The perceptual wink model of non-switching attentional blink tasks

Patrice Rusconi¹ · David E. Huber²

Published online: 13 October 2017
© Psychonomic Society, Inc. 2017

Abstract The attentional blink (AB) is a temporary deficit for a second target (T2) when that target appears after a first target (T1). Although sophisticated models have been developed to explain the substantial AB literature in isolation, the current study considers how the AB relates to perceptual dynamics more broadly. We show that the time-course of the AB is closely related to the time course of the transition from positive to negative repetition priming effects in perceptual identification. Many AB tasks involve a switch between a T1 defined in one manner and a T2 defined in a different manner. Other AB tasks are non-switching, with all targets belonging to the same well-known category (e.g., letter targets versus number distractors) or sharing the same perceptual feature. We propose that these non-switching AB tasks reflect perceptual habituation for the target-defining attribute; thus, a ‘perceptual wink’, with perception of one attribute (target identity) undisturbed while perception of another (target detection) is impaired. On this account, the immediate benefit following T1 (lag-1 sparing) reflects positive repetition priming and the subsequent deficit (the blink) reflects negative repetition priming for the realization that a target occurred. In developing the perceptual wink model, we extended the nROUSE model of perceptual priming to explain the results of two new experiments combining the AB and identity repetitions. This

establishes important connections between non-switching AB tasks and perceptual dynamics.

Keywords Attentional blink · Priming · Neural network modeling · Visual perception

Introduction

The attentional blink (AB; Chun & Potter, 1995; Raymond, Shapiro, & Arnell, 1992), is a temporary deficit in reporting the identity of a second target (T2) after presentation of a first target (T1), when the items are presented in rapid succession (e.g., 100 ms per item). It is one of the most reliable and well-studied tasks in the study of cognition, and a great deal of effort has gone into understanding the mechanisms underlying this task. In the current study, rather than focusing on the AB task as something to explain in isolation, we consider relations between the AB and other cognitive phenomena (e.g., perceptual priming), which leads us to propose a novel account of many AB tasks based on perceptual dynamics. After reviewing the literature that led us to develop this ‘perceptual wink’ model, we present the details of the model and then test the model with two experiments examining interactions between the AB and priming (both positive and negative priming). In the discussion, we present additional simulations of other results in the AB literature and discuss the relation between this model and other formal models of the AB.

Perceptual priming and the nROUSE model

It takes time to form a stable percept (Gorea, 2015). For example, images appear to blend together and overlap when they are presented very briefly in quick succession (Hogben & Di

Electronic supplementary material The online version of this article (<https://doi.org/10.3758/s13423-017-1385-6>) contains supplementary material, which is available to authorized users.

✉ David E. Huber
dehuber@psych.umass.edu

¹ School of Psychology, University of Surrey, AD Building, Guildford GU2 7XH, UK

² Department of Psychological and Brain Sciences, University of Massachusetts, 135 Hicks Way, Amherst, MA 01003, USA

Lollo, 1974). This blending also occurs for high-level attributes such as the orthography or meaning of words, a phenomenon referred to as ‘priming’ (Evetts & Humphreys, 1981). In some circumstances, immediate repetitions produce deficits rather than benefits (Humphreys, Besner, & Quinlan, 1988). The neural responding optimally with unknown sources of evidence (nROUSE) model (Huber & O’Reilly, 2003) assumes that priming deficits reflect neural habituation (Tsodyks & Markram, 1997). Habituation reduces source confusion from prior presentations, minimizing blending, but this produces a form of ‘repetition blindness’ (Kanwisher, 1987); lingering habituation makes it difficult to reactivate the same perceptual representation. The nROUSE model successfully explained many perceptual phenomena, including repetition and semantic priming of words (Huber, 2008), repetition priming of faces (Rieth & Huber, 2010), repetition priming of episodic recognition judgments (Huber, Clark, Curran, & Winkielman, 2008), semantic satiation (Tian & Huber, 2010), repetition priming of same/different judgments (Davelaar, Tian, Weidemann, & Huber, 2011), the priming of speeded affective valence responses (Irwin, Huber, & Winkielman, 2010), as well as the time course of evoked electrophysiological potentials for several of these paradigms (Huber, Tian, Curran, O’Reilly, & Woroch, 2008; Tian & Huber, 2013).

Figure 1 shows the repetition priming results of Huber (2008). In the target-primed condition, short duration primes increased accuracy whereas long duration primes decreased accuracy. However, in the foil-primed condition (i.e., a prime identical to the incorrect answer), this pattern was reversed. These results are compared to a baseline condition presenting an unrelated prime word (neither-primed), which produced a u-shaped deficit with prime duration. According to the nROUSE model, there is a rise and fall of neural activation in response to the prime word that is the inverse of the pattern seen in the neither primed condition: Prime activation reaches a peak level between 100 and 200 ms and this prime duration produces the largest deficit for an unrelated prime word, owing to inhibition between competing representations (i.e., a forward masking effect). At the same time, this rise and fall has important consequences for words that are identical to the prime. More specifically, lingering prime activation blends with a subsequent identical target or foil, with this activation boosting performance in the target primed condition, but driving performance below the chance level of 50% in the foil primed condition (i.e., the prime is mistaken for the target). For longer prime durations, habituation lessens this prime activation, reducing the overall priming effect. Even longer prime durations produce a priming deficit because it becomes difficult to reactivate perceptual representations in response to a repeated target word or when a choice word repeats the prime; lingering habituation more than offsets the benefit of any lingering activation from the prime. Simulating prime activation on a millisecond by millisecond basis, the

nROUSE model provided an accurate fit of all three of these conditions based on the rising and falling time course of prime activation.

The perceptual wink model

Chun and Potter (1995) presented a sequence of characters one at a time for 100 ms each, asking observers to report any letters in the sequence, but not numbers. The AB effect they reported is replotted in Figure 1. As seen in the figure, the AB has the same time course as the neither-primed condition from Huber (2008), suggesting that the rise and fall of perceptual activation may underlie both effects. In the case of the Chun and Potter AB paradigm, one can imagine a ‘letter detector’ that is activated in response to targets, and a ‘number detector’ that is activated in response to distractors. With highly practiced categories such as these, category can be identified pre-attentively (automatically), resulting in visual pop-out for detection of the category regardless of the number of distractors (Schneider & Shiffrin, 1977). This notion of categorical pop-out supports the claim that these category representations (e.g., a number or letter detector) exist separate from identity representations (e.g., an ‘L’ or ‘M’ detector); that is, the fact that there can be categorical pop-out for practiced sets suggests that detecting that a character is a letter may be a separate process from identifying which letter it is. Thus, the benefit for T2 when it appears immediately after T1 (aka, lag-1 sparing) may reflect positive priming of the letter detector (if letters are targets), whereas the subsequent blink deficit may reflect negative priming for the letter detector. Because the letter detector and number detector are in competition with each other, the ISI between T1 and T2 reflects the rise and fall of the number detector (if numbers are distractors), explaining the similarity between the neither primed condition and the AB.

On this perceptual wink account (see Figure 2), performance in the AB task requires attentional processes for loading identifies into short-term memory (STM), but the blink itself is not caused by a deficit in these attentional processes. Instead, the blink is a failure of perception to trigger the requisite attentional response; i.e., a failure to identify that a second target occurred because ‘targetness’ (i.e., whatever it is that defines targets as different from distractors) is habituated at the time that T2 occurs. Even though the identity of T2 is undisturbed, the identity of T2 is not loaded into memory because the observer fails to identify that T2 was in fact a target. Thus, the attentional blink can instead be thought of as a ‘perceptual wink’: during the deficit, the identity of T2 is seen (one eye open), while the target-defining feature is missed (one eye closed). This proposal finds support from a recent finding that while the AB paradigm affects conscious awareness of T2 (i.e., a failure to realize that there was a T2),

Fig. 1 A comparison between the word priming results of Huber (2008), Experiment 1, and the attentional blink (AB) results of Chun and Potter (1995), Experiment 1. For word priming, as tested with two-alternative forced choice (2AFC) testing, three conditions are plotted, depending on whether the target, foil, or neither choice word was a repetition of the prime. The prime was doubled-up to avoid presenting the prime in exactly the same position as the target, although similar results are found with a single lower-case prime. The AB task presented a sequence of numbers, shown for 100 ms each, with two target letters embedded in the sequence. The plotted AB results show the probability of reporting the second letter

(T2) given successful report of the first letter (T1). As seen in the figure, the neither primed condition has the same time course as the T1/T2 inter-stimulus interval (ISI) for the AB. The nROUSE model accurately fits all three priming conditions based on increasing prime activation for short prime durations followed by decreasing prime activation for long prime durations, owing to perceptual habituation. The question considered in the current study is whether this transition from positive to negative priming explains the transition from a lag-1 sparing benefit to a blink in the AB task.

perceptual integration of T2 is intact during the blink (Fahrenfort, van Leeuwen, Olivers, & Hogendoorn, 2017).

Here, we briefly describe the perceptual wink model's behavior as applied to the current experiments, which utilize upper-case targets and lower-case distractors. The complete mathematical details of the model are reported in Appendix A. The loading of letter identities into STM is dictated by the current activation of an attentional 'gate node', which is driven by the task demands of searching the sequence for upper-case letters while ignoring lower-case letters, implemented through +1 and -1 weights to upper- (the target detector) and lower-case (the distractor detector) nodes, respectively. Critically, the gate node does not habituate, and its moment by moment activation value is purely driven by the fluctuating dynamics of the target and distractor detector nodes. Thus, although attention experiences a transient deficit (a blink),

the cause of this blink is a transient deficit in perception that a target occurred. For this task, T1 activates the upper-case node and this activation can carry over to the time when T2 is presented, if T2 immediately follows T1. Thus, lag-1 sparing is a positive priming effect because the target detector (the upper-case node in Fig. 2) is pre-activated by the presentation of T1. If a distractor or blank screen appears after T1, the target detector loses its activation, eliminating the priming advantage. By itself, this loss of activation does not cause a blink. However, habituation makes it difficult to reactivate the upper-case node in response to T2, so there is a repetition blindness to detect that a second target occurred (this is not the same as identity repetition blindness—see Chun, 1997b). Because the observer experiences difficulty detecting that the T2 letter was an upper-case letter, the attentional gate node does not become fully active, and the identity of T2 is not

Fig. 2 The perceptual wink model as applied to an AB task with upper-case targets and lower-case distractors. Each circle in the figure is a 'node', describing average activity of many neurons with similar inputs and outputs. All nodes contained within each rectangle inhibit each other, producing masking/interference effects. Four free parameters were used in fitting the results of Experiment 1 and those parameters were then fixed to make predictions. Dotted trapezoids indicate comparison processes, with winner-take-all (WTA) divisive-normalization from letter identity

to short-term memory ($\delta = 1$ for the most active letter identity and 0 for other identities) and linear divisive-normalization (Luce choice) from short-term memory (STM) to response probabilities. STM encoding strength is dictated by α , which is a non-habituating 'gate node' driven by the difference between the upper and lower-case nodes. For a standard blink, the lower-case node is habituated by the first target, and is difficult to re-activate (a perceptual wink), resulting in weak encoding into STM of the second target's identity (even though that identity is highly active).

loaded into STM as fully as it otherwise would even though that identity (the corresponding Letter Identity node in Fig. 2) is highly active.

In summary, the perceptual wink model explains the attentional blink as reflecting the multifaceted nature of perception, considering the separate perceptual representations that support the identification of a visually presented character (e.g., which letter?) versus detection of the target-defining attribute (e.g., is it upper-case?). In the model, performance is determined by the degree to which each identity is loaded into STM; if T2 target detection is poor, then the identity of T2 is weakly loaded into STM even if that identity is highly active (e.g., the perceptual system knows which letter it is, but remains unsure if the letter is upper-case). In absolute terms, this is a modest effect, but because T2 is brief and followed by a mask (i.e., the next character), the magnitude of the identity of T2 in STM may be weakened to the point that the observer reports something other than T2. Critically, if T2 is not followed by a mask, the identity of T2 in STM accumulates to a high degree owing to an extended duration of target detection. Indeed, elimination of the post-T2 mask eliminates the attentional blink by virtue of a ceiling effect (Giesbrecht & Di Lollo, 1998). If instead performance is limited by a mask that visually overlaps with T2 (i.e., integration masking) rather than post-T2 masking, then the limit on performance is not the degree of accumulation for the identity of T2 in STM, but rather whether the identity that is being accumulated is the correct identity (e.g., the mask visually corrupts the stimulus, resulting in report of a visually similar letter). Indeed, integration masking appears to

eliminate the attentional blink even though T2 performance remains below ceiling (Giesbrecht & Di Lollo, 1998). These different T2 masking effects highlight the distinction between perceptual target detection and perceptual target identification, with the perceptual wink model assuming that the attentional blink reflects a temporary deficit in target detection but not target identification.

Switching versus non-switching attentional blink tasks

The original task termed an 'attentional blink' was developed by Raymond et al. (1992). In their task, observers reported a white letter (T1) followed by a switch to monitoring for a black X (T2). The perceptual wink model appears to be at odds with this AB task, as well as with other AB tasks that involved a switch from searching for a T1 defined in one manner to a T2 defined in a different manner. Without a consistent target feature or target category, a deficit for T2 cannot reflect perceptual habituation to detect a target. However, we consider the possibility that there are different variants of attentional blink tasks: (1) switching, whenever T1 and T2 are defined differently (e.g., a white letter followed by a black X) or appear in different spatial locations; and (2) non-switching, whenever T1 and T2, both appearing in the same location, are examples of the same class of stimuli, where class might be defined categorically (e.g., letter targets versus number distractors) or based on visual features (e.g., white letter targets versus black letter distractors). As detailed above, perceptual priming exhibits the same time course as the AB and it is

notable that task-switching phenomena, such as the psychological refractory period (Pashler, 1994), also exhibit this same time course. Because perceptual priming, task switching, and the attentional blink, have existed as three separate literatures, separate formal theories have been developed to explain each respective literature. Given the similar time course of transient deficits in each of these literatures, we consider a more parsimonious account in which perceptual dynamics explain perceptual effects as well as non-switching AB tasks while task-switching dynamics explain task-switching effects, as well as switching AB tasks (support for the switching/non-switching distinction is reviewed in the general discussion).

In the current study, we take an important first step in this division of the AB literature by extending the nROUSE model of perceptual dynamics to non-switching AB tasks. The success of the nROUSE model in explaining AB effects will be of interest to AB researchers, and in the general discussion we review and contrast this account with other AB models. Our contribution to the literature is the demonstration that a model of perceptual dynamics can explain non-switching AB tasks. Thus, the ‘playing field’ for model comparison is greatly enlarged, challenging other AB models to explain the results of perceptual paradigms. While proponents of other AB models may cite the inability of the perceptual wink to explain switching AB tasks, we equivalently cite the inability of existing AB models to explain the wide variety of perceptual paradigms also explained by the nROUSE model. The question then becomes where to place the theoretical divide; should it be placed between switching versus non-switching AB tasks, or between all tasks that have been given the label of ‘attentional blink’ versus perceptual paradigms that reveal the same biphasic time course as seen in the AB?

The final form of a switching AB model should be an extension of the already existing formal theories of switching dynamics (Reeves & Sperling, 1986). Indeed, the extension of the threaded cognition model to the attentional blink is an example of taking this approach to the AB literature (Taategen, Juvina, Schipper, Borst, & Martens, 2009). Because the threaded cognition model (Salvucci & Taategen, 2008) was originally developed to explain the psychological refractory period and other task-switching effects, it made specific predictions for the manner in which the AB should change when including a secondary task. In a similar manner, by extending a model of perceptual dynamics to the AB literature, the perceptual wink model makes specific predictions for the manner in which the AB should change with perceptual priming.

Overview of the current study

The current study is a first step in relating the AB literature to the perceptual dynamics literature, asking whether an existing

model of perceptual dynamics can be extended to explain non-switching AB tasks. If perceptual deficits underlie non-switching AB tasks, the dynamics that explain priming should explain the AB. We tested the perceptual wink model in two new AB experiments with upper-case target letters embedded in a sequence of lower-case distractor letters. We assume that the perception of each letter is broken into perception of identity (regardless of visual form) and perception of case (see Fig. 2). Because these are both orthographic representations, the corresponding identification dynamics were set to those that best-fit the word priming results.

Beyond demonstrating the adequacy of the model in handling the basic blink pattern, the perceptual wink account makes specific predictions when combining the AB with repetition priming. More specifically, because T1 primes target detection, which is the assumed cause of lag-1 sparing, the post-T1 item is encoded into STM even if that item is a distractor (i.e., even if it is a lower-case letter). However, if the post-T1 item happens to have the same identity as T2, this will increase T2 accuracy (pre-T2 priming), and if that item happens to have the same identity as T1, this will increase T1 accuracy (post-T1 priming). In contrast, placement of a prime in other positions within the stream will not produce these effects owing to a lack of target detection. These predictions were tested in Experiment 1. The perceptual wink model also predicts that with sufficient exposure to a particular letter, the ability to identify that letter will be reduced owing to habituation. We tested this in Experiment 2 by including a condition that used a single repeated distractor letter, predicting that this would reduce the magnitude of the blink (i.e., repetition blindness for the distractor makes it easier to identify and detect targets).

No other AB theory assumes a perceptual basis for the blink (for review, see the general discussion). In the perceptual wink model, the attentional gate node rises and falls with T1-T2 lag (the α node in Fig. 2), but this is caused by a rise and fall in perception (the upper and lower-case nodes in Fig. 2). Unlike other theories, this account makes specific predictions for the relation between the AB and perceptual priming. Priming has been examined in many AB experiments (for review, see the general discussion), but our design is unique, providing a comprehensive exploration of priming of targets (Experiment 1) and priming of distractors (Experiment 2) as a function of lag. The question posed here is whether the perceptual dynamics that explain positive priming of the target (Experiment 1) and negative priming of a constantly repeated distractor (Experiment 2) can also explain the time course of the AB. Other models of the AB could possibly explain these results if augmented with positive and negative priming mechanisms. In contrast, the perceptual wink model is highly constrained to explain the rise and fall of priming and the rise and fall of the AB with the same perceptual dynamics (albeit as applied to different perceptual representations).

Experiment 1: T1/T2 repetition priming and the AB

Method

Participants

Thirty undergraduate students (five male) from the University of Milano-Bicocca participated in the study in exchange for course credit. They ranged in age from 19 years to 48 years ($M = 24.2$ years, $SD = 6.8$). This sample size was chosen based on prior AB experiments reported in the literature and based on pilot versions of this experiment.

Apparatus

The experiment was programmed and run using E-Prime™ (Psychology Software Tools, Pittsburgh, PA), on PCs with CRT monitors and 75 Hz display refresh rate.

Materials and Design

All Rapid Serial Visual Presentation (RSVP) streams contained two targets, which were upper-case letters, and 12 distractors, which were lower-case letters. Targets and distractors were selected from the same pool of 14 letters (“a”, “b”, “d”, “e”, “g”, “h”, “m”, “n”, “q”, “r”, “j”, “i”, “f”, “t”). The font was bold Courier New, size 18, in black against a gray background.

There were 50 conditions in a fully within-subjects design, representing all combinations of trial type (standard blink, post-T1 priming, or pre-T2 priming), T1-T2 lag (2, 3, or 6, plus lag 1 only for the standard blink trial type), and T1 position within the sequence (3–7). T1 position was varied so participants would not know when to expect the first target. The reported results collapsed over this variable, producing ten conditions of interest, representing the combinations of trial type and T1–T2 lag (see Fig. 3). For the standard blink trial type, none of the distractors were the same letter as the target. For the post-T1 priming trials type, the distractor immediately following T1 was a lower-case version of T1. Finally, for the pre-T2 priming trial type, the distractor immediately before T2 was a lower-case version of T2. The 50 conditions were presented three times each in random order across the 150 trials of the experiment.

Procedure

Participants were tested individually in an experimental session that lasted approximately 14 min. Each stream of items began with a fixation point “+”, which appeared for 400 ms, followed by 14 items presented for 100 ms each. Finally, a mask “####” appeared for 200 ms. Participants were fully informed that there were two targets on every trial and they were

instructed to guess the target identities in any order. If their second response was the same as their first response, they received the instructions: “You were wrong: You typed in the same upper-case letter”, and they were then required to provide a different response for their second response. Before the experimental trials, there were 15 practice trials based on the standard blink trial type. For these 15 practice trials, lags of 1, 3, and 6, were presented using each of the five different T1 positions.

Results

The key accuracy measure is shown in Fig. 3, which plots the probability of providing T2 as one of the two responses given that T1 was the other response. A within-subjects analysis of variance (ANOVA) with trial type (standard blink, post-T1 priming, pre-T2 priming) and lag (2, 3, and 6) as factors showed a two-way interaction, $F(4,116) = 4.68$, $P = .002$, $\eta^2 = .029$, indicating that priming effects differed as a function of lag, warranting further statistical tests probing the nature of these priming differences. In addition, there was a significant main effect of trial type, $F(2,58) = 5.99$, $P = .004$, $\eta^2 = .018$, with accuracy lower in the standard blink trials ($M = .5$, $SDE = .02$) than the post-T1 priming trials ($M = .56$, $SDE = .03$), $P = .006$, and the pre-T2 priming trials ($M = .56$, $SDE = .03$), $P = .001$. Collapsing across lags, there was no significant difference between post-T1 priming and pre-T2 priming trials, $P = .829$. There was also a significant main effect of lag, $F(2,58) = 54.76$, $P = .001$, $\eta^2 = .449$. Accuracy was higher at lag 6 ($M = .73$, $SDE = .02$) than lag 2 ($M = .48$, $SDE = .03$) or lag 3 ($M = .42$, $SDE = .04$), all $ps < .001$. There was a significant difference in performance between lags 2 and 3, $P = .033$.

The pre-T2 priming advantage (i.e., higher accuracy for pre-T2 than standard blink) when collapsing across lag is qualified by a significant interaction between lag and the pre-T2 versus standard blink trial types, $F(2, 58) = 8.20$, $P = .001$, $\eta^2 = .042$. More specifically, there was a large pre-T2 priming effect at lag 2 ($M_{pre-T2-standard} = .17$, $SD_{pre-T2-standard} = .19$), $t(29) = 5.09$, $P < .001$, $d = .93$, but no reliable pre-T2 priming effects at lag 3 ($M_{pre-T2-standard} = -.02$, $SD_{pre-T2-standard} = .19$), $t(29) = -.59$, $P = .563$, $d = -.11$, or lag 6 ($M_{pre-T2-standard} = .03$, $SD_{pre-T2-standard} = .17$), $t(29) = .86$, $P = .395$, $d = .16$. In contrast to pre-T2 priming, the interaction between lag and the post-T1 versus standard blink trial types was unreliable, $F(2, 58) = 2.78$, $P = .070$, $\eta^2 = .013$, suggesting that the post-T1 priming advantage occurred regardless of lag (despite what appears to be a selective absence of post-T1 priming only at lag 3).

In summary, the observed data revealed significant post-T1 priming that did not interact with lag and pre-T2 priming only at lag 2. As seen in Fig. 3, the nROUSE model accounted for both of these patterns. With four free parameters, the nROUSE model was fit to the four joint probability values

Fig. 3 Experimental paradigm, modeling results with best-fit parameters, and observed data for Experiment 1. For each of the three trial types at each T1-T2 lag, the graph shows the probability of identifying the second target (T2) given that the first target (T1) was identified. The experimental paradigm (overlaid top box) shows the three trials types with example lag 2 trials. With four free parameters, the nROUSE model was fit to the 40 joint probability of identifying both T1 and T2, only T1, only T2, or neither T1 nor T2 as seen in Fig. 4. These joint probabilities were then

used to calculate the conditional probabilities shown in this figure. Error bars are 95% confidence intervals. According to the perceptual wink account, Pre-T2 priming occurred only at lag-2 because at this lag the T2 prime (e.g., lower-case j) is encoded into memory owing to upper-case priming from T1. Furthermore, Post-T1 priming occurred at all lags because post-T1 encoding of the prime (e.g., lower-case h) eliminates interference that would have occurred in the standard blink condition owing to post-T1 encoding of the post-T1 distractor (e.g., lower-case m).

(identification of T1 and T2, only T1, only T2, and neither T1 nor T2), for each of the ten conditions. As seen in Fig. 4, the model captured the data across these 40 joint probabilities (95.6% of the variance accounted for).

Experiment 2: Distractor repetition priming and the AB

Method

Except as noted, all methods were identical to Experiment 1. Experiment 2 examined repetition priming of the distractors by comparing a standard blink distractor type to one where the same distractor was used throughout the RSVP sequence (the repeated distractor trial type). In addition, this experiment manipulated the visual similarity of the distractors as compared to the targets, with targets and distractors drawn from a pool of either curvy or straight letters (see Appendix Table B1). As reported in the Appendix B analyses, performance was better when distractors were dissimilar from targets (e.g., curvy distractors paired with straight targets), but this effect was weak ($\eta^2 = .009$, corresponding to a change of 4%), and did not interact with any of the other variables. Therefore, we collapsed over distractor similarity.

Participants

Forty-three undergraduate students (eight males), ranging in age from 19 years to 41 years ($M = 22.7$ years, $SD = 3.5$) participated in exchange for course credit. However, only 41 participants provided a full data set, with one lost because of a black-out during the experiment, and another lost because of a corrupted data file.

Materials and design

Targets and distractors were drawn from the pools of letters reported in Table 1. The main factors of the design were distractor similarity (similar vs. dissimilar), and distractor type (repeated distractor vs. standard blink). A particular level of distractor similarity was created through two combinations of different letters types. For instance, the similar distractor type consisted of curvy targets paired with curvy distractors, as well as straight targets paired with straight distractors (see Appendix Table B2 for the full design). Unlike Experiment 1, all factors were tested at lag 1, 3, and 6. As in Experiment 1, T1 position varied from 3rd to 7th. Combining all factors, there were 120 unique conditions, which occurred three times across 360 experimental trials.

Fig. 4 Modeling results and observed joint probabilities for Experiment 1. With four free parameters, the model captured 95.6% of the variance. As seen in the figure, the model and data show a marked difference in Pre-T2 priming as compared to the standard blink. According to the model, this occurs only at lag-2 because only in this condition does upper-case priming from T1 cause encoding of the T2 prime. In contrast, post-T1 priming produced an overall upward shift in the probabilities of responding with both T1 and T2 across all lags. According to the model, this occurred because of a reduction in the probability of giving a response other than T1 or T2 (more specifically, the elimination of responding with the post-T1 distractor, such as occurs in the standard blink condition).

Procedure

Participants were tested individually in an experimental session that lasted approximately 20 min. Unlike Experiment 1, each item appeared for 67 ms, with blank screens of 25 ms occurring between each item (i.e., stimulus onset asynchrony = 92 ms, as in Olivers & Meeter, 2008). The 15 practice trials included a sample from all manipulations (i.e., during

practice, participants experienced all levels of all factors, but not all combinations of all levels).

Results

The probability of providing T2 as one of the two responses given that T1 was the other response is shown in Fig. 5, with the results broken down by lag and distractor type. A within subjects analysis of variance of these conditional probabilities, with distractor type (repeated distractor vs. standard blink) and lag (1, 3, and 6) showed a main effect of distractor type, $F(1,40) = 171.84$, $P < .001$, $\eta^2 = .297$, with greater accuracy for repeated distractor trials ($M = .72$, $SDE = .03$) than standard blink trials ($M = .5$, $SDE = .02$). There was also a main effect of lag, $F(2,80) = 101.5$, $P < .001$, $\eta^2 = .366$, with higher accuracy at lag 1 ($M = .73$, $SDE = .02$) than lag 6 ($M = .64$, $SDE = .03$) and lag 3 ($M = .45$, $SDE = .02$), all P s $< .001$. Performance at lag 3 was significantly lower than lag 6, $P < .001$. Finally, there was a significant distractor type by lag interaction, $F(2,80) = 40.31$, $P < .001$, $\eta^2 = .062$. For the repeated distractor trials, there was no reliable difference between performance at lag 1 ($M = .77$, $SDE = .03$) versus lag 6 ($M = .79$, $SDE = .03$), $P = .340$ whereas for the standard blink trials, accuracy was greater at lag 1 ($M = .69$, $SDE = .02$) than lag 6 ($M = .49$, $SDE = .03$).

As seen in Fig. 5, the nROUSE model predicted these results without any free parameters. As seen in Fig. S1 in the Supplementary material, when freeing up the same four parameters as used in Experiment 1, the model captured the data across the 24 joint probabilities, including the observation of greater T2 accuracy than T1 in the lag-1 condition (this T2 advantage was marginal in Experiment 1, but robustly present in Experiment 2, and in general it is found in the AB literature).

Discussion

In the perceptual wink model, the same perceptual dynamics underlie repetition priming and the AB, with these effects differing only because repetition priming affects identity, whereas a prior target affects detection that a target occurred (perception of upper-case in the current situation). This phenomenon is described as a wink because one perceptual attribute exhibits a deficit (e.g., the ability to detect that T2 was upper-case) while other perceptual representations are undisturbed (e.g., the ability to identify the T2 letter). Thus, the blink in these experiments reflects positive priming (lag-1 sparing) and negative priming (the blink) for the perception of letter-case. Lingering activation (i.e., positive priming) for the target attribute of letter-case causes the encoding into short-term memory of anything appearing immediately after T1, explaining lag-1 sparing and interactions with primes placed in the post-T1 position. However, after an intervening

Fig. 5 Experimental paradigm, *a priori* model predictions (no free parameters, using the parameter values determined by Experiment 1), and observed data for Experiment 2. For each of the two distractor types at each T1-T2 lag, the graph shows the probability of identifying the second target (T2) given that the first target (T1) was identified. The experimental paradigm (overlaid *top box*) shows the two distractor types with example lag 2 trials (although the experiment did not include lag 2). Blank screens were inserted between letters to give an appearance of a sequence even for the repeated distractor condition. The data and model predictions show a reduced difference between the distractor type

conditions only at lag-1. According to the model, T2 encoding was boosted in the lag-1 condition because the presentation of T1 primes the upper-case node, benefiting the post-T1 item. With the standard blink, even for a recovered upper-case node at lag 6, T2 encoding does not achieve the same degree as compared to upper-case priming at lag-1. However, with a repeated distractor, T2 encoding is high even in the absence of upper-case priming because of a relative lack of lower-case activation owing to an habituated visual response to the repeated distractor. Error bars are 95% confidence intervals.

distractor (or a blank screen), this positive priming of upper-case detection is lost, and lingering habituation for upper-case detection produces negative priming and thus the AB deficit (see also Fig. S4).

In developing the perceptual wink model, the nROUSE model of perceptual priming was extended to the AB paradigm. Using perceptual dynamics fixed by the word priming results, the model explains the basic blink pattern and interactions between the blink and priming of targets (Experiment 1) and priming of distractors (Experiment 2). Critically, both kinds of priming produced accuracy benefits, rather than deficits. In other words, priming of distractors was in truth a negative priming effect in the sense that the repeated presentation of the same distractor caused habituation for that letter identity, making it easier to identify the target letter identities and easier to detect when a target occurred. If priming of distractors had been a positive effect in the sense of boosting distractor perception, then a repeated distractor would have been a stronger competitor to targets, and performance would have been lower rather than higher. In the perceptual wink model, the same habituation dynamics that produced this benefit from priming of distractors also produced the AB deficit and the benefit from priming of targets. The success of the

model supports the claim that the root cause of the blink in non-switching AB tasks is perception, which in turn affects attention.

Upper-case priming for whatever appears in the lag-1 position

According to the perceptual wink model, at the time of any lag-1 presentation, the target detector (upper-case) is still active, even if the lag-1 item is a distractor (lower-case). This lingering activity for the upper-case detector triggers encoding of the post-T1 item regardless of its identity. Upper-case priming explains lag-1 sparing in the standard blink and it also explains: (1) a pre-T2 priming benefit that only occurred in the lag-2 condition; (2) a post-T1 priming benefit that was invariant with lag; and (3) the asymmetry in the degree of recovery from the blink when comparing the standard blink condition to the repeated distractor condition. Each of these effects was explained by the model as emerging from complex dynamics. Next, we attempt intuitive explanations of these three effects based on our analyses of the model's behavior (see also Figs. S4–S8).

Regarding pre-T2 priming, the benefit of priming is by virtue of a mistake. Only for lag-2 is the pre-T2 prime placed in the post-T1 position. Thus, only in this condition is the prime encoded into STM, boosting T2 performance. In other words, the observer might give an answer of T2 not because of accurate encoding of T2, but rather because of the mistaken encoding of the prime (a prime that happens to have the same identity as T2).

Regarding post-T1 priming, the benefit of priming is through reduced interference. In the standard blink condition, the post-T1 distractor is partially encoded, causing response interference (increasing the probability of giving a response other than T1 or T2). However, with post-T1 priming, the post-T1 item has the same identity as T1, eliminating this response interference (increasing the probability of responding with both T1 and T2 by eliminating the probability of responding with the post-T1 distractor).

Regarding the priming of distractors, the benefit of priming is again through reduced interference, but in this case it is a twofold effect. First, because the repeated distractor's identity is habituated, there is overall less response interference, boosting performance in general. Second, the recovery from the blink appears to be more complete (lag-6 performance is equal to lag-1 performance), whereas this is not true for the standard blink. According to the model, recovery from the blink is nearly complete at lag-6 in the standard blink. Nonetheless, lag-1 performance is better than lag-6 performance because in the lag-1 condition, T2 has an extra benefit from upper-case priming (i.e., lag-1 sparing is more than just preserving performance—it is a net benefit). This upper-case priming also occurs in the repeated distractor condition, but T2 encoding is already at ceiling regardless of this boost considering that the habituated distractor produces minimal lower-case detector activation (i.e., it is easy to achieve upper-case activation even without upper-case priming from T1).

A comparison with other priming-AB studies

These are not the first experiments to investigate interactions between priming and the AB. Indeed, the first example of priming in an AB task predates the adoption of the term 'attentional blink'. In the study of Broadbent and Broadbent (1987), upper-case target words were presented with lower-case distractor words in an RSVP stream, and semantic priming was examined by comparing target pairs that were unrelated to those that were semantic associates. No semantic priming benefit was found except in the specific circumstance of a failure to identify T1 in the lag-1 condition. In contrast, Maki, Frigen and Paulson (1997) found robust semantic priming between targets that was largely invariant with lag. Furthermore, they found semantic priming from a distractor

placed in the pre-T2 position, regardless of T1–T2 lag (in contrast to our Experiment 1). Additional experiments found semantic priming of T2 from a distractor placed in a variety of positions between T1 and T2 (both post-T1 and pre-T2). Although these results differ from our experiments, this may reflect the difference between semantic priming and repetition priming; according to the nROUSE model, while habituation of orthography readily eliminates priming or even produces negative repetition priming, semantic priming only produces benefits (Rieth & Huber, 2017). Thus, it is expected that semantic priming would produce benefits for a wider variety of positions within the RSVP stream.

Aside from using semantic priming rather than repetition priming, the Maki et al. (1997) study potentially differs from the current study by using a paradigm that might be a switching AB task; in their task, T1 was light green and T2 was light red. Thus, observers might have been searching for any colored word (a non-switch task), or they might first search for a red word and then switch to searching for a green word (a switch task). Indeed, most studies of priming in the AB have used switching AB tasks. For instance, Potter et al. (2005) examined semantic priming in an AB task that presented simultaneous upper/lower RSVP streams of distractor masks and target words. The two targets were always presented in different streams, so the T2 deficit may have reflected a spatial attention switch (e.g., the time needed to disengage from the T1 stream and switch spatial attention to the T2 stream). In any event, this paradigm produced results similar to Maki et al. (1997), revealing semantic priming of T2 across a range of delays between the T1 prime and the T2 target. In addition, Maki et al. (1997) examined backward priming, finding that a semantically related T2 produced a boost to T1 accuracy only when T2 immediately followed T1 (see Rieth & Huber, 2017, for an application of the nROUSE model to both forward and backward semantic priming).

Repetition priming has been examined in AB tasks, although most studies have used a switching paradigm. Shapiro et al. (1997) had people report a white digit (T1), followed by switch to a lag-3 upper-case black letter (T2), and then another switch to a lag-6 (relative to T1) lower-case black letter (T3). Observers were told to report not only the identity of these subsequent targets, but also the letter case, and half of the time the T3 letter was the same identity as T2 (i.e., repetition priming). In keeping with the nROUSE model and the repetition priming results of Huber et al. (2002), they found a negative repetition priming effect of T2 on T3 when T2 was identified, but a positive repetition priming effect when T2 was not identified. Thus, the observed identity priming effects were fully compatible with the perceptual dynamics of the perceptual wink model. A second experiment used words instead of letters, with T1 being the only white word, T2 the only red word, and T3 the only yellow word (distractors were black words). In keeping with the

nROUSE model when comparing semantic priming and repetition priming, this study produced only positive priming despite using the same lags as the repetition priming experiment.

The study of Chua, Goh and Hon (2001) used a switching AB task, examining the effect of letter repetition priming (differing in letter case), with a prime of T2 placed in a variety of different locations. In a series of six experiments, they found sizable priming only when the prime was in the lag-1 position. When the prime was prior to T1, in a later lag, or in the lag-1 position when using a task that did not require a response to T1, priming was eliminated. Thus, it appears that some level of attention to the prime is needed for repetition priming (i.e., presentation of a task relevant T1 produces greater processing of the prime in the lag-1 position, producing greater priming). This finding is similar to results of Experiment 1 in which the pre-T2 prime was only effective when placed in the lag-1 position (i.e., the lag-2 condition) whereas the post-T1 prime was effective regardless of lag (although note that in the current experiment, the post-T1 prime was a prime of T1 rather than T2).

Similar to the study of Chua et al. (2001), Dux and Marois (2008) examined repetition priming from letter distractors presented during the blink. It is not entirely clear whether their paradigm was switching or non-switching considering that distractors were white letters whereas the two targets were colored, but of different colors from each other. In any case, the prime always appeared in the lag-2 position and they examined T2 performance at lag-4 versus lag-10 to measure the AB. On average, they found small positive priming at lag-4 and no priming at lag-10, but analyses of individual differences revealed that some observers produced negative priming at lag-4, whereas others produced large positive priming at this lag. Furthermore, these individual differences reliably correlated with the AB magnitude (those with smaller ABs had negative priming) and T1 performance (those with high T1 accuracy had negative priming). The authors viewed attention as the causal factor of these individual differences, with differences in distractor inhibition during the blink causing priming differences (however, see Elliott & Giesbrecht, 2015 for a failure to replicate this correlation result and a computer simulation suggesting that the previously reported correlation may have been an artifact when using the same data twice). If this correlation exists, the perceptual wink model provides a qualitatively different account, with differences in perception as the causal factor: individuals with more rapid orthographic perception more easily perceive T1, more readily experience letter habituation (i.e., negative priming) and more rapidly recover the ability to detect the existence of a second target (i.e., revealed as a small AB magnitude). The results of Slagter and Georgopoulou (2013) support this alternative explanation. Their study replicated Dux and Marois, but also included a lag-2 condition. They found that priming correlated

with the blink magnitude when comparing lag-4 to lag-10, but not when comparing lag-2 to lag-10, suggesting that these individual differences reflect the speed of AB recovery rather than blink magnitude.

The priming study of Chun (1997b) is highly relevant to the current experiments. In a series of experiments, Chun found a double dissociation between the AB and repetition blindness (i.e., negative repetition priming). Like the current experiments, these were non-switching AB tasks (e.g., letter targets with digit or symbol distractors) and, like the current experiments, repetition priming was introduced through a case change. However, unlike the current experiments, priming was between T1 and T2 rather than from a distractor. This study found that repetition blindness followed a faster time course than the AB (compare for instance the target-primed condition of Figure 1 to the AB results in Figure 1). Furthermore, some manipulations alleviated repetition blindness but not the AB (a condition in which the two targets differed in color), while other manipulations alleviated the AB but not repetition blindness (searching for letter targets with symbol distractors). In the perceptual wink model, these dissociations reflect the difference between the perceptual representation of identity (which underlie repetition blindness) versus the perceptual representation of targetness (which underlie the attentional blink). Furthermore, the AB has a slightly different time course than repetition blindness considering that the repetition blindness condition entails all of the deficit from habituated target detection, as well as the deficit from habituated visual line segments of the repeated character, providing a weakened input to target detection. In a study currently in preparation, we replicated Chun's findings while using forced choice testing to rule out alternative explanations based on strategic guessing. Application of the perceptual wink model demonstrated that the different time course for each effect can be explained by the same perceptual dynamics.

In summary, most priming studies of the AB used switching AB tasks, and the perceptual wink model might not apply to the observed AB deficits. Nevertheless, the observed priming effects were largely consistent with priming studies in the perception literature. Given the success of the nROUSE model in explaining perceptual priming effects, it is likely that the perceptual wink model, which incorporates nROUSE's perceptual dynamics, could accommodate many of these interactions between priming and the AB.

Applying the perceptual wink model to other AB results

The AB literature is large and cannot be addressed in its entirety in the current paper, although we consider some of the more prominent findings. A key discovery was the 'spread of sparing' that occurs with multiple back-to-back targets (Di

Lollo, Kawahara, Ghorashi, & Enns, 2005). This is naturally explained by the perceptual wink model as reflecting continued priming of the target detector, which maintains STM encoding with an uninterrupted series of targets. After fully developing the perceptual wink model and applying it to Experiment 1, we learned of a study reporting that the spread of sparing reverses with the insertion of a blank screen (Chen & Zhou, 2015). If the model can account for these results, this would simultaneously demonstrate that the model handles the spread of sparing, as well as the challenging result that the spread of sparing is undone by insertion of a blank screen.

Chen and Zhou (2015) used a bare-bones version of the AB, in which the task was to report numbers rather than letters. Rather than placing the target numbers in a long sequence, there were just two targets, with a single letter distractor after T1 and a single letter distractor after T2. Rather than manipulating lag by changing the relative position of T2 in the sequence, a blank screen of different durations was presented after the post-T1 distractor, prior to T2 (see Fig. 6). This target-distractor-target (TDT) condition was run in one block of trials and observers knew to attempt report of two target numbers. In a different block of trials, the procedure was modified by replacing the intervening distractor with an intervening target instead. For this target-target-target (TTT) condition, observers knew to attempt report of three numbers. In applying the perceptual wink model to the results of this study, the simulation did not use any free parameters, fixing all parameters to the values dictated by the fit of Experiment 1. As seen in Figure 6, the *a priori* prediction (*a priori* in the sense that the model was developed and parameterized without knowing of this result) of the perceptual wink model was qualitatively accurate, capturing both the spread of sparing (TTT > TDT) without a blank screen as well the reversal of this effect (TTT < TDT) with the insertion of a blank screen. As reported in Fig. S2, the model quantitatively captures these results if allowed the same four free parameters as in Experiment 1.

The model produces a spread of sparing owing to continued priming of target detection, but this spread of sparing is reversed by insertion of a blank screen because target detection priming fades during the blank interval, revealing the underlying habituation for the ability to detect a target (see also Figs. S9–S10).

The study of Chen and Zhou (2015) presented each character only for 50 ms, rather than 100 ms, which is more typical in the AB literature. As such, the ability of the perceptual wink model to capture this result does not necessarily indicate that the perceptual wink model can explain the spread of sparing reported in other studies. To explore this situation, we again kept all parameters fixed to the values that best-fit Experiment 1, simulating the Chen and Zhou paradigm as if the procedures had presented each character for 100 ms rather than 50 ms. As reported in Fig. S3, the model's predictions were

nearly identical to those seen in Fig. 6 (i.e., the model produces a spread of sparing with 100 ms presentations). However, the model is not able to maintain the spread of sparing indefinitely (in keeping with prior studies), and the perceptual representation common to all targets will eventually fatigue and produce worse performance for a string of targets as compared to a situation that gives the target representation a rest by inserting distractors. To demonstrate this, we once again simulated the Chen and Zhou procedure, but this time we assumed that each character was presented for 200 ms. In this case, the spread of sparing was eliminated even without the insertion of a blank screen. In other words, after 400 ms of constant presentation of targets (i.e., the first two targets at 200 ms per target), the target detector was sufficiently habituated as to make it difficult to encode another target.

Beyond the spread of sparing, and its reversal with a blank screen, we also applied the perceptual wink model to other challenging results in the literature. For these simulations, the best-fitting parameter values from Experiment 1 produced similar results, although the results of Experiment 2 were more typical of the literature, producing a T2 advantage at lag-1 as compared to T1 accuracy (the parameter values from Experiment 2 indicate a more sluggish attentional gate node, which enhances the lag-1 boost). Thus, the parameters from Experiment 2 were used in the simulation results shown in Fig. 7 (henceforth, these will serve as 'default' parameters). The standard blink condition shown in Fig. 7 was based on the stimulus onset asynchrony experiment reported by Bowman and Wyble (2007), which involved eight different potential letter targets in a sequence of letter targets versus eight different potential digit distractors.

Bowman and Wyble (2007) found that the time course of the blink is invariant with the presentation rate. In other words, a blink that lasts for 600 ms is found regardless of whether that 600 ms is filled with five intervening distractors presented for 100 ms each versus 11 presented for 50 ms each. In the simulation shown in Fig. 7, the standard blink condition was run at 100 ms/item whereas the rate doubled condition was 50 ms/item. As such, a T1–T2 stimulus onset asynchrony of 100 in the rate doubled condition placed the T2 in the lag-2 position rather than lag-1. In line with the findings of Bowman and Wyble, the model produced lag-2 sparing in this rate doubled situation. The perceptual wink model produces lag-2 sparing because the letter detector is still active 100 ms after T1 (a 50-ms intervening distractor is insufficient to eliminate letter-detector priming).

Another challenging finding in the AB literature is that, in the midst of the blink, presentation of a distractor that shares an attribute with targets greatly reduces the blink magnitude. For instance, Nieuwenstein, Chun, van der Lubbe, & Hooge (2005) used an AB task with red number targets and black letter distractors. Thus, while the basic task was to report numbers but not letters, font color served a redundant cue. In

Fig. 6 Experimental paradigm, *a priori* model predictions (no free parameters, using the parameter values determined from Experiment 1), and observed data for the no noise conditions reported by Chen and Zhou (2015). These are *a priori* predictions in the sense that the perceptual wink model was developed and fit to Experiment 1 prior to learning of this study. As seen in the overlaid box showing the experimental paradigm, this experiment used a minimal version of the attentional blink with two (TDT) or three (TTT) target numbers and either one or two distractor letters. The final item was a distractor and the item between the first and last target was either a target or a distractor (in the TTT condition, the last target is labeled T2 for comparison to the TDT condition). A blank

interval of different durations was interposed between the intervening item and the final target. For each condition, the graph shows the probability of identifying the final target (T2) given that the first target (T1) was identified. The perceptual wink model explains the ‘spread of sparing’ (TTT > TDT) because the number detector is kept active (i.e., primed) by an uninterrupted series of targets. However, if interrupted by a blank screen, this activation is lost, revealing the underlying habituation of the number detector. This habituation is greater than a standard blink because of two rather than one prior targets, and performance in the TTT condition is worse than the TDT condition. Thus, the spread of sparing is reversed by insertion of a blank screen.

some conditions, T2 was pre-cued by a red letter distractor in the pre-T2 position. In the perceptual wink model, even habituated representations can be made active with sufficient input. Thus, pre-cueing the target detector should allow some degree of T2 encoding in the midst of the blink. To simulate this pre-cueing effect, the pre-T2 distractor was connected to the target detector node rather than the distractor node. This corresponds to the assumption that observers used color rather than number/letter to determine what to encode into short-term memory. As seen in Fig. 7, and as reported by Nieuwenstein et al. (2005), this pre-cue increased T2 performance as compared to the uncued situation (standard blink), particularly in the midst of the blink (e.g., stimulus onset asynchrony = 300).

A surprising result is that there is a similar time course of the blink regardless of whether the interval between T1 and T2 is filled with distractors or a blank screen (Nieuwenstein, Potter, & Theeuwes, 2009). In Fig. 7, the “no intervening distractors” condition was designed to mimic the conditions of Nieuwenstein et al. (2005), which compared a standard letter/number blink to one in which the distractors between

T1 and T2 were replaced with a blank screen. The perceptual wink model produces this effect because letter detector activation fades with a blank screen, revealing the underlying perceptual habituation (i.e., it is difficult to re-activate the letter-detector).

Finally, because the perceptual wink has separate perceptual representations for target detection and target identity, it could potentially explain distractor intrusions in AB tasks where targets are indicated by a separate visual feature (e.g., an outline box or circle). In these tasks, observers often report distractors preceding and following each target. Modifying the procedure developed by Chun (1997a), the study of Vul, Nieuwenstein and Kanwisher (2008) presented RSVP streams of letters (separated by blanks) where targets were indicated by a surrounding circle. At the end of each sequence, observers were asked to report each of the two targets in the order that they appeared. First-report (T1 question) intrusion distributions were analyzed separately from second-report (T2 question) distributions across ten different lags (see Fig. 8 for a replotting of their results). In the figure, item position is

Fig. 7 Application of the perceptual wink model to several other prominent results in the AB literature using the best-fitting parameter values from Experiment 2. Similar results are found with the best-fitting parameter values from Experiment 1, although the Experiment 2 values produce better T2 performance than T1 in the lag-1 condition, as is often found in the literature. The rate-doubled condition presented all stimuli for 50 ms rather than 100 ms and thus the 100 ms stimulus onset

asynchrony was a lag-2 situation (i.e., lag-2 sparing with rate doubling). The T2 pre-cued condition presented a stimulus in the pre-T2 position that carried the target defining attribute (i.e., priming the target detector ameliorates the blink). The no intervening distractors condition replaced all distractors between T1 and T2 with a blank screen (i.e., the blink does not require a distractor). See text for further details.

plotted relative to the position of T1. The colored curves show the distribution of intrusions and the black dashed lines highlight the T2 correct responses. Analyses revealed that in the midst of the blink, the second report distribution is suppressed (i.e., worse T2 accuracy during the blink), delayed (more intrusions after T2 than before), and diffused (i.e., intrusions across a wide range of positions, both before and after T2).

Rieth and Vul (2012) developed an ideal observer model of the second report data, although their model did not explain the first report data and did not capture the bimodal nature of the intrusion distribution for some lags (see for instance the second report distribution in the lag-4 condition, which shows a first mode at the lag-1 position and a second mode at lag-5). Indeed, a recent re-analysis of these data using a mixture model (Goodbourn et al., 2016) concluded that lags 1 and 2 were best captured by a single attentional episode (i.e., one that encompassed both targets) whereas lags 3 and above were best captured by separate attentional episodes (i.e., one centered on T1 and one centered on T2). Because this mixture modeling was descriptive, measuring rather than predicting the number of attentional episodes, it remained an open question whether a fully specified process model of the AB could account for these intrusion distributions.

The perceptual wink model as applied to these data was unchanged from the one fitting the Experiment 2 data and producing the simulations shown in Fig. 7 (the four free parameters from Experiment 2 were fixed for this simulation). This paradigm involved two aspects that differed

from the other paradigms, necessitating two new parameters. First, because the signal indicating the presence of a target (i.e., the circle surrounding the target) was peripheral and separate from the visual input of the target letter's identity, a new parameter specified the strength of visual input from the target signal. Second, unlike previous analyses, this application of the perceptual wink model required consideration of report order. To capture report order, we assumed that relative temporal position information was available for pairs of items that received some degree of encoding into STM. This was implemented with a logistic function based on relative position difference between the two items, with a free parameter for the growth rate. This relative temporal distinctiveness assumption is borrowed from the SIMPLE memory model (Brown, Neath, & Chater, 2007), which has been applied to serial recall tasks.

As seen in Fig. 8, the model captures all of the qualitative trends in the data and provides a reasonable account of the quantitative trends based on a fit with two free parameters.

Support for the switching/non-switching distinction

The perceptual wink model cannot explain temporary deficits in AB tasks that clearly involve a task switch from a T1 defined in one manner to a T2 defined in a different manner. Therefore, in extending the nROUSE model to the AB literature by only considering the switching AB tasks, it is important to consider the evidence in support of the switch/non-

Fig. 8 Application of the perceptual wink model to intrusion data of Vul et al. (2008). This paradigm presented a stream of letters, with each letter presented for 58 ms, followed by a 25 ms blank screen. Targets were enclosed by a circle. *Upper graphs:* Data and model results for the first item reported at the end of the sequence (i.e., an attempt to produce T1). *Lower graphs:* Data and model results for the second item reported at the end of the sequence (i.e., an attempt to produce T2). The *x-axis* shows position in the sequence relative to the T1 position. The *colored lines* show intrusion distributions for each of the ten lags tested in this study. The *black dashed line* shows T2 accuracy. In applying the perceptual

wink to these data, the parameters from Experiment 2 were used (the same parameters used in Fig. 7). Two new parameters were allowed to freely vary to capture the data. The first was the strength of visual input for the circle surrounding the target and the second was a logistic growth rate parameter specifying the probability of order reversals as a function of position differences (e.g., if the simulated trial indicated a report of items from positions 4 and 6, the probability of mistakenly reporting the position 6 item first was calculated using a logistic function based on a position difference of 2).

switch distinction. Although the magnitude of the blink in switching and non-switching AB tasks is correlated (Dale & Arnell, 2013), this does not necessarily mean that the two kinds of AB tasks reflect entirely identical underlying mechanisms; a formal explanation of individual differences needs to be developed and tested before such correlations can be interpreted. For instance, individual differences in the ability to stay motivated throughout a lengthy testing session are likely to impact both kinds of AB tasks in the same manner, although such differences may be unrelated to the root cause(s) of the blink in each task. Other results support the switching/non-switching distinction by observing dissociations between the two types of AB tasks. Some of this evidence comes from neural measures, implicating different neural substrates for the deficit with each type of AB task (Brissou, 2015; Vachon & Jolicoeur, 2011). Other evidence comes from behavioral dissociations. For instance, switching AB tasks produce cross-modal T2 deficits whereas non-switching AB tasks do not (Potter, Chun, Banks, &

Muckenhoupt, 1998). Also, blink magnitudes on one non-switching AB task correlate more strongly with those of different non-switching AB tasks as compared to the strength of correlation with those of switching AB tasks (Kelly & Dux, 2011). Dale, Dux, and Arnell (2013) performed a factor loading analysis of AB correlations, concluding that the “switch-versus-no-switch distinction may have influenced the strength of the relationships observed” (p. 466). Finally, Visser, Bischof, and Di Lollo (1999) reviewed AB results, concluding that lag-1 sparing (i.e., the lack of a T2 deficit immediately following T1) is robust for non-switching AB tasks, but often missing with switching AB tasks.

Even though the perceptual wink model does not readily explain deficits in switching AB tasks, its assumptions vis-à-vis target detection priming explain the finding that lag-1 sparing is less robust in switching tasks. When T2 appears as the post-T1 item, lingering perceptual responses to T1 trigger attentional mechanisms and encoding of T2 to the extent that the observer has not yet adopted the attentional set appropriate to

T2. For instance, in the Raymond et al. (1992) paradigm, observers might still be detecting the presence of white lettering at the time when the black X appears in the lag-1 position, causing them to partially encode the black X precisely because they have not yet switched to the T2 attentional set. However, this lag-1 sparing is less robust than in non-switching tasks because the T2 does not match the current attentional set. In contrast, in a non-switching task, at the time of the lag-1 T2, the target detector is primed by T1 and the T2 matches the attentional set.

A clear example of lag-1 sparing in a switching task comes from the “two glimpses” result of Weichselgartner and Sperling (1987). In that study, observers viewed a sequence of numbers and were instructed to report the four numbers that appeared subsequent/simultaneous with the appearance of a square around one of the numbers. This involved a task-switch from monitoring for a square to encoding numbers, and yet there was a form of lag-1 sparing (the first glimpse) such that, with high reliability, observers reported the number inside the square as well as the next number. This was followed by a failure to report subsequent numbers, and then a re-emergence of number reporting (a second glimpse) after the task switch. According to the perceptual wink model, if lag-1 sparing occurs in switching AB tasks, it may arise from the same mechanism as lag-1 sparing in non-switching AB tasks (i.e., detection of T1 activates the attentional gate, and this gate is still active at the time of the post T1 item). However, the deficit for switching AB tasks should relate to similar deficits in the task switching literature and the deficit for non-switching AB tasks should relate to similar deficits in the perceptual literature. This situation explains why switching and non-switching tasks correlate with each other to some degree, but two different switching or two different non-switching tasks correlate more strongly with each other than they do across the switching/non-switching distinction.

Finally, we consider the intriguing possibility that after a great deal of training, some switching AB tasks may become non-switching tasks, akin to the learning of new categories reported by Shiffrin and Schneider (1977). For instance, another way to state the rule in the task of Raymond et al. (1992) is to report any white letter or black X, regardless of which comes first (their task never presented black Xs prior to the white letter). With sufficient practice, observers may have learned this unique conjunction of attributes as the perceptual definition of targets.

A comparison with other AB models

The AB literature has sparked the development of many theoretical models (for a review, see Dux & Marois, 2009). However, there are three important features shared amongst the existing AB models that differ from the perceptual wink

model proposed here: (1) nearly all of these models are verbal (e.g., Chun & Potter, 1995; Di Lollo et al., 2005; Nieuwenstein et al., 2005; Raymond et al., 1992; K. L. Shapiro, Raymond, & Arnell, 1994; Ward, Duncan, & Shapiro, 1996) rather than formal theories specified with mathematical formulae and/or computer simulations; (2) nearly all of the models were developed specifically to explain the AB literature rather than attempting to bridge to related attentional or perceptual tasks, with the notable exception of the threaded cognition model of Taatgen et al. (2009); and (3) none of the models place the root cause of the blink in the dynamics of perception. Nevertheless, the perceptual wink model is similar to existing AB models in many ways. Besides reviewing these similarities and differences, we consider in detail how the perceptual wink model compares to the two most well-developed formal models of the AB: the episodic simultaneous type/serial token model (Wyble, Nieuwenstein, & Bowman, 2009) and the boost and bounce model (Olivers & Meeter, 2008).

Like all models of the AB, the perceptual wink model contains fluctuating attention dynamics, but unlike other theories, these attention dynamics passively follow perceptual dynamics. For this reason, the perceptual wink model only explains deficits in non-switching AB tasks. Because other theories place the root cause of the AB within attentional dynamics, they explain switching AB tasks in the same manner as non-switching AB tasks. This could be viewed as a limitation of the perceptual wink model. However, we view this as a strength; by placing the root cause of the blink in perceptual dynamics, the perceptual wink model necessarily predicts interactions of the AB with perceptual priming. Other theories may be able to capture these priming effects in a post-hoc manner through additional assumptions, but they do not necessarily make predictions about perceptual priming.

In addition to predicting interactions between the AB and priming, the perceptual wink model predicts that the AB should be explainable by well-established perceptual dynamics. The reported simulations specified perceptual dynamics according to the nROUSE model of perceptual priming as previously applied to word/letter identification. Aside from a freely varying inhibition parameter, which was necessary to capture the lack of visual crowding for singly presented letters as opposed to letters in the context of words (Pelli, Palomares, & Majaj, 2004), the only free parameters used to explain Experiment 1 concerned the addition of an attentional gate node, requiring two new parameters (an integration time constant and an activation threshold), and a parameter capturing the probability of guessing among the possible answers, implemented as residual activation within STM for all possible responses. The success of the model, despite using the previously established perceptual dynamics, served as a critical sufficiency check; if the model had been unable to explain

the attentional blink, this would have falsified the claim that perceptual dynamics underlie the attention blink.

Although the placement of the blink dynamics within perceptual representations is unique, the functional form of these dynamics is not unique, appearing in several other AB models. It is important to note that this does not reflect a “borrowing” of these dynamics from other models—the equations dictating the dynamics of perceptual wink model were first published in 2003 (Huber & O'Reilly, 2003). Nevertheless, it is instructive to consider the similarities between these models. For instance, similar to the Locus Coeruleus AB model of Nieuwenhuis, Gilzenrat, Holmes, and Cohen (2005), the perceptual wink model assumes that the deficit is a kind of neural refractory period. Similar to the episodic simultaneous type/serial token (eSTST) model of Wyble et al. (2009), the perceptual wink model does not impose any resource limitations within the attentional system. Similar to the boost and bounce (BB) model of Olivers and Meeter (2008), the perceptual wink model assumes that T1 initiates an attentional response that carries over and enhances encoding of the post-T1 item (a boost).

The BB model (Olivers & Meeter, 2008) was developed to explain the spread of sparing in which the onset of the blink can be delayed by a continual sequence of targets. In the BB model, the continued presence of targets keeps the attentional gate active, allowing the encoding of items into STM. This is equally true in the perceptual wink model, although there are limits to this spread in the perceptual wink model considering that the target detector continues to accrue greater habituation over the course of several targets. In contrast to these models, the original simultaneous type/serial token (STST) model of Bowman and Wyble (2007) could not explain the spread of sparing. In the STST model, encoding into STM occurs via a ‘blaster’ that is triggered by T1. The blaster causes attentional enhancement, producing sufficient input to trigger type processing (e.g., which letter) of the target and thus the binding of the type into an episodic token (e.g., an instance of that letter). The initiation of binding triggers inhibitory feedback to the blaster, shutting down attention to reduce interference between competing types. This inhibition produces the blink. However, in the original STST model, the onset of inhibition has a fixed time course. Lag-1 sparing was explained as the T2 binding to the same token as T1, but this process could not be extended to multiple targets in succession. The eSTST model remedied this situation by allowing flexibility in the dynamics of the blaster, delaying inhibition in the presence of continued targets, with each target receiving its own token.

In the BB model, the blink occurs as a result of the post-T1 distractor. T1 causes attentional enhancement (boosting the attentional gate neurons), with this boost being delayed somewhat so as to also affect the post-T1 item, producing lag-1 sparing. However, if the post-T1 item is a distractor, the mistaken encoding of that distractor triggers inhibition (a bounce),

and thus an attentional blink. In other words, the BB model requires a distractor to initiate the blink. This can be contrasted with the eSTST model, which assumes that the blink is initiated by encoding T1 into STM, with the caveat that additional targets can delay the onset of this inhibition. Simply put, in the BB model, the blink is triggered by a target followed by a distractor, whereas in the eSTST model the blink is triggered by a target followed by the absence of a subsequent target. In this regard, the perceptual wink model sits somewhere between these alternatives. When distractors are presented between T1 and T2, the blink in the perceptual wink model reflects both lingering habituation for the target detector (akin to inhibition of the blaster in the eSTST model) and inhibition of the target detector owing to competing activation of the distractor detector (akin to the bounce in the BB model). If blank screens are presented between T1 and T2, the perceptual wink model still produces a blink (as does the eSTST model) because of lingering habituation, although the magnitude of the blink is reduced owing to the lack of inhibition from the distractor detector (as well as weaker forward masking of T2).

In light of the different mechanisms for triggering the blink in each of these models, the results of Chen and Zhou (2015) are particularly informative. The TTT condition produced a spread of sparing, which the original STST model cannot handle, but the BB, eSTST, and perceptual wink models can handle. However, the reversal of the spread of sparing by inserting a blank screen between the second and third Ts cannot be handled by the BB model because in the absence of an intervening distractor, no blink is initiated. The eSTST model does not handle this result with its default parameters, but can handle this with a change of its parameter values. In contrast, the perceptual wink model predicts this result regardless of its parameter values. In the perceptual wink model, two initial targets provide greater habituation of the target detector as compared to one initial target (i.e., the TT of the TTT condition, as compared to the TD of the TDT condition). However, this habituation is more than offset by the priming advantage of the target detector being kept continually active and by the absence of a competing distractor, with these two factors producing the spread of sparing for a third target. In the event of an intervening blank screen, the target detector loses its activation (no target detector priming) and the distractor detector also loses its activation (no competing distractor). Thus, after a sufficiently long blank screen between the second and third targets, the TTT condition is necessarily worse than the TDT condition (regardless of parameters) because of greater habituation following two initial targets. Critically, the perceptual wink model was developed without its authors knowing of this result, and this was a true prediction.

In terms of non-switching AB tasks, the eSTST and perceptual wink models seem to fair equally well, but do so for

completely different reasons (perceptual habituation versus attentional inhibition). Because these two models assume different root causes of the blink, the perceptual wink model cannot explain switching AB tasks (unless it was augmented with task-switching dynamics) whereas the eSTST model cannot explain the transition from positive to negative perceptual priming (unless it was augmented with perceptual dynamics). Additional work on the switching/non-switching distinction and the interplay between perceptual dynamics and the attentional blink is needed to adjudicate between these models.

Conclusions

In conclusion, the reported experiments and modeling suggest that non-switching AB tasks are something of a misnomer. Although the blink reflects the waxing and waning of selective attention, the root cause of this waxing and waning does not lie within the attentional system. Instead, attention simply follows the time course of perceptual activation and habituation. Thus, attention is fully prepared and capable of accepting a second target, but fails to do so because perceptual processes fail to identify that a second target occurred (see Schneider & Shiffrin, 1977 Experiment 3a for a target detection task that produced an apparent blink). Similar to the proposal made by Awh et al. (2004), this perceptual wink explanation follows from an appreciation for the multifaceted nature of perception, with the blink reflecting a deficit for one perceptual attribute (e.g., the ability to detect the target) which leaves intact other perceptual attributes (e.g., the ability to identify the target).

Acknowledgments We thank Carlo Toneatto of the University of Milano-Bicocca for help with the implementation of the experiments. We also thank Brad Wyble, Lucas Huszar, Barry Giesbrecht, and one anonymous reviewers for constructive comments on draft versions of this manuscript. This work was supported by NSF BCS-0843773 and NSF BCS-1431147.

Appendix A: Model details

Figure 2 depicts the connectivity of the perceptual wink model as applied to a letter-case AB task (Layer 1 simulating low-level visual responses and Layer 2 simulating orthographic responses). The equations reported here are equivalent to the ones reported in the original nROUSE paper (Huber & O'Reilly, 2003). However, rather than presenting their most general form, we clarify the operations of the model by unpacking them into separate equations for separate layers/inputs.

When the preferred visual stimulus is displayed, the dynamic activation of node i in the visual layer is

captured by the average membrane potential, v_i , which is updated by the amount, Δv_i , every millisecond according to Equation A1.

$$\frac{\Delta v_i}{S_1} = (1-v_i)-v_i \left\{ L + I_1 \sum_{\forall j} o_j \right\} \quad (\text{A1})$$

S_1 is the integration time constant for visual nodes (layer 1), L is the magnitude of leak current, I_1 is an inhibition constant that scales the output, o_j , from each visual node j of items presented in the same screen position (including self-inhibition). Because action potentials only occur above a certain threshold (θ), output is defined by Equation A2, which applies equally to all nodes j in all layers of the model.

$$o_j = \begin{cases} (v_j - \theta) a_j & \text{if } v_j > \theta \\ 0 & \text{otherwise} \end{cases} \quad (\text{A2})$$

Equation A1 drives activation towards the maximum value of 1.0 when the preferred stimulus is presented. However, when that stimulus is removed, Equation A1 mathematically reduces to Equation A3, and activation approaches zero quickly, if other visual nodes are active (i.e., visual masking through inhibition), or slowly (i.e., iconic persistence), owing to leak current L .

$$\frac{\Delta v_i}{S_1} = -v_i \left\{ L + I_1 \sum_{\forall j} o_j \right\} \quad (\text{A3})$$

The variable a_j in Equation A2 is the current level of synaptic resources (e.g., neurotransmitter) available to node j for effective signaling of receiving cells. Thus, Equation A2 is the probability of an action potential multiplied by the post-synaptic effect of each action potential. Every node j has a unique time-varying level of available resources as determined by Equation A4.

$$\frac{\Delta a_j}{S_n} = R(1-a_j) - Do_j \quad (\text{A4})$$

As with Equation A2, Equation A4 applies equally to all nodes in all layers of the model, although the integration time constant, S_n , is set to the appropriate value for layer n in which node j resides (e.g., for the nodes of the visual layer, $S_n = S_1$).

Equations A1 and A3 describe membrane potential update for Layer 1 based on the presence or absence of visual input, respectively. In the case of Layer 2, Equation A5 is used to update the membrane potential of node k based on the sum of all output values j from nodes in Layer 1, each multiplied by the connection weight, w_{kj} , between the sending node j and the receiving node k .

$$\frac{\Delta v_k}{S_2} = (1-v_k) \sum_{\forall j} w_{kj} o_j - v_k \left\{ L + I_2 \sum_{\forall l} o_l \right\} \quad (\text{A5})$$

Connection weights are set to either 1 or 0, depending on whether each pair of nodes is or is not connected (as indicated by the arrows in Figure 2). Inhibition for layer 2 nodes is based on the sum over the output of all layer 2 nodes, I , that are engaged in the same type of identification. In other words, as indicated in Fig. 2, the letter identities mutually inhibit each other (i.e., they compete to determine which letter is currently present) and, separately, upper versus lower-case mutually inhibit each other (i.e., these two nodes compete to determine the case of the currently presented letter).

This describes the basic nROUSE dynamics used in all previous publications of the model. The model parameters were set to ‘default’ values that were developed to explain the Experiment 1 word repetition priming results reported by Huber (2008), and, as seen in Fig. 1. A recent study (Rieth & Huber, 2017) replicated these results and provided a more accurate fit of the word priming results by modifying the I , S_1 and S_2 parameters, and these values are used here. More specifically, these parameters were: $L = .15$, $I = .9844$, $\theta = .15$, $D = .324$, $R = .022$, and $S_1 = .0294$, $S_2 = .0609$. The inhibition parameter was set to its default value for the orthographic layer ($I_2 = .9844$) but it was a free parameter for the visual layer (I_1). The reason for this change is that presentation of a word necessarily involves visual crowding between adjacent letters (Pelli et al., 2004). Thus, the strength of visual inhibition determined from the word priming data likely overestimates the degree of inhibition for a single letter presented without any flanking letters. Indeed, as compared to word priming value of .9844, the best-fitting visual inhibition parameter was estimated to be .1128 for the AB results of Experiment 1.

Previous applications of the nROUSE model did not need to address the selection of items allowed into STM, but this is a crucial component of any AB model, considering that most AB tasks use more items than can be maintained in STM. We assume that STM encoding is dictated by an attentional gate node, which is labeled α in Fig. 2. Because the attentional-set for this AB task was to report upper-case but not lower-case letters, the connection strengths between the attentional gate node and the upper/lower-case nodes were +1 and −1, respectively. The attentional gate node is a new addition to the nROUSE model requiring two new free parameters. As previously mentioned, it was assumed that this node does not experience habituation. This is implemented by fixing the synaptic resources (a) at 1.0. With this assumption, the cause of the AB must lie in the dynamics of the upper/lower-case nodes. How rapidly the attentional gate node responds to input from the letter-case nodes is dictated by the integration time constant (S_3), which was a free parameter dictating the timing delay between when the upper-case node activates and when letter identities are encoded

into STM. Unlike other nodes, which use the default firing threshold ($\theta = .15$), the threshold for the attentional gate node was a free parameter, serving to rescale the magnitude (but not the dynamics) of the blink. In summary, these two free parameters can shift the letter-case dynamics (delays) and can rescale the letter-case dynamics, but they cannot change the rapidity with which a blink is initiated or the rapidity with which the AB recovers—instead, these dynamics were fixed *a priori* by prior applications of the nROUSE model to the word priming data.

A ubiquitous finding in neuroscience is divisive normalization whereby the activity of one neuron is normalized (i.e., inhibited) by the activity of other neurons in the same region (e.g., Barlow, 1953). The inhibition parameters in the nROUSE model implement divisive normalization within perceptual responses, and we assume that divisive normalization also occurs for the encoding of letter identities into STM and for the decoding of letter identities out of STM when making overt responses. In an abstract sense, divisive normalization can be thought of as producing a probability distribution over the possible items, as described by Equation A6.

$$p(i) = \frac{e_i^\gamma}{\sum_j e_j^\gamma} \quad (\text{A6})$$

In this equation, the response probability of a particular alternative i is the evidence e in favor of that alternative raised to the power γ , as compared to the sum over all possible responses j . If the power constant γ is set to infinity, divisive normalization is winner-take-all (WTA) and if the power constant is set to 1, this equation becomes Luce’s (1959) choice rule. In the case of STM encoding, we assume WTA divisive normalization (e.g., only one item can be encoded into STM at any particular moment). Thus, if two or more letter identities are simultaneously active, only the most active is encoded, with the strength of encoding dictated by the attentional gate node at that point in time (α_t). This is captured by the summation equation shown in Fig. 2, in which δ is the delta function, with $\delta = 1$ for the most active letter identity at time t and $\delta = 0$ for the other letters. The summation is over all time steps (the model is run millisecond by millisecond) during the RSVP sequence, and over all possible letter identities (i). Thus, this summation provides a total evidence (e) accumulation in STM for each possible letters identity. However, because of the actions of attentional gating, only four letter identities tend to be encoded to any degree: T1, the post-T1 distractor (in conditions other than lag-1), T2, and the post-T2 distractor. Nevertheless, it is assumed that all possible letters have some residual degree of evidence accumulation at the start

of the trial, perhaps through proactive interference from previous trials (Wickens, Born, & Allen, 1963). This residual STM activation is the a third STM free parameter.

The process of giving two (or three) responses at the end of the RSVP trial is explicitly modeled. The Luce choice rule is applied to give a first response distribution. Next, for each possible first response, a conditional second response distribution is calculated by setting to zero the evidence (e_i) for the appropriate first response before re-applying the Luce choice rule. This prevents the model from producing the same response twice (which is not allowed by the experimental procedures). In the case of three responses, this procedure is applied recursively to find conditional third response distributions for each possible sequence of first two responses. However, the data were too sparse to allow a full analysis of response order and/or a full analysis over all specific letter identities (e.g., the probability of responding with the first distractor in the RSVP stream, the second distractor, etc.). Therefore, these first and second responses were collapsed to provide the probability of giving both T1 and T2 (regardless of response order), T1 but not T2, T2 but not T1, and neither T1 nor T2. Thus, even though full response distributions are not analyzed, the probability of responding with letter identities other than T1 and T2 provides an important constraint on the model, serving to increase or decrease the probability that neither T1 nor T2 is identified.

In summary, four free parameters were used to explain the 40 joint probabilities of Experiment 1 (i.e., 10 conditions, each with 4 joint probabilities of T1/T2, representing 30 degrees of freedom in the observed data). The best-fitting parameter value for S_3 was the maximum allowed value of 1.0 (above 1, the model becomes unstable and may oscillate), indicating very rapid attentional gating based on the activity of the upper/lower-case nodes (as reported in the legend to Fig. S1, when fitting the data of Experiment 2, this parameter took on a smaller value, with the attentional gate delay capturing the observation that T2 performance was higher than T1 performance in the lag-1 condition). The firing threshold for the gate node (θ_3) was .2903 and residual STM activation was .3074. Finally, visual inhibition (I_1) was .1128. With these parameter values, the model captured 95.6% of the variance across the 40 joint probabilities, with a chi-square value of 140.5. The best-fitting parameter values and goodness of fit metrics for the Chen and Zhou study are reported in the figure caption of Fig. S2.

Application of the model to the study of Vul et al. (2008) required two additional parameters. The first reflected the strength of visual input from the circle that surrounded targets. Because this symbol was larger than the characters, it may have provided a more salient signal for activation of the target detector as compared to

paradigms that triggered the target detector through an orthographic property of the target character (e.g., number/letter category or upper/lower case). In the simulation, the visual node corresponding to the surrounding circle had the same dynamics as all of the other visual nodes, but the best-fitting connection strength to the target detector was 55.6, rather than the default value of 1.0. It is critical to note two things about this value. First, this does not mean that the target detector was 55.6 times more active considering that membrane potential is bounded at 1.0. In fact, this value scarcely changes the maximum activation of the target detector, but it does result in the target detector being activated more quickly. Second, because the dynamics of the visual node corresponding to the circle were unchanged, this greater connection to the target detector should not be considered a greater connection weight (which would have resulted in faster habituation owing to more rapid loss of synaptic resources). Instead, this reflects a greater number of neurons representing the surrounding circle (55.6 times as many visual neurons were activated by the surrounding circle, as compared to a target letter).

The second new parameter concerned the judgment of temporal order. If both of items sampled from STM (i.e., the two items slated as answers, according to two implementations of the Luce choice rule in Equation A6) had encoding values greater than the residual STM activation, then it was assumed that temporal position information was available. In this case, the logistic equation shown in Equation A7 was used.

$$p(i \text{ before } j) = \frac{1}{1 + e^{t(pos\ i - pos\ j)}} \quad (A7)$$

This equation gives the probability of producing item i before item j based on their positions in the sequence and the logistic growth rate parameter, t . For instance, if the position of i is much later than j , this equation approaches 0 (i.e., j is given as the first response), whereas if the position of i is much earlier than j , this equation approach 1 (i.e., i is given as the first response). If the two items have the same position (which is not possible in this paradigm because items were presented one at a time), then it would be a toss-up as to which would be given as the first response. The best-fitting growth rate parameter was 1.08. In the event that just one of items was encoded above the residual STM activation value, or if neither of the items were encoded, then response order was a toss-up.

With these two free parameters, the 520 separate response categories were fit (i.e., 26 possible letter responses, with each one having a unique position relative to the position of T1, for each of the first and second reports at each of 10 lags) capturing 81% of the variance, with a chi-square of 3954 (488 degrees of freedom).

Appendix B: Experiment 2 similarity tables and analyses

Table B1 Targets and distractors used in Experiment 2

Distractors		Targets	
Curvy	Straight	Curvy	Straight
a	h	B	A
b	r	D	E
d	j	Q	H
q	k	G	K
e	f	P	F
g	t	R	T
p	y	J	Y

Table B2 The eight conditions of Experiment 2 rising from the 2×2 design factoring similarity and distractor type

Condition	Distractors	Targets	Distractor similarity	Distractor type
1	Curvy	Curvy	Similar	Repeated
2	Curvy	Curvy	Similar	Standard
3	Curvy	Straight	Dissimilar	Repeated
4	Curvy	Straight	Dissimilar	Standard
5	Straight	Curvy	Dissimilar	Repeated
6	Straight	Curvy	Dissimilar	Standard
7	Straight	Straight	Similar	Repeated
8	Straight	Straight	Similar	Standard

We ran a 2 (distractor type: repeated vs. unrepeated) $\times$ 2 (similarity: similar vs. dissimilar) $\times$ 3 (Lag: lag 1, lag 3, lag 6) within-subjects repeated-measures ANOVA on the mean proportion of T2 | T1 report. The analysis yielded a main effect of distractor type, $F(1,40) = 171.84$, $P < .001$, $\eta^2 = .261$, due to the greater participants' accuracy when the distractors were repeated ($M = .72$, $SDE = .03$), than when they were different ($M = .5$, $SDE = .02$). There was also a main effect of similarity, $F(1, 40) = 15.71$, $p < .001$, $\eta^2 = .009$, indicating that participants were more accurate when targets and distractors were dissimilar ($M = .63$, $SDE = .02$) than when they were similar ($M = .59$, $SDE = .02$). Finally, there was a main effect of Lag, $F(2,80) = 101.5$, $P < .001$, $\eta^2 = .321$, indicating that performance at Lag 1 ($M = .73$, $SDE = .02$) was higher than accuracy at both Lag 6 ($M = .64$, $SDE = .03$), and Lag 3 ($M = .45$, $SDE = .02$), all P s $< .001$. The difference between performance at Lag 3 and Lag 6 was significant, $P < .001$.

The interaction between distractor type and similarity was not significant, $F(1,40) = .1$, $P = .748$, $\eta^2 = .000$, as well as the one between similarity and Lag, $F(1, 40) = .73$, $P = .398$, $\eta^2 = .001$ [lower-bound correction], and the three-way interaction among distractor type, similarity, and Lag, $F(2,80) = .4$, $P = .675$, $\eta^2 = .000$. The only significant interaction was the one between distractor type and Lag, $F(2,80) = 40.31$, $P < .001$, $\eta^2 = .054$. Specifically, in the trials in which the distractor was repeated there was not a significant difference between performance at Lag 1 ($M = .77$, $SDE = .03$) and at Lag 6 ($M = .79$, $SDE = .03$), $P = .340$, and performance at Lag 3 ($M = .58$, $SDE = .03$) was significantly worse than both at Lag 1 and at Lag 6, all P s $< .001$. When the distractors were different, accuracy was greater at Lag 1 ($M = .69$, $SDE = .02$) than both at Lag 6 ($M = .49$, $SDE = .03$) and at Lag 3 ($M = .32$, $SDE = .02$), all P s $< .001$. Also the performance at Lag 3 and at Lag 6 was significantly different, $P < .001$.

References

- Awh, E., Serences, J., Laurey, P., Dhaliwal, H., van der Jagt, T., & Dassonville, P. (2004). Evidence against a central bottleneck during the attentional blink: Multiple channels for configural and featural processing. *Cognitive Psychology*, 48(1), 95–126. doi: [https://doi.org/10.1016/S0010-0285\(03\)00116-6](https://doi.org/10.1016/S0010-0285(03)00116-6)
- Barlow, H. B. (1953). Action potentials from the frogs retina. *Journal of Physiology-London*, 119(1), 58–68.
- Bowman, H., & Wyble, B. (2007). The simultaneous type, serial token model of temporal attention and working memory. *Psychological Review*, 114(1), 38–70.
- Brisson, B. (2015). Task switching mediates direct interference of intertarget distractors in the attentional blink: An event-related potential study. *Psychophysiology*, 52(7), 919–932.
- 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.
- Brown, G. D. A., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, 114(3), 539–576.
- Chen, X., & Zhou, X. L. (2015). Revisiting the spread of sparing in the attentional blink. *Attention Perception & Psychophysics*, 77(5), 1596–1607. doi: <https://doi.org/10.3758/s13414-015-0886-5>
- Chua, F. K., Goh, J., & Hon, N. (2001). Nature of codes extracted during the attentional blink. *Journal of Experimental Psychology-Human Perception and Performance*, 27(5), 1229–1242. doi: <https://doi.org/10.1037/0096-1523.27.5.1229>
- Chun, M. M. (1997a). Temporal binding errors are redistributed by the attentional blink. *Perception & Psychophysics*, 59(8), 1191–1199.
- Chun, M. M. (1997b). Types and tokens in visual processing: A double dissociation between the attentional blink and repetition blindness. *Journal of Experimental Psychology-Human Perception and Performance*, 23(3), 738–755. doi: <https://doi.org/10.1037/0096-1523.23.3.738>
- Chun, M. M., & Potter, M. C. (1995). A 2-stage model for multiple-target detection in rapid serial visual presentation. *Journal of Experimental Psychology-Human Perception and Performance*, 21(1), 109–127.
- Dale, G., & Arnell, K. M. (2013). How reliable is the attentional blink? Examining the relationships within and between attentional blink tasks over time. *Psychological Research-Psychologische Forschung*, 77(2), 99–105.

- Dale, G., Dux, P. E., & Arnell, K. M. (2013). Individual differences within and across attentional blink tasks revisited. *Attention Perception & Psychophysics*, 75(3), 456–467.
- Davelaar, E. J., Tian, X., Weidemann, C. T., & Huber, D. E. (2011). A habituation account of change detection in same/different judgments. *Cognitive Affective & Behavioral Neuroscience*, 11(4), 608–626. doi: <https://doi.org/10.3758/s13415-011-0056-8>
- Di Lollo, V., Kawahara, J., Ghorashi, S. M. S., & Enns, J. T. (2005). The attentional blink: Resource depletion or temporary loss of control? *Psychological Research-Psychologische Forschung*, 69(3), 191–200. doi: <https://doi.org/10.1007/s00426-004-0173-x>
- Dux, P. E., & Marois, R. (2008). Distractor inhibition predicts individual differences in the attentional blink. *Plos One*, 3(10), e3330.
- Dux, P. E., & Marois, R. (2009). The attentional blink: A review of data and theory. *Attention Perception & Psychophysics*, 71(8), 1683–1700. doi: <https://doi.org/10.3758/APP.71.8.1683>
- Elliott, J. C., & Giesbrecht, B. (2015). Distractor suppression when attention fails: Behavioral evidence for a flexible selective attention mechanism. *Plos One*, 10(4), e0126203.
- Evett, L. J., & Humphreys, G. W. (1981). The use of abstract graphemic information in lexical access. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 33A(4), 325–350. doi: <https://doi.org/10.1080/14640748108400797>
- Fahrenfort, J. J., van Leeuwen, J., Olivers, C. N. L., & Hogendoorn, H. (2017). Perceptual integration without conscious access. *Proceedings of the National Academy of Sciences of the United States of America* 114(14), 3744–3774. *Proceedings of the 25th Annual Conference of the Cognitive Science Society*.
- Giesbrecht, B., & Di Lollo, V. (1998). Beyond the attentional blink: Visual masking by object substitution. *Journal of Experimental Psychology-Human Perception and Performance*, 24(5), 1454–1466.
- Goodbourn, P. T., Martini, P., Barnett-Cowan, M., Harris, I. M., Livesey, E. J., & Holcombe, A. O. (2016). Reconsidering temporal selection in the attentional blink. *Psychological Science*, 27(8), 1146–1156.
- Gorea, A. (2015). A refresher of the original Bloch's law paper (Bloch, July 1885). *i-Perception*, 6(4), 1–6. doi: <https://doi.org/10.1177/2041669515593043>
- Hogben, J. H., & Di Lollo, V. (1974). Perceptual intergration and perceptual segregation of brief visual stimuli. *Vision Research*, 14, 1059–1069. doi: [https://doi.org/10.1016/0042-6989\(74\)90202-8](https://doi.org/10.1016/0042-6989(74)90202-8)
- Huber, D. E. (2008). Immediate priming and cognitive aftereffects. *Journal of Experimental Psychology-General*, 137(2), 324–347. doi: <https://doi.org/10.1037/0096-3445.137.2.324>
- Huber, D. E., Clark, T. F., Curran, T., & Winkielman, P. (2008a). Effects of repetition priming on recognition memory: Testing a perceptual fluency-disfluency model. *Journal of Experimental Psychology-Learning Memory and Cognition*, 34(6), 1305–1324. doi: <https://doi.org/10.1037/a0013370>
- Huber, D. E., & O'Reilly, R. C. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: Temporal segregation through synaptic depression. *Cognitive Science*, 27(3), 403–430. doi: https://doi.org/10.1207/s15516709cog2703_4
- Huber, D. E., Shiffrin, R. M., Quach, R., & Lyle, K. B. (2002). Mechanisms of source confusion and discounting in short-term priming: 1. Effects of prime duration and prime recognition. *Memory & Cognition*, 30(5), 745–757.
- Huber, D. E., Tian, X., Curran, T., O'Reilly, R. C., & Woroch, B. (2008b). The dynamics of integration and separation: ERP, MEG, and neural network studies of immediate repetition effects. *Journal of Experimental Psychology-Human Perception and Performance*, 34(6), 1389–1416. doi: <https://doi.org/10.1037/a0013625>
- Humphreys, G. W., Besner, D., & Quinlan, P. T. (1988). Event perception and the word repetition effect. *Journal of Experimental Psychology: General*, 117(1), 51–67. doi: <https://doi.org/10.1037/0096-3445.117.1.51>
- Irwin, K. R., Huber, D. E., & Winkielman, P. (2010). Automatic affective dynamics: An activation-habituation model of affective assimilation and contrast. *Modeling Machine Emotions for Realizing Intelligence*, 1, 17–34. doi: <https://doi.org/10.1007/978-3-642-12604-8>
- Kanwisher, N. G. (1987). Repetition blindness—type recognition without token individuation. *Cognition*, 27(2), 117–143. doi: [https://doi.org/10.1016/0010-0277\(87\)90016-3](https://doi.org/10.1016/0010-0277(87)90016-3)
- Kelly, A. J., & Dux, P. E. (2011). Different attentional blink tasks reflect distinct information processing limitations: An individual differences approach. *Journal of Experimental Psychology-Human Perception and Performance*, 37(6), 1867–1873.
- Luce, R. D. (1959). *Individual choice behavior; a theoretical analysis*. New York: Wiley.
- Maki, W. S., Frigen, K., & Paulson, K. (1997). Associative priming by targets and distractors during rapid serial visual presentation: Does word meaning survive the attentional blink? *Journal of Experimental Psychology-Human Perception and Performance*, 23(4), 1014–1034. doi: <https://doi.org/10.1037/0096-1523.23.4.1014>
- 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), 291–307. doi: <https://doi.org/10.1037/0096-3445.134.3.291>
- 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., & Theeuwes, J. (2009). Unmasking the attentional blink. *Journal of Experimental Psychology-Human Perception and Performance*, 35(1), 159–169.
- Olivers, C. N. L., & Meeter, M. (2008). A boost and bounce theory of temporal attention. *Psychological Review*, 115(4), 836–863. doi: <https://doi.org/10.1037/a0013395>
- Pashler, H. (1994). Dual-task interference in simple tasks—data and theory. *Psychological Bulletin*, 116(2), 220–244. doi: <https://doi.org/10.1037/0033-2909.116.2.220>
- Pelli, D. G., Palomares, M., & Majaj, N. J. (2004). Crowding is unlike ordinary masking: Distinguishing feature integration from detection. *Journal of Vision*, 4(12), 1136–1169. doi: <https://doi.org/10.1167/4.12.12>
- 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., Dell'Acqua, R., Pesciarelli, F., Job, R., Peressotti, F., & O'Connor, D. H. (2005). Bidirectional semantic priming in the attentional blink. *Psychonomic Bulletin & Review*, 12(3), 460–465.
- 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. doi: <https://doi.org/10.1037/0096-1523.18.3.849>
- Reeves, A., & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review*, 93(2), 180–206.
- Rieth, C. A., & Huber, D. E. (2010). Priming and habituation for faces: Individual differences and inversion effects. *Journal of Experimental Psychology-Human Perception and Performance*, 36(3), 596–618. doi: <https://doi.org/10.1037/a0018737>
- Rieth, C. A., & Huber, D. E. (2017). Comparing different kinds of words and word-word relations to test an habituation model of priming. *Cognitive Psychology*, 95, 79–104. doi: <https://doi.org/10.1016/j.cogpsych.2017.04.002>
- Rieth, C. A., & Vul, E. (2012). Expectations about the temporal structure of the world result in the attentional blink and repetition blindness.

- Paper presented at the The 34th Annual Meeting of the Cognitive Science Society, Sapporo, Japan.
- Salvucci, D. D., & Taatgen, N. A. (2008). Threaded cognition: An integrated theory of concurrent multitasking. *Psychological Review*, 115(1), 101–130. doi: <https://doi.org/10.1037/0033-295X.115.1.101>
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information-processing. I. Detection, search, and attention. *Psychological Review*, 84(1), 1–66. doi: <https://doi.org/10.1037/0033-295X.84.1.1>
- Shapiro, K., 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: <https://doi.org/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), 357–371.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information-processing .2. perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84(2), 127–190.
- Slagter, H. A., & Georgopoulou, K. (2013). Distractor inhibition predicts individual differences in recovery from the attentional blink. *Plos One*, 8(5), e64681.
- Taatgen, N. A., Juvina, I., Schipper, M., Borst, J. P., & Martens, S. (2009). Too much control can hurt: A threaded cognition model of the attentional blink. *Cognitive Psychology*, 59(1), 1–29.
- Tian, X., & Huber, D. E. (2010). Testing an associative account of semantic satiation. *Cognitive Psychology*, 60(4), 267–290. doi: <https://doi.org/10.1016/j.cogpsych.2010.01.003>
- Tian, X., & Huber, D. E. (2013). Playing "Duck Duck Goose" with neurons: Change detection through connectivity reduction. *Psychological Science*, 24(6), 819–827 <https://doi.org/10.1177/0956797612459765>
- Tsodyks, M. V., & Markram, H. (1997). The neural code between neo-cortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences of the United States of America*, 94(2), 719–723. doi: <https://doi.org/10.1073/pnas.94.2.719>
- Vachon, F., & Jolicoeur, P. (2011). Impaired semantic processing during task-set switching: Evidence from the N400 in rapid serial visual presentation. *Psychophysiology*, 48(1), 102–111.
- Visser, T. A. W., Bischof, W. F., & Di Lollo, V. (1999). Attentional switching in spatial and nonspatial domains: Evidence from the attentional blink. *Psychological Bulletin*, 125(4), 458–469.
- 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. (1996). The slow time-course of visual attention. *Cognitive Psychology*, 30(1), 79–109.
- Weichselgartner, E., & Sperling, G. (1987). Dynamics of automatic and controlled visual-attention. *Science*, 238(4828), 778–780. doi: <https://doi.org/10.1126/science.3672124>
- Wickens, D. D., Born, D. G., & Allen, C. K. (1963). Proactive-inhibition and item similarity in short-term-memory. *Journal of Verbal Learning and Verbal Behavior*, 2(5–6), 440–445. [https://doi.org/10.1016/S0022-5371\(63\)80045-6](https://doi.org/10.1016/S0022-5371(63)80045-6)
- Wyble, B., Nieuwenstein, M., & Bowman, H. (2009). The attentional blink provides episodic distinctiveness: Sparing at a cost. *Journal of Experimental Psychology-Human Perception and Performance*, 35(3), 787–807. doi: <https://doi.org/10.1037/a0013902>