Chapter 4 Supplement

Examples of the Design of ERP Experiments

In this supplement, I will discuss three experiments in which ERPs were used to answer a significant question in cognitive neuroscience. In each case, the experiments were designed in such a manner that the conclusions did not depend on identifying specific ERP components, making the results much stronger than is otherwise possible.

Example 1: Auditory Selective Attention

Background

Given that I did my graduate work with Steve Hillyard, I was imprinted at an early age with the experimental design that he and his colleagues developed in 1973 to address the classic "locus-of-selection" question. This question asks whether attention operates at an early stage of processing, allowing only selected inputs to be perceived (which is the "early selection" hypothesis of investigators such as Broadbent [1958] and Treisman [1969]), or whether attention instead operates at a late stage such that all incoming sensory events receive equal perceptual processing but only selected stimuli reach decision processes, memory, and behavioral output systems (which is the "late selection" hypothesis of investigators such as Deutsch and Deutsch [1963] and Norman [1968]).

Two factors made this a difficult question to address with traditional behavioral measures. First, it is difficult to assess the processing of an ignored stimulus without asking the subjects to respond to it, in which case the stimulus may become attended rather than ignored. Second, if responses to ignored stimuli are slower or less accurate than responses to attended stimuli, it is difficult to determine whether this reflects an impairment in sensory processes or an impairment in higher-level decision, memory, or response processes. ERPs are particularly well suited for solving both of these problems. First, ERPs are easily recorded in the absence of an overt response, making them ideal for monitoring the processing of an ignored stimulus. Second, ERPs provide precise information about the time course of processing, making them ideal for answering questions about the stage of processing at which an experimental effect occurs. In the case of attention, for example, it is possible to determine whether the early sensory-evoked ERP components are suppressed for ignored stimuli relative to attended stimuli, which would be consistent with the early selection hypothesis, or whether attention influences only the later ERP components, which would be consistent with the late selection hypothesis.

Experimental Design

Several studies in the 1960s and early 1970s used this logic to address the locus-of-selection question, but these experiments had various methodological shortcomings that made them difficult to interpret. Hillyard, Hink, Schwent, and Picton (1973) reported an experiment that solved these problems and provided strong evidence for early selection. This experiment is illustrated in figure S4.1A. Subjects were instructed to attend to the left ear in some trial blocks and the right ear in others. A rapid sequence of tone pips was then presented, with half of the stimuli presented in each ear. To make the discrimination between the two ears very easy, the tones were presented at a different pitch in each ear. Subjects were instructed to monitor the attended ear and press a button whenever a slightly higher-pitched "deviant" target tone was detected in that ear, which

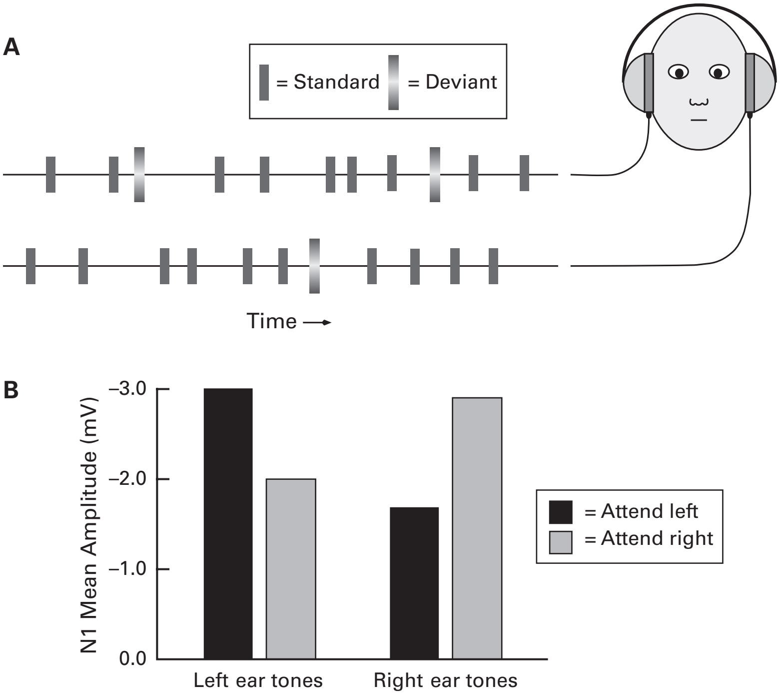

Figure S4.1
Experimental paradigm (A) and results (B) from the study of Hillyard et al. (1973). Subjects listened to streams of tone pips in the left ear and right ear. Most of the tones were a standard frequency (800 Hz in the left ear and 1500 Hz in the right ear), but occasional deviant tones were presented at a slightly higher frequency (840 Hz left ear; 1560 Hz right ear). Subjects were instructed to attend to the left ear for some trial blocks and to the right ear for others, and they counted the number of deviant tones in the attended ear. The average N1 amplitude was measured for the standard tones and used as an index of sensory processing. Left ear tones elicited a larger N1 wave when attention was directed to the left ear than when attention was directed to the right ear; conversely, right ear tones elicited a larger N1 wave when attention was directed to the right ear than when attention was directed to the left ear. Reproduced by permission from Luck (1998b). Copyright 1998 Psychology Press.

occurred infrequently and unpredictably. Higher-pitched tones were also presented occasionally in the ignored ear, but subjects were instructed not to respond to these “ignored deviants.”

In some prior ERP experiments, subjects were required to respond to or count all attended stimuli and withhold responses to all ignored stimuli (including internal responses such as counting). As a result, any differences in the ERPs evoked by attended and ignored stimuli could have been due to response-related or counting-related activity that was present for the attended ERPs but absent for ignored ERPs (see tip 4 in chapter 4). Hillyard et al. (1973) avoided this problem by presenting both target and nontarget stimuli in both the attended and ignored ears and asking the subjects to count the targets in the attended ear. The analyses were focused on the nontargets, to which subjects made neither an external response like a button press nor an internal response like counting. To ensure that subjects attended to the nontargets in the attended ear, even though no response was required for them, the targets and nontargets within an ear were difficult to discriminate from each other (but easy to discriminate from the stimuli in the other ear). Thus, subjects were motivated to focus attention onto all of the stimuli presented in one ear and ignore all stimuli within the other ear. The main experimental question was whether the early sensory ERP components evoked by a nontarget stimulus presented in the attended ear would be larger than those evoked by a nontarget stimulus presented in the ignored ear.

The sensory ERP components are highly sensitive to the physical characteristics of the evoking stimulus. As a result, one cannot legitimately compare the ERP evoked by an attended tone in the left ear with an ignored tone in the right ear: any differences between these ERPs could be due to differences between the two ears that have nothing to do with attention. The design used by Hillyard et al. (1973) circumvents this problem by allowing the ERP elicited by the same physical stimulus to be compared under different psychological conditions. For example, the ERP evoked by a left nontarget during attend-left blocks can be compared with the ERP evoked by the same left nontarget during attend-right blocks. Because the same stimulus is used in both cases, any differences in the ERPs between the attend-left and attend-right conditions must be due to differences in attentional processing.

In some attention experiments, the investigators compare an “active” condition in which the subject responds to the stimuli with a “passive” condition in which the subject makes no response to the stimuli and perhaps engages in a distracting activity such as reading a book. Frequently, however, the task in the active condition is much more demanding than the distraction task in the passive condition, leading to greater overall arousal during the active condition. If we compare the ERPs in the two conditions, any differences might be due to these global arousal differences rather than selective changes in stimulus processing. To ensure that differences in global arousal would not interfere with their study, Hillyard et al. compared ERPs evoked during equally difficult attend-left and attend-right conditions rather than active and passive conditions.

Results

Now that we have discussed the logic behind this study, let’s consider the results. As shown in figure S4.1B, the N1 component was found to be larger for attended stimuli than for ignored stimuli. Specifically, the N1 elicited by left ear tones was larger when the left ear was attended than when the right ear was attended, and the N1 elicited by right ear tones was larger when the right ear was attended than when the left ear was attended. These effects began approximately

60–70 ms after stimulus onset and peaked at approximately 100 ms poststimulus. In this manner, Hillyard et al. (1973) were able to demonstrate that attention can influence the processing of a stimulus within the first 100 ms after stimulus onset, which is consistent with the early-selection hypothesis. If you are interested in reading more about this experimental paradigm, see the review by Luck and Kappenman (2012a).

Larger Issues

I would like to emphasize three aspects of this study. First, it was specifically designed to address an existing and significant question that had previously eluded investigators. This contrasts with many ERP experiments in which the researchers simply take an interesting cognitive paradigm and run it while recording ERPs to “see what happens.” This latter approach, while sometimes a useful first step, rarely leads to important conclusions about the mind and brain.

Second, this study uses a component-independent design that does not rely on identifying specific ERP components. The ERPs elicited by the attended and ignored stimuli diverged around 60–70 ms poststimulus, and it is this timing information that was the crucial result of the study, and not the fact that the effect occurred in the latency range of the N1 component. In fact, there has been some controversy about whether attention influences the N1 component per se, but the finding of an effect within 100 ms of stimulus onset will continue to be important regardless of the outcome of this dispute.

A third important aspect of this study is that it used ERPs to assess the processing of stimuli for which subjects made no overt response. This is one of the main advantages of the ERP technique over behavioral measures, and many of the most significant ERP experiments have exploited this ability of ERPs to provide “covert monitoring” of processing.

Example 2: Partial Information Transmission

Background

Most early models of cognition assumed that simple cognitive tasks were accomplished by means of a sequence of discrete processing stages. These models were challenged in the late 1970s by models in which different processes could occur in parallel, such as McClelland’s (1979) cascade model and Eriksen and Schultz’s (1979) continuous flow model. At the most fundamental level, these new models differed from the traditional serial models in that they postulated that partial results from one process could be transmitted to another process, such that the second process could begin working before the first process was complete. Traditional discrete-stage models, in contrast, assumed that a process worked until it achieved its result, which was then passed on to the next stage. This is a crucial distinction, but it is very difficult to test without a direct means of observing the processes that are thought to overlap. This is exactly the sort of issue that ERPs can easily address.

By coincidence, two different laboratories conducted similar ERP studies of this issue at about the same time, and both used the lateralized readiness potential (LRP) as a means of assessing whether response systems can become activated before stimulus identification is complete (Miller & Hackley, 1992; Osman, Bashore, Coles, Donchin, & Meyer, 1992). Both were excellent studies, but here I will focus on the one conducted by Jeff Miller and Steve Hackley. These investigators tested the specific hypothesis that subjects will begin to prepare a response to a stimulus based on the most salient aspects of the stimulus, even if they later withheld that response because of additional information extracted from the stimulus. In other words, they predicted that motor processing may sometimes begin before perceptual processing is complete, which would be incompatible with traditional discrete-stage models of cognition.

Experimental Design

Miller and Hackley presented subjects with one of four stimuli on each trial: a large S; a small S; a large T; or a small T. Subjects responded with one hand for S and with the other hand for T, but they responded only to one of the two sizes; no response was given for the other size (half of the subjects responded to large stimuli and half to small). Thus, this paradigm was a hybrid of a go/no-go design (go for one size and no-go for the other) and a two-alternative forced-choice design (one response for S, a different response for T). The shape difference was very salient, but the size difference was relatively difficult to discriminate. Consequently, subjects could begin to prepare a given hand to respond as soon as they discriminated the shape of the letter, and they could later choose to emit or withhold this response when the size of the letter was eventually discriminated.

Miller and Hackley used the LRP component—which reflects response preparation—to determine if the subjects prepared a specific hand for response on the basis of letter identity, even on trials when no response was ultimately made because the letter was the wrong size. Miller and Hackley's (1992) paper provides a wonderful review of the LRP literature of the time, and I highly recommend reading it (for a more recent and more comprehensive review, see also Smulders & Miller, 2012). Here I will mention two essential aspects of the LRP that were essential for testing their hypothesis. First, the LRP begins well before a response and can occur even if there is no response, indicating that it reflects response preparation rather than response execution. Second, the LRP is larger over the hemisphere contralateral to the movement being prepared, and the LRP is isolated from other components by means of a contralateral-minus-ipsilateral difference wave (see figure 3.15 in chapter 3).

The presence of a nonzero voltage in this difference wave is virtually absolute proof that the brain has begun to determine which hand should make the button-press response for the current stimulus. In other words, there is no way to get a consistently larger response over the hemisphere contralateral to a given hand unless the brain has figured out that this is the hand that should be prepared. Thus, Miller and Hackley predicted that an LRP would be observed contralateral to the hand indicated by the shape of the letter, even if the size of the letter indicated that no response should be made (note that EMG recordings were used to ensure that absolutely no motor activity was present on these trials). When the size of the letter was finally perceived, the LRP would then be terminated if it indicated that no response should be made.

Results

The results of this study are shown in figure S4.2. Panel A shows the averaged ERP waveforms from the left- and right-hemisphere electrode sites for left- and right-hand responses when the stimulus was the appropriate size for a response ("Go Trials"). The electrode sites used here are labeled C3′ and C4′ to indicate premotor sites just lateral to the standard C3 and C4 sites in the left and right hemispheres, respectively. The ERP waveform at the left-hemisphere site was more

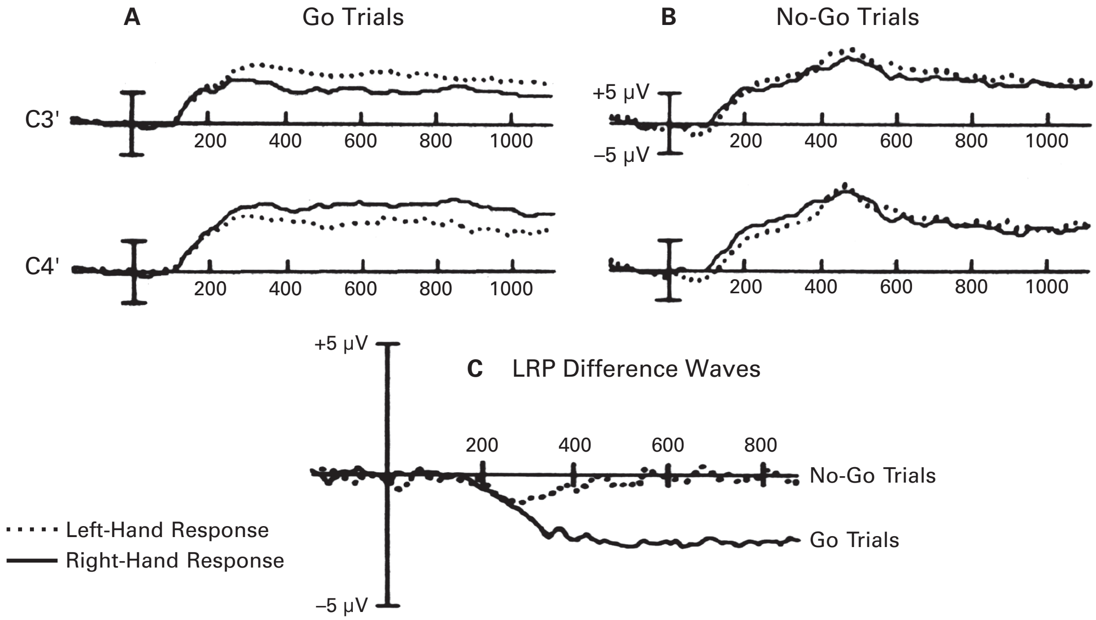

Figure S4.2
Grand average ERP waveforms from the study of Miller and Hackley (1992). The waveforms in panels A and B reflect trials on which the size of the letter indicated that a response should be made (“Go Trials,” panel A) or should not be made (“No-Go Trials,” panel B). These waveforms were recorded from left central (C3′) and right central (C4′) electrode sites. Broken lines reflect trials on which the shape of the letter indicated a left-hand response, and solid lines reflect trials on which the shape of the letter indicated a right-hand response. Panel C shows LRP difference waves for the go and no-go trials. Note that a brief LRP deflection was present on no-go trials, even though no response (or EMG activity) was present on these trials. Adapted with permission from Miller and Hackley (1992). Copyright 1992 American Psychological Association.

negative on trials with a right-hand response than on trials with a left-hand response, and the complementary pattern was observed at the right-hemisphere site. This effect began approximately 200 ms poststimulus and continued for at least 800 ms, which is typical for the LRP component.

Panel B of figure S4.2 shows the responses observed when the size of the stimulus indicated that no response should be made (“No-Go Trials”). From approximately 200 to 400 ms poststimulus, the ERPs were more negative over the left hemisphere when the shape of the stimulus was consistent with the right-hand response than with the left-hand response, and the complementary pattern was observed at the right-hemisphere site. After 400 ms, however, the ERP waveform was slightly more negative in both hemispheres for right-hand responses relative to left-hand responses.

Panel C of figure S4.2 shows the waveforms for the go and no-go trials after the LRP was isolated by means of a contralateral-minus-ipsilateral difference wave. For both go and no-go trials, the LRP began to deviate from baseline at approximately 200 ms. On go trials, the LRP continued until the end of the recording epoch. On no-go trials, in contrast, the LRP returned to baseline at approximately 400 ms. Thus, the brain began to prepare the response indicated by the shape of the letter on both go and no-go trials, even though no response was executed on no-go trials. This provides what I believe to be ironclad evidence that, at least under some conditions, response systems receive partial information about a stimulus before the stimulus has been completely identified.

Larger Issues

There are many unjustified conclusions that you might be tempted to draw from the waveforms shown in figure S4.2. First, you might suppose that subjects typically began preparing the response at approximately 200 ms poststimulus, the time point at which the LRP began to deviate from baseline. However, the onset of a response in an averaged ERP waveform reflects the earliest onset times, not the average onset times (see figure 2.6 and rule 5 in chapter 2). Thus, the 200-ms onset latency of the LRP in figure S4.2C reflects the fastest trials from the fastest subjects. Similarly, you might be tempted to assume that the LRP was only about half as large on no-go trials as on go trials. However, it is possible that the single-trial LRPs were just as large on no-go trials as on go trials but were present on only 50% of trials.

Fortunately, the main conclusions from this experiment do not depend on any unjustifiable conclusions. In fact, it doesn't even matter whether or not the waveforms shown in figure S4.2C reflect the same LRP component that has been observed in previous experiments. The simple fact that the distribution of voltage over the scalp was different when the stimulus signaled a contralateral response rather than an ipsilateral response is sufficient to provide solid evidence that the brain had begun to determine which response was associated with the letter. So it doesn't matter if the effect reflects an ipsilaterally larger P3 component, a contralaterally larger N400 component, or some new component that has never been observed before. Thus, this was a component-independent experimental design. That is, although the experiment was designed around the known properties of the LRP, it does not matter whether the effects observed in the experiment were due to this component. The mere finding of a contralateral-ipsilateral difference on the no-go trials—regardless of what underlying component produced it—was sufficient to conclude that partial information was extracted from the stimulus and used to prepare the response before the stimulus was fully identified.

I would like to make one additional observation about this experiment: The data shown in figure S4.2 are extremely clean. In the parent ERP waveforms (panels A and B), the waveforms are almost completely noise-free during the prestimulus interval. Even in the difference waves (panel C), which are plotted at a higher magnification, the prestimulus noise level is very small compared to the size of the LRP effects. Clean data lead to much greater confidence and much stronger conclusions: Even if the *p* value of an experimental effect passes the magical 0.05 criterion, noisy-looking waveforms make a poor impression on the reader and make it difficult for you and for your audience to have confidence in the results.

Example 3: Dual-Task Performance

Background

It is often difficult for a person to perform two tasks at the same time. For example, it is difficult to discuss the principles of ERP experimental design while driving a BMW 325xi at high speeds on a winding road in a snowstorm. Dual-task interference can also be seen in much simpler tasks, and many researchers have studied dual-task interference in the *attentional blink* paradigm (for a review, see Shapiro, Arnell, & Raymond, 1997). In this paradigm, a very rapid stream of approximately 20 stimuli is presented at fixation on each trial, and the subject is required to detect two targets (called T1 and T2) from among the nontarget stimuli. A typical experiment is illustrated in figure S4.3A. In this experiment, T1 and T2 are digits, and the nontarget stimuli are letters. At the end of each trial, subjects report the identities of the two digits. While subjects are processing T1, they may be unable to process T2 effectively, leading to errors in identifying T2. This would be expected to occur primarily when T2 is presented shortly after T1, and to assess the time course of the interference between T1 and T2, the lag between T1 and T2 is varied.

Figure S4.3B shows the pattern of results that is typically observed in this paradigm. T2 accuracy is severely impaired when T2 occurs two to four items after T1, but T2 accuracy is quite high at a lag of one item or at lags of five or more items. In contrast, T1 accuracy is typically quite good at all lags. This pattern of results is called the *attentional blink* because it is analogous to the impairment in accuracy that would result from a T1-triggered eyeblink. This is an inter-

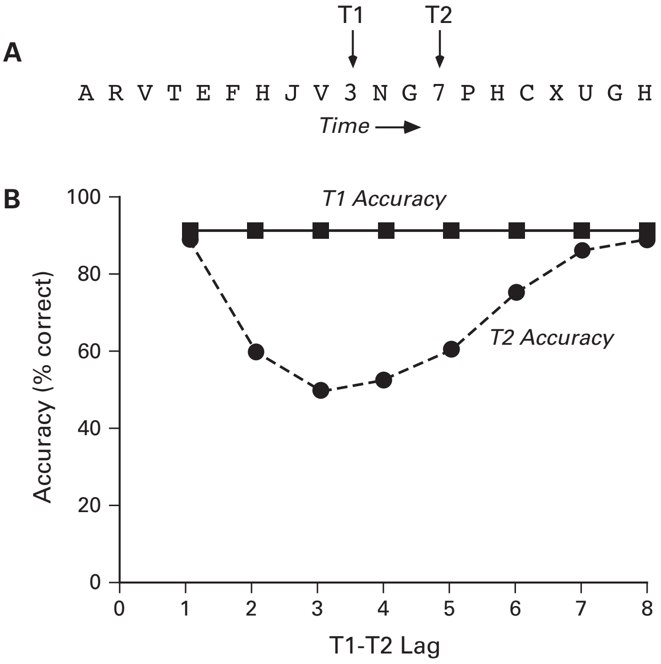

Figure S4.3
Experimental paradigm (A) and idealized results (B) from a typical attentional blink paradigm (based on the experiments of Chun & Potter, 1995). The stimuli are presented sequentially at fixation at a rate of 10 per second. T1 and T2 are digits, and the other stimuli are letters; subjects are required to report the identities of T1 and T2 at the end of the trial. The lag between T1 and T2 is varied, and although T1 accuracy is generally found to be independent of lag, T2 accuracy typically drops significantly at lags 2–4.

esting example of dual-task interference, and it has generated a great deal of research. One of the most fundamental questions is whether the attentional blink reflects a failure to identify T2 (Raymond, Shapiro, & Arnell, 1992) or whether T2 is identified but is not stored in a durable form in working memory (Shapiro, Raymond, & Arnell, 1994; Chun & Potter, 1995). A long time ago, Potter (1976) demonstrated that stimuli can be identified more rapidly than they can be stored in working memory, and so it seemed plausible that both T1 and T2 could be identified even though only T1 could be stored in memory when T2 occurred shortly after T1.

To test this hypothesis, Ed Vogel, Kim Shapiro, and I conducted several experiments in which we examined the P1, N1, P3, and N400 components in variants of the attentional blink paradigm (Luck, Vogel, & Shapiro, 1996; Vogel, Luck, & Shapiro, 1998). Here, I will discuss only the most definitive experiment, in which we focused on the N400 component. As discussed in chapter 3, the N400 component is typically elicited by words that mismatch a previously established semantic context. For example, a large N400 would be elicited by the last word of the sentence, “I opened the dishwasher and pulled out a clean eyebrow,” but a small N400 would be elicited if this same sentence ended with the word “plate.” The N400 can also be elicited with simple word pairs, such that the second word in PIG–HAT will elicit a large N400, whereas the second word in COAT–HAT will elicit a small N400. The N400 component is well suited for determining whether a word has been identified, because a semantically mismatching word cannot elicit a larger response unless that word has been identified to the point of semantic (or at least lexical) access. Thus, if a given word elicits a larger N400 when it mismatches the semantic context than when it matches the semantic context, this can be taken as strong evidence that the word was identified to a fairly high level. We therefore designed an experiment to determine whether the N400 component would be suppressed for words presented during the attentional blink period. Many previous experiments have examined the effects of semantic mismatch, so we were confident that we could adapt this approach to the attentional blink context.

Experimental Design

The stimuli and task for this experiment are illustrated in figure S4.4A. Each trial began with the presentation of a 1000-ms “context word” that established a semantic context for that trial. After a 1000-ms delay, a rapid stream of stimuli was presented at fixation. Each stimulus was seven characters wide. Distractor stimuli were seven-letter strings of randomly selected consonants. T1 was a digit that was repeated seven times to form a seven-character stimulus string. T2 was a word, presented in red, that was either semantically related or semantically unrelated to the context word that had been presented at the beginning of the trial. The T2 word was flanked by Xs, if necessary, to ensure that it was seven characters long.[1] At the end of each trial, the subjects made two unspeeded responses, one to indicate whether T1 was an odd digit or an even digit, and another to indicate whether T2 was semantically related or unrelated to the context word (e.g., PURSE–DOG or CAT–DOG). Related and unrelated T2 words occurred equally often, as did odd and even T1 digits. Each T2 word was presented twice for each subject (in widely separated trial blocks), appearing once as a related word and once as an unrelated word. This made it possible to be certain that any differences in the ERPs elicited by related and unrelated words were due to their semantic relationship with the context word rather than any peculiarity of the words themselves. The lag between T1 and T2 was either 1, 3, or 7; this restricted set of lags was necessary so that a sufficient number of trials could be obtained at each lag.

As illustrated in figure S4.4B, the rapid presentation of stimuli in the attentional blink paradigm leads to an overlap problem when ERPs are recorded. Specifically, the response to a given stimulus in the sequence lasts for hundreds of milliseconds, overlapping the responses to several of the subsequent stimuli. This makes it difficult to isolate the ERP elicited by T2 from the ERPs elicited by the preceding and subsequent stimuli. To overcome this problem, we computed difference waves in which we subtracted the response elicited by T2 when it was semantically related to the context word from the response elicited by T2 when it was an unrelated word. The responses to the other items in the stimulus stream should be essentially identical on related-T2 and unrelated-T2 trials, and this difference wave therefore provides a relatively pure measure of the brain's differential response as a function of the semantic relatedness of T2 relative to the context word. A large difference between related-T2 and unrelated-T2 trials can therefore be used as evidence that the T2 word was identified sufficiently to determine its semantic relationship to the context word.

Results

The results of this experiment are shown in panels C and D of figure S4.4. T2 accuracy (reporting whether T2 was semantically related or unrelated to the context word) was highly impaired at lag 3 relative to lags 1 and 7; this is the usual attentional blink pattern. In contrast, the N400 component elicited by T2 was equally large at all three lags. Thus, although the subjects could not accurately report whether T2 was related or unrelated to the context word at lag 3, the N400 wave differentiated between related and unrelated trials, indicating that the brain made this discrimination quite accurately. This result indicates that stimuli are fully identified during the attentional blink, but are not reported accurately because they are not stored in a durable form in working memory.

This conclusion relies on a hidden assumption, namely that the N400 component would be significantly smaller if the perceptual processing of T2 had been impaired at lag 3. This is a potentially problematic assumption because it is difficult to know what the relationship is between the amplitude of an ERP component and the speed or accuracy of a cognitive process. To avoid relying on this assumption, we conducted a control experiment to determine whether a decrease in the perceptibility of a word would cause a significant decrease in N400 amplitude (see experiment 3 in Vogel et al., 1998). In this experiment, we simply added random visual noise of varying intensity to the words. As the intensity of the noise increased, both behavioral accuracy and N400 amplitude declined in a roughly linear manner. Moreover, a change in accuracy that was comparable to the impaired accuracy observed at lag 3 in the main experiment led to a large and statistically significant decline in N400 amplitude. Thus, the absence of a decline in N400 amplitude at lag 3 in the main experiment provides strong evidence that the behavioral errors at lag 3 reflected an impairment that followed word identification.

These results also suggest that a great deal of processing can occur in the absence of awareness, because the brain apparently identified the T2 word at lag 3 even though subjects could not report the word (although it is possible that subjects were briefly aware of the T2 word even though they could not report its semantic relationship moments later).

Larger Issues

I would like to draw attention to three aspects of this attentional blink experiment. First, this experiment used ERPs to monitor a psychological process—word identification—that we sus-

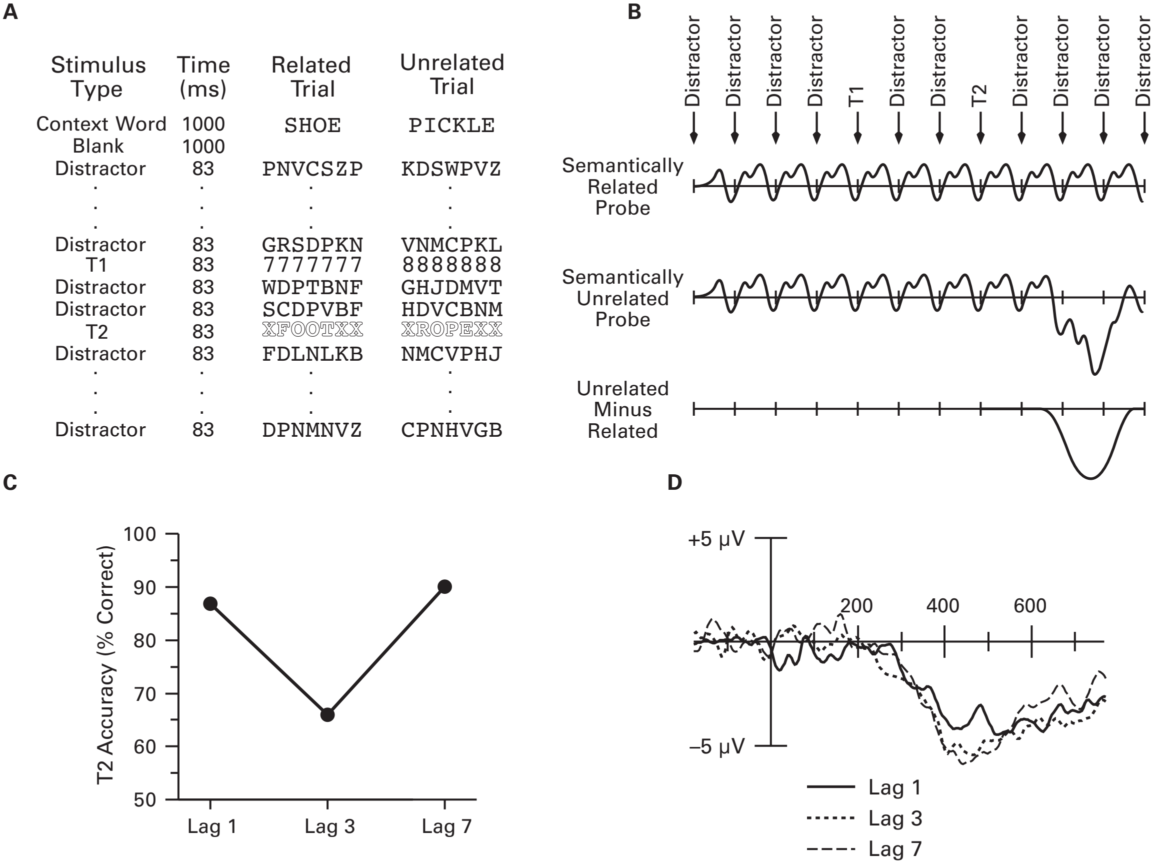

Figure S4.4
Paradigm and results from the study of Vogel et al. (1998). (A) Example stimuli. (B) Subtraction method used to overcome the overlap problem. (C) Mean discrimination accuracy for T2 as a function of lag. (D) Grand average ERP difference waveforms from the Cz electrode site, formed by subtracting related-T2 trials from unrelated-T2 trials.

pected was present but could not be observed in the subjects' overt behavior. Many ERP experiments are designed to show correspondences between behavioral results and ERPs, but it is often more informative to demonstrate an interesting pattern of both similarities and differences. I should note that we ran an additional attentional blink experiment that focused on the P3 wave, and we found that the P3 wave elicited by an oddball was completely suppressed at lag 3, consistent with the proposal that the attentional blink reflects an impairment in working memory. In this manner, we were able to show a theoretically sensible pattern of similarities and differences between overt behavior and ERPs.

A second notable aspect of this experiment is that we used the N400 component as an index of word identification even though the neural/psychological process that generates the N400 component is probably not directly related to word identification. Because word identification was a necessary antecedent to the presence of an N400 in this experiment, we could use the N400 as an indirect index of word identification. Similarly, the P3 wave can be used as an index of the time required to categorize a stimulus even though the neural/psychological process that generates the P3 component is probably unrelated to categorization (as discussed in the section on the P3 wave in chapter 3).

A third notable aspect of this experiment is that difference waves were used to isolate both the activity elicited by a single stimulus and a specific ERP component elicited by that stimulus. It is sometimes necessary to present stimuli in such close temporal proximity that the ERPs elicited by each stimulus will overlap each other in time, and difference waves can often be used to circumvent this problem (for additional examples of this approach, see Luck, Fan, & Hillyard, 1993; Luck & Hillyard, 1995; Luck, 1998a). Care must be used, however, because you cannot simply assume that the difference wave provides a pure estimate of the size of the component of interest. In the attentional blink experiment, for example, the overall N400 may have been smaller at lag 3 than at lags 1 and 7, even though the difference in N400 amplitude between related and unrelated T2 words was the same at all three lags. In this experiment, the key question was whether the brain could differentiate between related and unrelated T2 words, and the size of the overall N400 was immaterial. In fact, it didn’t even matter whether the activity in the unrelated-minus-related difference waves was the N400, the P3, or some other component; no matter what component it was, it demonstrated that the brain could distinguish between related and unrelated words. Thus, difference waves can be very useful, but they must be interpreted carefully.

The main body of chapter 4 describes eight strategies for dealing with the problem of isolating specific ERP components, and this experiment embodies all of them. First, it was focused on a specific component. If you read the whole series of four experiments (see Vogel et al., 1998), you will see that we examined several different components (P1, N1, N400, and P3), but different experiments were focused on different components, and each experiment was designed to isolate only one specific component. Second, we focused on a large component (N400) in the experiment presented here. The first experiment in the series actually focused on the P1 wave, but it was much smaller than the N400, and the conclusions were not as strong. Third, the experiment described here hijacked a component from a different domain. That is, although the N400 component is ordinarily associated with language research, our goal in using it was to ask how attention influences object identification. Fourth, we used a well-studied experimental manipulation (semantically related versus unrelated word pairs), and we factorially combined this with the manipulation of interest (the lag in the attentional blink paradigm). Because so many previous studies had used this manipulation, we knew we would get a large N400 effect. Fifth, we focused on difference waves, which allowed us to deal with the overlap problem and made it possible to demonstrate that the brain was differentially processing related and unrelated words during the attentional blink period. Sixth, we focused on a component—N400—that was easy to isolate. Seventh, we used this component to isolate the processes that necessarily precede it. That is, although the N400 does not directly reflect the process of word identification, an N400 difference between related and unrelated words implies that the words were identified. Eighth, this was a component-independent experimental design, because the conclusions depended only on the presence of a large related-minus-unrelated difference during the attentional blink period, irrespective of which underlying ERP component was responsible for this difference.

Of course, it's not usually possible to use all eight strategies in a single experiment. But if you use as many of these strategies as possible, your conclusions will be much stronger.

Note

1. We initially allowed each stimulus string to vary between four and seven characters. However, when presented at 10 stimuli per second, these varying lengths made the stimulus sequence very unpleasant to watch. We therefore made every stimulus seven characters wide. We couldn't easily find enough seven-character words that met a variety of linguistic criteria (e.g., relatively high frequency of occurrence), so we used words that were between four and seven characters and flanked the words with Xs when necessary to make them seven characters wide.