lag interaction, $F(3, 78) = 4.93$, MSE = .9.27, p = .007, $\eta^2 p$ = .16. Mean AB magnitude was 31.6% for non-blinkers and 49.0% for blinkers (t(26) = 3.53, SE = 4.95, p = .002). These findings suggest that both the blinkers as well as the non-blinkers showed a sizeable AB, but that it was substantially smaller in the non-blinkers than in the blinkers.

A positive Pearson product-moment correlation was found between individual AB magnitudes in Experiments 1a and 1b, r = .42, p = .027. A similar correlation was found for
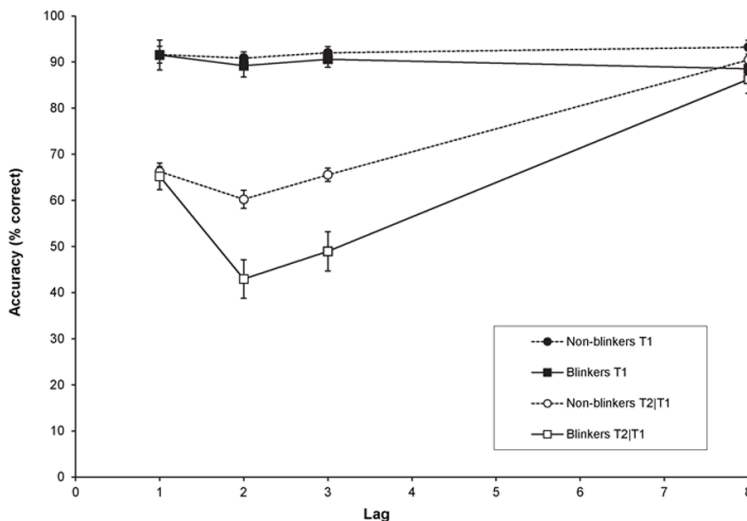

Figure 3.2. Target accuracy in Experiment 1b. Mean percentage correct report of T1 and T2 given correct report of T1 as a function of lag, for non-blinkers and blinkers. Error bars reflect standard error of the mean.

T2|T1 performance, r = .44, p = .019, but not for T1 performance (p = .14). These findings suggest that although AB magnitude was generally larger in Experiment 1b than in Experiment 1a, individuals with a relatively small or large AB in Experiment 1a continued to show a relatively small or large AB in Experiment 1b, respectively.

Suppression

We estimated the efficacy of selection (A) as the proportion of trials during which an item was reported from a 7-item window around the target (spanning three items before to three items after the target) as follows:

$$A = \sum_{i=ks}^{ke} P_i \, ,$$

where P_i is the probability (i.e., empirical frequency) of reporting an item from serial position i relative to the target position (i = 0), and ks and k_e are the lower and upper bounds, respectively, of the window used to compute the measure (in this case, −3 and 3, respectively). Thus, we calculated how frequent each participant reported a letter from the 7-item window surrounding T1 or T2 to indicate the availability of the distractors around the target. In contrast to the previous analyses, order reversals were counted as incorrect, because for these and the
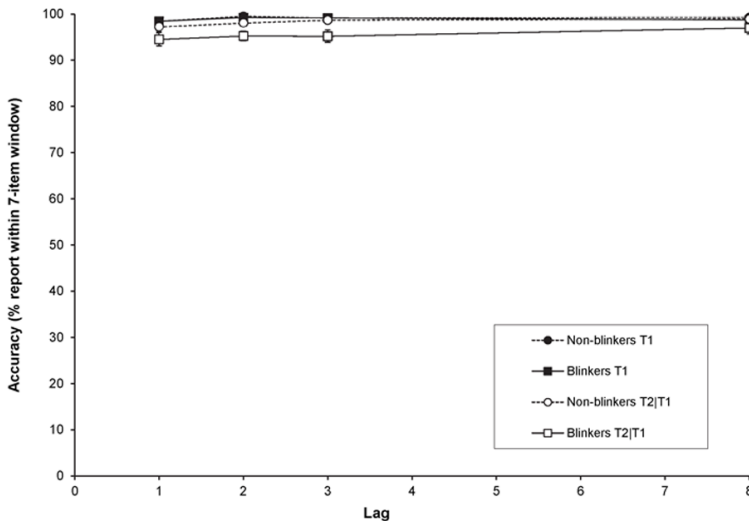

Figure 3.3. Suppression in Experiment 1b. Suppression of the temporal selection process expressed as the accuracy of reporting an item within the 7-item window around a given target as a function of lag, for blinkers and non-blinkers.
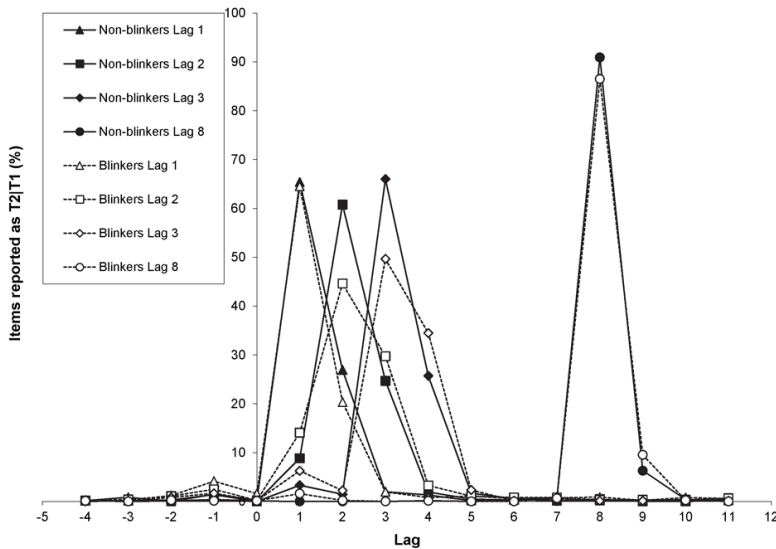
Figure 3.4. Distribution of T2|T1 reports in Experiment 1b. The percentage of letters at a particular position in the RSVP stream that were reported as T2 given correct report of T1 as a function of lag, for blinkers and non-blinkers.

following analyses we were interested in the exact serial location of the reported letters.

As shown in Figure 3.3, performance within the 7-item window was close to or at ceiling for both blinkers and non-blinkers. Given that 17 different letters could be presented within the stream, the chance to randomly select a letter within the 7-item window was 7/17 (i.e., 42%). A paired t-test revealed that the accuracy of reporting an item within the 7-item window differed significantly from the level of chance, $t(27) = 60.0$, SE = .9, $p < .001$; $t(27) = 74.29$, SE = .74, $p < .001$; $t(27) = 61.0$, SE = .91, $p < .001$; $t(27) = 68.5$, SE = .82, $p < .001$ for lags 1, 2, 3, and 8, respectively. A RM-ANOVA of T1 showed an effect of lag, $F(3, 78) = 4.1$, MSE = 1.09, $p = .009$, $\eta^2p = .14$, but both the Group x Lag interaction ($p = .66$), as the effect of group ($p = .62$) were non-significant.

A RM-ANOVA of T2|T1 revealed a significant effect of group, $F(1, 26) = 9.87$, MSE = 21.05, $p = .004$, $\eta^2p = .28$, whereas neither the effect of lag ($p = .27$) nor the Group $\times$ Lag interaction ($p = .91$) was significant. These findings suggest that overall, little or no suppression seemed to be present, and that the AB did not induce any suppression as a function of lag in this study. Given that many theoretical and computational models of the AB assume that the AB is caused by the suppression that is induced by T1 and/or the distractor that immediately follows T1 (Dux & Marois, 2008; Olivers, 2007; Vogel et al., 1998; Vul et al., 2008), it is striking to find no evidence for an AB-induced suppression effect for T2, which would otherwise be reflected in a sizeable drop in performance during lags 2 and 3. However, it is important to note that because performance in the current experiment was close to ceiling, such an effect might be concealed. Figure 3.4 provides a more detailed picture regarding the distribution of T2|T1 reports, revealing that participants tend to report either the letter preceding or following the second target when making intrusion errors. We will discuss this pattern of intrusions further in the section below on 'relative T2+3 intrusions'.
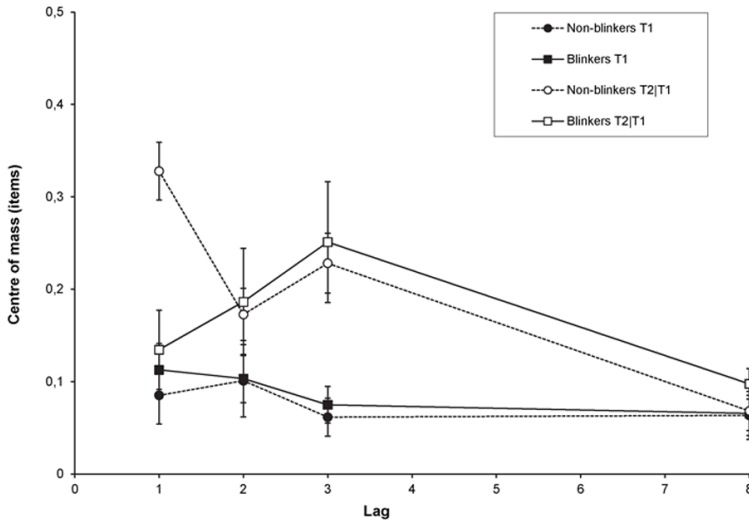
Figure 3.5. Delay in Experiment 1b. Delay of the temporal selection process expressed as the centre of mass of reports in the selection window around a given target as a function of lag, for blinkers and non-blinkers.

Delay

In order to measure the latency of these intrusion errors in a similar manner as (Chun, 1997; Vul et al., 2008) did, we calculated the center of mass (C) of reports in the window around a given target as follows:

$$C = \frac{\sum_{i=ks}^{ke} P_i * i}{A}.$$

Originally employed by Chun (Chun, 1997), the center of mass corresponds to the average reported serial position relative to the target. A positive center of mass indicates that participants are more likely to report items following the target, whereas a negative center of mass would indicate a bias to report items preceding the target. If the center of mass is more positive for T2 than for T1, this means that selection is delayed for T2 relative to T1. Order reversals were counted as incorrect in this analysis.

Figure 3.5 shows the measure of delay for T1 and T2 as a function of lag, for blinkers and non-blinkers. A RM-ANOVA of the center of mass for T1 only revealed a significant main effect of lag, $F(3, 78) = 3.03$, MSE = .004, $p = .045$, $\eta^2p = .1$. For T2|T1 we found an effect of lag, $F(2.1, 54.64) = 6.48$, MSE = .02, $p = .003$, $\eta^2p = .20$; no main effect of group ($p = .35$); and a Group × Lag interaction, $F(2.1, 54.64) = 3.63$, MSE = .02, $p = .03$, $\eta^2p = .12$. The non-blinkers show a delay that is particularly pronounced at lag 1, whereas for blinkers the strongest delay is observed at lag 3. Independent samples t-tests revealed a significant diffe-
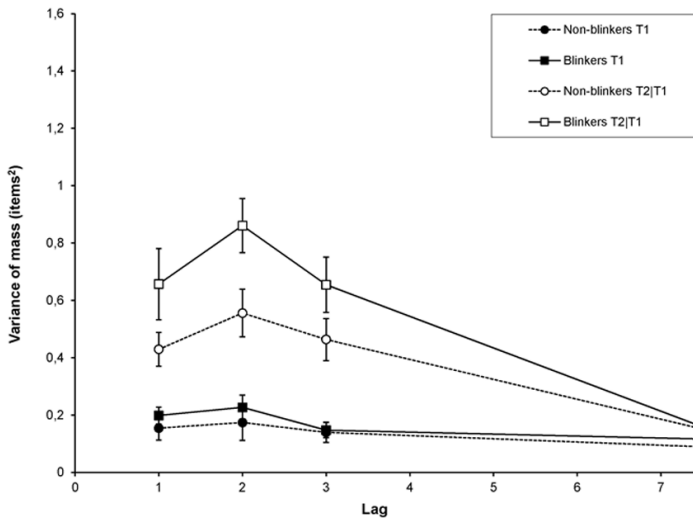
Figure 3.6. Diffusion in Experiment 1b. Diffusion of the temporal selection process expressed as the variance of the centre of mass in the selection window around T1 or T2 as a function of lag, for blinkers and non-blinkers.

rence between non-blinkers and blinkers at lag 1 only, $t(26) = 3.88$, $SE = .51$, $p = .001$. This might reflect a difference in the use of letters following the second target for the two groups, however, it must be noted that this could also reflect a difference in the binding of letter identity and color, which is discussed more extensively in the general discussion.

Diffusion

Similarly to Vul et al. (Vul et al., 2008), we estimated the precision of selection around the center of mass (see Figure 3.6) by calculating the variance of the center of mass (V), as follows:

$$V = \frac{\sum_{i=ks}^{ke} P_i * (i - C)^2}{A} .$$

Here, the variance of the center of mass reveals to which extent the reports of the letters are diffused around the center of mass, reflecting the spread of selection. Again, order reversals were counted as incorrect.

For T1, we only found a significant effect of lag, $F(1.82, 47.36) = 6.41$, $MSE = .01$, $p = .004$, $\eta^2 p = .2$; whereas for T2|T1 we found a significant effect of group, $F(1, 26) = 4.29$, $MSE = .21$, $p = .048$, $\eta^2 p = .14$; and lag, $F(2.2, 57.2) = 33.01$, $MSE = .08$, $p < .001$, $\eta^2 p = .56$; but no significant Group × Lag interaction ($p = .11$). These results clearly reflect that—compared to non-blinkers—blinkers are less precise in selecting the second but not the first target.
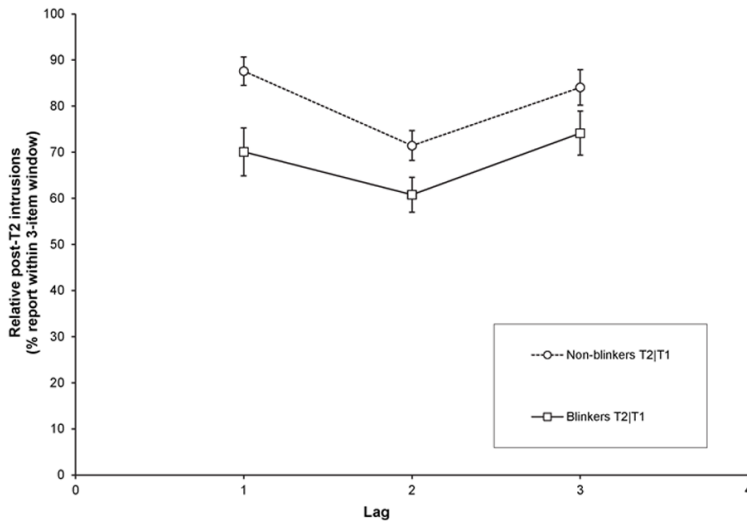
Figure 3.7. Intrusion errors in Experiment 1b. Percentage of erroneously selected letters (relative to all trials with an incorrect T2 response) presented 1-3 serial positions following T2 as a function of lag, for blinkers and non-blinkers.

Relative T2+3 intrusions

The relatively high performance within the 7-item window reveals that response errors were far from random, as illustrated in Figure 3.3 and Figure 3.4. The latter figure indicates that for lags 2 and 3, blinkers show more post-target intrusions than non-blinkers do. However, blinkers show more errors overall, so a more meaningful comparison would be to determine the pattern of relative intrusion errors, controlling for differences in the total error rate. To that end, we examined the percentage of erroneously selected letters presented at one to three serial positions following a target, relative to all errors on a given lag. Order reversals were counted as incorrect. For T1, as well as for T2 at lag 8, the number of post-target intrusions was insufficient to allow for a meaningful analysis. Therefore, this analysis was restricted to T2|T1 at lags 1 to 3 only. For this analysis, the average number of trials over participants available in blinkers was 16.3, 25.3, and 27.0 for lags 1, 2, and 3, respectively. In non-blinkers this was 21.2, 19.8, and 20.4 for lags 1, 2, and 3, respectively.

In Figure 3.7 the percentage T2+3 intrusions relative to all errors on a given trial are plotted as a function of lag. A RM-ANOVA of the T2+3 intrusions with lag (1, 2, and 3) as a within-subjects factor and group (non-blinkers and blinkers) as a between-subjects factor revealed significant effects for lag, $F(1.46, 37.91) = 7.97$, MSE = 264.53, $p = .003$, $\eta^2p = .24$; and group, $F(1, 26) = 9.93$, MSE = 339.4, $p = .004$, $\eta^2p = .28$; but a significant Group × Lag interaction was not found ($p = .48$). Thus, compared to blinkers, when a selection error was made, the T2 response of non-blinkers more frequently matched one of the items following the second target. In contrast to the pattern of absolute intrusion rates (see Figure 3.4), the current analysis of relative post-target intrusions shows that this was not only the case at lag 1, but also at lags 2 and 3 (see Figure 3.7).
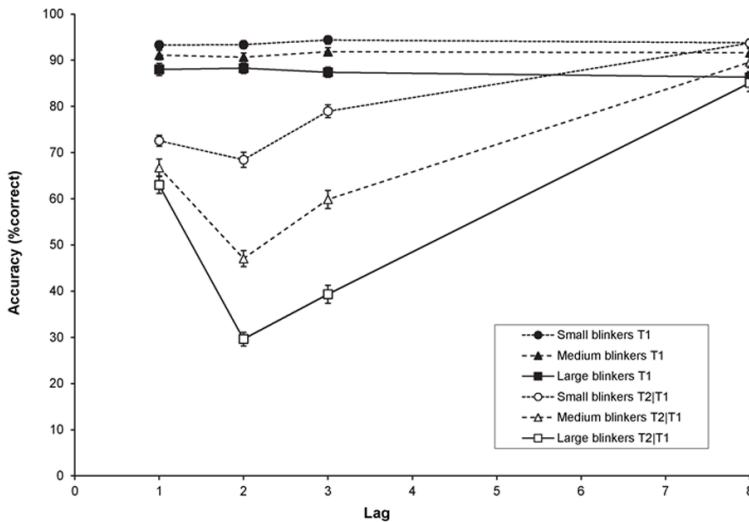
Figure 3.8. Target accuracy in Experiment 2. Mean percentage correct report of T1 (black symbols) and T2 given correct report of T1 (white symbols) as a function of lag, for small blinkers (circles), medium blinkers (triangles), and large blinkers (squares). Error bars reflect standard error of the mean.

Order reversals

The percentage of order reversals for trials during which T1 and T2 were both correct was 8.3%, .3%, .5%, and .2% at lags 1, 2, 3, and 8, respectively. A significant main effect of lag reflected the decrease of order reversals as a function of lag, $F(1.1, 28.39) = 23.83$, $MSE = 49.52$, $p < .001$, $\eta^2p = .48$. No effect of group ($p = .6$) or an interaction effect between group and lag ($p = .54$) was found, suggesting no difference in order reversals between non-blinkers and blinkers. Given that AB magnitude was larger in Experiment 1b than in Experiment 1a for both groups, it is perhaps surprising that there were substantially more order reversals in Experiment 1a. An explanation might at least partially lie in the fact that the SOA was much shorter in Experiment 1a (80 ms) than in Experiment 1b (120 ms).

Experiment 2

After initial analysis, 21 students were excluded from further analyses due to insufficient identification performance of T1 (<70%). In total, 111 participants remained for further analyses. Given that Experiment 2 featured a wide range of AB magnitudes, we treated AB magnitude in the analyses of Experiment 2 as a continuous variable. However, for the sake of clarity, figures for Experiment 2 feature three subgroups, based on individuals' AB magnitude in the first block of the experiment. Mean AB magnitude was 15.9% (range = 1.3–27.0%) for the group of 'small blinkers', 39.1% (range = 27.0–47.5%) for the group of 'medium blinkers', and 60.3% (range = 48.2–92.8%) for the group of 'large blinkers'.
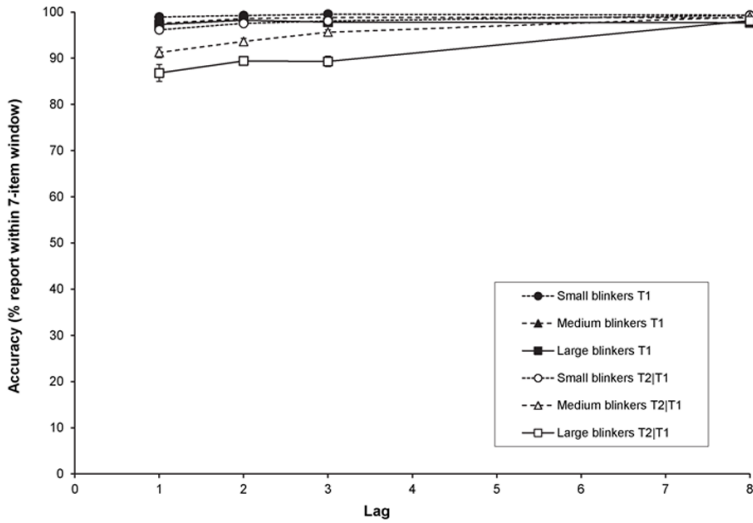
Figure 3.9. Suppression in Experiment 2. Suppression of the temporal selection process expressed as the accuracy of reporting an item within the 7-item window around a given target as a function of lag, for small, medium, and large blinkers.
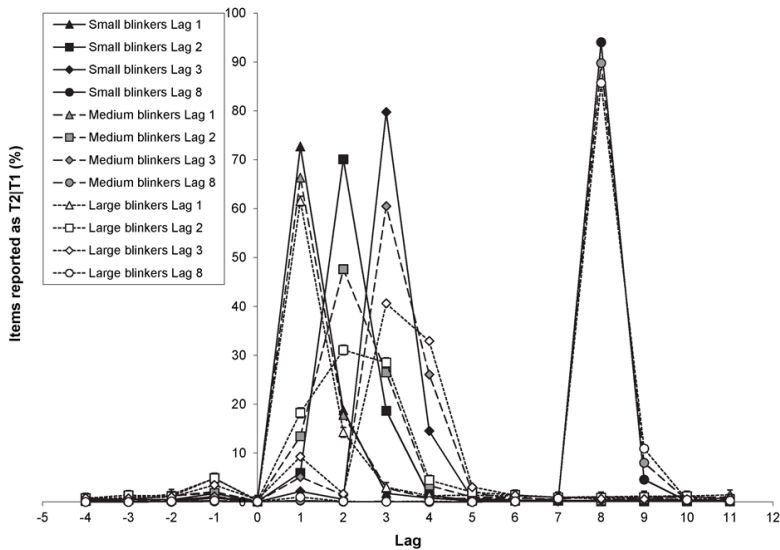
Figure 3.10. Distribution of T2|T1 reports in Experiment 2. The percentage of letters at a particular position in the RSVP stream that were reported as T2 given correct report of T1 as a function of lag, for small, medium, and large blinkers.
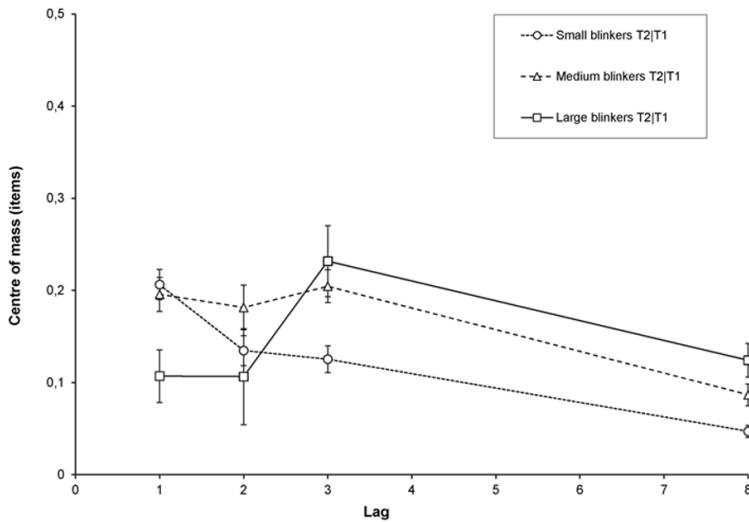
Figure 3.11. Delay in Experiment 2. Delay of the temporal selection process expressed as the centre of mass of reports in the selection window around a given target as a function of lag, for small, medium, and large blinkers.

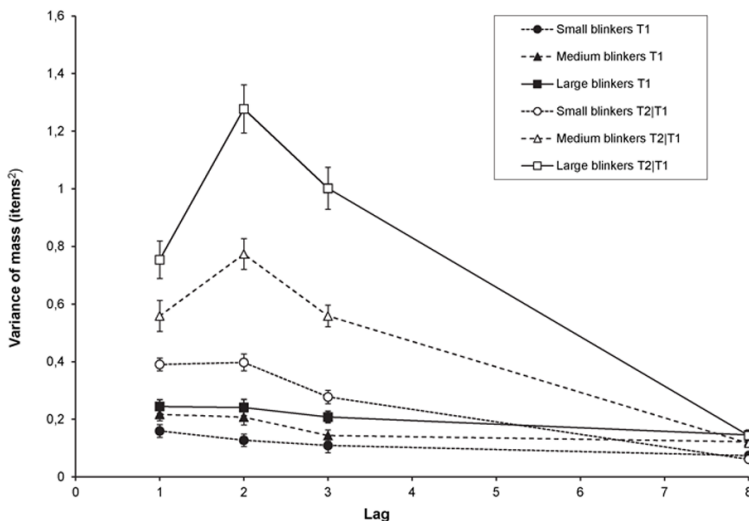**Figure 3.12. Diffusion in Experiment 2.** Diffusion of the temporal selection process expressed as the variance of the centre of mass in the selection window around T1 or T2 as a function of lag, for small, medium, and large blinkers.

In Figure 3.8, T1 accuracy and T2|T1 accuracy are plotted as a function of lag (1, 2, 3, and 8), for the small blinkers (circle symbols), the medium blinkers (triangle symbols), and the large blinkers (square symbols). A RM-ANCOVA of T1 performance with lag (1, 2, 3, and 8) as a within-subjects factor and AB magnitude as a continuous between-subjects factor (i.e., covariate) revealed no effect of lag (p = .07), but there was a main effect of AB magnitude, $F(1, 109)$ = 22.37, MSE = 116.0, p<.001, η^2p = .17, and a significant AB magnitude × Lag interaction, $F(3, 327)$ = 3.3, MSE = 10.63, p = .022, η^2p= .03.

A RM-ANCOVA of T2|T1 revealed an effect of lag, $F(3, 327)$ = 40.93, MSE = 63.83, p<.001, η^2p = .27; AB magnitude, $F(1, 109)$ = 365.59, MSE = 134.9, p<.001, η^2p = .77; and a significant AB magnitude × Lag interaction, $F(3, 327)$ = 90.39, MSE = 63.83, p < .001, η^2p = .45. These results confirm the presence of clear individual differences in AB magnitude, as illustrated in Figure 3.8.

Suppression

The amount of suppression was calculated in the same manner as in Experiment 1b. Again, a paired t-test revealed that the accuracy within the 7-item window differed significantly from the level of chance, $t(110)$ = 61.47, SE = .82, p<.001; t (110) = 99.74, SE = .52, p < .001; t(110) = 95.43, SE = .56, p < .001; t(110) = 199.97, SE = .56, p < .001 for lags 1, 2, 3, and 8, respectively.

Figure 3.9 shows the accuracy within a 7-item window for T1 and T2|T1 as a function of lag, for the different groups. A RM-ANCOVA of T1 showed an effect of AB magnitude, $F(1, 109)$ = 23.42, MSE = 6.45, p<.001, η^2p = .18; but no significant effect of lag (p = .45) or an AB magnitude × Lag interaction (p = .45).

A RM-ANCOVA of T2|T1 revealed no effect of lag (p = .30), but there was an effect of AB magnitude, $F(1, 109)$ = 88.29, MSE = 46.76, p<.001, η^2p = .45; and an AB magnitude × Lag interaction, $F(1.94, 211.65)$ = 14.81, MSE = 29.47, p<.001, η^2p = .12. Thus, as can be seen in Figure 3.9, little or no suppression occurred in small blinkers, whereas suppression of distractors as a function of lag clearly occurred in large blinkers. However it must be noted that, as in Experiment 1b, the ceiling effect might be a restrictive factor here.

The distribution of T2|T1 reports can be found in Figure 3.10. Here it can be seen that, again, the main contributors of the high accuracy in the 7-item window are the reports of the targets either preceding or following the target, plus the reports of the target itself.

Delay

The amount of delay during the temporal selection process was calculated as in Experiment 1b. The results for T2|T1 as a function of lag are plotted in Figure 3.11. For the sake of clarity, T1 is not plotted. A RM-ANCOVA of T1 showed an effect of lag,$F(3, 327)$ = 4.18, MSE = .003, p = .006, η^2p = .04; and AB magnitude, $F(1, 109)$ = 7.99, MSE = .02, p = .006, η^2p = .07; but no significant AB magnitude × Lag interaction was found (p = .66).

For T2|T1, a RM-ANCOVA showed an effect of lag, $F(3, 327)$ = 12.81, MSE = .02, p<.001, η^2p = .11; no main effect of AB magnitude (p = .33); but a significant AB magnitude × Lag interaction, $F(3, 327)$ = 12.54, MSE = .02, p<.001, η^2p = .10. As shown in Figure 3.11,

consistent with our findings in Experiment 1b, there was a remarkable delay at lag 1 for small blinkers, whereas for large blinkers the delay was most pronounced at lag 3.

Diffusion

Shown in Figure 3.12, diffusion during the temporal selection process was calculated as in Experiment 1b. A RM-ANCOVA of T1 revealed a main effect of AB magnitude, $F(1, 109) = 7.55$, $MSE = .09$, $p = .007$, $\eta^2 p = .07$; but no significant effect was found of lag ($p = .24$) or AB magnitude × Lag interaction ($p = .76$).

For T2|T1 we found a significant effect of lag, $F(3, 327) = 9.64$, $MSE = .05$, $p < .001$, $\eta^2 p = .08$; AB magnitude, $F(1, 109) = 164.85$, $MSE = .16$, $p < .001$, $\eta^2 p = .60$; and also an AB magnitude × Lag interaction, $F(3, 327) = 60.16$, $MSE = .05$, $p < .001$, $\eta^2 p = .36$. These results clearly confirm the results of Experiment 1b, namely that the temporal selection process of small blinkers is more precise than that of large blinkers. The significant interaction with lag as observed in the current experiment indicates that this is especially the case during the AB interval.

Relative T2+3 intrusions

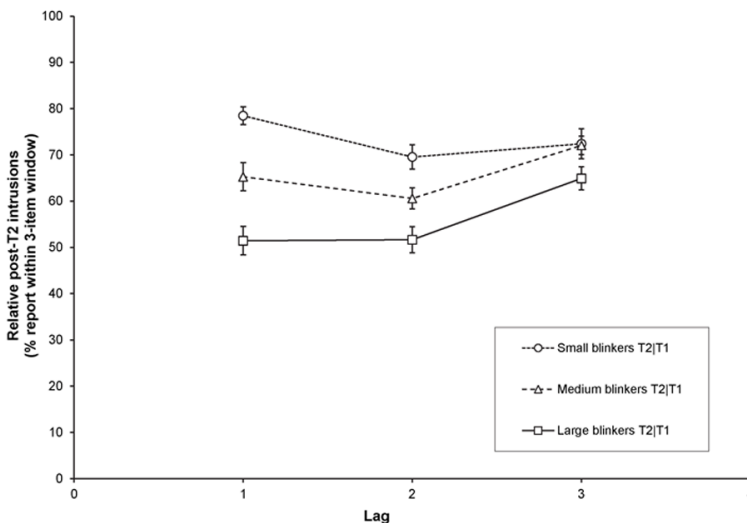Focusing on lags 1 to 3, we examined the percentage of erroneously selected letters presented

Figure 3.13. Intrusion errors in Experiment 2. Percentage of erroneously selected letters (relative to all trials with an incorrect T2 response) presented 1-3 serial positions following T2 as a function of lag, for small, medium, and large blinkers.

one to three serial positions following T2 relative to all errors on a given lag, as shown in Figure 3.13. For this analysis, the average number of trials over participants available was 14.4, 21.1, and 19.9 for lags 1, 2, and 3, respectively.

A RM-ANCOVA revealed a significant effect of lag, $F(2, 218) = 4.6$, MSE = 177.7, p = .011, $\eta^2p = .04$; AB magnitude, $F(1, 109) = 61.52$, MSE = 368.31, p<.001, $\eta^2p = .36$; and AB magnitude × Lag, $F(2, 218) = 10.09$, MSE = 177.7, p<.001, $\eta^2p = .09$, such that small blinkers made relatively more post-target intrusions than large blinkers did, particularly at the shorter lags (see Figure 3.13). Thus, besides making fewer mistakes, small blinkers made more educated guesses with the T2 response frequently matching with one of the subsequent items in the RSVP stream.

Order reversals

As in the former experiments, we calculated the percentage of order reversals for trials during which T1 and T2 were both reported correctly. Here, we found no effect of lag (p= .065), but there was a significant effect of AB magnitude, $F(1, 109) = 24.38$, MSE = 15.18, p<.001, $\eta^2p = .18$; and a significant AB magnitude × Lag interaction, $F(1.14, 124.34) = 19.09$, MSE = 10.8, p<.001, $\eta^2p = .15$, such that large blinkers had more order reversals than small blinkers did, particularly at the short lags. These results suggest that a small or absent AB does not come at a cost for temporal order information, and is better preserved for small blinkers than for large blinkers.

General Discussion

The aim of this study was threefold. Previously, we found that some individuals show little or no AB when required to identify two target letters presented in a sequential stream of non-target digits. Our first goal was to investigate whether these 'non-blinkers' would continue to show no AB when required to identify two red target letters amongst a stream of black non-target letters, thus testing the generality of their remarkable ability in avoiding an AB. Earlier, it was found that they failed to do so when targets had to be selected based on rotation or semantic features (Martens et al., 2010; Martens et al., 2010). After replicating the differential performance between blinkers and non-blinkers in a standard alphanumeric AB task, we found that when targets and distractors could only be distinguished on the basis of color, a substantial AB occurred in both groups. Though color is a stimulus property that is available relatively early in the processing pathway (Rotte et al., 1997; Smid & Heinze, 1997; Wijers, 1989), apparently early target selection was not possible to the extent that non-blinkers failed to avoid the occurrence of an AB. Combined with the previous observation of an AB in non-blinkers when alphanumeric stimuli were presented in the auditory modality (Martens et al., 2009), the current results seem to suggest that the non-blinkers' ability might indeed be quite task-specific, requiring the presence of visual alphanumeric category information. However, given that AB magnitude in our colored targets task remained smaller in non-blinkers than in blinkers, there must be more to the story.

Interestingly, the colored targets paradigm as employed here allowed us to study individual differences in target selection efficiency in more detail. More specifically, our second

aim was to study possible differences in the temporal profile of blinkers and non-blinkers by examining the amount of suppression, delay, and diffusion of the temporal selection process during the AB (Vul et al., 2008). We expected to find differences in these three dissociable dimensions of temporal selection, because even in the colored target task clear differences in AB magnitude were observed.

Suppression

Surprisingly, little suppression was observed in both Experiments 1b and 2; the efficacy of selection, measured as the percentage of trials during which an item was reported from a 7-item window around either T1 or T2 (i.e., spanning three items before to three items after the target), was generally high. In Experiment 1b, a significant difference between blinkers and non-blinkers in the amount of suppression for T2 was found, which, however, was not modulated by lag. This finding is similar to what was reported by Popple and Levi (2007). It must be noted though that in their study, as well as in the current one, patterns of AB-induced suppression may have been obscured by ceiling effects.

In Experiment 2, employing a larger sample of subjects and thus a wider range of AB magnitudes, the interaction of AB magnitude and lag reflected signs of suppression of T2 and the surrounding distractors at the shortest lags for large blinkers, whereas small blinkers continued to show no suppression whatsoever. Although the finding of suppression as a function of lag corresponds with findings from previous studies (Botella et al., 2011; Chun, 1997; Vul et al., 2008), all of these papers reported substantially more suppression.

An explanation for these differential findings might lie in differences in methods, stimuli, and overall task difficulty. Whereas both our study and that of Popple and Levi (2007) employed integral dimensions of the stimuli as the relevant features (color and shape), Vul et al. (2008) as well as Chun (1997) used composed targets (a letter surrounded by an annulus or colored frame). Although the study by Botella and colleagues (2011) did use color as an integrated target feature, they introduced a task-switch by varying the color of the two targets, and possibly reduced the effectiveness of color as a target-specific feature by also varying the color of each distractor in the stream. It is thus not inconceivable that the latter studies introduced additional factors into the AB task that further complicated the binding and subsequent selection of targets. In addition, the level of overall performance in (Vul et al., 2008) was dramatically low (~10–50%), making comparisons with other AB studies—that typically feature much higher performance—difficult.

Another notable finding pertains to the individual differences in the amount of suppression. In multiple studies it has been suggested that the AB is due to a failure to effectively suppress distractors (Dux et al., 2006; Dux & Harris, 2007; Dux & Marois, 2008; Harris et al., 2010). Specifically, based on findings in their priming study, Dux and Marois (2008) suggested that large blinkers in particular fail to suppress the processing of irrelevant distractors, whereas small blinkers frequently manage to avoid an AB by successful suppression of these distractors. If that would indeed be the case, however, one would expect to see strong suppression in non-blinkers and little or no suppression in large blinkers, exactly opposite to the pattern of findings reported here.

Instead, we propose that non-blinkers are somehow able to select targets at an earlier processing stage than blinkers do, to some extent even when targets are not defined by alpha-

numeric category. Consequently, compared to blinkers, non-blinkers may have little need to suppress distractors, as stable target representations can more readily and easily be formed. The less effective this early selection, the stronger the need for suppression at a later stage of processing, a pattern that is indeed in line with the levels of suppression that we observed in small, medium, and large blinkers, respectively (see Figure 3.9). However, it must be noted that given the relatively modest amount of suppression observed in the current study, it is hard to conceive that suppression alone can account for the significant AB that was obtained in the majority of participants. Moreover, it remains puzzling why the strongest suppression tended to occur at lag 1, whereas the strongest AB was consistently found at lag 2.

Delay

Another surprising finding emerged in the latency measure of the intrusion errors. Following Vul et al. (2008) and Chun (1997), the center of mass was calculated as a measure of delay. Whereas for large blinkers, the maximal delay was consistently found at lag 3, for small blinkers the maximum in both experiment 1b and 2 was observed at lag 1. This latter finding, however, may at least partly reflect an artifact of the T2 center of mass calculation, and at first sight does not seem to be very meaningful. That is, the small blinkers' seemingly large delay at lag 1 may be the simple consequence of a) the fact that the diffusion of responses was substantially smaller for small blinkers than for large blinkers (who made intrusions from a wider window; see section below), b) the fact that small blinkers made relatively more post-target intrusions than blinkers did (see Figure 3.7 and Figure 3.13), and c) the fact that correct T1 responses are excluded from the calculation. The combination of these factors at lag 1 may thus be responsible for an inflated center of mass for small blinkers, and a center of mass that is close to zero for large blinkers. However, given that the results found here correspond to the pattern of relative post-target intrusion errors (further discussed below), they may nevertheless reflect a genuine difference between small and large blinkers.

The pattern of results is quite different from that reported by Vul et al. (2008) and Chun (1997), who both reported finding a negative center of mass at the shortest lags. Again, an explanation might lie in differences in methods, stimuli, and overall task difficulty, as well as the fact that their participants showed more suppression than the individuals in the current study did.

Diffusion

Perhaps the most telling and straightforward finding is provided by the measure of diffusion, expressing the precision of selection for each group of individuals. Calculated as the variance of the center of mass, the amount of diffusion showed a consistent pattern that matched closely with that of the AB, reaching the lowest temporal precision at lag 2. Although the amount of overall diffusion was much lower than that reported by Vul et al. (2008), the pattern of diffusion as a function of time between the targets is very similar. In addition, our current findings clearly showed that, compared to small blinkers, large blinkers were less precise in selecting the second but not the first target.

This pattern of diffusion fits with the idea that non-blinkers are able to select targets

at an earlier processing stage than blinkers do. Early target selection may reduce interference from distractors, allowing subsequent processing of the targets to proceed faster and more accurately in non-blinkers than in blinkers, reflected in earlier P3s (Martens et al., 2006; Martens et al., 2010) and less diffusion.

Relative intrusion errors

In addition to these three dimensions of temporal selection, we analyzed the percentage of erroneously selected letters presented one to three serial positions following T2 relative to all errors on a given lag (see Figure 3.7 and Figure 3.13). Errors in the temporal selection process have been studied before (Botella et al., 2011; Chun, 1997; Popple & Levi, 2007), but individual differences were not considered and differences in the total number of errors were not controlled for. Given that intrusions of items following T2 are inherently related to the total number of errors made, we studied the relative number of intrusions, allowing comparisons between blinkers and non-blinkers in the type of intrusions irrespective of the total rate of response errors. In both Experiments 1b and 2, we found that non-blinkers and small blinkers made relatively more post-T2 intrusions than blinkers did. In Experiment 2, within the group of small blinkers, most post-T2 intrusions were made at lag 1, whereas within the group of large blinkers most of these intrusions occurred at lag 3. This pattern matches quite well with the differences in delay that we observed for the different groups, but poses a challenge in terms of interpretation. Although we argued that the latter differences might at least partly be due to the way in which the center of mass was calculated, the significant interaction between group and lag in the relative post-T2 intrusions does indicate systematic differences in the selection process employed by blinkers and non-blinkers, especially at lag 1.

Note however, that some caution is generally required in the interpretation of what a shift in the center of mass as well as the number of relative post-target intrusion errors actually reflect. Given that the particular task employed in the current study required the binding of a color to a particular letter, the delay that is associated with a positive shift in the center of mass or an increase in post-target intrusions may be due to non-blinkers and blinkers having differential processing speeds in either the color, letter, or the binding of features (or a combination thereof). Future research is needed to isolate these different components of the temporal selection process.

Relative order reversals

In response to the proposition that the AB reflects a cognitive strategy of enforcing an episodic distinction between successive stimuli of Wyble, Bowman, and Nieuwenstein (2009), our third and final aim was to determine whether avoiding an AB comes at a cost. Given the non-blinkers' ability to largely avoid the occurrence of an AB, information concerning temporal order and the correct binding of features into targets might be compromised in non-blinkers. If that were indeed the case, non-blinkers should show relatively more order reversals, compared to large blinkers. However, while correcting for differences in target accuracy, the opposite pattern of results was observed. Although no significant difference in relative order reversals was found between blinkers and non-blinkers in Experiment 1b, individuals with

little or no AB showed fewer rather than more order reversals than large blinkers as showed in Experiment 1a and 2. Even though the AB may have a functional role in providing episodic distinctiveness, our results suggest that avoiding an AB does not come at a cost for temporal order information.

Conclusions

By studying individual differences in response errors, we found that only a modest amount of suppression of T2 and surrounding distractors was present in blinkers. In addition, lower accuracy was closely accompanied by reduced precision during target selection in blinkers. In comparison, the temporal selection process seems to be faster and more precise in non-blinkers, and we found no evidence of suppression. Non-blinkers did show a sizeable AB when target selection was based on color features rather than alphanumeric category, but continued to outperform blinkers. Finally, we found that non-blinkers did not lack episodic distinctiveness; temporal order information was actually preserved better in individuals with a small rather than a large AB. Intriguingly, non-blinkers showed most intrusions as well as a selection delay at lag 1, a finding that deserves further investigation.

Acknowledgments

Author Contributions

Conceived and designed the experiments: CW SMW EV SM. Performed the experiments: EV. Analyzed the data: CW SMW SM. Wrote the paper: CW SMW SM.

Chapter 4: Word frequency and the attentional blink

The effects of target difficulty on retrieval and consolidation processes

In which we show that difficult words can outlast the attentional blink

Chapter was previous published in:

Stefan M. Wierda

Niels Taatgen

Hedderik van Rijn

Sander Martens

Abstract

Background

When a second target (T2) is presented in close succession of a first target (T1) within a stream of non-targets, people often fail to detect T2—a deficit known as the attentional blink (AB). Two types of theories can be distinguished that have tried to account for this phenomenon. Whereas attentional-control theories suggest that protection of consolidation processes induces the AB, limited-resource theories claim that the AB is caused by a lack of resources. According to the latter type of theories, increasing difficulty of one or both targets should increase the magnitude of the AB. Similarly, attentional-control theories predict that a difficult T1 increases the AB due to prolonged processing. However, the prediction for T2 is not as straightforward. Prolonged processing of T2 could cause conflicts and increase the AB. However, if consolidation of T2 is postponed without loss of identity, the AB might be attenuated.

Methodology/Principal Findings

Participants performed an AB task that consisted of a stream of distractor non-words and two target words. Difficulty of T1 and T2 was manipulated by varying word-frequency. Overall performance for high-frequency words was better than for low-frequency words. When T1 was highly frequent, the AB was reduced. The opposite effect was found for T2. When T2 was highly frequent, performance during the AB period was relatively worse than for a low-frequency T2. A threaded-cognition model of the AB was presented that simulated the observed pattern of behavior by taking changes in the time-course of retrieval and consolidation processes into account. Our results were replicated in a subsequent ERP study.

Conclusions/Significance

The finding that a difficult low-frequency T2 reduces the magnitude of the AB is at odds with limited-resource accounts of the AB. However, it was successfully accounted for by the threaded-cognition model, thus providing an explanation in terms of attentional control.

Introduction

It is well known that the human mind is limited in the conscious processing of relevant stimuli (e.g., letters) when presented in close temporal proximity in a sequential stream of irrelevant stimuli (e.g., digits). Most people show a reduced ability to successfully report a second target (T2) when presented within 200-500 ms of a first (T1), a phenomenon known as the attentional blink (AB; Martens & Wyble, 2010; Raymond et al., 1992). Although there are a diversity of models and theories of this phenomenon, they can roughly be divided in two types: limited-resource accounts (Chun & Potter, 1995; Jolicœur & Dell'Acqua, 1998; Shapiro et al., 1994) and attentional-control accounts (Di Lollo et al., 2005; Olivers & Meeter, 2008; Taatgen et al., 2009; Wyble et al., 2009).

In limited-resource accounts of the AB, the common assumption is that there is a pool of resources available for processing targets and that this pool is limited. Whenever a target must be stored for later report, resources are drawn from the resource-pool in order to consolidate that target. Because this pool of resources is limited, there is a chance that the pool is still depleted due to the ongoing consolidation of T1 at the moment that T2 is encountered. Because there are not enough resources available for the processing of T2, an AB occurs. Thus, in these theories, a capacity-limitation of the attentional system underlies the phenomenon of the attentional blink.

On the other hand, there are theories that advocate an attentional-control account of the AB. The common theme in these theories is that processing of T1 is being protected by an attentional-control mechanism. Whenever a distractor is encountered, some kind of protection mechanism is trigged, preventing incoming information to be consolidated into working memory, effectively protecting the consolidation of T1. Because T2 is presented while T1 is being consolidated, the protection mechanism prevents T2 from being consolidated. Thus, whereas limited-resource accounts contribute the AB to a limited pool of resources, attentional-control accounts attribute the AB to some process actively suppressing the consolidation of new information.

In attempts to test and contrast these theories, several studies have been conducted that manipulated the difficulty of T1. Two types of difficulty manipulations can be distinguished: data-limited and resource-limited manipulations. Following the definitions of Norman and Bobrow (1975), data-limited manipulations affect the physical characteristics of the stimuli (e.g., contrast), whereas resource-limited manipulations affect the difficulty of a task (e.g., number of candidate targets). Most AB studies that varied T1 difficulty employed data-limited manipulations, but the results have been mixed, with some studies finding an increased AB (Christmann & Leuthold, 2004), an attenuated AB (Chua, 2005; Chun & Potter, 1995; Raymond et al., 1992), or no effect (McLaughlin et al., 2001; Ward, Duncan, & Shapiro, 1997). Others have reported data-driven difficulty effects on the AB, but only when T1 was not masked (Visser & Ohan, 2007; Visser, 2007). It is known that if the distractors following T1 are replaced by blanks, essentially removing the mask on T1, the AB is clearly attenuated; the longer the duration of the blank interval, the smaller the AB (Chun & Potter, 1995; Raymond et al., 1992). However, in the majority of AB studies, targets are typically masked by a subsequent distractor, which is often considered as a requirement to induce an AB (but see Nieuwenstein, Potter, & Theeuwes, 2009). The mixed results by the studies described here make it hard to find conclusive evidence for either limited-resource or control-process accounts of the AB.

Although fewer in number, studies employing resource-limited manipulations show a

more consistent pattern of results. Tasks that increased the informational load associated with T1 encoding typically produced a larger AB. Shapiro and colleagues (Shapiro et al., 1994) first showed that increasing the set-size from which a T1 could be drawn from 3 to 25 increased the AB. However, it should be noted that this was tested between rather than within subjects (10 in each group), and that data from the difficult condition (set size 25) was obtained from a different study (experiment 2 from Raymond et al., 1992) using a slightly different procedure. An alternative explanation, for instance in terms of individual differences between groups, can thus not be ruled out.

In another study that manipulated difficulty to affect the AB (Ouimet & Jolicoeur, 2007), a T1 was used that consisted of five digits. The digits 0-4 were presented either in ordered (i.e., '01234') or shuffled (e.g., '04231') sequence. Participants had to report whether and in what sequence the target item occurred in the RSVP. The task for T2 was to identify a single digit represented by a 5-digit number (e.g., '33333'). It was found that the ordered sequence produced a minimal AB compared to the AB produced by the shuffled digits. However, one should be cautious interpreting these results, for the ordered (low load) task could be seen as a recognition task (merely remember whether an ordered sequence was presented), whereas identification is required for the shuffled condition (report the full sequence). Furthermore, as the tasks for T1 and T2 were different, the effect of a task-switch potentially confounded the results (Potter et al., 1998; Potter, Wyble, Pandav, & Olejarczyk, 2010).

A third example of a resource-limited manipulation is provided by Martens and colleagues (Martens et al., 2006). In their study, T1 difficulty was manipulated by changing the probability of occurrence associated with the identity of T1 (i.e., one of the candidate targets occurred more often than the other target items). It was found that an infrequently reoccurring T1 target letter induced a larger AB magnitude than a frequently reoccurring T1 letter.

In a fourth study, difficulty of T1 was manipulated by varying the word frequency. Burt, Howard, and Falconer (2011) showed that the AB is attenuated by word frequency. Participants had to identify two color-marked words in a stream of irrelevant pseudo-words. They found that high-frequency words induced a smaller blink than low-frequency words. The T2 word was always medium frequent. According to the authors, the T1 difficulty effects are more readily accounted for by limited-resource than by attentional-control theories.

Another line of evidence comes from event-related potential (ERP) studies. A late parietal component—the P300—has been associated with the AB (Kranczioch, Debener, & Engel, 2003; McArthur et al., 1999; Vogel et al., 1998), and is thought to reflect processes involved in the consolidation of targets into working memory (Isreal, Wickens, Chesney, & Donchin, 1980; Polich, 2007). During the AB critical period, the P300 is suppressed for the second target (Vogel et al., 1998). However, earlier components associated to perceptual processing and the relatively late N400 (associated to semantic processing) can still be found (Luck, Vogel, & Shapiro, 1996; Rolke, Heil, Streb, & Hennighausen, 2001; Vogel et al., 1998). These findings indicate that—to some extent—targets are being processed up to the semantic level, and are presumably accessed in memory, but are nevertheless not available for consciousness report. Therefore, the impairment seems to be at a post-perceptual stage of processing specifically related to the consolidation of a target for later report.

In addition, effects of the P300 found in AB studies could be taken as evidence in favor of resource-depletion theories, because manipulations that cause targets to elicit larger P300 amplitudes are generally found to increase AB magnitude, which suggest some kind of trade-off between the amount of processing and the probability that a target is detected

(McArthur et al., 1999). For example, when a secondary task has to be performed next to the AB task, both P300 amplitude and AB magnitude decrease (Wierda et al., 2010) (although resource-limited theories would have some issues explaining why a secondary task increases performance on the primary task). Indeed, some argue that the P300 can be used to index the allocation of resources (Isreal, Chesney, Wickens, & Donchin, 1980; Isreal et al., 1980; Kranczioch & Bryant, 2011), but one should be cautious to interpret the amplitude of the P300 as a direct index of resource allocation. For example, whereas high-frequency words are easier to detect than and induce a smaller AB than low-frequency words (Burt et al., 2011), they elicit a larger P300 amplitude than low-frequency words (Polich & Donchin, 1988; see Kok, 2001 for a review of manipulation effects on the amplitude of the P300).

Several of the abovementioned studies have revealed evidence that the difficulty of the AB inducing task can influence the magnitude of the AB, but very few manipulated T1 difficulty within subjects without adding additional stimuli (Burt et al., 2011; Martens et al., 2006). As mentioned above, Burt et al. (2011) argued that their findings support limited-resource rather than attentional-control accounts of the AB. The goal of the current study was to replicate their findings and further investigate whether the AB is caused by a limitation in resources or by attention-control processes. Limited-resource accounts predict that performance should decrease when T2 is made more difficult. Because T1 and T2 are supposed to draw resources from the same limited-resource pool, the difficulty (in terms of frequency) of both T1 and T2 should affect the magnitude of the AB in a similar fashion. Predictions made by attentional-control accounts are more subtle. Whereas attentional-control theories also predict that a difficult T1 would increase the magnitude of the AB due to prolonged processing of T1, the predictions made for T2 are not as straightforward. Prolonged processing of T2 would affect the AB only at lag 2, when the protection mechanism for consolidation of T1 is trigged, and the effect could go both ways. Either the prolonged processing of T2 directly competes with both the processing and protection of T1, leaving no room for processing T2 and thus decreasing the probability of T2 to survive the AB period, or the prolonged processing could carry the target beyond the duration of the consolidation of T1 and its protection mechanism, increasing the probability of T2 to be consolidated. Thus, its prediction relies on the subtle timing of the target identification and consolidation processes.

Similar to the study of Burt et al. (2011), a natural manipulation was employed in the current study by using words that intrinsically varied in frequency of usage outside the context of the experiment. It is known that high-frequency words are processed faster and identified with greater accuracy than low-frequency words (Van Rijn & Anderson, 2003). And because high-frequency words induce a smaller AB than low-frequency words, we assume high-frequency words to be easier targets than low-frequency words.

In the current study, targets consisted of words within a stream of unpronounceable non-words (Experiment 1) or within a stream of digits (Experiment 2). Target difficulty varied as a function of word frequency, without the need for stimulus degradation or other perceptual manipulations. Resource depletion theories predicted that overall identification performance for low-frequency targets would be lower than for high-frequency words, and more importantly, that a low-frequency T1 would induce a larger AB effect on a subsequent T2. Furthermore, a low-frequency T2 would require more attention or resources. According to most resource-limited theories, the largest AB was thus likely to occur for a low-frequency T2 following a low-frequency T1. Whereas attentional-control theories also predict a negative impact of a difficult T1 on the AB due to the prolonged duration of processes needed to

identify and consolidate T1, the predictions made by attentional-control theories on the effect of difficulty of T2 are less straightforward. Difficulty affects the timing of different parallel processes, and as such can have either a positive or negative effect on the AB, depending on the onsets, offsets, and duration of cognitive processes during the critical AB period in which consolidation of T1 is being protected.

As described below, an attenuated AB was observed when T2 became more difficult, which is hard to explain with any resource-depletion theory. An extension of our threaded cognition model of the AB (Taatgen et al., 2009) is therefore presented, providing an explanation in terms of attentional control for this somewhat surprising finding. To confirm the finding that a low-frequency T2 is relatively easier to detect than a high-frequency T2 during the AB critical period, an ERP experiment (Experiment 2) was conducted in which only the word-frequency of T2 was manipulated. In line with the result of Experiment 1, and consistent with the model, again a relatively small AB was observed for low-frequency words when compared to high-frequency words. Also, smaller P300 amplitudes were found for low-frequency words when presented at long lags—consistent with findings that low-frequency words induce a smaller P300 than high-frequency words—but no difference between high-frequency and low-frequency words was observed during short lags.

Experiment 1

Methods

Participants

Twenty native German speaking psychology students (aged 18-25, mean = 20.4, with normal or corrected-to-normal visual acuity) from the University of Groningen were recruited via an online sign-up program, and received course credits for participating in the experiment. Informed consent was obtained prior to the experiment. The Ethical Committee Psychology of the University of Groningen approved the experiment.

Stimuli and Apparatus

E-Prime 1.2 software was used to generate stimuli and to collect responses, running under Windows XP on a PC with a 17-inch 100-Hz CRT monitor. In total, 576 high-frequency (HF; Mannheim frequency 63 to 6413) and low-frequency (LF; Mannheim frequency 9 to 19) German words (four to six letters in length) were pseudo-randomly picked from the German word forms CELEX corpus (Baayen, Piepenbroek, & Van Rijn, 1993). The target words were balanced for word length and word frequency. Distractor stimuli were pseudo-randomly generated strings of consonants, consisting of the same number of characters as the targets on a given trial. The first letter of each word and non-word was presented in uppercase. The remaining letters were presented in lowercase. All stimuli were presented in black, Courier New font, size 18, on a white background at a viewing distance of ~50 cm. The monitor's resolution

during the experiment was set at 1,024 x 768 pixels.

Procedure

The experiment consisted of one practice block and three testing blocks, with a short break between each testing block. The practice block contained 9 trials and each testing block contained 288 trials.

The participants' task was to identify two words (the targets) presented amongst a rapid serial visual presentation (RSVP) stream of non-words (the distractors). Participants were instructed to fixate on a cross in the middle of the screen. After pressing the spacebar, the fixation cross remained on the screen for 750 ms, followed by a blank screen. After 100 ms, the stream was presented, consisting of 22 stimuli. Each stimulus in the stream was presented for a duration of 150 ms without inter stimulus interval. T1 was always presented on the fifth temporal position within the stream. T2 was presented on the first, second, or seventh position after T1 (i.e., lag 1, 2, or 7). Within each block, each combination of lag, T1 and T2 word frequency (HF-HF, LF-HF, HF-LF, and LF-LF), and word length (4 to 6) was presented equally often. A specific word was never presented twice on the same trial.

At the end of the stream, a question mark appeared, prompting participants to verbally report T1 and T2 to the experimenter. The correct answers were presented to the experimenter on a second display. Using the numeric keypad on a keyboard, the experimenter typed a "0" if a response matched with T1, a "1" if it matched with T2, a "2" if no response was given, and a "3" if it matched with neither of the targets. Responses were accepted and counted correct regardless of the order in which they were reported.

Data Analysis

Following Wierda et al. (2010) and Wierda, van Rijn, Taatgen, and Martens (2012), accuracy scores were analyzed using binominal mixed effects models. Given that our hypothesis predicted a different number of observations per cell, mixed effects models are preferred over methods that assume an equal number of observations per cell. Analyses were performed using the lme4 package (version 0.999375-31; Bates & Sarkar, 2008). Lag, T1 word frequency, and T2 word frequency were entered as fixed factors in each model. For both word-frequency factors, the natural logarithm of the Mannheim word frequencies was entered in each model as continuous predictor. Subject was entered as random factor in each model.

Results and Discussion

Figure 4.1 shows identification performance of T1 as a function of lag. Identification perfomance of T2 given that T1 is correctly identified is shown in Figure 4.2.
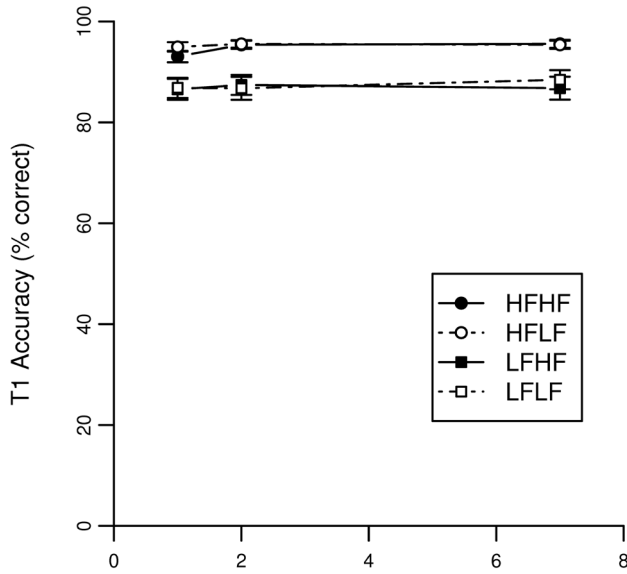
Figure 4.1. Accuracy scores of the AB task for T1 in Experiment 1. The lag corresponds to the temporal location of T2 relative to T1.

Table 4.1. The estimates and z-values of the mixed-effects model for T1 accuracy.

	Mixed-effects model T1			
	Estimate β	Standard Error	z-value	p-value
(Intercept)	2.204	0.208	9.778	0.000
Word frequency T1	0.308	0.064	4.805	0.000
Word frequency T2	-0.054	0.048	-1.128	0.259
Lag 1	-0.188	0.178	-1.055	0.291
Lag 2	-0.288	0.180	-1.598	0.110
Word frequency T1, Word frequency T2	0.028	0.025	1.151	0.250
Word frequency T1, Lag 1	0.093	0.088	1.059	0.289
Word frequency T1, Lag 2	0.130	0.092	1.418	0.156
Word frequency T2, Lag 1	0.054	0.065	0.823	0.411
Word frequency T2, Lag 2	0.832	0.068	1.227	0.220
Word frequency T1, Word frequency T2, Lag 1	-0.063	0.032	-1.981	0.048
Word frequency T1, Word frequency T2, Lag 2	-0.039	0.034	-1.119	0.263

T1 accuracy

A binominal mixed effects model was fitted on the accuracy of T1. Table 4.1 lists the statistics for the model's factors. Here we will focus on the three significant estimates. The lag 7 condition was used as baseline and is reflected in the intercept. This factor indicates that a word with a natural logarithmic frequency of 0 would be responded to correctly in X% of all trials. The model revealed that T1 frequency predicts T1 accuracy, such that performance increases with T1 frequency (β=0.308). The three-way interaction between Lag 1, T1 word frequency, and T2 word frequency (β=-0.063) indicates that if the two targets immediately follow each other, the positive influence of the word frequency of T1 on the accuracy is adjusted downwards as a function of the frequency of T2. This suggests that at short lags, the frequency of the second word might interfere with the processing of the first word.

T2 accuracy

A binomial mixed effects model was fitted on T2 accuracy for trials with a correct T1 response. Table 4.2 shows statistics for each fixed factor. A marginally significant effect was found for the word frequency of T1, indicating that there is an overall long-lasting frequency effect (β=0.073, p=0.060) of T1. Interestingly, this marginally significant effect is positive, suggesting that a higher frequency for T1 is associated with better performance on T2. The main effect of T2 word frequency (β=0.297) is similar to the effect of word frequency on T1 (β=0.308),

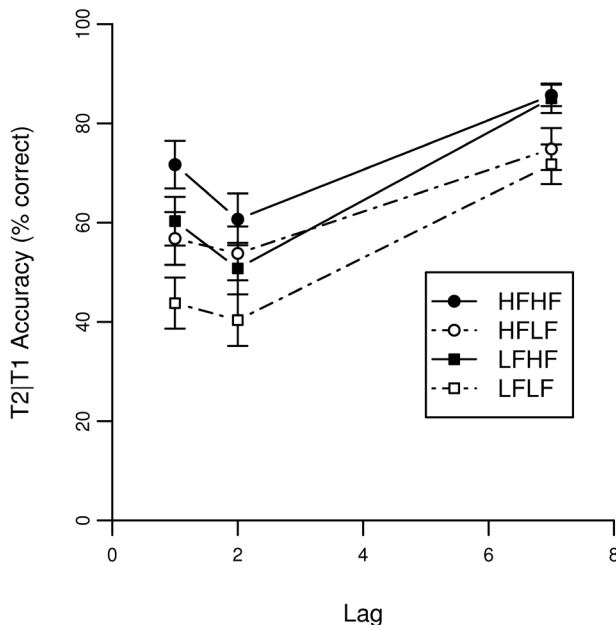**Figure 4.2. Accuracy scores of the AB task on lag 1, 2 and 7 for T2 given T1 correct in Experiment 1. The lag corresponds to the temporal location of T2 relative to T1.**

Table 4.2. The estimates and z-values of the mixed-effects model for T2|T1 accuracy.

	Mixed-effects model T2|T1			
	Estimate β	Standard Error	z-value	p-value
(Intercept)	0.880	0.260	3.381	0.001
Word frequency T1	0.073	0.039	1.878	0.060
Word frequency T2	0.297	0.044	6.687	0.001
Lag 1	-1.582	0.135	-11.679	0.001
Lag 2	-1.600	0.135	-11.840	0.001
Word frequency T1, Word frequency T2	-0.009	0.016	-0.532	0.595
Word frequency T1, Lag 1	0.162	0.051	3.196	0.001
Word frequency T1, Lag 2	0.144	0.051	2.850	0.004
Word frequency T2, Lag 1	0.017	0.056	0.299	0.765
Word frequency T2, Lag 2	-0.130	0.055	-2.348	0.019
Word frequency T1, Word frequency T2, Lag 1	-0.008	0.021	-0.379	0.704
Word frequency T1, Word frequency T2, Lag 2	-0.000	0.020	-0.011	0.990

indicating that accuracy on T2 increases with higher natural-logarithmic word-frequencies in a similar manner as for T1. The negative estimate of lag 1 and lag 2 reflects the AB, showing that during the AB critical period performance is lower than outside the AB critical period (at lag 7). Furthermore, the interaction between T1 word frequency and both lag 1 and lag 2 demonstrates that the AB is modulated by T1 frequency. The positive estimate implies that the AB is larger when T1 is low frequent. Finally, an interaction between T2 frequency and lag 2 was found. The negative estimate indicates that the AB is relatively larger when T2 is highly frequent. This latter finding is somewhat surprising, as the effect is only found at lag 2, and one might expect a larger AB when T2 is difficult rather than easy, following limited-resource theories on the AB. However, in the next section, we describe a computational model of the AB that provides an explanation for this effect.

Model

In order to explain the patterns in the data, in particular the finding that a high-frequency T2 leads to a larger rather than smaller AB, we modified the threaded cognition (TC) model of the AB by Taatgen et al. (2009) to fit the current task. The TC model, which is based on the ACT-R cognitive architecture (Anderson, 2007) assumes that several cognitive modules are involved in the AB task. More in particular, a visual module is needed to perceive the input, a declarative memory module is necessary to assess the category of an input (e.g., target versus distractor), and an imaginal module is used to consolidate targets (comparable to working memory). Finally, procedural memory coordinates the flow of information (Figure 4.3). The
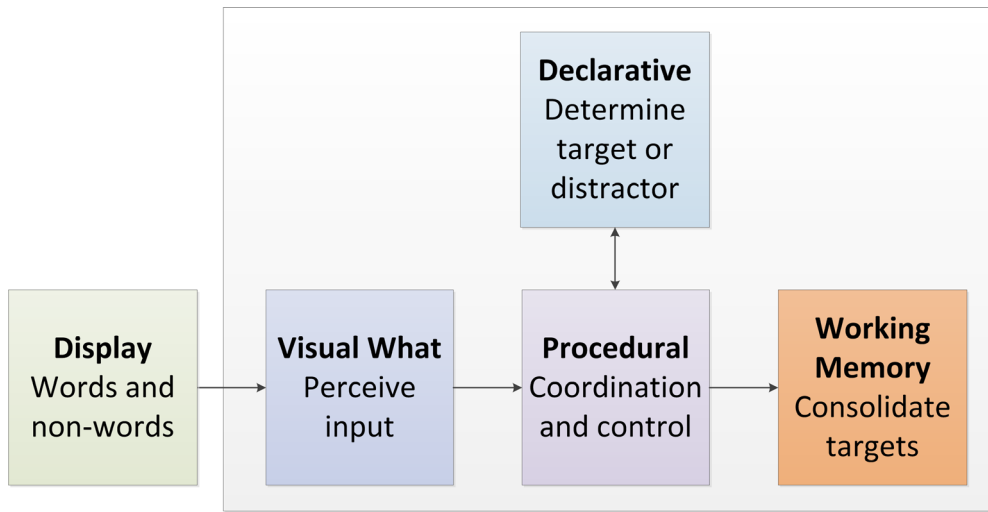
Figure 4.3. An overview of the modules and their role in the TC AB model.

TC assumption is that all modules can operate in parallel, but that a single module can only do one thing at a time. In the TC model, the AB is explained by a (procedural) control strategy that blocks the scanning for targets during memory consolidation. This control strategy is employed when a distractor is encountered. This explanation has similarities with those offered by some other models, in particular the Boost and Bounce model (Olivers & Meeter, 2008) and the eSTST model (Wyble et al., 2009). Specific about the TC model is that this control strategy has to compete with other processes, which enables it to explain why the AB is reduced in cases where there is distraction or a secondary task (Olivers & Nieuwenhuis, 2005; Taatgen et al., 2009; Wierda et al., 2010).

In order to fit the model to the current experiment, we changed the timing of the model to comply with the current experiment, and slightly changed the function of declarative memory. In the standard model, declarative memory was mainly used to determine the category of the stimulus, but now it is used to retrieve the representation of the word so that it can be reported later on. The assumption of the model is that the retrieval time of a low-frequency word is longer than that of a high-frequency word, and that the accuracy of identifying a word is also slightly lower. This is consistent with previous ACT-R models of lexical decision (Van Rijn & Anderson, 2003). Furthermore, a second assumption is that it takes slightly longer to consolidate a low-frequency word in memory than a high-frequency word. Although an intervening distractor causes the AB in the model, performance on lag 1 (i.e., no intervening distractor) is almost as low as performance on lag 2 (i.e., during the AB period). Whereas the low performance on lag 2 is explained by the control strategy to protect T1 consolidation, performance on lag 1 is due to the direct competition between processes needed to consolidate T1 and T2. It is important to note that there was no difference in performance or fit between the modified model as presented in this study and the original model as reported by Taatgen et al. (2009).

The crucial aspect of the model that can explain why the AB is relatively smaller in the
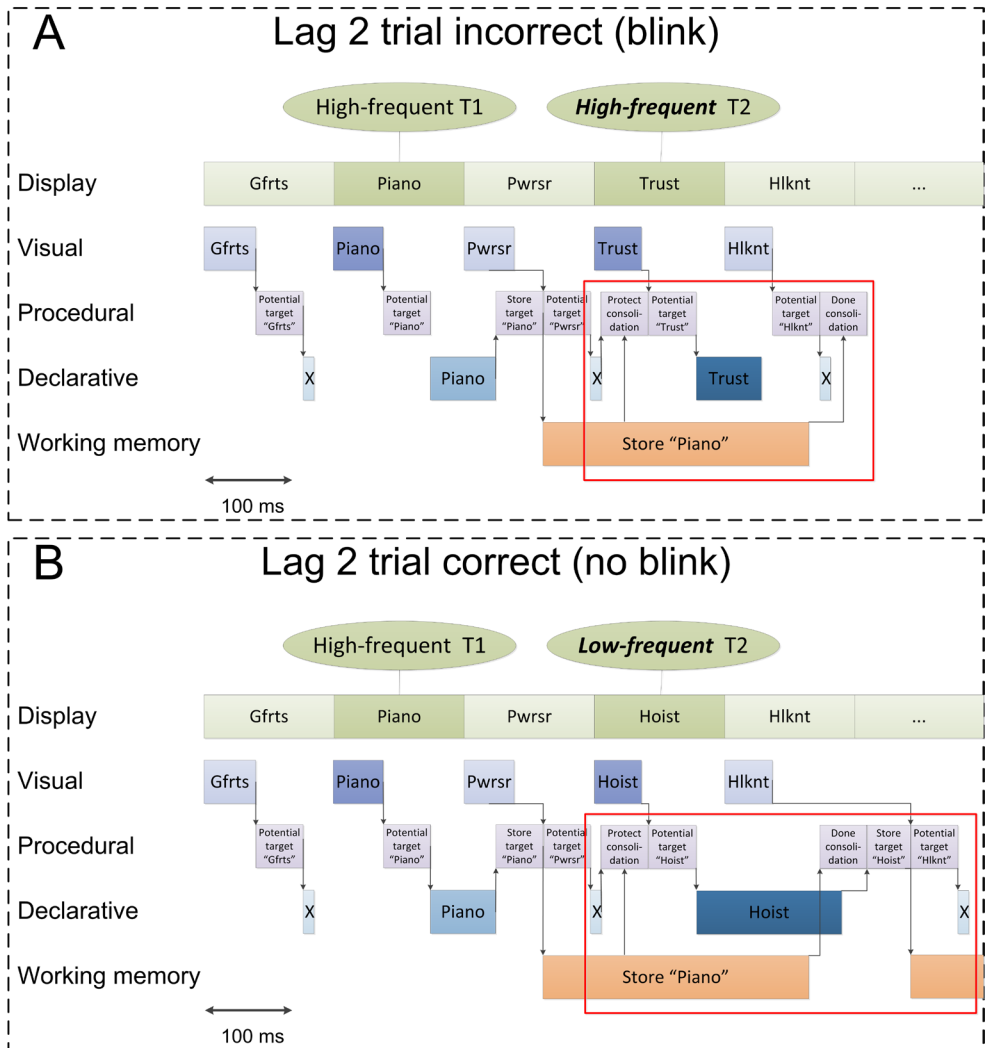
Figure 4.4. Examples of the model traces for the HFHF-condition (a) and the HFLF-condition (b). The second target was presented at lag 2.

cases where the T2 is of low frequency is that retrieving that word sometimes extends beyond the consolidation of T1, surpassing the strategic protection of consolidation. This is illustrated in Figure 4.4, where the activity of the four modules (along with a row representing the input) is displayed. Figure 4.4A illustrates a HF-HF trial in which there is an AB. After the word piano has been detected, the "Protect Consolidation" step in the procedural module temporarily prohibits targets from being consolidated, resulting in an AB. In the HF-LF example in Figure 4.4B, on the other hand, retrieval of the word hoist extends beyond the consolidation of T1, and therefore does not result in an AB. Because at Lag 1, no intervening distractor triggers the protection of T1 consolidation, the effect is absent for lag 1. The results of the model are
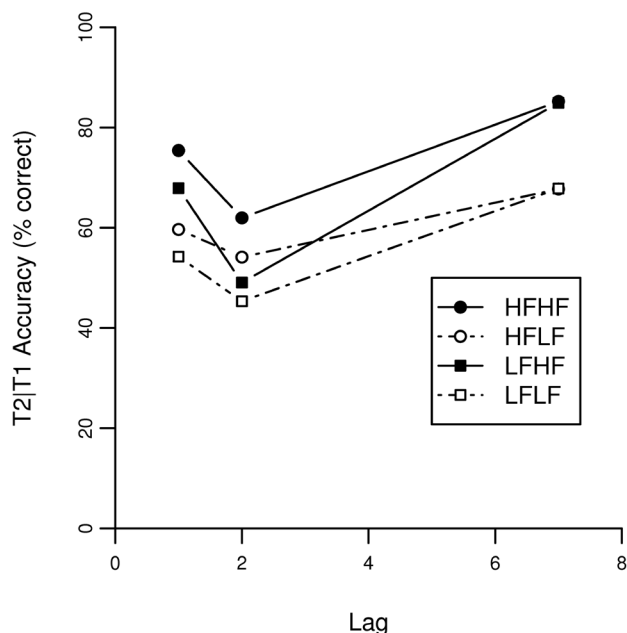
Figure 4.5. Accuracy scores of the AB task as produced by the model on lag 1, 2, and 7 for T2 given T1 correct. The lag corresponds to the temporal location of T2 relative to T1.

shown in Figure 4.5, and fit the overall patterns in the data quite well.

In order to verify the results of Experiment 1 and test the hypothesis that indeed late processes were affected by word frequency, we set up an ERP experiment and focused on the P300 component, which is associated with late-stage processing of targets and is strongly related to the AB phenomenon (e.g., Kranczioch, Debener, Schwarzbach, Goebel, & Engel, 2005; Martens et al., 2006; McArthur et al., 1999; Vogel et al., 1998; Vogel & Luck, 2002; Wierda et al., 2010). Providing converging evidence for the observed patterns of behavior, we expected to find ERP differences associated with late-stage processing of LF words relative to HF words. The amplitude of the P300 for a LF word was expected to be lower and the peak was expected to be later than those of a HF word. However, in line with our behavioral results and data from our computational model, these frequency-induced differences are expected to at least partially cancel out during the AB interval due to the fact that the late-stage processing of a LF T2 word extends beyond the period of T1 interference, escaping the AB more often than a HF T2.

Experiment 2

Methods

Participants

Twenty-one native German speaking psychology students (aged 19-24, mean = 21.0, with normal or corrected-to-normal visual acuity) from the University of Groningen were recruited via an online sign-up program, and received course credits for participating in the experiment. Informed consent was obtained prior to the experiment. The Ethical Committee Psychology of the University of Groningen approved the experiment.

Stimuli and Apparatus

E-Prime 2.0 software was used to generate stimuli and to collect responses, running under Windows XP on a PC with a 17-inch 100-Hz CRT monitor. In total, 190 high-frequency (HF; Mannheim frequency 85 to 1,425), 380 medium-frequency (MF; Mannheim frequency 17 to 76), and 190 low-frequency (LF; Mannheim frequency 9 to 16) German words (four to eight letters in length) were pseudo-randomly picked from the German word forms CELEX corpus (Baayen et al., 1993). The first target word was always a MF word; the second target word was either a HF word or a LF word. Target-words were enclosed by 'X''s such that every stimulus had a length of twelve characters (e.g., the word BERGBAU would be presented as XXBERGBAUXXX). Distractor stimuli were pseudo-randomly generated strings of digits, also consisting of twelve characters. Each word was presented in uppercase. All stimuli were presented in black, Courier New font, size 27, on a white background at a viewing distance of ~50 cm. The monitor's resolution during the experiment was set at 1,024 x 768 pixels.

Procedure

Similar to Experiment 1, the current experiment also consisted of one practice block and three testing blocks, with a short break between each testing block. The practice block contained 20 trials and each testing block contained 120 trials.

The participants' task was to identify two words (the targets) presented amongst a rapid serial visual presentation (RSVP) stream of digit-strings (the distractors). Participants were instructed to fixate on a cross in the middle of the screen. After pressing the spacebar, the fixation cross remained on the screen for 500 ms. After the fixation cross disappeared, the stream was presented, consisting of 18 stimuli. Each stimulus in the stream was presented for a duration of 120 ms without inter stimulus interval. T1 was always presented on the fourth temporal position within the stream. T2 was presented on the second, seventh, or eighth position after T1 (i.e., lag 2, 7, or 8). Within each block, each combination of lag and T2 word frequency (HF and LF) was presented equally often. A specific word was never presented twice in the experiment.

At the end of the stream, a question appeared, prompting participants to verbally report T1 and T2 to the experimenter. The correct answers were presented to the experimenter on a second display. Using the numeric keypad on a keyboard, the experimenter typed a "0" if a response matched with T1, a "1" if it matched with T2, a "2" if no response was given, and a "3" if it matched with neither of the targets. Responses were accepted and counted correct regardless of the order in which they were reported.

EEG recording

During the experiment, the EEG signal was recorded using a 64-channel electro-cap with tin electrodes (the organization of the electrode adhered to the international 10/20 system) connected to an REFA 8-64 average reference amplifier. Impedance was reduced to less than 10kΩ for all electrodes. The data was sampled with a frequency of 2kHz and digitally reduced to 500Hz. The vertical electrooculogram (EOG) was measured from two tin electrodes placed approximately 3 cm below the left eye and 1 cm above the brow of the left eye. The horizontal EOG was recorded from tin electrodes attached approximately 2 cm to the outside corner of each eye. Two tin electrodes attached to the two mastoids served as an offline reference. Brain Vision Recorder (Brain Products GmbH, Munich, Germany) was used to control the data acquisition.

Data Analysis

Preprocessing of the EEG data was done using Brain Vision Analyzer. Accuracy scores were analyzed using binominal mixed effects models. EEG data were analyzed using permutation tests and mixed effects models. Lag and T2 word frequency were entered as fixed factors in each mixed effects model. As in Experiment 1, the natural logarithm of the Mannheim word frequencies was entered in each model as continuous predictor. Subject was entered as random factor in each model. The p-values reported for the non-binominal models of the EEG data were calculated by performing 10000 Markov Chain Monte Carlo (MCMC) samplings. The permutation tests were used to determine the time-windows to be tested in the

Table 4.3. The estimates and z-values of the mixed-effects model for T1 accuracy.

	Mixed-effects model T1			
	Estimate β	Standard Error	z-value	p-value
(Intercept)	1.501	0.252	5.968	0.001
Word frequency T2	0.009	0.039	0.226	0.821
Lag 2	0.044	0.227	0.196	0.845
Lag 7	0.117	0.230	0.510	0.610
Word frequency T2, Lag 2	-0.018	0.054	-0.335	0.738
Word frequency T2, Lag 7	-0.015	0.055	-0.277	0.782
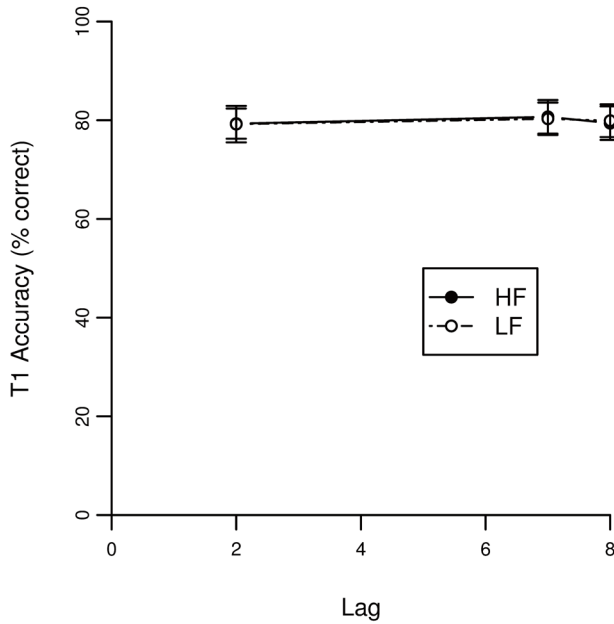
Figure 4.6. Accuracy scores of the AB task for T1 in Experiment 2. The lag corresponds to the temporal location of T2 relative to T1.

Table 4.4. The estimates and z-values of the mixed-effects model for T2|T1 accuracy.

	Mixed-effects model T2|T1			
	Estimate β	Standard Error	z-value	p-value
(Intercept)	0.397	0.254	1.563	0.118
Word frequency T2	0.226	0.043	5.209	0.001
Lag 2	-1.027	0.226	-4.543	0.001
Lag 7	-0.322	0.238	-1.353	0.176
Word frequency T2, Lag 2	-0.153	0.056	-2.722	0.007
Word frequency T2, Lag 7	0.052	0.061	0.856	0.392

mixed effects models (as an alternative for visual inspection of the EEG grand-averages). Analyses were performed using the lmer and pvals.fnc functions in the lme4 (version 0.999375-31; Bates & Sarkar, 2008) and languageR packages for the statistical software R.

Behavioral Results and Discussion

Figure 4.6 shows accuracy of T1 as a function of lag. Performance of T2 given that T1 is correctly identified is depicted in Figure 4.7.
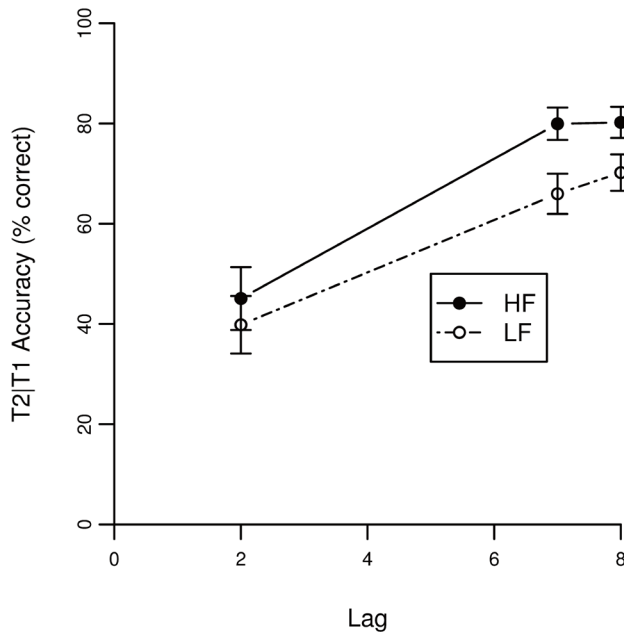
Figure 4.7. Accuracy scores of the AB task on lag 1, 2, and 7 for T2 given T1 correct in Experiment 2. The lag corresponds to the temporal location of T2 relative to T1.

T1 accuracy

A binominal mixed effects model was fitted on the accuracy of T1. Table 4.3 lists the statistics for the model's factors. None of the factors (i.e., T2 word frequency and lag) significantly predicted T1 accuracy (note that T1 word frequency was not manipulated and thus not tested). These results are in accordance with the findings of Experiment 1 presented above.

T2 accuracy

A binomial mixed effects model was fitted on T2 accuracy for trials with a correct T1 response. Table 4.4 shows the statistics for each fixed factor. Again, a main effect of T2 word frequency is found ($\beta = 0.226$). The negative estimate of lag 2 ($\beta = -1.027$) again reflects the AB, showing that at the early lag performance is lower than at later lags. Confirming the results found in Experiment 1, an interaction between T2 frequency and lag 2 was found ($\beta = -0.153$). Again, the negative estimate indicates that the AB is relatively larger when T2 is highly frequent.

Electrophysiological Results and Discussion

The EEG data was rereferenced to the mastoid electrodes. In order to remove noise, the data was filtered using a high-pass filter with a cutoff frequency of 1 Hz (24dB/oct) and a low-pass filter with a cutoff frequency of 40 Hz (24 dB/oct). As we were interested in the activity at the

parietal sites, data were then pooled over the parietal electrodes CPz, P1, P2, POz, and Pz. Next, the data were divided in T2-timelocked segments of one second (-200 ms to 800 ms). The 200ms before onset of T2 served as baseline activity. Segments containing eye-blinks were excluded from analysis. Also, if the difference in voltage between the minimum and maximum data-point in a segments exceeded 100 μVolt, the segment was excluded. The final exclusion criterion was when the difference between two successive data-points exceeded 50 μVolt. In total, three segments were excluded. The grand averages for lag 2, 7, and 8 are shown in Figure 4.8A, Figure 4.8B, Figure 4.8C, respectively.

Time-window determination

To avoid using visual inspection to determine the time-window of interest, permutation tests were used instead. For purpose of finding the window of interest, the data were binned in bins of 50ms, resulting in a total of 20 bins. In the first permutation test, all lags where averaged together and the difference between high-frequency and low-frequency words were tested for each bin. To correct for multiple comparisons, the null-distribution was constructed from the maximum and minimum t-statistics across all bins of each permutation (Nichols & Holmes, 2002). The null-distribution was constructed of 5000 randomly generated permutations. If the t-statistic was smaller or larger than the 0.025 or 0.975 quantile of the null-distribution, respectively, the bin was marked as a time-window of interest. In the first permutation test, a time-window from 200ms to 250ms (consisting of one bin, $p < 0.0446$) and a time-window from 500ms to 600ms (consisting of two bins, $p < 0.0072$ and $p < 0.0001$) were found. The latter time-window was also found when the permutation test was performed on data from lag 7 ($p = 0.0062$ and $p = 0.0068$) and lag 8 ($p < 0.0230$ and $p < 0.0020$), but not at lag 2. The time-window from 200ms to 250ms was not found for the separate lags. A mixed-effect model for the first time-window did not reveal any effects of lag or word-frequency. Also, analyses on peak latencies did not reveal any evidence for latency shifts within the windows of interest (also, inspection of Figure 4.8 shows no indication of latency shifts within the windows of interest). The results of the mixed-effects model on amplitude differences for the time-window from 500ms to 600ms are discussed below.

Table 4.5. The estimates and z-values of the mixed-effects model for P300 amplitude (only correct trials are included).

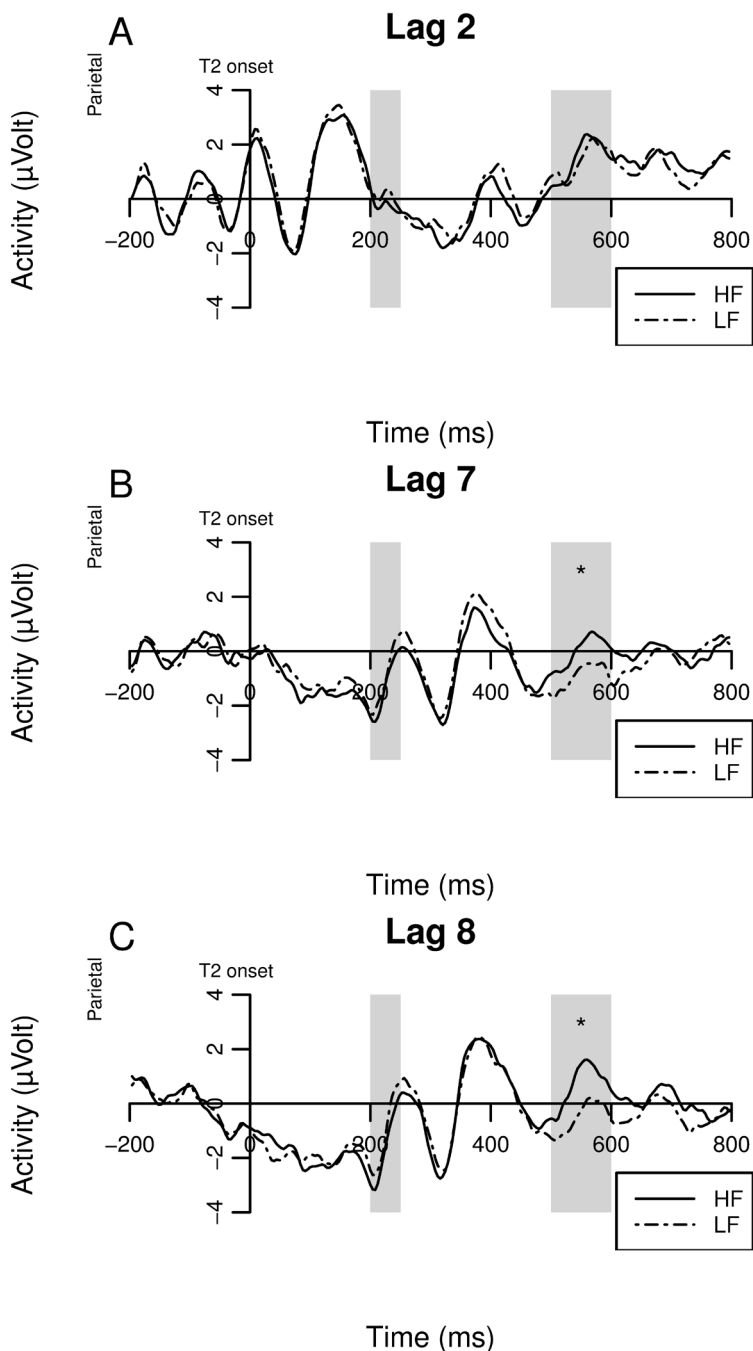|	Mixed-effects model P300 amplitude			
	Estimate β	Standard Error	z-value	p-value
(Intercept)	-1.145	0.532	-2.152	0.031
Word frequency T2	0.388	0.099	3.937	0.001
Lag 2	2.666	0.701	3.804	0.001
Lag 7	-0.231	0.603	-0.383	0.702
Word frequency T2, Lag 2	-0.398	0.165	-2.418	0.016
Word frequency T2, Lag 7	-0.078	0.141	-0.554	0.580

Figure 4.8. Grand averages for the ERPs of high-frequency and low-frequency words at lag 2 (A), 7 (B), and 8 (C). The ERPs are time-locked on target onset. Only correct trials were analyzed.

Parietal late-positivity

As mentioned above, a mixed-effects model was used to analyze the mean activity in a time-window from 500ms to 600ms at the pooled parietal electrodes. We assumed that the late parietal activity is a late P300 component, as the time-course is similar to the P300 time-course found in literature (Polich & Donchin, 1988; note that the time-course is also similar to the P600 found in morphosyntactic-violation tasks, but this component seems to be distinct from the P300, see Frisch, Kotz, von Cramon, & Friederici, 2003). The statistics of the model are shown in Table 4.5. The main effect of T2 frequency shows that parietal activity increases as word frequency increases ($\beta = 0.388$, $p < 0.001$). Also, in comparison to lag 8, increased activity was found at lag 2 ($\beta = 2.667$, $p < 0.001$). This is not surprising, as T1 related activity was likely to be present in the time-window at lag 2, but was absent in lags 7 and 8. Interestingly, an interaction effect of T2 frequency at lag 2 was found ($\beta = -0.398$, $p = 0.016$). Note that the size of the estimate is almost equal to that of the main effect of T2 frequency, but in the opposite direction (0.388 vs. -0.398), suggesting that the effect of word frequency was absent at lag 2. This indicates that the post-perceptual difference caused by word-frequency in successfully reported words disappears during the AB critical period.

Although we did not find the expected effects of latency, we did find amplitude differences in the P300. Outside the AB critical period, the low-frequency words were more likely to induce a relatively smaller P300 compared to high-frequency words. However, during the AB critical period, this difference was not observed. A likely explanation is that particularly words that were retrieved relatively quickly, mostly HF words that typically induce the largest P300 response, were more likely to be blinked. The net result is that the frequency-induced difference in P300 amplitude that was observed at late lags was absent at the early lag.

General Discussion

Previous studies have shown that manipulating T1 difficulty can modulate the AB. However, so-called 'data-driven manipulation' studies in which the physical target properties (e.g., contrast) were varied have often produced mixed results, or required T1 to remain unmasked. In addition, 'resource-driven manipulation' studies that changed the processing load rather than perceptual properties of T1 are both sparse and sometimes allow alternative explanations due to various methodological problems, including the presence of a task-switch, differing target-templates, or the use of small groups with between- rather than within-subject manipulations.

To address these issues, we manipulated the difficulty of both targets within subjects by presenting high- and low-frequency words as targets within a stream of distractor non-words. By virtue of the different frequencies that words have within a language, the difficulty of our word stimuli intrinsically varied in a more natural way than previous resource-driven difficulty manipulations. Based on findings from lexical decision studies (e.g., Van Rijn & Anderson, 2003), low-frequency words were assumed to be more difficult than high-frequency words. In addition, it was predicted that a low-frequency T1 should induce a greater AB effect.

Consistent with the study of Burt et al. (2011), we found that a low-frequency T1 produced a larger AB than a high-frequency T1. Unexpectedly though, an easy high-frequency T2 produced a relatively larger AB than a more difficult low-frequency T2, when compared to

performance at lag 7. This finding is at odds with limited-resource explanations of the AB, but can be accounted for in terms of attentional control and our computational model.

In our model, this relatively smaller AB observed for low-frequency T2 targets is attributed to the longer retrieval times of a low-frequency word from declarative memory. This is in line with findings by Polich and Donchin (1988), who showed that the P300—an electrophysiological component associated with working memory consolidation—is delayed and its amplitude is decreased when a word has a low rather than high frequency. However, if the retrieval of T2 takes long enough so that it completes after the consolidation of T1 has completed, then T2 will be consolidated, reflected in a relatively smaller AB for low-frequency T2s. The combination of these orthogonal effects (a relatively small P300 for low-frequent words versus a relatively larger P300 due to a relatively smaller AB) may have led to the absence of significant word-frequency-related differences in P300 amplitude at lag 2. The explicit distinction in our model between the unconscious recognition of a target (the retrieval from declarative memory) and the conscious recognition of the target (the consolidation process) fits well with previous findings of post-perceptual semantic processing of blinked items (e.g., Martens, Wolters, & van Raamsdonk, 2002; Shapiro, Driver, Ward, & Sorensen, 1997; Vogel et al., 1998). Based on this explicit distinction, we predict that manipulations that shorten rather than lengthen the retrieval-time of a T2—by increasing its activation in declarative memory—might cause T2's subsequent consolidation process to be blocked due to overlap with T1's consolidation process, paradoxically leading to an increase in AB magnitude. For example, if one would conceptually prime a target T2 word by showing its pictorial counterpart beforehand, we predict that AB magnitude increases (for some initial evidence that this might indeed be the case, see Koelewijn, Van der Burg, Bronkhorst, & Theeuwes, 2008).

In summary, word frequency can be used to manipulate the difficulty of targets presented in RSVP in a resource-limited manner without introducing any sort of task-switching cost (Potter et al., 1998) or perceptual degradation. A low-frequency T1 word is more difficult to process and consequently increases the AB for T2, as reflected in the present results. However, the data indicate that if consolidation is delayed by a difficult T2–in our case through a prolonged retrieval from memory–there is a higher chance that the item will be successfully consolidated and reported. Simulations show that if the processing time prior to the consolidation stage increases, a T2 is indeed less likely to be blinked. The behavioral findings, computational model, and electrophysiological results presented here strongly support an attention-control rather than limited-resource account of the AB.

Acknowledgements

Chapter 5: Pupil dilation deconvolution

reveals the dynamics of attention at high temporal resolution

In which we present a new method to analyze pupil dilation, allowing us to study the time-course of attention

Chapter was previous published in:

PNAS, 2012

Stefan M. Wierda

Hedderik van Rijn

Niels Taatgen

Sander Martens

Abstract

The size of the human pupil increases as a function of mental effort. However, this response is slow, and therefore its use is thought to be limited to measurements of slow tasks or tasks in which meaningful events are temporally well separated. Here we show that high-temporal-resolution tracking of attention and cognitive processes can be obtained from the slow pupillary response. Using automated dilation deconvolution, we isolated and tracked the dynamics of attention in a fast-paced temporal attention task, allowing us to uncover the amount of mental activity that is critical for conscious perception of relevant stimuli. We thus found evidence for specific temporal expectancy effects in attention that have eluded detection using neuroimaging methods such as EEG. Combining this approach with other neuroimaging techniques can open many research opportunities to study the temporal dynamics of the mind's inner eye in great detail.

Introduction

The size of the human eye pupil often is used as a measure of mental effort because it is assumed that the pupil size is related to the amount of cognitive control (Jepma & Nieuwenhuis, 2011), attention (Verney, Granholm, & Marshall, 2004), and cognitive processing (Moresi et al., 2008) required by a given task. However, because the pupillary response is slow—pupil size increases slowly in response to a relevant event and peaks after approximately 1 s—measuring effort by assessing pupil dilation traditionally was reserved for slow tasks or tasks in which meaningful events are well separated in time.

Here we show that high-temporal-resolution (~10 Hz) tracking of attention and cognitive processes can be obtained from the slow pupillary response (~1 Hz). Using automated dilation deconvolution, based on the quantitative analysis of the pupillary response (Hoeks & Levelt, 1993), we isolated and tracked the dynamics of attention in a fast-paced temporal attention task, allowing us to uncover the amount of mental activity that is critical for conscious perception of relevant stimuli.

We modeled the pupillary response as a function of a series of cognitive events, extending the approach of Hoeks and Levelt (1993). In their model, each cognitive event is associated with an attentional pulse, which is assumed to trigger a dilation of the pupil as a function of that attentional pulse's strength. The number of pulses, the temporal location of pulses, and the strength of each pulse that add up to a dilation of the pupil can be set at specific values or can be free to vary. Given the additive nature of the pupillary response (Hoeks & Levelt, 1993), a prediction for the pupillary response pattern evoked by a task can be derived by convolving the attentional pulses with a pupillary response function, similar to the convolution process in functional MRI (fMRI) analyses. This pupillary response function is described as an Erlang gamma function, and its constants have been determined empirically (Hoeks & Levelt, 1993). Apart from predicting a pupillary response, this method also can be used to derive a pattern of pulses that underlies an observed pupillary response by means of a deconvolution process. However, the method described by Hoeks and Levelt (1993) is limited to an isolated dilation during a relatively short interval. Our extensions allow deconvolving longer intervals in which multiple independent cognitive events might take place. An optimization algorithm takes a vector of the temporal locations of the attentional pulses and determines the strength associated with each pulse by minimizing the mismatch between the normalized observed and the predicted pupillary response. Concurrent with the estimation of the strength of the pulses, the slope accounting for linear drifts in the data is estimated. Normalizing the pupil dilation and accounting for the slow drift allows the analysis of longer intervals of pupillary data, up to complete experimental trials with multiple cognitive events.

Because the temporal resolution is determined solely by the temporal locations defined in the attentional pulses vector, this deconvolution method allows the use of pupil dilation to assess the involvement of the cognitive system at a much higher temporal resolution than the slow pupillary response seemingly would allow. To demonstrate the benefits of extracting information with high temporal resolution from pupillary response patterns, we recorded pupil size during an attentional blink (AB) task (1992). The AB is a deficit in reporting the second of two targets presented in close temporal succession (~150–500 ms) within a stream of nontargets (i.e., distractors), reflecting temporal limitations of attention. To assess the involvement of the cognitive system at all stages of the AB task, the onset of all stimuli, both distractors and targets, are represented in the attentional pulses vector.
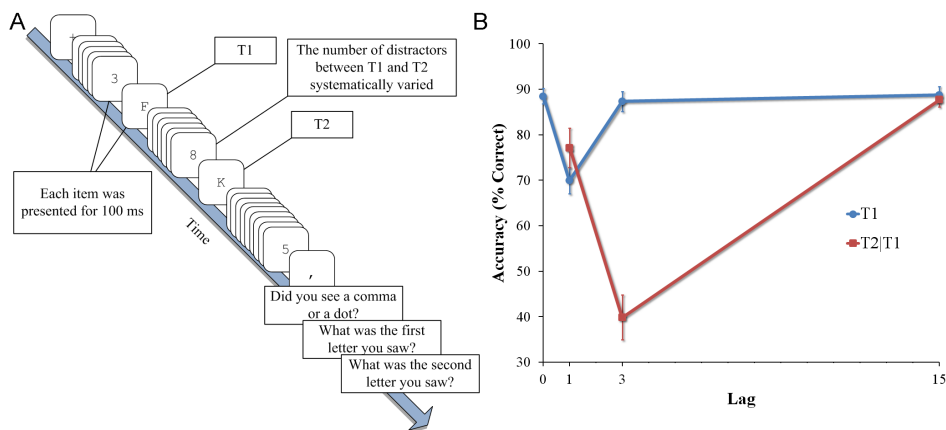
Figure 5.1. The AB task. Panel A shows an example of a dual-target trial, which required two unspecified target letters (T1 and T2) to be reported in a stream of distractors (digits). At the end of the trial, a comma or dot appeared to keep attention focused on the stream. The task was to first report whether a comma or dot had been presented. Secondly, the targets had to be reported in order of appearance. However, responses in either order were counted as correct in the analyses. Panel B depicts the accuracy scores of T1 and T2|T1 (i.e., T2 accuracy on trials where T1 is correctly reported) as a function of the temporal interval between the targets. A binominal mixed-effects model revealed that there was a substantial decrement in performance at lag 3 when compared to performance at lag 15 (p < 0.001), revealing a robust AB effect. Lag 0 refers to performance in the one-target condition.

The AB paradigm is particularly well suited to test our method for the following three reasons. First, it has been shown that a single target presented in a rapid stream of nontargets elicits a pupillary response (Privitera, Renninger, Carney, Klein, & Aguilar, 2010). Second, because of the fast-paced nature of the task, with stimuli typically presented every 100 ms, the pupillary responses to two closely succeeding targets are bound to overlap. Thus, the paradigm allowed us to test whether our method can identify the cognitive involvement of targets presented at a frequency higher than a single pupillary response. Third, the AB is one of the most intensively studied phenomena in attention research in the past two decades and has been investigated with various neuroimaging techniques, including EEG and magnetoencephalography (MEG), both of which excel in the temporal domain (Martens & Wyble, 2010). Given this extensive literature, we can compare and validate our results.

Results

Behavioral Results

An example of a trial and the behavioral results of the AB task are shown in Figure 5.1. Participants were required to detect and identify zero, one, or two unspecified target letters that were
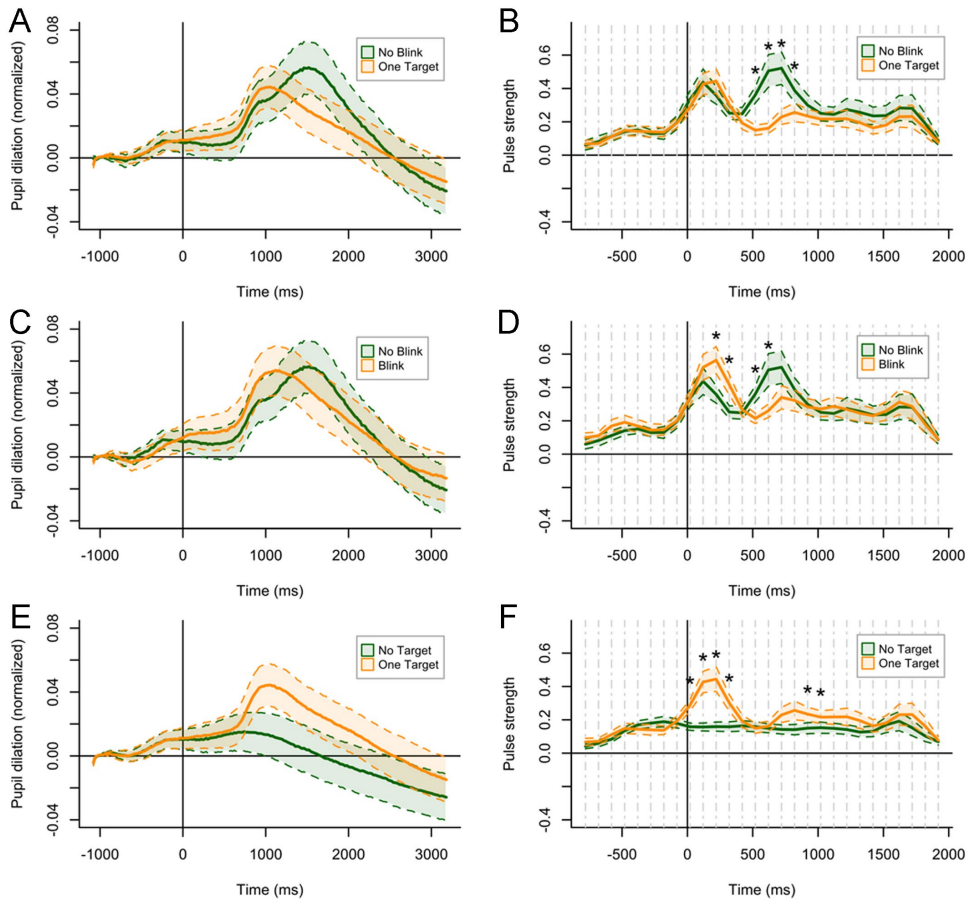
Figure 5.2. Pupil dilation averages and the underlying attentional pulses. The left panels show the normalized pupil dilation, the right panels show the strength of the deconvolved attentional pulses. The strengths of the attentional pulses for each condition of interest were compared using permutation tests. In each panel, the data depicted is time-locked on the presentation of the first target. Panel A and B show the comparison between no-blink and one-target trials, panel C and D show the comparison between no-blink and blink trials, and panel E and F show the comparison between the one-target and no-target trials. In the blink and no-blink trials, the second target's onset was always at 300 ms following T1. To correct for multiple comparisons, a single threshold test was used in each permutation test (Nichols & Holmes, 2002). The single threshold statistic t for panel B, D, and F was 2.654, 2.682, and 2.677, respectively. Significant differences (α = 0.05) in the strength of attentional pulses are denoted by an asterisk (*).

presented within a sequential stream of nontarget digits (i.e., distractors). In trials in which the stream contained two targets, the second target was presented either in the AB critical period at 300 ms after the first target or outside the critical period. Trials were presented in a random order. In the remainder of this paper, we focus on no-target trials, one-target trials, and dual-target trials in which the second target was presented 300 ms after the first target.

Isolating Temporally Proximal Stimuli

To test whether it was possible to isolate the attentional pulses corresponding to two distinct but temporally proximal stimuli, we compared the one-target trials (in which only a single meaningful stimulus was presented and correctly identified) with the no-blink trials (the condition in which both targets were correctly identified). The orange line in Figure 5.2A represents one-target trials; as expected, only a single pupillary response is visible. In contrast, the green line shows that the dilation response to the first and second target is not clearly separable in the no-blink trials, because the observed normalized dilation consists of one complex dilation pattern. However, dilation deconvolution clearly identified two distinct groups of attentional pulses, as is shown in Figure 5.2B. Similar activity has been reported in EEG studies (8, 9), but rarely are the effects in EEG studies as clearly separable as shown in Figure 5.2B.

Cognitive Involvement of the First Target

To examine the relationship between the cognitive involvement associated with the first target and the occurrence of a subsequent AB, we contrasted the no-blink trials vs. the blink trials (trials in which a second target was presented but not identified correctly). The pupil size corresponding to these two types of trials, which feature identical stimuli and instructions but have different behavioral outcomes, is shown in Figure 5.2C. As expected, the strengths of the attentional pulses associated with the second target were lower during blink trials (Figure 5.2D). This finding is in line with EEG studies in which target-specific activity time-locked to the second target typically is absent (Craston, Wyble, Chennu, & Bowman, 2009; Kranczioch et al., 2003; Martens et al., 2006; Vogel et al., 1998). Interestingly, the estimated strength for the first target was higher for blink trials than for no-blink trials. Thus, the current study replicates the finding that processing demands of the first target are critical for the occurrence of the AB (Akyürek, Leszczyński, & Schubö, 2010; Martens et al., 2006; Ouimet & Jolicoeur, 2007; Wierda et al., 2010). However, whereas previous studies involved experimental manipulations such as changing the difficulty of the first target, this study demonstrates this phenomenon without any experimental manipulation, mirroring the results of an MEG study (Shapiro et al., 2006). This result suggests that our method is highly sensitive and can account for subtle fluctuations of attention that have remained elusive in fMRI or EEG measurements.

Expectancy Effects in Attention

Because of the structure of the task, participants might have come to expect a second target if

a first target was perceived. To test whether we could find any expectancy effects in temporal attention, we compared one-target trials with no-target trials. Figure 5.2E shows the normalized pupil dilation for both types of trials and shows a fairly stable pupil dilation over the whole trial for the no-target trials, whereas the presentation of a target in the one-target trials results in an increased dilation of the pupil that peaks about 1,000 ms after presentation. Figure 5.2F depicts the attentional pulses derived from the dilation deconvolution method. Up to 500 ms, the plot closely resembles what might be extracted from the normalized pupil dilation plot, because the increased strength for the initial pulses corresponds with the processing of the first stimulus. Although the normalized pupil dilation in Figure 5.2E does not show any salient effects after the initial dilation peak, the strengths of the attentional pulses show a second increase for the one-target trials observed at about 1,000 ms. Because this time frame is associated with the effect of a second target in dual-target trials, this increase in strength for the attentional pulses most likely reflects an expectancy effect. Indeed, more general effects of expectancy have been found in attentional processing (Correa, Lupiáñez, Madrid, & Tudela, 2006). However, the expectancy effect in the AB reported here might easily have eluded discovery because of its fairly small effect size in comparison with the effects observed with the normal processing of target stimuli. For example, the much stronger effects of the processing of the first target might obscure the expectancy effect in EEG studies, as it did in the raw pupillary patterns shown in Figure 5.2E.

Discussion

As shown empirically in this study, dilation deconvolution can provide valuable information regarding the occurrence and timing of attentional processes that underlie human cognition. The results of the experiment show that dilation deconvolution can track and isolate attentional processing of multiple events at close temporal proximity, thus revealing the temporal dynamics of the mind's eye at a surprisingly high resolution.

Although the dilation deconvolution presented here is based on the pupillary response function of Hoeks and Levelt (1993), other models of the pupillary response could be implemented as well. For example, the more complex bimodal pupillary response function as described by O'Neill and Zimmerman (2000) could be implemented in our pupil deconvolution method. Although in theory distinguishing between sympathetic and parasympathetic influences on pupil dilation might provide additional information, implementing this response function would have doubled the number of parameters (or even tripled them, depending on whether the distance between sympathetic and parasympathetic pulses is fixed). We therefore opted for the more parsimonious pupillary response function of Hoeks and Levelt.

The reliability of the present pupil deconvolution method is strengthened by the consistency of our results with those in previous neuroimaging studies. Our results also are consistent with AB theories (Dux & Marois, 2009; Martens & Wyble, 2010). In addition, evidence of subtle expectancy effects in temporal attention was obtained, as was direct evidence for the crucial role of the processing demands of the first target. This pattern of results supports the notion that both pupil dilation and the AB might be linked closely to the activity of the no-radrenergic system (Aston-Jones & Cohen, 2005; Nieuwenhuis, Gilzenrat, Holmes, & Cohen, 2005; Nieuwenhuis, De Geus, & Aston-Jones, 2011). Finally, it is worth noting that a practical advantage of using our pupil dilation approach is that, in comparison with other neuroima-

ging techniques, it is a relatively cost-effective and accessible method to measure cognitive workload and attention.

In conclusion, our method allows the analysis of the processing of stimuli that are presented in close temporal succession, both within a trial (as demonstrated here) and between trials of tasks in which an induced pupil dilation on a given trial overlaps with the pupil dilation of a subsequent trial (Gilzenrat, Nieuwenhuis, Jepma, & Cohen, 2010). Combining this method with other behavioral and neuroimaging methods can open a wide range of interesting research opportunities.

Methods

The study was approved by the Psychology Ethical Committee of the University of Groningen. In total, 20 students of the University of Groningen participated in the study. The students received course credits in return for their participation. Informed consent was obtained from each participant. Two participants were excluded from the analyses because of artifacts and poor performance. After exclusion, nine male and nine female participants with normal vision, age 18–28 y, remained.

AB Task

Participants performed a standard AB task, requiring unspeeded identification of zero to two target letters (uppercase consonants excluding Q, V, and Y) presented in a sequential stream of nontarget digits (excluding 0 and 1). Each item was presented in black on a white background in a bold 18-point Courier New font in the middle of a 17-inch computer screen set at a resolution of 1,024 × 786 pixels with a refresh rate of 100 Hz. Viewing distance was 50 cm. The stream consisted of 50 items presented at a rate of 100 ms per item. The experiment contained one practice block of 32 trials and two experimental blocks of 160 trials each, with a short break between the two blocks. Each trial contained zero, one, or two targets. The first target (T1) always was the tenth item presented in the stream. The second target (T2) was the 1st, 3rd, 8th, 15th, or 30th item after T1 (i.e., lag 1, lag 3, lag 8, lag 15, or lag 30, respectively). Each block contained 30 trials for lag 1, lag 3, and lag 15, because these were the lags of interest, and five trials for each of the two filler lags (i.e., lag 8 and 30). Within a block, these dual-target trials were intermixed randomly with 30 no-target and 30 one-target trials. Each trial started with the presentation of a fixation-cross in the middle of the screen. Participants were instructed to press the spacebar to initiate the trial. The fixation-cross disappeared 900 ms after the space bar was pressed, and the first item appeared on the screen. After all targets and distractors in a trial had been presented, a comma or a dot was shown on the screen for 100 ms.

Before reporting the targets, participants had to indicate whether the last character was a comma or a dot. This comma/dot task was included to encourage the participants to remain fixated to the center of the screen throughout stimulus presentation. Following the comma/dot task, participants were prompted to report which letters were presented by pressing the corresponding letters on the keyboard in the order in which the letters had appeared or to press the spacebar if no target was observed. The accuracy data were analyzed using

binominal mixed-effects models available in the lme4 package of the statistical software package R (www.r-project.org). In the analysis of T2 accuracy, only trials in which T1 was reported correctly were analyzed.

Preprocessing of Pupil Size Data

Pupil size was measured using the EyeLink 1000 eye-tracker (www.sr-research.com). Data were sampled at 250 Hz and down-sampled to 50 Hz. Because we were interested primarily in the pupil dilation in response to the first target and the second target presented at lag 3 and wanted to avoid artifacts caused by the eye blinks that frequently occur during the last phase of the trial, the data were segmented in segments of 4,300 ms and time locked to 200 ms before stream onset. Segments containing eye blinks were excluded from the analysis. The average pupil size of the interval from −200 ms to 0 ms to stream onset was used as a baseline. Data were normalized by calculating the percentage increase of the pupil size compared with the baseline for each data point

$$x_{norm} = \frac{x_{data} - baseline}{baseline}.$$

Pupillary Response Model

The deconvolution of the dilation patterns is driven by an optimization algorithm. By calculating the misfit between an observed pupil dilation pattern and a pupil dilation pattern derived from a sequence of attentional pulses that is slightly changed per iteration, the optimization algorithm effectively performs the deconvolution by repeated convolutions. In total, 34 attentional pulses were modeled corresponding to the first 34 items in the AB stream. We limited our analyses to these pulses, because the development of the pupillary response from pulse onset to maximum peak dilation falls within the 4,300-ms segment for these pulses. The remaining 16 items were not modeled, because a large part of the pupillary response to these items exceeds the segment and would thus not be estimated properly by the model.

The predicted pupil dilation pattern was calculated by convolving the attentional input

$$i = \{ w_1, w_2, w_3, ..., w_{32}, w_{33}, w_{34} \}$$

with the Erlang gamma function

$$h = s \cdot (t^n) \cdot e^{(\frac{-n \cdot t}{t_{max}})};$$

where w_l is the strength of the attentional pulses, l is the position of each pulse in vector i, n is the number of layers set at 10.1, $t_{max} = 930$ is the position of the maximum response, and s = 1/1027 is a constant to scale the pupillary response function. Hoeks and Levelt (1993) empirically determined the constants n and tmax. The temporal location of each pulse was fixed

at the onset of each stimulus (i.e., 100 ms between each pulse). The predicted pupil dilation then was obtained by

$$x = l \cdot b + (i * h);$$

where b is the slope parameter that accounts for drifts in the data. The strengths wl were fit by minimizing the mean square error
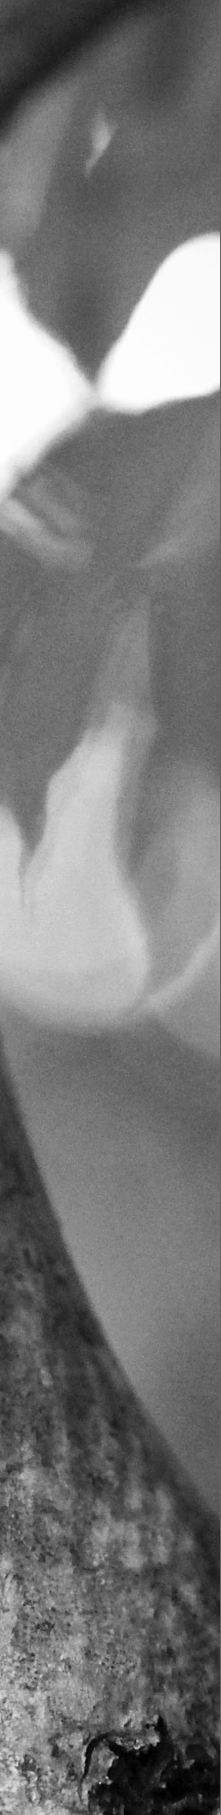
$$f = \sum |x_p - x_m|^2$$

using a linearly constrained optimization algorithm, where x_p is the predicted pupil dilation and x_m is the measured data. The initial values of the strengths for the optimization algorithm were chosen randomly from a range of 0.2 to 0.5 for each pulse. Strengths were constrained so that values below zero were not allowed. Because this optimization method could result in local minima, this process was repeated 200 times per combination of participant and condition to obtain stable strength averages of the attentional pulses. An implementation of the pupil dilation deconvolution method as described above can be found in appendix of this thesis (the preprocessed data can be found at the website of PNAS in the SI Text).

Analysis of Attentional Pulses

The attentional pulses generated by the model were tested using nonparametric permutation tests that are used commonly in fMRI studies. Permutation tests have the advantage over parametric tests of being more conservative and relying on fewer assumptions (Holmes, Blair, Watson, & Ford, 1996; Nichols & Holmes, 2002). In the analyses, the first three and last three pulses were excluded. In each test, two conditions were compared. The permutation distribution was obtained by randomly labeling the conditions for each pulse per participant 10,000 times. The null hypothesis stated that there would be no difference between condition A and condition B; thus the difference score (i.e., A − B) would equal zero for a pulse if there were no effect. The statistic used in the test was the standard t-statistic. To account for multiple comparisons within the permutation test, a single threshold test was performed per permutation test. Thus, each pulse with a t-value exceeding the single threshold would reject the null hypothesis. In total, three permutation tests were done. In the first test, the one-target trials were compared with the trials in which T2 was presented at lag 3 and correctly identified (i.e., no-blink trials). In the second test, lag-3 trials with correct T2 responses were compared with lag-3 trials with incorrect T2 responses (i.e., blink trials). Finally, no-target trials were compared with one-target trials.

In which I summarize the results of this thesis and discuss the implications following these results

In this chapter, I will briefly summarize the empirical chapters, and I will address the issues raised in these chapters in a broader context. This chapter ends with concluding remarks and implications of the studies reported in this thesis. Detailed discussion of the results of each specific chapter can be found at the end of each chapter.

EEG shows that less effort increases performance

In Chapter 2, we investigated the paradoxical effect of improvement when individuals are distracted during an AB task. Participants performed a red dot detection task concurrent with the AB task—a dot circled the stream of letters and digits of the AB task, and one had to report whether or not the dot had turned red for a brief period. During this task, brain activity was measured by means of an electrophysiological encephalogram (EEG). Although we were mainly interested in the effects of distraction on the EEG, we also wanted to compare the results to a previous EEG study that compared the brain-activity of blinkers to that of non-blinkers (Martens et al., 2006).

More specifically, we focused on a component in the EEG signal that is associated with working memory consolidation (the P300 that peaks at the parietal electrode sites). We found that this target-specific brain activity was reduced for target items in the red dot detection—concurrent to the AB task—condition compared to the condition in which only a standard AB task was performed (i.e., without red dot detection). In addition, occipital electrode sites showed less distractor-related brain-activity. The fact that both the P300 and the distractor-related occipital activity were reduced during the red-dot task suggests that when people invest less attention in the AB task, performance improves, possibly because task-irrelevant stimuli (the distractors) and T1 (that directly competes with T2) receive less attention. In contrast to the study that compared blinkers to non-blinkers (Martens et al., 2006), no latency shifts or change in frontal activity (reflected in the FSP) were found. Although individuals seemed to invest less effort in the primary AB task when distracted, resulting in a smaller AB, most individuals still performed worse than the non-blinkers. Thus, differences in both brain activity and behavior continued to exist between distracted blinkers and natural non-blinkers. Perhaps non-blinkers perform the task in a fundamentally different way than blinkers do, by using for example categorical information to select targets from non-targets at an early processing stage, as suggested by an earlier study (Martens et al., 2010).

Differences in selection errors between blinkers and non-blinkers

In Chapter 3, the differences between blinkers and non-blinkers were further explored. By studying the temporal dynamics of the AB task in terms of suppression (i.e., the amount of erroneous reports of distractors in a window around the actual target), delay (i.e., the point in time at which most items are reported), and diffusion (i.e., the variance in delay), a temporal profile was made for both blinkers and non-blinkers. In a first experiment, blinkers and non-blinkers were selected based on their performance on an alphanumeric AB task. In the second experiment, blinkers and non-blinkers performed an AB task requiring the identification of red letters (i.e., the targets) amongst a stream of black letters (i.e., the distractors). Thus, targets were defined by color, rather than alphanumeric category. In this task, non-blin-

kers did show an AB, but they still outperformed blinkers. The fact that non-blinkers did blink on this task allowed us to investigate their temporal profile. Performance was compared to that of a new group of individuals that also performed the same AB task (with red target letters amongst black distractor letters). Again, a temporal profile was constructed for each individual.

Only a modest amount of suppression was found amongst blinkers, which is surprisingly different from what was found in previous research. The differences found between this study and previous research might be due to subtle differences in the AB task (e.g., distractors, target feautures, and instructions), individual differences, or ceiling effects. However, most theories explain the AB by some sort of mechanism that suppresses the incoming stimuli. If that would indeed be the case, much stronger suppression effects should have been observed, which we clearly did not. In fact, our non-blinkers did not show any sign of suppression. In addition, they were more precise than blinkers in selecting the second target, and they made less order reversals compared to the blinkers. However, during lag-1 sparing, non-blinkers did make relatively more order reversals and they showed a delay in selecting the second target. The results from this chapter show that individuals with a small AB or no AB at all are faster—except for lag 1 as mentioned above—and more precise in the target selection process during the period of the AB than large blinkers are. Again, the findings point toward a difference in the strategy blinkers and on-blinkers use when performing an AB task.

Difficult words can outlast the AB

Chapter 4 explored the underlying mechanism of the AB by manipulating the target-difficulty in a purely resource-driven manner. In an AB with words as targets and non-words as distractors, difficulty of targets was manipulated by using words as targets that varied in word-frequency. Words that occur often in our natural language (high-frequency words) are faster and easier to retrieve from memory than words that are rare (low-frequency words). Thus, if the AB is resource-driven, this resource-driven manipulation should yield straightforward results; the more resources a target takes, the larger the AB. However, if there is more to the AB than a resource-limited bottleneck, the effect of such a manipulation may be quite different, as was indeed the case as described in this chapter.

Whenever an easy, high-frequency word was presented as T1, performance on the AB task increased compared to when a low-frequency words was presented as a T1—as predicted by resource-driven accounts of the AB. In contrast to and unpredicted by the latter type of theories, when an easy, high-frequency word was presented as T2, relative performance went down (when corrected for the baseline difference between easy, high-frequency words and hard, low-frequency words). This paradoxical phenomenon was explained in terms of the threaded cognition AB model by assuming that retrieval of a word from memory always takes place, even during the AB period. This assumption is supported by the evidence presented in Chapter 3, showing that there is almost no effect of suppression during the AB period. Then, if retrieval takes long enough, which is the case for low-frequency words that need a relatively long time to be retrieved, the retrieval process carries the word out of this critical period, making it available for report.

A follow-up EEG experiment described in this chapter verified the results of the first experiment, and showed that late parietal activity (i.e., the P300 that is not only associated

with working memory consolidation, but also is sensitive to word-frequency effects) discriminates between high-frequency words and low-frequency words outside the period of the AB, but not during the AB period. This lack of a frequency-induced difference is again explained in terms of the model. Words with a high frequency have a high P300 and have a short retrieval time. Because words with a short retrieval time are more likely to decay during the AB when presented as T2, fewer words with a high P300 survive the AB. Thus, by blinking the words with a high P300, the difference between high-frequency words and low-frequency words gets smaller (or, in this case, seems to disappear). The results challenge resource accounts of the AB in which interference and depletion of resources play a major role, whereas an attentional control strategy account of the AB (such as threaded cognition AB model) can explain the results.

Pupil dilation deconvolution reveals the dynamics of temporal attention

In Chapter 5, a novel method called pupil dilation deconvolution was presented, which can be used to investigate high-paced tasks such as the AB task, by deconvolving the measured pupil dilation with the slow pupillary response. The pupillary response is the manner in which the pupil dilates in response to one or more cognitive events that takes place in one's mind in response to external stimuli. By assuming that every stimulus elicits such a response, and that this response follows a characteristic pattern, one can trace back the individual responses belonging to each stimulus, which in turn says something about the amount of workload and attention devoted to the processing of that stimulus. This process is called deconvolution—we estimate the pupillary responses to each stimulus that together add up to the measured pupil dilation. These estimated weights can then be interpreted as a measure of the strength of the cognitive response to the corresponding stimulus.

The method that we presented successfully reveals differences in workload and attention between blink and no-blink trials. Furthermore, subtle expectancy effects were found in one-target trials at the temporal position where a second target was to be expected, which have never been reported before. Despite the relatively long response time of the pupil—approximately one second—the method of pupil dilation deconvolution enhances the temporal resolution to the extent that it allows the analysis of the processing of multiple stimuli, even when presented at a rate of 10 items per second. Combining this method with other behavioral and neuroimaging methods can open a wide range of interesting research opportunities. Fast-paced tasks such as the AB task might be further investigated using this method, but one could also think of more practical or clinical applications. For example, Stoll, Chatelle, Carter, Koch, Laureys, and Einhäuser (2013) found that communcation is possible through the pupillary response in locked-in syndrome patients. The proposed method, together with neuroimaging techniques, could perhaps improve communcation with these patients. Furthermore, the method could be used in human factors research, in order to study interfaces and the amount of cognitive effort and attention they require.

Concluding remarks

Throughout the chapters of this thesis, seemingly paradoxical effects were reported. They are hard to reconile when the AB is assumed to be due to a depletion of a limited pool of central resources. In contrast, when the AB is considered to be the result of the use of a particular strategy to control attention, the results can be more readily explained, as described in this thesis. Of course, structural limitations of our cognitive system should still be taken into account when devising a theory of the AB. Thus, although the AB does not seem to arise from a cognitive bottleneck itself, the strategy that is deployed to control attention, together with structural limitations of our cognitive system, are likely to add up to the AB.

An effective way to consider structural limitations is by framing a theory that explains the AB within an integrated framework of cognition (Taatgen et al., 2009), such as the ACT-R cognitive architecture (Anderson, 2007). By doing so, the theory and model of the AB inherit the structural limitations of the architecture, which are constructed by decades of research in other domains of cognition. We do not only constrain the number of possible models this way, but the use of a general framework to explain the AB also gives us the opportunity to model and study temporal attention in relation to other tasks and domains. Specific examples are given in Chapter 2, in which we study the relation between the AB task and a secondary red dot, and in Chapter 4 where the role of declarative memory in the AB is shown.

In addition, by conducting behavioral, neuroimaging, and psychophysiological experiments the number of possible strategies can be reduced. For example, the studies presented in this thesis show that both distractors and targets are processed up to a late stage. Clear evidence for late stage processing is presented in Chapter 4, which shows that even blinked targets are retrieved from declarative memory. This is supported by the results presented in Chapter 3, that show only a modest amount of suppression during the AB period and also that distractors are processed up to a stage in which they are confused with targets. Finally, the expectancy effect found in Chapter 5 shows that even though a target is blinked, a certain amount of attention is still deployed during the AB.

An important lesson learned from the AB for cognitive research in general is that when studying cognitive phenomena and aspects of cognition, one should always keep in mind individual differences. As with the AB phenomenon, it could be that the brain is capable of doing much more, but that it restricts itself for some reason. The task strategy used by individuals has a huge impact on the outcome of an experiment, and these strategies might well be influenced through training or instructions. Bad performance on a task does not necessarily imply that the human brain is limited in doing that task, as is illustrated by the AB phenomenon. Thus, one should be cautious to generalize results from specific tasks to general theories or real life situations.

This thesis leaves several important questions open. To name a few, what exactly separates the blinkers from the non-blinkers? Why do non-blinkers show an AB when the paradigm is only slightly changed? And perhaps the most intriguing question: Why are individuals so strongly inclined to deploy a seemingly suboptimal strategy that causes attention to blink? If indeed a certain control strategy is responsible for the AB, it makes little sense that most of us deploy a counterproductive strategy. As demonstrated by non-blinkers, we are in principle capable of deploying strategies that are more effective. In addition, following an hour of training with a salient T2, the AB can be "unlearned" (Choi et al., 2012).
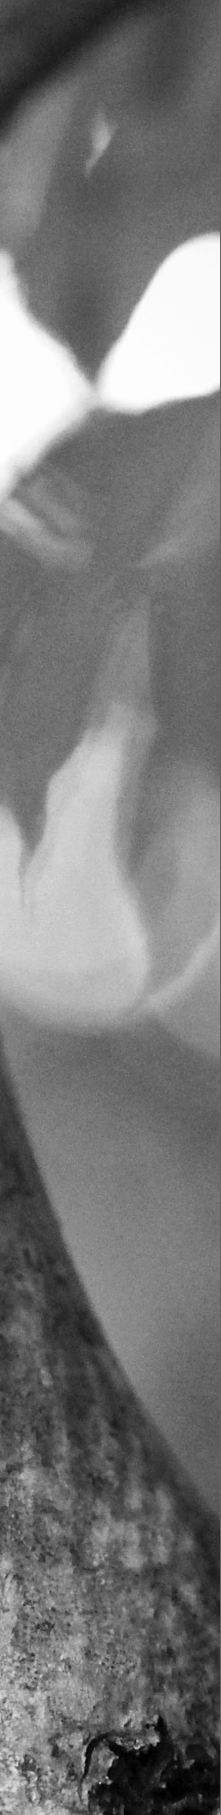
Perhaps the manner in which we select and consolidate information is effective and beneficial in the majority of every-day situations, but not so much in RSVP tasks—tasks that are rarely encountered outside the context of the psychophysiological lab (although there are plenty of fast-paced tasks in real life, such as driving a car, watching television, playing action video games, or piloting an airplane). If there is a functional benefit to the AB, then could there be a cost for non-blinkers on some other task? That is, the particular strategy that they employ might be optimal for an AB task, but might be quite problematic when used in a different task or context. To support this point, one would need to devise an experiment in which being a non-blinker will actually lead to deteriorated performance, and in which being a blinker is beneficial to performance. To my knowledge, such a task is still to be found and the functional purpose of the AB—if it exists—remains a mystery to be unraveled.

To conclude, the neuroimaging data, behavioral data, and computational model presented here in this thesis all point towards the AB as a strategic bottleneck. Framing the AB in a broader, more general theory of cognition and attention will help to figure out why most people employ this strategic bottleneck. Corny but true: although a bit closer, we are still a long way from fully understanding how the brain goes about in parsing the world around us.

The references to literature used throughout the thesis can be found here

Akyürek, E. G., Leszczyński, M., & Schubö, A. (2010). The temporal locus of the interaction between working memory consolidation and the attentional blink. Psychophysiology, 47(6), 1134-1141. doi:10.1111/j.1469-8986.2010.01033.x

Anderson, J. R. (2007). How can the human mind occur in the physical universe? United States: Oxford University Press.

Arend, I., Johnston, S., & Shapiro, K. (2006). Task-irrelevant visual motion and flicker attenuate the attentional blink. Psychonomic Bulletin & Review, 13(4), 600-607.

Arnell, K. M., Howe, A. E., Joanisse, M. F., & Klein, R. M. (2006). Relationships between attentional blink magnitude, RSVP target accuracy, and performance on other cognitive tasks. Memory & Cognition, 34(7), 1472-1483.

Arnell, K. M., Stokes, K. A., Maclean, M. H., & Gicante, C. (2010). Executive control processes of working memory predict attentional blink magnitude over and above storage capacity. Psychological Research, 74(1), 1-11. doi:10.1007/s00426-008-0200-4

Arnell, K. M., & Stubitz, S. M. (2010). Attentional blink magnitude is predicted by the ability to keep irrelevant material out of working memory. Psychological Research, 74(5), 457-467. doi:10.1007/s00426-009-0265-8

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(1), 403-450. doi:10.1146/annurev.neuro.28.061604.135709

Baayen, R. H., Davidson, D., & Bates, D. (2008). Mixed-effects modeling with crossed random effects for subjects and items. Journal of Memory and Language, 59(4), 390-412. doi:10.1016/j.jml.2007.12.005

Baayen, R. H., Piepenbroek, R., & Van Rijn, H. (1993). The CELEX lexical database [CD-ROM] Philadelphia: Linguistic Data Consortium, University of Pennsylvania, PA.

Barr, D. (2008). Analyzing 'visual world' eyetracking data using multilevel logistic regression. Journal of Memory and Language, 59(4), 457-474. doi:10.1016/j.jml.2007.09.002

Bates, D., & Sarkar, D. (2008). Lme4: Linear mixed-effects models using S4 classes (R package version 0.999375-28 ed.)

Botella, J., Privado, J., de Liano, B. G., & Suero, M. (2011). Illusory conjunctions reflect the time course of the attentional blink. Attention, Perception & Psychophysics, 73(5), 1361-1373. doi:10.3758/s13414-011-0112-z; 10.3758/s13414-011-0112-z

Bowman, H., & Wyble, B. (2007). The simultaneous type, serial token model of temporal attention and working memory. Psychological Review, 114(1), 38-70.

Broadbent, D. E., & Broadbent, M. H. P. (1987). From detection to identification: Response to multiple targets in rapid serial visual presentation. Perception & Psychophysics, 42(2), 105-113.

Burt, J., Howard, S., & Falconer, E. (2011). T1 difficulty affects the AB: Manipulating T1 word frequency and T1 orthographic neighbor frequency. Attention, Perception, & Psychophysics, 73(3), 751-765. doi:10.3758/s13414-010-0054-x

Choi, H., Chang, L. H., Shibata, K., Sasaki, Y., & Watanabe, T. (2012). Resetting capacity limitations revealed by long-lasting elimination of attentional blink through training. Proceedings of the National Academy of Sciences of the United States of America, 109(30), 12242-12247. doi:10.1073/pnas.1203972109; 10.1073/pnas.1203972109

Christmann, C., & Leuthold, H. (2004). The attentional blink is susceptible to concurrent perceptual processing demands. The Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology, 57(2), 357-381.

Chua, F. K. (2005). The effect of target contrast on the attentional blink. Perception & Psychophysics, 67(5), 770-788.

Chun, M. M. (1997). Temporal binding errors are redistributed by the attentional blink. Perception & Psychophysics, 59(8), 1191-1199.

Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. Journal of Experimental Psychology: Human Perception and Performance, 21(1), 109-127.

Colzato, L. S., Hommel, B., & Shapiro, K. L. (2010). Religion and the attentional blink: Depth of faith predicts depth of the blink. Frontiers in Psychology, 1, 147. doi:10.3389/fpsyg.2010.00147; 10.3389/fpsyg.2010.00147

Colzato, L. S., Bajo, M. T., van den Wildenberg, W., Paolieri, D., Nieuwenhuis, S., La Heij, W., . . . Trento, I. (2008). How does bilingualism improve executive control? A comparison of active and reactive inhibition mechanisms. Journal of Experimental Psychology: Learning, Memory, and Cognition, 34(2), 302.

Colzato, L. S., Spapè, M., Pannebakker, M. M., & Hommel, B. (2007). Working memory and the attentional blink: Blink size is predicted by individual differences in operation span. Psychonomic Bulletin & Review, 14(6), 1051-1057.

Correa, Á, 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(1), 116-128. doi:10.1016/j.brainres.2005.11.074

Craston, P., Wyble, B., Chennu, S., & Bowman, H. (2009). The attentional blink reveals serial working memory encoding: Evidence from virtual and human event-related potentials. Journal of Cognitive Neuroscience, 21(3), 550-566.

Dale, G., & Arnell, K. M. (2010). Individual differences in dispositional focus of attention predict attentional blink magnitude. Attention, Perception, & Psychophysics, 72(3), 602-606.

Dehaene, S., Sergent, C., & Changeux, J. P. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. Proceedings of the National Academy of Sciences, 100(14), 8520-8525.

Dell'Acqua, R., Dux, P. E., Wyble, B., & Jolicœur, P. (2012). Sparing from the attentional blink is not spared from structural limitations. Psychonomic Bulletin & Review, 19(2), 232-238. doi:10.3758/s13423-011-0209-3

Dell'acqua, R., Jolicoeur, P., Luria, R., & Pluchino, P. (2009). Reevaluating encoding-capacity limitations as a cause of the attentional blink. Journal of Experimental Psychology.Human Perception and Performance, 35(2), 338-351. doi:10.1037/a0013555; 10.1037/a0013555

Dell'Acqua, R., Sessa, P., Jolicoeur, P., & Robitaille, N. (2006). Spatial attention freezes during the attention blink. Psychophysiology, 43(4), 394-400.

Di Lollo, V., Kawahara, J., Shahab Ghorashi, S. M., & Enns, J. T. (2005). The attentional blink: Resource depletion or temporary loss of control? Psychological Research, 69(3), 191-200.

Dux, P. E., Asplund, C. L., & Marois, R. (2008). An attentional blink for sequentially presented targets: Evidence in favor of resource depletion accounts. Psychonomic Bulletin & Review, 15(4), 809-813.

Dux, P. E., Asplund, C. L., & Marois, R. (2009). Both exogenous and endogenous target salience manipulations support resource depletion accounts of the attentional blink: A reply to olivers, spalek, kawahara, and di lollo (2009). Psychonomic Bulletin & Review, 16(1), 219-224.

Dux, P. E., Coltheart, V., & Harris, I. M. (2006). On the fate of distractor stimuli in rapid serial visual presentation. Cognition, 99(3), 355-382. doi:10.1016/j.cognition.2005.04.002

Dux, P. E., & Harris, I. M. (2007). On the failure of distractor inhibition in the attentional blink. Psychonomic Bulletin & Review, 14(4), 723.

Dux, P. E., & Marois, R. (2008). Distractor inhibition predicts individual differences in the attentional blink. PLoS One, 3(10), e3330. doi:doi: 10.1371/journal.pone.0003330

Dux, P. E., & Marois, R. (2009). The attentional blink: A review of data and theory. Attention, Perception & Psychophysics, 71(8), 1683.

Feinstein, J. S., Stein, M. B., Castillo, G. N., & Paulus, M. P. (2004). From sensory processes to conscious perception. Consciousness and Cognition, 13(2), 323-335.

Ferlazzo, F., Lucido, S., Di Nocera, F., Fagioli, S., & Sdoia, S. (2007). Switching between goals mediates the attentional blink effect. Experimental Psychology, 54(2), 89-98.

Frisch, S., Kotz, S. A., von Cramon, D. Y., & Friederici, A. D. (2003). Why the P600 is not just a P300: The role of the basal ganglia. Clinical Neurophysiology : Official Journal of the International Federation of Clinical Neurophysiology, 114(2), 336-340.

Gilzenrat, M., Nieuwenhuis, S., Jepma, M., & Cohen, J. (2010). Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function. Cognitive, Affective, & Behavioral Neuroscience, 10(2),

252-269. doi:10.3758/CABN.10.2.252

Green, C. S., & Bavelier, D. (2003). Action video game modifies visual selective attention. Nature, 423(6939), 534-537.

Handy, T. C., Soltani, M., & Mangun, G. R. (2001). Perceptual load and visuocortical processing: Event-related potentials reveal sensory-level selection. Psychological Science, 12(3), 213-218.

Harris, I. M., Benito, C. T., & Dux, P. E. (2010). Priming from distractors in rapid serial visual presentation is modulated by image properties and attention. Journal of Experimental Psychology.Human Perception and Performance, 36(6), 1595-1608. doi:10.1037/a0019218; 10.1037/a0019218

Hoeks, B., & Levelt, W. (1993). Pupillary dilation as a measure of attention: A quantitative system analysis. Behavior Research Methods, 25(1), 16-26. doi:10.3758/BF03204445

Holmes, A. P., Blair, R. C., Watson, G., & Ford, I. (1996). Nonparametric analysis of statistic images from functional mapping experiments. Journal of Cerebral Blood Flow and Metabolism: Official Journal of the International Society of Cerebral Blood Flow and Metabolism, 16(1), 7-22.

Hommel, B., Kessler, K., Schmitz, F., Gross, J., Akyürek, E., Shapiro, K., & Schnitzler, A. (2006). How the brain blinks: Towards a neurocognitive model of the attentional blink. Psychological Research, 70(6), 425-435.

Isreal, J. B., Chesney, G. L., Wickens, C. D., & Donchin, E. (1980). P300 and tracking difficulty: Evidence for multiple resources in dual-task performance. Psychophysiology, 17, 259-273.

Isreal, J. B., Wickens, C. D., Chesney, G. L., & Donchin, E. (1980). The event-related brain potential as an index of display-monitoring workload. Human Factors, 22(2), 211-224.

Jepma, M., & Nieuwenhuis, S. (2011). Pupil diameter predicts changes in the Exploration–Exploitation trade-off: Evidence for the adaptive gain theory. Journal of Cognitive Neuroscience, 23(7), 1587-1596. doi:10.1162/jocn.2010.21548

Jolicœur, P., Sessa, P., Dell'Acqua, R., & Robitaille, N. (2006). On the control of visual spatial attention: Evidence from human electrophysiology. Psychological Research, 70(6), 414-424.

Jolicoeur, P., Sessa, P., Dell'Acqua, R., & Robitaille, N. (2006). Attentional control and capture in the attentional blink paradigm: Evidence from human electrophysiology. European Journal of Cognitive Psychology, 18(4), 560-578.

Jolicœur, P., & Dell'Acqua, R. (1998). The demonstration of short-term consolidation. Cognitive Psychology, 36(2), 138-202. doi:10.1006/cogp.1998.0684

Kihara, K., Yagi, Y., Takeda, Y., & Kawahara, J. I. (2011). Distractor devaluation effect in the attentional blink: Direct evidence for distractor inhibition. Journal of Experimental Psychology.Human Perception and Performance, 37(1), 168-179. doi:10.1037/a0019948; 10.1037/a0019948

Kliegl, R., Masson, M. E. J., & Richter, E. M. (2009). A linear mixed model analysis of masked repetition priming. Visual Cognition, 18(5), 655-681.

Koelewijn, T., Van der Burg, E., Bronkhorst, A., & Theeuwes, J. (2008). Priming T2 in a visual and auditory attentional blink task. Perception & Psychophysics, 70(4), 658-666.

Kok, A. (2001). On the utility of P3 amplitude as a measure of processing capacity. Psychophysiology, 38(3), 557-577.

Kranczioch, C., & Bryant, D. (2011). Attentional awakening, resource allocation and the focus of temporal attention. Neuroreport, 22(4), 161-165. doi:10.1097/WNR.0b013e3283438b76; 10.1097/WNR.0b013e3283438b76

Kranczioch, C., Debener, S., & Engel, A. K. (2003). Event-related potential correlates of the attentional blink phenomenon. Cognitive Brain Research, 17(1), 177-187.

Kranczioch, C., Debener, S., Schwarzbach, J., Goebel, R., & Engel, A. K. (2005). Neural correlates of conscious perception in the attentional blink. Neuroimage, 24(3), 704-714.

Lavie, N., & Tsal, Y. (1994). Perceptual load as a major determinant of the locus of selection in visual attention. Perception & Psychophysics, 56(2), 183-197.

Loach, D., & Marí-Beffa, P. (2003). Post-target inhibition: A temporal binding mechanism? Visual Cognition, 10(5), 513-526.

Luck, S. J., Vogel, E. K., & Shapiro, K. L. (1996). Word meanings can be accessed but not reported during the attentional blink. Nature, 383(6601), 616-618.

Maclean, M. H., & Arnell, K. M. (2010). Personality predicts temporal attention costs in the attentional blink paradigm. Psychonomic Bulletin & Review, 17(4), 556-562. doi:10.3758/PBR.17.4.556; 10.3758/PBR.17.4.556

Maki, W. S., & Padmanabhan, G. (1994). Transient suppression of processing during rapid serial visual presentation: Acquired distinctiveness of probes modulates the attentional blink. Psychonomic Bulletin & Review, 1(4), 499-504.

Mangun, G. R., & Hillyard, S. A. (1991). Modulations of sensory-evoked brain potentials indicate changes in perceptual processing during visual-spatial priming. Journal of Experimental Psychology: Human Perception and Performance, 17(4), 1057-1074.

Martens, S., Dun, M., Wyble, B., & Potter, M. C. (2010). A quick mind with letters can be a slow mind with natural scenes: Individual differences in attentional selection. PLoS One, 5(10), e13562. doi:10.1371/journal.pone.0013562

Martens, S., Elmallah, K., London, R., & Johnson, A. (2006). Cuing and stimulus probability effects on the P3 and the AB. Acta Psychologica, 123(3), 204-218.

Martens, S., & Johnson, A. (2005). Timing attention: Cuing target onset interval attenuates the attentional blink. Memory and Cognition, 33(2), 234-240.

Martens, S., & Johnson, A. (2009). Working memory capacity, intelligence, and the magnitude of the attentional blink revisited. Experimental Brain Research, 129, 43-52.

Martens, S., Johnson, A., Bolle, M., & Borst, J. (2009). A quick visual mind can be a slow auditory mind: Individual differences in attentional selection across modalities. Experimental Psychology, 56(1), 33-40.

Martens, S., Kandula, M., & Duncan, J. (2010). Restricted attentional capacity within but not between sensory modalities: An individual differences approach. PLoS One, 5(12), 473-490.

Martens, S., Korucuoglu, O., Smid, H. G. O. M., & Nieuwenstein, M. R. (2010). Quick minds slowed down: Effects of rotation and stimulus category on the attentional blink. PLoS One, 5(10) doi:10.1371/journal.pone.0013509

Martens, S., Munneke, J., Smid, H., & Johnson, A. (2006). Quick minds don't blink: Electrophysiological correlates of individual differences in attentional selection. Journal of Cognitive Neuroscience, 18(9), 1423-1438.

Martens, S., & Valchev, N. (2009). Individual differences in the attentional blink: The important role of irrelevant information. Experimental Psychology, 56(1), 18-26.

Martens, S., Wolters, G., & van Raamsdonk, M. (2002). Blinks of the mind: Memory effects of attentional processes. Journal of Experimental Psychology: Human Perception and Performance, 28(6), 1275-1287.

Martens, S., & Wyble, B. (2010). The attentional blink: Past, present, and future of a blind spot in perceptual awareness. Neuroscience & Biobehavioral Reviews, 34, 947-957. doi:doi:10.1016/j.neubiorev.2009.12.005

McArthur, G., Budd, T., & Michie, P. (1999). The attentional blink and P300. Neuroreport, 10(17), 3691-3695.

McLaughlin, E. N., Shore, D. I., & Klein, R. M. (2001). The attentional blink is immune to masking-induced data limits. The Quarterly Journal of Experimental Psychology Section A, 54(1), 169-196.

Moresi, S., Adam, J. J., Rijcken, J., Van Gerven, P. W. M., Kuipers, H., & Jolles, J. (2008). Pupil dilation in response preparation. International Journal of Psychophysiology, 67(2), 124-130. doi:10.1016/j.ijpsycho.2007.10.011

Nichols, T. E., & Holmes, A. P. (2002). Nonparametric permutation tests for functional neuroimaging: A primer with examples. Human Brain Mapping, 15(1), 1-25. doi:10.1002/hbm.1058

Nieuwenhuis, S., Gilzenrat, M. S., Holmes, B. D., & Cohen, J. D. (2005). The role of the locus coeruleus in mediating the attentional blink: A neurocomputational theory. Journal of Experimental Psychology: General, 134(3), 291.

Nieuwenhuis, S., De Geus, E. J., & Aston-Jones, G. (2011). The anatomical and functional relationship between the P3 and autonomic components of the orienting response. Psychophysiology, 48(2), 162-175. doi:10.1111/j.1469-8986.2010.01057.x

Nieuwenstein, M. R. (2006). Top-down controlled, delayed selection in the attentional blink. Journal of Experimental Psychology: Human Perception and Performance, 32(4), 973-985.

Nieuwenstein, M. R., Chun, M. M., van der Lubbe, R. H. J., & Hooge, I. T. C. (2005). Delayed attentional engagement in the attentional blink. Journal of Experimental Psychology: Human Perception and Performance, 31(6), 1463-1475.

Nieuwenstein, M. R., & Potter, M. C. (2006). Temporal limits of selection and memory encoding: A comparison of whole versus partial report in rapid serial visual presentation. Psychological Science, 17(6), 471-475. doi:10.1111/j.1467-9280.2006.01730.x

Nieuwenstein, M. R., Potter, M. C., & Theeuwes, J. (2009). Unmasking the attentional blink. Journal of Experimental Psychology: Human Perception and Performance, 35(1), 159-169.

Norman, D. A., & Bobrow, D. G. (1975). On data-limited and resource-limited processes. Cognitive Psychology, 7(1), 44-64. doi:DOI: 10.1016/0010-0285(75)90004-3

Olivers, C. N. L., & Meeter, M. (2008). A boost and bounce theory of temporal attention. Psychological Review, 115(4), 836-863.

Olivers, C. N. L., & Nieuwenhuis, S. (2005). The beneficial effect of concurrent task-irrelevant mental activity on temporal attention. Psychological Science, 16(4), 265-269.

Olivers, C. N. L., & Nieuwenhuis, S. (2006). The beneficial effects of additional task load, positive affect, and instruction on the attentional blink. Journal of Experimental Psychology: Human Perception and Performance, 32(2), 364-79.

Olivers, C. N. L., van der Stigchel, S., & Hulleman, J. (2007). Spreading the sparing: Against a limited-capacity account of the attentional blink. Psychological Research, 71(2), 126-139.

Olivers, C. N. L., & Watson, D. G. (2006). Input control processes in rapid serial visual presentations: Target selection and distractor inhibition. Journal of Experimental Psychology: Human Perception and Performance, 32(5), 1083-1092.

Olivers, C. N. L. (2007). The time course of attention: It is better than we thought. Current Directions in Psychological Science, 16(1), 11-15. doi:10.1111/j.1467-8721.2007.00466.x

Olivers, C. N. L., & Watson, D. G. (2008). Subitizing requires attention. Visual Cognition, 16(4), 439-462. doi:10.1080/13506280701825861

O'Neill, W. D., & Zimmermann, S. (2000). Neurologial interpretations and the information in the cognitive pupillary response. Methods of Information in Medicine-Methodik Der Information in Der Medizin, 39(2), 122-124.

Ouimet, C., & Jolicoeur, P. (2007). Beyond task 1 difficulty: The duration of T1 encoding modulates the attentional blink. Visual Cognition, 15(3), 290-304. doi:10.1080/13506280600693741

Polich, J. (2007). Updating P300: An integrative theory of P3a and P3b. Clinical Neurophysiology, 118(10), 2128-2148.

Polich, J., & Donchin, E. (1988). P300 and the word frequency effect. Electroencephalography and Clinical Neurophysiology, 70(1), 33. doi:DOI: 10.1016/0013-4694(88)90192-7

Popple, A. V., & Levi, D. M. (2007). Attentional blinks as errors in temporal binding. Vision Research, 47(23), 2973-2981. doi:10.1016/j.visres.2007.06.022

Potter, M. C., Chun, M. M., Banks, B. S., & Muckenhoupt, M. (1998). Two attentional deficits in serial target search: The visual attentional blink and an amodal task-switch deficit. Journal of Experimental Psychology: Learning, Memory, and Cognition, 24(4), 979-992.

Potter, M. C., Staub, A., & O Connor, D. H. (2002). The time course of competition for attention: Attention is initially labile. Journal of Experimental Psychology: Human Perception and Performance, 28(5), 1149-1162.

Potter, M. C., Wyble, B., Pandav, R., & Olejarczyk, J. (2010). Picture detection in rapid serial visual presentation: Features or identity? Journal of Experimental Psychology: Human Perception and Performance, 36, 1486-1494. doi:10.1037/a0018730

Privitera, C. M., Renninger, L. W., Carney, T., Klein, S., & Aguilar, M. (2010). Pupil dilation during visual target detection. Journal of Vision, 10(10) doi:10.1167/10.10.3

Quené, H., & van den Bergh, H. (2008). Examples of mixed-effects modeling with crossed random effects and with binomial data. Journal of Memory and Language, 59(4), 413-425.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink. Journal of Experimental Psychology: Human Perception and Performance, 18(3), 849-860.

Rolke, B., Heil, M., Streb, J., & Hennighausen, E. (2001). Missed prime words within the attentional blink evoke an N400 semantic priming effect. Psychophysiology, 38(2), 165-174.

Rotte, M., Heinze, H. J., & Smid, H. (1997). Selective attention to conjunctions of color and shape of alphanumeric versus non-alphanumeric stimuli: A comparative electrophysiological study. Biological Psychology, 46(3), 199-221.

Shapiro, K. L., Schmitz, F., Martens, S., Hommel, B., & Schnitzler, A. (2006). Resource sharing in the attentional blink. Neuroreport, 17(2), 163-166.

Shapiro, K. L., Driver, J., Ward, R., & Sorensen, R. E. (1997). Priming from the attentional blink: A failure to extract visual tokens but not visual types. Psychological Science, 8(2), 95-100. doi:10.1111/j.1467-9280.1997.tb00689.x

Shapiro, K. L., Raymond, J. E., & Arnell, K. M. (1994). Attention to visual pattern information produces the attentional blink in rapid serial visual presentation. Journal of Experimental Psychology: Human Perception and Performance, 20(2), 357.

Shih, S. I. (2008). The attention cascade model and attentional blink. Cognitive Psychology, 56(3), 210-236.

Slagter, H. A., Johnstone, T., Beets, I. A., & Davidson, R. J. (2010). Neural competition for conscious representation across time: An fMRI study. PLoS One, 5(5), e10556. doi:10.1371/journal.pone.0010556; 10.1371/journal.pone.0010556

Slagter, H. A., Lutz, A., Greischar, L. L., Francis, A. D., Nieuwenhuis, S., Davis, J. M., & Davidson, R. J. (2007). Mental training affects distribution of limited brain resources. PLoS Biology, 5(6), e138.

Slagter, H. A., Tomer, R., Christian, B. T., Fox, A. S., Colzato, L. S., King, C. R., . . . Davidson, R. J. (2012). PET evidence for a role for striatal dopamine in the attentional blink: Functional implications. Journal of Cognitive Neuroscience, 24(9), 1932-1940. doi:10.1162/jocn_a_00255; 10.1162/jocn_a_00255

Smid, H. G. O. M., & Heinze, H. J. (1997). An electrophysiological study of the selection of the color and shape of alphanumeric characters in response choice. Biological Psychology, 44(3), 161-185.

Stol, J., Chatelle, C., Carter, O., Koch, C., Laureys, S., & Einhäuser, W. (2013). Pupil responses allow communication in locked-in syndrome patients. Current Biology, 23(15), R647-648. doi: 10.1016/j.cub.2013.06.011

Strayer, D. L., & Johnston, W. A. (2001). Driven to distraction: Dual-task studies of simulated driving and conversing on a cellular telephone. Psychological Science, 12(6), 462-466.

Taatgen, N. A., Juvina, I., Schipper, M., Borst, J., & Martens, S. (2009). Too much control can hurt: A threaded cognition model of the attentional blink. Cognitive Psychology, 59, 1-29.

Tang, M. F., Badcock, D. R., & Visser, T. A. (In Press). Training and the attentional blink: Limits overcome or expectations raised? Psychonomic Bulletin & Review, doi:10.3758/s13423-013-0491-3

Team, R. D. C. (2009). R: A language and environment for statistical computing. Viennia, Austria, , R Foundation for Statistical Computing.

Van Rijn, H., & Anderson, J. R. (2003). Modeling lexical decision as ordinary retrieval. In F. Detje, D. Doerner, & H. Schaub (Ed.), Proceedings of the 5th international conference on cognitive modeling (pp. 207-212) Bamberg: Universitaetsverlag Bamberg.

Verney, S. P., Granholm, E., & Marshall, S. P. (2004). Pupillary responses on the visual backward masking task reflect general cognitive ability. Int. J. Psychophysiol., 52(1), 23-36. doi:10.1016/j.ijpsycho.2003.12.003

Visser, T. A. W. (2007). Masking T1 difficulty: Processing time and the attentional blink. Journal of Experimental Psychology: Human Perception and Performance, 33(2), 285.

Visser, T. A. W., & Ohan, J. L. (2007). Data-limited manipulations of T1 difficulty influence the attentional blink. Canadian Journal of Experimental Psychology, 61(2), 102-108.

Vogel, E. K., & Luck, S. J. (2002). Delayed working memory consolidation during the attentional blink. Psychonomic Bulletin & Review, 9(4), 739-743.

Vogel, E. K., Luck, S. J., & Shapiro, K. L. (1998). Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. Journal of Experimental Psychology: Human Perception and Performance, 24(6), 1656-1674.

Vul, E., Hanus, D., & Kanwisher, N. (2008). Delay of selective attention during the attentional blink. Vision Research, 48(18), 1902-1909. doi:10.1016/j.visres.2008.06.009; 10.1016/j.visres.2008.06.009

Vul, E., Nieuwenstein, M., & Kanwisher, N. (2008). Temporal selection is suppressed, delayed, and diffused during the attentional blink. Psychological Science, 19(1), 55-61.

Ward, R., Duncan, J., & Shapiro, K. (1997). Effects of similarity, difficulty, and nontarget presentation on the time course of visual attention. Perception and Psychophysics, 59(4), 593-600.

Wierda, S. M., van Rijn, H., Taatgen, N. A., & Martens, S. (2010). Distracting the mind improves performance: An ERP study. PLoS One, 5(11) doi:10.1371/journal.pone.0015024

Wierda, S. M., van Rijn, H., Taatgen, N. A., & Martens, S. (2012). Pupil dilation deconvolution reveals the dynamics of attention at high temporal resolution. Proceedings of the National Academy of Sciences, 109(22), 8456-8460. doi:10.1073/pnas.1201858109

Wijers, A. A. (1989). Visual selective attention: An electrophysiological approach. unpublished doctoral dissertation. Groningen: University of Groningen.

Willems, C., Wierda, S. M., van Viegen, E., & Martens, S. (2013). Individual differences in the attentional blink: The temporal profile of blinkers and non-blinkers. PLoS One, 8(6), e66185. doi:10.1371/journal.pone.0066185; 10.1371/journal.pone.0066185

Wyble, B., Bowman, H., & Nieuwenstein, M. (2009). The attentional blink provides episodic distinctiveness: sparing at a cost. Journal of Experimental Psychology: Human Perception and Performance, 35(3), 787-807.

Appendix

The implementation of the pupil dilation deconvolution method of Chapter 5 is presented in this appendix

Sourcecode of pdd_functions.R

```
# file: pdd_functions.R
#
# Pupil dilation deconvolution method as described in:
#
# Wierda, Van Rijn, Taatgen & Martens (submitted) Pupil dilation
# deconvolution reveals the dynamics of attention at high
# temporal resolution.
#
# Please cite this work if using the deconvolution method and/or any
# code
# presented in these scripts.
#
# Method developed by Stefan Wierda & Hedderik van Rijn.
#
# Source code and updates can be found at
# https://sourceforge.net/p/wierdapdd/
#
# Script has been tested with Project R v.2.14.2 (http://www.r-pro
# ject.org/)
# and R Studio v.0.95.263 (http://rstudio.org/)


cat("Loading pupipl dilation deconvolution functions\n\n")
cat("Please cite Wierda, Van Rijn, Taatgen & Martens (submitted) in
any work based on\nor derived from these functions.\n\n")

se <- function(X) {
  sd(X)/(sqrt(length(X)))
}

####################################################################
########
########          Hoeks & Levelt pupillary response function
########
####################################################################


h_pupil <- function(t,n=10.1,t_max=930,f=1/(10^27) )
  # n+1 = number of laters
  # t_max = response maximum
  # f = scaling factor
{
  h<-f*(t^n)*exp(-n*t/t_max)
  h[0] <- 0
  h
}
```

```
####################################################################
########
########   Functions for fitting the strength & slope
########
####################################################################

create.matrix.ones <- function(n)
{
  o <- c()
  for(i in 1:n)
  {
    tmp <- rep(0,n)
    tmp[i] <- 1
    o <- rbind(o,tmp)
  }
  o
}

pupil.model.strength.slope  <-  function(pars,locations=NULL,input.
length=145,output.text = FALSE,h=NULL)
{
  if(!is.null(data.x))
  {
    length_data <- length(data.x)
    # The number of spikes to be fitted
    n <- (length(pars))  - 1
    # Each subject has fixed locations for all conditions

    locations <- locations
    strengths <- pars[1:n]
    slope <- pars[n+1]

    if(output.text)
    {
      cat("Values:\n")
      cat("\tLocations:",locations,"\n")
      cat("\tStrength:",strengths,"\n")
      cat("\tSlope:",slope,"\n")
    }

    locations <- floor(locations)
    i <- rep(0,input.length)
    loc_check <- (( length(unique(locations)) == length(locations)) &&
(max(locations) < length(i) )) && !any(diff(locations)<0)

    if(output.text)
    {
      cat("\tLocation check:",loc_check,"\n")
    }
```

```
    if(loc_check)
    {
      i[locations] <- strengths
      o <- convolve(i,rev(h),type="open")
      o <- o + (1:length(o))*slope
    }
    else
    {
      o <- rep(1000,length_data)
    }
  }
  else
  {
    cat("ERROR: This function requires a global h and data.x variable
containing the h-functions and pupil-data to be fitted\n")
    o <- NULL
  }
  o[1:length_data]
}


pupil.fit.strength.slope <- function(pars,locations=locations,input.
length=145,output.text = FALSE,h=NULL)
{
  tmp_data <- data.x
     tmp_pred <- pupil.model.strength.slope(pars,locations,input.
length,output.text,h)
  o <- sum((tmp_pred[1:length(tmp_data)] - tmp_data)^2)
}

pupil.optim.strength.slope <- function(x,n=4,locations=NULL,input.
length=145,output.text=FALSE,h=NULL)
{
  data.x <<- x
  ui <- create.matrix.ones(n+1)
  strength_ci <- rep(0,n)

  slope_ci <- -5

  # Constrains
  ci <- c(strength_ci,slope_ci)
  # Starting values
  strength_theta <- runif(n,0.01,0.25)
  slope_theta <- -2
  theta <- c(strength_theta,slope_theta)

  if(output.text)
  {
    tmp <- ui %*% theta - ci
    if(prod(tmp) > 0)
    {
```

```
      o <- TRUE
    }
    else
    {
      o <- FALSE
    }
    o
    cat("Initial value check:",o,"\n")
    cat("\tCall to Optim function\n")
  }

   fit_pars <- constrOptim(f=pupil.fit.strength.slope,grad=NULL,theta=
theta,ui=ui,ci=ci,outer.iterations  =  10000,locations=locations,in-
put.length=input.length,output.text=output.text,h=h)
  fit_pars

}


pddeconvolution <-function(size,time,pulse_locations)
{
  h <- h_pupil(time)
  h
  pulses <- 1:length(pulse_locations)
  model_pred <- pupil.optim.strength.slope(size,n=length(pulses),lo-
cations=pulse_locations,input.length=max(pulse_locations)+1,h=h)
  model_pred
  output <- c()
  output$pulses <- model_pred$par[pulses]
  output$slope <- model_pred$par[length(model_pred$par)]
  output
}

pdconvolution <-function(pulses,slope,time,pulse_locations)
{
  h <- h_pupil(time)
  output <- c()
  output <- data.frame(time=time,x=pupil.model.strength.slope(c(pul-
ses,slope),locations=pulse_locations,input.length=max(pulse_locati-
ons)+1,h=h,output.text=F))
  output
}

line_green <-function(x,y,y_err)
{
  lines(x,y,lwd=2,col="darkgreen")

  err1 <- y - y_err
  err2 <- y + y_err

  lines(x,err1,type="l",col=c("darkgreen"),lwd=1,lty=2)
```

```
  lines(x,err2,type="l",col=c("darkgreen"),lwd=1,lty=2)
          polygon(c(x,rev(x)),c(err1,rev(err2)),col=rgb(0,100,0,30,-
maxColorValue=255),border=NA)
}

line_orange <-function(x,y,y_err)
{
  lines(x,y,lwd=2,col="darkorange")

  err1 <- y - y_err
  err2 <- y + y_err

  lines(x,err1,type="l",col=c("darkorange"),lwd=1,lty=2)
  lines(x,err2,type="l",col=c("darkorange"),lwd=1,lty=2)
   polygon(c(x,rev(x)),c(err1,rev(err2)),col=rgb(255,140,0,30,maxCo-
lorValue=255),border=NA)
}
```

Sourcode of main.R

```
# file: main.R
#
# Example code for pupil dilation deconvolution method as described
# in:
#
# Wierda, Van Rijn, Taatgen & Martens (submitted) Pupil dilation
# deconvolution reveals the
# dynamics of attention at high temporal resolution.
#
# Please cite this work if using the deconvolution method and/or any
# code
# presented in
# these scripts.
#
# The example code deconvolutes Lag 3 trials in which the second tar
# get is
# identified
# correctly (i.e., no-blink trials). Note that this example code runs
# the deconvolution
# routine only once. The above mentioned paper reports estimated
# effects
# based on the
# average of 200 runs.
#
# Method developed by Stefan Wierda & Hedderik van Rijn.
#

# Source code and updates can be found at
# https://sourceforge.net/p/wierdapdd/
#
# Script has been tested with Project R v.2.14.2 (http://www.r-pro-
ject.org/)
# and R Studio v.0.95.263 (http://rstudio.org/)
#

# Source the files needed for the deconvolution
source("pdd_functions.R")

# Load the data from the experiment of Wierda et al. (submitted)

dat <- read.table(file="data.txt")

# This is the dataframe we are going to store our pulses data in
model.data <- NULL
model.prediction <- NULL

for(i in unique(dat$Subject))
```

```
{
  # Get the subset of data: Lag 3, T1 and T2 both accurately responded
  tmp_dat <- dat[dat$Subject==i & dat$Lag==3 & dat$T1acc==1 & dat$T2acc
== 1,]

  # Now average the size of the pupil over each trial
  eyeDat <- with(tmp_dat,aggregate(size,list(time=time),mean))

  # Set the location of the pulses
  pulse_locations <- seq(1,170,5) # This adds up to 34 pulses.
  real_locations <- eyeDat$time[pulse_locations] # This vector gives
the onsets of the pulses

  # perform deconvolution for this subject and store output in frame
  tmp <- pddeconvolution(eyeDat$x,eyeDat$time,pulse_locations)

  # Perform convolution to obtain the predicted response for later use
  tmp2 <- pdconvolution(tmp$pulses,tmp$slope,eyeDat$time,pulse_loca-
tions)

  # Add the obtained data to the data-frame
  if(is.null(model.data))
  {
          model.data  <-  data.frame(subject=i,slope=tmp$slope,pul-
se=1:length(pulse_locations), pulse_size=tmp$pulses,locations=pulse_
locations,locations_ms=real_locations)
    model.prediction <- data.frame(subject=i,time=tmp2$time,x=tmp2$x)
  }
  else
  {
      model.data  <-  rbind(model.data,data.frame(subject=i,slope=t-
mp$slope,pulse=1:length(  pulse_locations),pulse_size=tmp$pulses,lo-
cations=pulse_locations,locations_ms= real_locations))
    model.prediction <- rbind(model.prediction, data.frame(subject=i,-
time=tmp2$time,x=tmp2$x))
  }
}

# Aggregate the data in order to plot
tmp_dat <- dat[dat$Lag==3 & dat$T1acc==1 & dat$T2acc == 1,]
tmp_dat <- with(tmp_dat,aggregate(size,list(Subject=Subject,time=ti-
me),mean))

data_se <- with(tmp_dat,aggregate(x,list(time=time),se))
data_mean <- with(tmp_dat,aggregate(x,list(time=time),mean))

model_se <- with(model.prediction,aggregate(x,list(time=time),se))
model_mean  <-  with(model.prediction,aggregate(x,list(time=time),-
mean))

pulses_se <- with(model.data,aggregate(pulse_size,list(time=locati-
```

```
ons_ms),se))
pulses_mean <-with(model.data,aggregate(pulse_size,list(time=locati-
ons_ms),mean))

# Now plot the result
par(mfrow=c(1,2))

plot(data_mean,type="n",xlab="Time                    (ms)",ylab="Size",y-
lim=c(-0.04,0.075),main="Data vs. Model")

line_green(data_mean$time,data_mean$x,data_se$x)
line_orange(model_mean$time,model_mean$x,model_se$x)

legend(data_mean$x[1],0.075,col=c("darkgreen","darkorange"),pt.
bg=c(rgb(0,100,0,30,maxColorValue=255),rgb(255,140,0,30,-
maxColorValue=255)),pch=c(22,22),legend=c("Data","Model"),pt.cex=2,-
box.col="darkgrey",bg="white",cex=.8)


plot(data_mean,type="n",xlab="Time                    (ms)",ylab="Strength",y-
lim=c(0,0.75),main="Pulses")

line_orange(pulses_mean$time[4:31],pulses_mean$x[4:31],pulses_
se$x[4:31])

legend(data_mean$x[1],0.75,col=c("darkorange"),pt.bg=c(rg-
b(255,140,0,30,maxColorValue=255)),pch=c(22),legend=c("Model"),pt.
cex=2,box.col="darkgrey",bg="white",cex=.8)
```

Nederlandse samenvatting van dit proefschrift

Dutch summary of this thesis

De voortdurende technologische ontwikkelingen zorgen ervoor dat er steeds meer informatie op ons af komt. Denk bijvoorbeeld aan de tablet en smartphone waarmee we sinds enkele jaren dag en nacht met elkaar in verbinding staan en een toenemende hoeveelheid informatie steeds sneller over en weer wordt uitgewisseld. Hoe bepalen we welke informatie relevant is en welke informatie niet? Om snel opeenvolgende relevante informatie te kunnen onderscheiden van irrelevante informatie gebruiken we een mechanisme in ons brein dat we selectieve temporele aandacht noemen. Selectieve temporele aandacht beschrijft de mogelijkheid om aandacht op het juiste moment te verdelen over specifieke stukjes relevante informatie. Door onze aandacht op bepaalde informatie te richten, kunnen we die informatie effectief verwerken en opslaan in ons geheugen of er direct op reageren, terwijl we de andere informatie negeren. Maar hoe werkt dat precies? Om daar een antwoord op te krijgen zoeken wetenschappers de grenzen van het menselijk denken op en zetten ze deze onder druk door in een hele korte periode heel veel informatie, zowel relevant als irrelevant, aan te bieden.

Een geschikte manier om selectieve temporele aandacht te testen en onder druk te zetten is gebruik te maken van de zogenaamde *attentional blink*-taak. De taak gaat als volgt: stel je voor dat je meedoet aan een spelprogramma waar je twee prijzen mee kunt winnen. De presentator legt uit dat hij een lijst van producten heeft gemaakt en dat er twee meubels in de lijst staan die je kunt winnen. De producten op de lijst zullen één voor één met een tiende van een seconde gepresenteerd worden op een groot scherm. Als je kunt onthouden welke meubels er op de lijst staan, mag je deze houden. Je weet echter niet op welke plek in de lijst beide meubels staan. Nu is het zo dat wanneer de twee meubels vlak achter elkaar op het scherm verschijnen, binnen een halve seconde, je vaak het tweede meubel mist. Op het moment dat je bezig bent met het eerste meubel in je geheugen op te slaan, mis je het tweede. Er treedt dus als het ware een knippering (*blink*) in je aandacht (*attention*) op (wanneer je met je ogen knippert zie je niets, wanneer je met aandacht knippert ook niet). Dit is het verschijnsel waar de *attentional blink* zijn naam aan te danken heeft. In het onderzoek dat wij gedaan hebben gebruikten we (in plaats van meubels en andere producten) cijfers, letters of woorden in een *attentional blink*-taak. In dit proefschrift wordt verslag gedaan van dit onderzoek, waarbij selectieve temporele aandacht en de *attentional blink*-taak centraal stonden.

De grote vraag die we wilden beantwoorden is waarom mensen een dergelijke *attentional blink* hebben. Het is toch veel efficiënter om je aandacht erbij te houden en beide stukjes relevante informatie te onthouden? De voornaamste theorieën om de *attentional blink* te verklaren vallen binnen twee stromingen: de capaciteitstheorieën en de controletheorieën. Onderzoekers die de capaciteitstheorie aanhangen denken dat de *attentional blink* ontstaat doordat we simpelweg niet genoeg capaciteit hebben om beide stukjes informatie te verwerken in een dergelijke korte tijdsperiode. Wanneer je de hersenen ziet als een computer en de vaardigheden en kennis die we hebben als software, dan zou dat betekenen dat de computer te langzaam is of te weinig geheugen heeft, waardoor je niet op tijd alle informatie kunt verwerken. Hoewel de software wellicht optimaal is, is dit probleem niet op te lossen, aangezien je de hardware, in dit geval de hersenen, niet zomaar kunt vervangen.

De hierboven beschreven theorie heeft echter een aantal zwakheden. Er zijn namelijk mensen die helemaal geen *attentional blink* laten zien ('non-blinkers', in tegenstelling tot 'blinkers' die wel een *attentional blink* laten zien). Het zou kunnen zijn dat de hersenen (hardware) van non-blinkers beter zijn dan die van blinkers en dat ze daarom alle relevante informatie kunnen verwerken. Maar het zou ook kunnen zijn dat deze non-blinkers de taak anders uitvoeren in hun hersenen dan blinkers. Er is ook bewijs gevonden dat de prestatie op de *atten-*

tional blink-taak vele malen verbeterd kan worden door blinkers op een specifieke manier te trainen op de *attentional blink*-taak. Tot slot blijkt dat wanneer je blinkers afleidt tijdens het doen van de *attentional blink*-taak, door hen bijvoorbeeld een tweede taak te laten doen, ze veel beter gaan presteren. De hierboven genoemde bevindingen wijzen erop dat de *attentional blink* niet verklaard kan worden als een puur capaciteitsprobleem.

Een andere theorie is dat we in principe de capaciteit wel hebben om alle relevante informatie te onthouden, maar dat we de taak op een dusdanige manier uitvoeren dat we niet al. onze capaciteit kunnen benutten. Als we ervoor kunnen zorgen dat we onze capaciteit wel goed in kunnen zetten, dan zouden we geen last hebben van de *attentional blink*. Ook hier kun je de hersenen zien als een soort computer, maar nu is die computer snel genoeg en heeft hij genoeg geheugen. De beperking zit dan in de software: de hardware kan de taak aan, maar de software niet. Dit zou op te lossen zijn door een andere, efficiëntere verwerkingsstrategie te hanteren, bijvoorbeeld door de taak anders te benaderen of op een andere manier aan te leren. Je vervangt de software dan, of past hem in feite aan. Omdat het in dit soort theorieën gaat om de controle, dat wil zeggen de manier waarop een taak gedaan wordt, noemen we deze theorieën controletheorieën. De manier waarop een taak uitgevoerd wordt is een controlestrategie.

In Hoofdstuk 2 onderzoeken we een dergelijke controlestrategie: wat gebeurt er in je brein als je afgeleid wordt tijdens het doen van een *attentional blink*-taak? Je zou verwachten dat het nog slechter zou gaan, maar verrassend genoeg bleek de prestatie op de taak te verbeteren in dit experiment.

Op het moment dat iemand een stukje relevante informatie opslaat, genereert dat een specifiek patroon van hersenactiviteit. Deze hersenactiviteit werd gemeten terwijl mensen de *attentional blink*-taak uitvoerden. Vervolgens hebben we de activiteit gemeten tijdens een soortgelijke taak, maar ditmaal terwijl de proefpersonen werden afgeleid. We constateerden dat de taak-specifieke hersenactiviteit tijdens de *attentional blink*-taak afnam op het moment dat mensen de tweede taak deden. Het lijkt er dus op dat wanneer mensen afgeleid worden door bijvoorbeeld een tweede taak, zij minder aandacht in de *attentional blink*-taak stoppen en dat de prestatie op die taak daar desondanks beter van wordt. Een mogelijke verklaring voor dit paradoxale verschijnsel is dat mensen te veel hun best doen om het eerste stukje informatie te verwerken, waardoor zij het tweede stukje informatie missen. Als ze vervolgens aandacht moeten stoppen in een tweede taak, en ze dus worden afgeleid van het verwerken van het eerste stukje informatie, hebben ze genoeg aandacht of capaciteit over voor het tweede stukje informatie. Hoewel de *attentional blink*-taak dus beter gaat tijdens afleiding, verdwijnt de *attentional blink* niet helemaal. Het gevonden verschil in hersenactiviteit is anders dan de verschillen die zijn gevonden tussen blinkers en non-blinkers.

Om het verschil tussen mensen die goed zijn in de taak, non-blinkers, in meer detail te vergelijken met mensen die slecht zijn in de taak, blinkers, kijken we in Hoofdstuk 3 naar het zogenaamde temporele profiel van deze mensen. Dit temporele profiel wordt opgesteld aan de hand van de fouten die men maakt in een *attentional blink*-taak. De mate waarin mensen stimuli onderdrukken, maar ook de vertraging die in de selectie van informatie zit, kan bekeken worden aan de hand van de gemaakte fouten. Verrassend genoeg blijkt dat blinkers minder hard stimuli onderdrukken dan verondersteld werd in voorgaande studies. Verder blijkt dat non-blinkers sneller zijn in het verwerken van informatie, hoewel ze juist langzamer zijn en vaker de relevante informatie door elkaar halen wanneer deze direct na elkaar volgt zonder tussenliggende irrelevante informatie. Blinkers en non-blinkers lijken dus een ander soort

strategie toe te passen in de manier waarop zij informatie selecteren.

In Hoofdstuk 4 wordt verder gekeken naar het onderliggende mechanisme van de *attentional blink*, ditmaal door de moeilijkheid van de binnenkomende informatie te vergroten in een *attentional blink*-taak met woorden als relevante informatie, en non-woorden als irrelevante informatie. De moeilijkheid wordt gevarieerd met behulp van de woordfrequentie van een woord. Woordfrequentie is een maat voor hoe vaak een woord voorkomt in een natuurlijke taal, bijvoorbeeld in boeken, de krant of op televisie. Woorden die vaak voorkomen zijn hoogfrequent en daardoor relatief eenvoudig te herkennen. Laagfrequente woorden komen minder vaak voor en zijn dus moeilijker te herkennen. Omdat het meer tijd en moeite zou kosten om die moeilijke woorden te verwerken, zou volgens de capaciteitstheorie dergelijke moeilijke informatie voor een grotere *attentional blink* moeten zorgen. Dit is inderdaad het geval wanneer het eerste stukje informatie moeilijk is (m.a.w., een laagfrequent woord). Maar wanneer het tweede stukje informatie moeilijk is, dan wordt de taak relatief gezien makkelijker.

Dit wederom verrassende effect verklaren we door te stellen dat je twee handelingen uitvoert met de relevante woorden die je ziet in deze taak. Eerst herken je het betreffende woord door het in je geheugen op te zoeken (herkennen), vervolgens sla je het woord in je geheugen op zodat je aan het eind van de taak weet welk woord je gezien hebt (opslaan). Zowel het herkennen heeft een bepaalde tijdsduur nodig als ook het opslaan van het woord. Verder veronderstellen we dat twee dezelfde handelingen niet tegelijkertijd uitgevoerd kunnen worden. Zo kun je twee woorden niet tegelijkertijd herkennen (dezelfde handeling), maar is het wel mogelijk om het ene woord op te slaan terwijl je het andere woord herkent (verschillende handelingen). Nu is het zo dat woorden die niet vaak in onze taal voorkomen (laagfrequente woorden) meer tijd en moeite vergen om herkend te worden, terwijl woorden die vaak voorkomen (hoogfrequente woorden) eenvoudiger en dus sneller herkend worden. Zolang we bezig zijn met het opslaan van het eerste relevante woord, zal het opslaan van het tweede woord dus moeten wachten. De kans dat het weede woord uiteindelijk succesvol gerapporteerd wordt zal kleiner zijn. Echter, als de herkenning van het tweede woord relatief veel tijd kost kan het voorkomen dat het opslaan van het eerste woord al is afgerond voordat het opslaan van het tweede woord begint. De kans op onderling conflicterende handelingen neemt in dat geval af. Daardoor is het mogelijk dat een laagfrequent (moeilijk) woord dat als tweede gepresenteerd wordt succesvol opgeslagen kan worden, terwijl een hoogfrequent (eenvoudig) tweede woord vaker in conflict komt met het opslaan van het voorgaande woord en dus verloren zal gaan.

In een vervolgexperiment waarbij we hersenactiviteit gemeten hebben, komt inderdaad naar voren dat de hoogfrequente woorden makkelijker verloren gaan, en dat juist de laagfrequente moeilijkere woorden een grotere kans hebben om gezien te worden. Simulaties met een computermodel laten eenzelfde patroon van resultaten zien. De bevindingen die in dit hoofdstuk gepresenteerd worden, zijn moeilijk te verklaren met capaciteitstheorieën, terwijl controletheorieën dit soort resultaten wel kunnen verklaren.

In hoofdstuk 5 presenteren we een nieuwe manier om te meten hoeveel aandacht iemand besteedt aan een bepaalde taak. Dit doen we aan de hand van de grootte van de pupil. Naast dat de pupil reageert op licht (als het donker is wordt je pupil groter, als het licht is kleiner), reageert de pupil ook als je ergens over nadenkt. Wanneer informatie verwerkt wordt, vergroot de pupil vrij langzaam en piekt na ongeveer een seconde. Doordat de *attentional blink*-taak een taak is waar informatie voor een tiende van een seconde aangeboden wordt en

de pupilreactie op die informatie een seconde nodig heeft om maximaal te worden, was het lastig (nagenoeg onmogelijk) om iets zinnigs te zeggen over veranderingen in de pupilgrootte tijdens een dergelijke snelle taak. Door de pupildeconvolutie methode uit dit hoofdstuk toe te passen (een nieuw ontwikkelde analysemethode), kon de pupilgrootte terug vertaald worden naar de onderliggende mentale activiteit die door gepresenteerde informatie werd opgewekt. Op deze manier kon er alsnog een behoorlijk nauwkeurig beeld worden verkregen van wanneer en in welke mate de aandacht tijdens het uitvoeren van de taak werd ingezet. Zo konden we met deze methode duidelijke verschillen in de aandacht aantonen wanneer informatie al dan niet succesvol gerapporteerd kon worden en vonden we subtiele aandachteffecten op het moment waarop het tweede stukje informatie vaak verscheen.

Conclusie

Door de hoofdstukken heen worden er in dit proefschrift steeds paradoxale effecten gerapporteerd. Deze effecten zijn lastig te verklaren als men uit gaat van capaciteitstheorieën. Wanneer we de *attentional blink* zien in het licht van controletheorieën en er dus van uit gaan dat de *attentional blink* ontstaat door de manier waarop mensen de taak uitvoeren, dan kunnen deze resultaten wel verklaard worden. Uiteraard is er—alle manipulaties ten spijt—wel een limiet aan de hoeveelheid informatie die het menselijk brein kort achter elkaar kan onthouden en zal er waarschijnlijk een combinatie van zowel controle- als capaciteitstheorieën nodig zijn om alle aspecten van het *attentional blink*-fenomeen goed te kunnen verklaren.

Enkele belangrijke vragen blijven nog open. Om er een aantal te noemen: wat is nu precies het verschil tussen blinkers en non-blinkers? Waarom gebruiken blinkers een strategie die niet optimaal is omdat deze tot de *attentional blink* leidt? Als het inderdaad zo is dat we de 'verkeerde' manier gebruiken om de taak te doen, waarom gebruikt het merendeel van de mensen dan toch deze manier? Wellicht dat de manier waarop we de relevante informatie verwerken in alledaagse taken juist optimaal is, maar dat deze in een taak zoals de *attentional blink*-taak—welke we niet vaak tegenkomen in het dagelijks leven—niet optimaal uitwerkt. Als de *attentional blink* inderdaad een nuttige functie heeft, is het dan zo dat non-blinkers misschien een andere taak niet goed kunnen uitvoeren? Dit zou betekenen dat er in principe een taak zou moeten bestaan die de non-blinkers slecht kunnen uitvoeren maar de blinkers juist goed. Tot op heden is een dergelijke taak niet gevonden, en blijft het nut en mechanisme achter de *attentional blink* helaas nog steeds een raadsel.

Ter afsluiting: de data gepresenteerd in dit proefschrift geeft evidentie voor het idee dat de *attentional blink* een probleem weerspiegelt in het aansturen van de aandacht, maar het moge duidelijk zijn dat er nog veel vragen onbeantwoord blijven. Hoewel dit proefschrift een stap in de goede richting is, zijn we nog steeds verre van het begrijpen hoe onze hersenen de wereld om ons heen een plaats geven in ons bewustzijn. Het zou, denk ik, in ieder geval goed zijn om de *attentional blink* in een bredere theorie over aandacht en het denken proberen in te passen, in de hoop op die manier uiteindelijk te achterhalen waarom de meeste mensen dit opmerkelijke fenomeen vertonen.

Acknowledgements

The acknowledgements or *dankwoord* for all of you who have supported me during the period that I worked on this thesis, and without whom I could not have gotten to this point

Na vier jaar is het eindelijk zover: mijn proefschrift is af! De afgelopen vier jaar heb ik hard gewerkt om dit resultaat te bereiken, maar zonder mijn promotores, co-promotores, collega's, studenten en de mensen die deel hebben genomen als proefpersonen zou ik dit wetenschappelijk stuk uiteraard nooit hebben kunnen schrijven. Ik ben dan ook erg dankbaar voor alle hulp en advies die ik de afgelopen jaren heb mogen ontvangen. De onderstaande mensen wil ik in het bijzonder bedanken.

Allereerst mijn dank aan de drie begeleiders met wie ik de afgelopen vier jaar intensief heb samengewerkt en waarmee we een goed team vormden.

Dr. S. Martens, beste Sander, als directe begeleider van mijn onderzoek wil ik jou als eerste bedanken. Bedankt voor de intensieve samenwerking de afgelopen vier jaar. Jouw enthousiasme voor het onderzoek en met name de *attentional blink* is erg aanstekelijk. Ik heb genoten van onze discussies en heb dan ook de wekelijkse bijeenkomsten altijd als prettig ervaren.

Dr. H. van Rijn, beste Hedderik, met name jouw kritische blik op het design van de studies die we samen hebben gedaan, plus het brainstormen over methodes om experimenten op een zo correct mogelijk manier te analyseren, hebben er toe geleid dat ik erg trots mag zijn op de publicaties in dit proefschrift.

Prof. dr. N. A. Taatgen, beste Niels, ik heb erg veel geleerd van en gehad aan jouw vermogen om zeer analytisch over de attentional blink na te kunnen denken.

Dr. A. Aleman, beste André, helaas hebben we tijdens mijn promotie niet heel veel samengewerkt, met name omdat het onderwerp van deze thesis iets verder van jouw onderzoeksveld af ligt. Ik wil je desalniettemin bedanken voor de kans die ik heb gekregen om mijn onderzoek in jouw groep te mogen doen.

Ook zou ik mijn twee andere co-auteurs willen bedanken. Eva, wie had ooit gedacht dat wij samen een artikel zouden publiceren? Charlotte, het afgelopen jaar heb ik de samenwerking met jou als zeer prettig ervaren en ik ben blij dat we samen een hoofdstuk hebben kunnen schrijven voor dit proefschrift. Uiteraard vind ik het ook erg fijn dat je mijn paranimf wilt zijn tijdens mijn verdediging.

Verder wil ik de andere studenten bedanken die betrokken zijn geweest bij de totstandkoming van dit onderzoek, waaronder Atser en Manasa.

Jonathan, bedankt voor de leuke discussies, gezellige feesten en mooie tijd bij de PhD Council. Tof dat ook jij mijn paranimf wilt zijn.

Amarins, Barbara, David, Emi, Funda, Hande en Vladimir, ook jullie bedankt voor jullie vriendschap en voor de leuke tijd bij de PhD Council.

De (oud-)leden van de CogMod groep, Ben, Burcu, Enkhbold, Florian, Fokie, Ioanna, Jacolien, Jason, Jelmer, Jennifer, Katja, Leendert, Marijke, Margreet, Marieke, Menno, Tadeusz, Trudy, Udo en alle andere mensen die de bijeenkomsten de afgelopen vier jaar bijgewoond hebben, bedankt voor alle input, discussie en inspiratie. Ik zie uit naar onze samenwerking tijdens mijn postdoc de komende twee jaar.

Uiteraard ook veel dank aan mijn (oud-)collega's van BCN en het NIC, Annerieke, Branislava, Claire, Diana, Doety, Edith, Edyta, Esther, Evelien, Frans, Hanneke, Hedwig, Heleen, Jacob, Jan-Bernard, Janine, Jelle, Jojanneke, Jorien, Katharina, Léon, Leonie, Linda, Lisette, Luca, Marie-José, Marieke, Marije, Mark, Marnix, Martijn, Michelle, Moniek, Natascha, Nicky, Nicolas, Nikola, Nynke, Rasa, Remco, Ronald, Ruggero, Ruud, Sanne, Shankar en Sima. Zoals jullie misschien wel weten ben ik niet al. te goed met namen, dus mocht ik je vergeten zijn in deze lijst te noemen, toch bedankt.

Naast het onderzoek voor dit proefschrift heb ik me ook bezig gehouden met onderzoek naar liefdesverdriet en liefdesrelaties samen met Dr. Love.

Prof. dr. G. ter Horst, beste Gert, ik vond het erg interessant om betrokken te zijn bij het onderzoek naar verliefdheid en ik hoop in de toekomst hier verder aan te werken samen met jou. Ook wil ik uiteraard de studenten bedanken met wie ik in het kader van dit onderwerp heb samengewerkt. Saif, René en Madelein, maar ook de andere studenten, bedankt voor jullie inzet voor dit onderzoek. Ik vond het erg prettig om met jullie samen te werken.

Aan het onderzoek doen en het schrijven van een proefschrift zit echter niet alleen een wetenschappelijke kant.

Naast Eva, die niet alleen een co-auteur, maar ook een vriendin van me is, wil ik ook mijn andere vrienden bedanken voor de leuke tijd die we de afgelopen jaren tijdens mijn studie en promotietraject hebben gehad. Arno, Bastiaan, Cas, Colin, Evert, Heike, Hielke, Inge, Jason, Jasper, Joas, Jolie, Stefan, Laurens, Maarten, Michel, Paul, Peter, Richard, Sander, Sander, Tarik, Tessa en Ying bedankt voor een onvergetelijke tijd.

I also want to thank Anne, Ben, Chris, Danielle, David, Erik, Gabe, Jessica, Kelley, Kristi, Luke, Martina, Raj, Ted, Valentinos, and Vick for the amazing time in Pittsburgh, just before I started my research for this thesis.

Ik wil mijn gehele familie bedanken voor de steun die ik heb mogen ontvangen de afgelopen jaren, maar toch vooral in het bijzonder voor het afgelopen jaar.

Pa, ik kan niet in woorden uitdrukken hoe erg ik het vind dat je dit proefschrift en dankwoord nooit zult kunnen lezen. Ik zal altijd van je blijven houden en je altijd blijven missen.

Moeders, jij hebt me grotendeels opgevoed en dankzij jou ben ik nu wie ik ben. Ik ben er trots op een zoon van je te zijn en ik zal uiteraard ook altijd van jou blijven houden.

Broers en zussen Eric, Inge, Jurjen en Laura, bedankt voor jullie steun en interesse. En hoewel ze zelf nog weinig besef hebben van alles wat er om hen heen gebeurt, wil ik bij deze Arthur en Joes hier nog noemen. Ook natuurlijk mijn dank aan Kinge, Richard, Roline en Marijn, die inmiddels ook als deel van de familie beschouwd kunnen worden.

Marion, ook jij, je dochters en je vrienden bedankt voor de gezelligheid en steun de afgelopen jaren, en in bijzonder natuurlijk het afgelopen jaar.

Cees, maar ook de rest van de familie en alles wat daarbij hoort, erg bedankt voor jullie steun en hulp het afgelopen jaar.
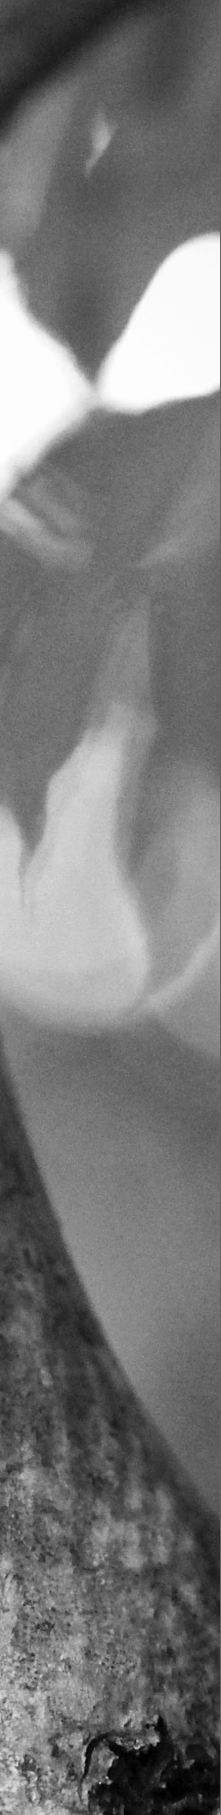
Ook mijn schoonfamilie mag natuurlijk niet ontbreken in dit dankwoord. Anton, Ria, Rob en Eda, Heleen en Freek, bedankt voor de leuke jaren en dat er nog veel leuke jaren mogen volgen.

Hoewel ik jullie tegenwoordig alleen nog sporadisch zie, wil ik ook jullie bedanken voor de rol die jullie hebben gespeeld in mijn jongere jaren, Teun, Jeanet, Anda, Mark, Pim en natuurlijk al jullie (klein)kinderen.

Cezanne en Casanova, jullie kunnen uiteraard ook niet in dit rijtje ontbreken, bedankt voor alle gezelligheid en knuffels!

Lieve Ingrid, ten slotte wil ik jou bedanken voor je steun en voor de jaren dat we al bij elkaar zijn. Je weet natuurlijk hoeveel ik van je hou, maar toch schrijf ik het hier nogmaals op. Ik hou héél veel van je laiffie.

A list of published peer-reviewed journal articles and conference abstracts

Journal Articles

1. **Wierda, S. M.**, Taatgen, N. A., van Rijn, H., & Martens, S. (in press). Word frequency and the attentional blink: The effects of target difficulty on retrieval and consolidation processes. PLoS One

2. Willems, C., **Wierda, S. M.**, van Viegen, E., & Martens, S. (2013). Individual differences in the attentional blink: The temporal profile of blinkers and non-blinkers. PloS One, 8(6), e66185

3. **Wierda, S. M.**, van Rijn, H., Taatgen, N. A., & Martens, S. (2012). Pupil dilation deconvolution reveals the dynamics of attention at high temporal resolution. Proceedings of the National Academy of Sciences, 109(22), 8456-8460.

4. **Wierda, S. M.**, van Rijn, H., Taatgen, N. A., & Martens, S. (2010). Distracting the mind improves performance: An ERP study. PLoS One, 5(11), e15024

Abstracts

1. **Wierda, S. M.**, van Rijn, H., Taatgen, N. A., & Martens, S. (2013). Pupil dilation deconvolution reveals the dynamics of attention at high temporal resolution. European Conference on Visual Perception. Bremen, Germany

2. **Wierda, S. M.**, van Rijn, H., Taatgen, N. A., & Martens, S. (2011). Pupil dilation deconvolution reveals the dynamics of attention at high temporal resolution. In Proceedings of the 13th Winter Conference of the Dutch Psychonomic Society

3. Cnossen, F., Grapengeter, M. Feenstra, L., **Wierda, S. M.** (2009). Learning flexible skills in anesthesiology: comparing list and context learning in a human patient simulator. In Proceedings of the the 31st Annual Conference of the Cognitive Science Society Cognitive Science Society

4. **Wierda, S. M.** & Taatgen, N. A. (2009). The Learning of Robust and Flexible Skills with Respect to System Failures in a Flight Management System. In Proceedings of the 12th Winter Conference of the Dutch Psychonomic Society
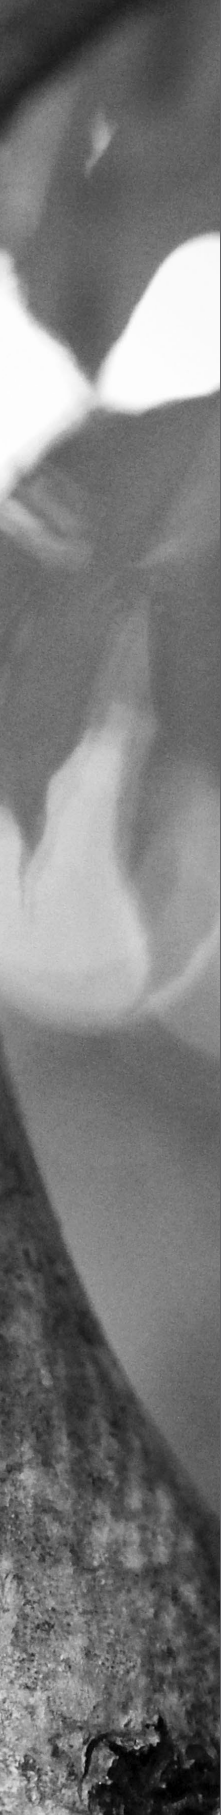
About the author of this thesis

Stefan M. Wierda was born in Steenwijk in 1984. From 1996 till 2002, he studied Atheneum at the RSG Steenwijk. In 2002, Stefan began studying Artificial Intelligence at the University of Groningen.

During this period, he was an active board member of the Groninger Studentenbond (GSb), Cognitievereniging (Cover), Centre for Information Sciences, and the council of the Faculty of Behavioral and Social Sciences. After receiving his bachelor degree, he studied Human Machine Communication at the University of Groningen. For his master thesis, he worked on a research project on the learning of robust and flexible skills at the Carnegie Mellon University in Pittsburgh (USA) for six months.

After graduating as a Master of Science in 2009, he started working on a PhD project at the Neuroimaging Center of the University Medical Center Groningen, which resulted in the current thesis. Supervised by Sander Martens, Hedderik van Rijn, and Niels Taatgen, he studied the constraints of temporal attention, and showed that the attentional blink phenomenon is best explained as a result of strategic restrictions in temporal attention rather than being caused by a fundamental cognitive bottleneck.

Currently, Stefan is working at the Artificial Intelligence Department of the University of Groningen as a postdoctoral fellow, focusing on Theory of Mind and the transfer of cognitive skills.