Journal of Experimental Psychology:
Learning, Memory, and Cognition
2005, Vol. 31, No. 3, 387–395

0278-7393/05/$12.00 DOI: 10.1037/0278-7393.31.3.387

Accessing Distant Premise Information: How Memory Feeds Reasoning

R. Brooke Lea
Macalester College

Elizabeth J. Mulligan
University of Colorado, Boulder

Jennifer Lee Walton
Bowdoin College

According to current psychological models of deduction, people can draw inferences on the basis of information that they receive from different sources at different times. In 3 reading-comprehension experiments, the authors demonstrated that premises that appear far apart in a text (distant) are not accessed and are therefore not used as a basis for logical inferences (Experiment 1), unless the premises are reinstated by a contextual cue (Experiment 2). In Experiment 3, the authors investigated whether these deductions are then integrated into the reader's situation model of the text. The results are interpreted in terms of a collaboration between memory-based text processing and higher level schema-driven logical reasoning.

Imagine that shortly after takeoff, an airline representative announces that you will have a choice of either fish or chicken for your in-flight meal. An hour later, a rumor spreads that the staff has run out of chicken. You can deduce that they will serve you fish. According to current mental-logic theories (e.g., Braine & O'Brien, 1998; Rips, 1994), you draw this inference rather effortlessly by combining the two premises (*chicken or fish* and *not chicken*) with the matching inference schema (or-elimination for this example: a or b; not a; therefore b) to produce the conclusion, in this case, that fish will be served. According to these psychological theories of deduction, readers make certain logical inferences when the necessary premises are simultaneously available in working memory (see Braine et al., 1995, or Lea, O'Brien, Fisch, Noveck, & Braine, 1990, for a full description of one such model). Indeed, one reason for suspecting that a mental logic would have evolved in preliterate humans is that it would have provided them with a means for integrating information arriving from different sources and at different times (Braine, 1990; Noveck, Lea, Davidson, & O'Brien, 1991). These models, however, are silent about how premises received at different times might end up in working memory simultaneously. In the example above, unless you recall what the meal choices were, learning that they are out of chicken cannot be used to produce new information. To be specific, it cannot yield a predictive inference about what you will be offered to eat. In the present research, we are concerned with the important role that low-level memory processes play in assembling information used by the inference engines in psychological theories of deduction.

Memory-Based Text Processing

Let us return to our hypothetical airline passenger. What conditions must exist for the logical inference to be made? Mental-logic models claim that a deduction will be made only when both premises are simultaneously available to the person. An hour has passed since the initial announcement, so how does the first premise (chicken or fish) arrive in working memory so that, together with the second premise, it can trigger the inference process? In several recent studies, researchers have directly addressed the question of how readers access information from long-term memory that is relevant to a text (e.g., Lea, Mason, Albrecht, Birch, & Myers, 1998; McKoon, Gerrig, & Greene, 1996; Myers & O'Brien, 1998). For example, Albrecht and Myers (1995) presented participants with passages in which a protagonist's goal was either satisfied or not satisfied in the early part of the passage. Following material that backgrounded the goal, a sentence appeared that was inconsistent with the unsatisfied goal condition but consistent with the satisfied goal condition. The authors found a slowdown in reading the target sentence in the unsatisfied goal condition, but only when a contextual cue in the pretarget sentence reminded the reader of the unsatisfied goal. For example, in a passage about someone whose goal was to make an airline reservation before midnight, the contextual cue was a *leather couch* on which she sat while trying to make the reservation. After the goal was backgrounded by an intervening episode, participants in the unsatisfied goal condition read sentences describing the protagonist sitting down and getting ready for bed (i.e., not fulfilling the goal)—participants showed no effect of the unsatisfied goal unless she sat down on the leather couch before

R. Brooke Lea, Department of Psychology, Macalester College; Elizabeth J. Mulligan, Psychology Department, University of Colorado, Boulder; Jennifer Lee Walton, Psychology Department, Bowdoin College.

We thank Ira Noveck, Barry Cohen, Jerome Myers, and Akiva Liberman for helpful conversations and suggestions regarding this research and Charles A. Weaver III and Debra L. Long for their comments on a draft of this article. We also thank Karin Cox, Aaron Mitchel, Julia Meyers, and Meghan Pedersen for assistance with data collection. This research was supported by a Keck Foundation Grant from Macalester College to R. Brooke Lea and a visiting professorship at the Center for Cognitive Sciences at the University of Minnesota.

Correspondence concerning this article should be addressed to R. Brooke Lea, Department of Psychology, Macalester College, 1600 Grand Avenue, Saint Paul, MN 55105. E-mail: lea@macalester.edu

getting ready for bed. That is, if the contextual cue *leather couch* was included, participants read the next two sentences significantly more slowly than they did if the cue was omitted. In short, the contextual cue reminded readers of the unsatisfied goal.

Albrecht and Myers (1995) discussed their results in terms of a resonance process that has been advanced by them and their colleagues (e.g., Lea et al., 1998; Myers & O'Brien, 1998; Myers, O'Brien, Albrecht, & Mason, 1994) and by McKoon and her colleagues (e.g., Gerrig & McKoon, 1998; McKoon et al., 1996). According to these researchers, the concepts and propositions derived from the sentence currently being processed signal to related elements in the discourse representation, triggering a spread of activation, or what others (e.g., Hintzman, 1986; Myers et al., 1994; Ratcliff, 1978) have labeled a resonance process. The more overlap between the information being read and information in memory, the greater the resonance. If the overlap is sufficiently high, then those elements will become more accessible and will be available for further processing. This is the heart of what McKoon et al. have labeled *memory-based text processing*. Work by Albrecht and Myers and McKoon et al. (1996) has shown that more than a single element from the discourse representation is often activated. For example, McKoon et al. found that elements from the entire introduction of the passage were reactivated after participants read a pronoun whose antecedent had last appeared in the introduction. Likewise, the contradiction effect reported by Albrecht and Myers relies on the contextual cue reactivating the entire goal from the first episode. It is worth noting that this memory-based text processing approach is consistent with the assumption by Ericsson and Kintsch (1995) that "If the episodic text memory that has been generated is coherent, text elements currently in focus of attention provide access to these structures" (p. 230).

Resonance, however, has its limitations. More constructive, schema-driven processes produce information that cannot simply be reactivated from memory. For example, our hypothetical airline passenger cannot "recall" that she will be served fish—that proposition has not been asserted. She can, however, deduce that she will be served fish by combining the information about her choices (fish or chicken) with the news that the chicken is finished. Thus, one cannot explain the inference that fish will be served by memory-based text processing theory alone.

We set out to examine the collaboration that must exist between low-level, bottom-up processes like resonance and the top-down, schema-driven processes they serve. In Experiment 1, we separated two logical premises in long texts so that readers were unlikely to have both premises in mind simultaneously. We found, not surprisingly, that readers did not make the logical inference when the premises were distant, though they did when the premises were presented close to each other in the text. In Experiment 2, we sought to reinstate the distant premises with contextual cues. In Experiment 3, we investigated whether these logical inferences are merely activated momentarily or are integrated into the reader's mental representation of the text.

Experiment 1

In Experiment 1, we presented readers with passages that contained two logical premises (e.g., *if it rains then we'll go to the museum* and *it rains*). In half of the passages, the two premises were separated by only 1 or 2 sentences (the *close* condition), whereas the premises in the other half of the passages were separated by an intervening episode that lasted for approximately 10 sentences (the *distant* condition). For example, consider the passage in the Appendix. In the introduction of this story, we meet Nathan, who is the best man at his best friend's wedding and who is being sent to buy more beer for the reception. In both versions of the story in Experiment 1, the premise in the introduction read: *if he did not get back before the food was served, then he would not be able to make his toast*. An intervening episode followed the first premise about Nathan having car trouble. In the close condition, the intervening episode is described in 1 sentence. In the distant condition, however, 10 sentences are devoted to Nathan's adventure. The purpose of the longer intervening episode is to background (i.e., clear from working memory) the first premise. Both conditions end with the same ending segment, in which the second premise is presented: *By the time he got back, people were halfway through the meal*. (The contextual cue, *messy*, was not presented in this experiment.) A sentence that was designed to be locally coherent unless the logical inference was made followed the second premise. We called this the *target* sentence. We examined reading times on this sentence in our analyses. In the example passage, the inference is that Nathan could not make his toast, and the target sentence stated that the groom's mother complimented Nathan on his toast.

If we succeeded in backgrounding the first premise in the distant condition, then readers would not have both premises in mind simultaneously, and no inference would be made. As a consequence, the target sentence would not contradict the readers' model of the text, and reading times would not be slowed. In the close condition, on the other hand, both premises would be in working memory and, as previous research has shown (e.g., Lea, 1995; Lea & Mulligan, 2002), the inference would be drawn. Because the inference concept is inconsistent with the target sentence, there should be a reading slowdown in the close condition.

Method

Participants. In return for course credit, 25 college students participated. All were native speakers of English.

Materials and design. We revised the experimental materials used in Albrecht and Myers (1995) to fit the current design. An example passage is presented in the Appendix. Each of the 18 experimental passages contained a conditional premise in the introduction (e.g., *if he did not get back before the food was served, then he would not be able to make his toast*). Following the introduction, an intervening episode introduced a new topic. In the close version of the story, the episode was resolved in 1 or 2 sentences. In the distant version, the intervening episode lasted for approximately 10 sentences. Both close and distant versions of the passage ended with an ending segment that contained the second premise that licensed the *modus ponens* inference (if p then q; p; therefore q), for example, *By the time he got back, people were halfway through the meal.* In asserting the antecedent of the conditional, we avoided using the same words as much as possible. For instance, the second premise in the example story stated that people were halfway through the meal, which conveys approximately the same information as *the food was served* without repeating any of the content words. This was important because we did not want the second premise itself to reactivate the first premise.

It is possible that the overlap of meaning information between the second premise and first premise would be sufficient to reactivate the first premise. If this is indeed the case, evidence that readers are making the inference should be found in both the close and distant conditions. In fact, Klin (1995) found that information from memory was reactivated by an overlap of traces that shared conceptual features but not surface features. However, several differences between Klin's materials and the present materials led us to believe that meaning overlap would not be sufficient to reactivate the first premise in the present study. As Klin pointed out, the strength of an item in memory determines its potential for resonance. The information that readers reactivated in her study in order to make causal inferences was elaborated extensively in the text, which should have resulted in a strong memory representation, and then was backgrounded with only 3 intervening sentences before the ending segment. In the present texts, the first premise is stated once with no elaboration and is then backgrounded with at least 10 intervening sentences before the ending segment. Because this format should produce a relatively weak memory representation for the first premise, we believe it unlikely that it will resonate strongly in response to meaning overlap from the second premise. Notably, Klin reported that in a prior study, she did not find evidence for inferences that depended on reactivating information from memory when the targeted information had not been elaborated in the text. Therefore, although we believe that resonance can be achieved through a variety of types of overlap with memory representations, the conditions in which each type of overlap will be sufficient to provide reactivation of the targeted information vary just as widely. Indeed, a central goal for memory-based text processing theorists is to account for these sources of variance. For the present materials, however, we did not predict that meaning overlap would be sufficient to reinstate the first premise.

The ending portion of the passages contained a target sentence that was locally coherent with the preceding text but inconsistent with the logical inference that could be made by integrating the two premises. The stories then ended with a final sentence. This final sentence also was inconsistent with the inference, and in previous research it was used as a second target sentence to investigate the time course of reactivation (e.g., Albrecht & Myers, 1995). Unlike this previous research, however, in the present study, we examined target sentence reading time to detect whether an inference had just been made (Experiments 1 and 2) or whether it was made and integrated into a mental representation of the text (Experiment 3). Thus, our hypotheses—and the subsequent data analyses—were limited to the first target sentence.

In addition to the 18 experimental passages (9 close and 9 distant), we designed 14 filler texts around the same form that contained no logical premises. Half of the fillers were approximately the length of the close passages, and half were similar in length to the distant ones. Each passage was followed by a comprehension question. Half of the comprehension statements were true, and the balance were false.

To summarize, participants read 32 passages, 18 of which permitted a logical inference and 14 of which did not. Participants read two practice stories before they began the experiment.

We constructed two stimuli lists such that one of the two versions of each experimental passage appeared on a list. Participants were randomly assigned to read one of the two lists; this resulted in two participant groups. The within-subjects factor (passage version) and between-subjects factor (participant group) yielded a 2 × 2 split-plot design. Participant group effects were not predicted.

Participants were seated before a computer screen in sound-insulated rooms. A keyboard, used to record participants' responses, was placed within easy reach.

Procedure. The stories were presented to participants in random order. Participants were asked to read the stories line by line at their own pace. When they had read and understood the sentence on the screen, they pressed the space bar and moved on to the next sentence. Sentences or parts of sentences replaced one another on the screen; no more than one line of text was present on the screen at a time. The *?/* button on the keyboard functioned as the "yes" button, and the *Zz* key served as the "no" button. The space bar was used to advance the text line by line. Participants were instructed to keep a finger on these three buttons at all times. At the end of each story, a sentence appeared on the screen, and participants were asked to indicate whether the sentence followed from, or was true of, the story they had just read. Participants' response to this comprehension statement advanced them to the next story.

Results and Discussion

For all of the experiments reported, F_1 refers to tests against an error term based on participant variability and F_2 refers to tests against an error term based on item variability. Results reported as significant had associated p values of less than .05. Outliers (response times that were more than three standard deviations from each participant's mean) were discarded, resulting in a loss of less than 5% of the data in each experiment.

As predicted, the target sentence reading times were significantly longer when the premises were close together (1,985 ms) in the text than when they were separated by a lengthy intervening episode (1,901 ms): $F_1(1, 23) = 5.79$, $MSE = 15{,}905$; $F_2(1, 16) = 5.61$, $MSE = 17{,}941$. Thus, participants appear not to have been making the target-sentence violating logical inferences when the premises were widely separated in the text (distant), but they were making the inferences when the premises were presented only a sentence or two apart (close). The latter result is consistent with previous results (Lea, 1995; Lea & Mulligan, 2002) in which premises were presented close together in texts. It is also consistent with our prediction that meaning overlap between the first and second premises would not be sufficient to reinstate the first premise. This experiment, however, lacked a no-inference baseline. Therefore it is possible that the inference was drawn in the distant condition, but for reasons that are not apparent, readers were not as affected by the inconsistency. In the remaining experiments, we addressed this possibility by including a no-inference baseline condition with distant passages. As we will see shortly, the results of Experiments 2 and 3 clearly indicate that readers do not make the logical inference in question when reading the distant passages used in Experiment 1.

The distant results suggest that, not surprisingly, if premises are widely separated in texts, readers do not recall the distant premise and therefore are unable to make the inference that follows. The question we wish to address concerns how distant premise information is reactivated so that the inference process is possible. In Experiment 2, we took up the question of reactivating distant premises.

Experiment 2

Albrecht and Myers (1995) used contextual cues to reactivate backgrounded goals. We used a similar approach by manipulating two factors in Experiment 2: cue (present or absent) and inference possibility (inference-permitting passage or no-inference permitting passage).

Cue Factor

In Experiment 1, we saw that meaning overlap between the first and second premises was not sufficient to reactivate the distant first premise so that an inference could be drawn. In this experiment, we look at whether surface feature overlap will be sufficient to reinstate the first premise. Certainly there are alternative methods for increasing the likelihood of resonance (e.g., elaborating the information in the first premise), but repeating a contextual cue near each premise has been shown experimentally to reinstate backgrounded information (e.g., Albrecht & Myers, 1995). We did not design the present experiments to test the boundary conditions for resonance; we simply wanted to create a situation in which resonance would be most likely to occur.

Using the distant passages from Experiment 1, we inserted a contextual cue (e.g., *messy trunk*) near the first premise at the beginning of the story and then manipulated whether that cue appeared when the second premise was presented toward the end of the passage. We used these specific noun phrases as cues because they have been shown to provide sufficient overlap to reactivate backgounded information (Albrecht & Myers, 1995). Our hypothesis was that if the contextual cue was sufficient to reinstate the first premise, then inferences would be made that contradict the target sentence, thus increasing reading times on those sentences.

Inference Factor

We also created no-inference versions of the passages in which two premises could not be combined to draw an inference. For example, for the passage presented in the Appendix, the first premise in the no-inference version was: *He thought about how cute the bride's sister looked in her bridesmaid's gown*. Note that this premise cannot be combined with the second premise (*By the time he got back, people were halfway through the meal.*) to draw a logical inference. The rest of the no-inference passage was identical to the inference version. These no-inference passages provided a baseline with which to compare reading times of the inference versions. Thus, if we found that the two inference versions were not different from each other, then (a) if they were different from the no-inference versions, we could conclude that inferences were made in both, whereas (b) if they were not different from the no-inference versions, we could conclude that inferences were made in neither.

Method

Participants. In return for course credit, 69 undergraduates participated. All participants were native speakers of English.

Materials, design, and procedure. We used the distant versions of the passages from Experiment 1, along with six additional passages for a total of 24 experimental passages. We developed four versions of each passage so that each passage either allowed an inference or did not (inference factor) and so that each passage contained a contextual cue that was either repeated in the ending segment or was not (cue factor). We crossed these two factors to create a 2 (inference) $\times$ 2 (cue) factorial design. The experimental passages were otherwise the same as those used in Experiment 1. We used 24 filler texts to break the correlations of patterns within the experimental passages.

We constructed four stimulus lists such that one of the four versions of each experimental passage appeared on a list. Participants were randomly assigned to read one of the four lists, thereby creating four participant groups. Thus, the two within-subjects factors (inference and cue) and the between-subjects factor (participant group) yielded a 2 $\times$ 2 $\times$ 4 split-plot design. Participant group effects were not predicted. There were no differences in procedure between Experiments 1 and 2.

Results and Discussion

Mean reading times for Experiment 2 are presented in Figure 1. In Experiment 1, we found no evidence for an inference effect with the distant (no-cue) versions of the passages, and that result was replicated in Experiment 2. There was a small and nonsignificant difference between the inference and no-inference conditions (M = 2,038 vs. M = 2,018, respectively) when the contextual cue was absent (both Fs $<$ 1). However, when the contextual cue was present, the difference between inference and no-inference became large (180 ms) and significant: M = 2,141 versus M = 1,961, F_1(1, 65) = 8.75, MSE = 97,964; F_2(1, 20) = 8.22, MSE = 36,139. Thus, with the cue present, there was a strong inference effect.

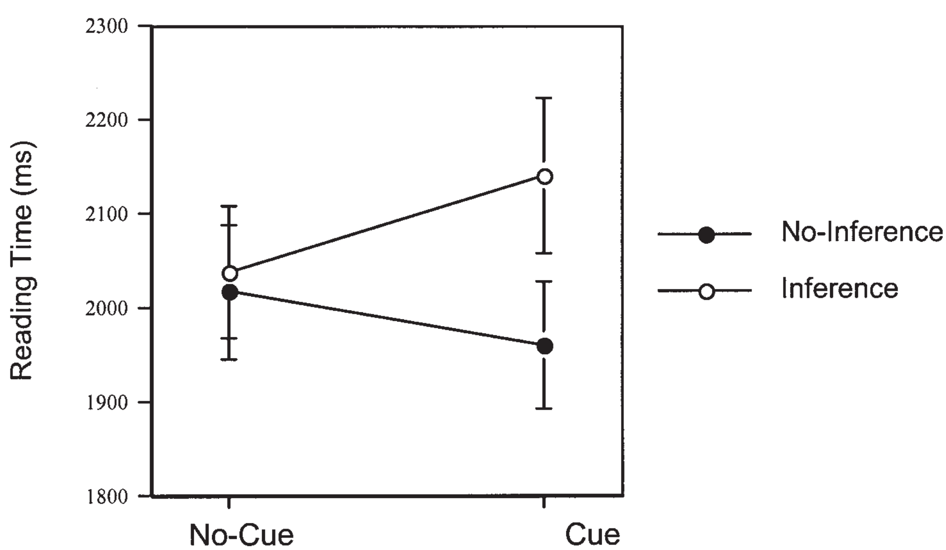

Figure 1. Mean target sentence reading times with standard error bars for Experiment 2. Solid circles indicate no-inference condition; open circles indicate inference condition.

Overall, the target sentences were read more slowly when they appeared in inference passages compared with when they were read in no-inference passages (2,090 vs. 1,990, respectively), though one can see from Figure 1 that the difference is due to differences in the cue conditions. An analysis of variance revealed a significant main effect for the inference factor by participants, $F_1(1, 65) = 5.94$, $MSE = 88{,}331$, though not by items, $F_2(1, 20) = 1.66$, $MSE = 60{,}476$, $p < .25$. More important for the present hypotheses, the inference/no-inference reading slowdown was much larger when the contextual cue was present (180 ms) than it was when the contextual cue was absent (20 ms). This interaction was significant, $F_1(1, 65) = 3.86$, $MSE = 88{,}482$; $F_2(1, 20) = 5.44$, $MSE = 37{,}847$. Thus, the presence of the cue led to a significantly bigger inference effect (reading slowdown) for the inference-permitting passages compared with the no-inference passages. Indeed, the presence of the cue in inference passages led to a significant 103-ms slowdown, on average: $M = 2{,}038$ versus $M = 2{,}141$, $F_1(1, 65) = 3.80$, $MSE = 90{,}801.8$; $F_2(1, 20) = 4.39$, $MSE = 28{,}908$. The comparable difference for no-inference passages (2,018 vs. 1,961) was in the opposite direction and was not significant (both Fs <1). Hence, the presence of the cue increased target sentence reading times in the inference passages without changing no-inference reading times appreciably.

The results of Experiment 1 suggested that readers made the logical inferences when the two premises were presented close together in the text but not when they were separated by a lengthy intervening episode. For our purposes, the critical difference between the two passage types was that the first premise was backgrounded in the distant versions of the texts and therefore was not in working memory when the second premise arrived. This effect was confirmed in Experiment 2, in which we added a no-inference control. In addition, in Experiment 2, we demonstrated that we could overcome the effects of textual distance by inserting a contextual cue near the second premise that reinstated the first premise, thus allowing the inference to be drawn.[1]

Several investigators (e.g., Corbett & Dosher, 1978; Klin, Murray, Levine, & Guzman, 1999; McKoon & Ratcliff, 1980) have distinguished between mere activation of information, on the one hand, and its integration into the discourse representation, on the other. In Experiments 1 and 2, we demonstrated that the logical inference was made. In Experiment 3, we addressed whether it is integrated into the reader's mental representation of the text.

Experiment 3

The results of Experiments 1 and 2, together with other experimental evidence (e.g., Lea, 1995; Lea & Mulligan, 2002) indicate that simple deductive inferences such as those specified in mental-logic models are routinely made during reading. Elsewhere, we have argued that these deductive inferences are elaborative in nature (Lea, 1995). That is, they are made even if they are not needed for maintaining local coherence. When the inferences concern what will happen at some future time in the story (as in Lea, 1995, and the present passages), they become predictive inferences. In these ways, the logical inferences under consideration here are comparable with other predictive inferences that have been studied extensively in the text processing literature.

After some initial disagreement over whether nonmandatory, predictive inferences were in fact made online (e.g., McKoon & Ratcliff, 1986; Potts, Keenan, & Golding, 1988), a consensus that inference concepts can be activated immediately after the predictive sentence, provided that the context is sufficiently predictive (e.g., Murray, Klin, & Myers, 1993; van den Broek, 1990), has formed. However, the question remains whether the explicit inference or only a general sense of the inference concept is incorporated into readers' mental representations of the text. For example, given a stimulus such as *The angry husband threw the delicate porcelain vase against the wall*, there is some debate about whether the specific inference *vase was broken* (Klin et al., 1999) or the more general inference *something bad happened to the vase* (Cook, Limber, & O'Brien, 2001) is incorporated into the representation.

One difference between logical inferences studied here and world-knowledge-based pragmatic inferences such as those described above is that the former lead to conclusions that were previously specified. For instance, in the airplane-food example presented earlier, the reader already knows what the choice of food is; if she reads that they have run out of chicken, she can deduce exactly what she will be offered, provided that the original choice is remembered. In Experiment 3, we investigated whether logical inferences lead to the sort of changes in a reader's situation model of a text that is comparable with a version of the text in which that information has been stated explicitly. We tested this by (a) including a condition in which the inference concept was presented explicitly in the text and (b) using a target sentence that was consistent with the inference concept and that was otherwise locally coherent. Thus, reading times should have been shorter when the target sentence already matched the reader's mental representation of the text (inference and explicit conditions) compared with when the sentence required integration into the representation. Using a consistency paradigm allowed us to examine in a more direct fashion the reader's mental representation of the text when the target information had been read explicitly compared with when it had been instantiated via inference.

In addition, to make a stronger statement about integration, we added three sentences between the second premise and the target sentence. If the inference concept is part of the reader's model of the text three sentences later, then the case for integration of the inference into the reader's mental representation is more secure.

Method

Participants. In return for course credit, 41 undergraduates participated. All participants were native speakers of English.

[1] A reviewer raised the possibility that our reinstatement cues were overspecified and therefore infelicitous in a way that might trigger processing anomalies like the repeated-name penalty (e.g., Gordon, Grosz, & Gilliom, 1993). The repeated-name penalty, however, occurs only when the referent is in discourse focus; when the referent is a nonfocused entity, there is no processing cost associated with referring to it with a repeated definite noun (Almor, 1999; Gordon et al., 1993). In the current texts, the contextual cue was out of discourse focus for at least 10 sentences before it appeared the second time. Therefore, no added cost should be associated with processing it, and it should not be particularly attention-grabbing. Furthermore, the texts contain many other adjective–noun pairs, some of which are repeated in the text. Therefore, as has been argued elsewhere (e.g., Albrecht & Myers, 1995, 1998) there is no reason to believe that there is something special about the cues that catches readers' attention.

Materials, design, and procedure. We modified the passages from Experiment 2 so that the target sentences no longer contradicted the logical inference that could be drawn from the second premise. Instead, the target sentences were written so that they included information that was consistent with the logical deduction. We used 18 experimental passages. We developed 3 of the 4 passage versions used in Experiment 2; we changed the no-cue/no-inference condition to an explicit condition in which the conclusion that could be inferred in the inference versions was presented explicitly in the text. Thus, the three conditions were explicit, inference/cue, and no-inference/cue. In the example passage presented in the Appendix, the logical inference is that Nathan will not make a toast; the explicit version of the passage stated that *he thought about how glad he was to be getting out of giving the best man's toast*. Accordingly, the target sentence was *As he put the champagne away, the groom walked over and accused him of not fulfilling his duties as best-man* (see Appendix). In addition, we added three sentences to separate the second premise and the target sentence. In the example passage presented in the Appendix, we inserted the sentences *The large reception hall was filled with the sound of chatter. The band was playing through a set of familiar songs. Everyone seemed to be enjoying themselves* before the target sentence. The passages, stimulus list preparation, and procedure were otherwise identical to those used in Experiment 2.

Results and Discussion

The crucial comparisons in Experiment 3 consisted of (a) target sentence reading times when we were most certain that readers are making the inferences in question (inference) compared with the reading times associated with passages in which the inference was explicitly stated (explicit) and (b) reading times for the inference passages versus reading times for the no-inference passages, in which the readers could not have been making the inference. Recall that for this experiment shorter, not longer, reading times are associated with readers making the inference in question. Taking the latter comparison first, the difference in reading times between inference and no-inference conditions was large (93 ms) and significant, $M = 1{,}686$ versus $M = 1{,}779$, $F_1(1, 40) = 10.81$, $MSE = 16{,}518$; $F_2(1, 14) = 6.414$, $MSE = 95{,}072$. This difference indicates that readers read the target sentences significantly faster when they appeared in passages that permitted a logical inference consistent with the target sentences compared with when they appeared in passages that did not sanction that inference. Furthermore, the reading time differences between the explicit and inference target sentences revealed that participants read the target sentences even faster in the inference condition than when the inference concept had been presented explicitly in the text ($M = 1{,}686$ vs. $M = 1{,}696$). This difference was not statistically significant (both Fs < 1). This comparison shows that participants read target sentences no slower when they had to infer information consistent with the sentences than when that information had been presented explicitly; if anything, they read the target sentences a little faster in the inference condition. Furthermore, the explicit condition reading times were significantly shorter than the no-inference reading times ($M = 1{,}696$ vs. $M = 1{,}779$): $F_1(1, 40) = 7.23$, $MSE = 19{,}479$; $F_2(1, 14) = 4.43$, $MSE = 71{,}367$. Thus, reading times were significantly shorter than control both when the target concept had been stated explicitly in the text and when it was inferred logically; reading times between the latter two were indistinguishable. This pattern of results is consistent with the conclusion that readers had incorporated the logical inference into their mental representation of the text by the time they were reading the target sentence; four sentences downstream, the inference concept remained part of the readers' mental representation of the text. Hence, these data support the conclusion that readers not only made the logical inference on presentation of the second premise but also integrated that inference into their situation model of the text.

Discussion

How are people able to logically integrate information that they receive from different sources at different times? The process can be broken down into three main steps. First, memory processes activate information from long-term memory so that relevant propositions can be considered together in working memory. In some cases, this may entail a conscious and deliberate search, though considerable research indicates that more automatic, bottom-up memory processes can easily account for the rapid retrieval of information that helps people understand what they are reading (e.g., Albrecht & Myers, 1995; Gerrig & McKoon, 1998; Lea et al., 1998; Long & Lea, in press; McKoon et al., 1996; Myers & O'Brien, 1998). Much of this work falls under the framework of a memory-based model of text processing (McKoon et al., 1996). According to this view, simple overlap between concepts can produce spontaneous reactivation of distant information through a process that can be described by the resonance metaphor. In the example text presented in the Appendix, the reintroduction of the contextual cue (messy trunk) in the ending segment was hypothesized to provide sufficient contextual overlap to reactivate propositions strongly linked to it in the introduction—such as the first premise. Although we did not measure activation per se in the present experiments, Experiments 1 and 2 together provided a compelling case that small, contextual cues did reactivate the distant premises.

In Experiment 1, we established that a first premise could be backgrounded by a lengthy intervening episode. Unlike passages in which the premises were presented close together in the text, the distant-premise versions of the texts showed participants reading smoothly through target sentences that would have appeared to be anomalous had the inference been made. In Experiment 2, we tested the contextual-cue hypothesis by adding the contextual cue just before the second premise in distant versions of the passages. We found that adding the cues reestablished the contradiction effect created by the inference. Remarkably, the only difference between the cue and no-cue versions of the stories was one or, at the most, two words that were quite unrelated (semantically) to the inference in question. Comparable reactivation effects with a similar set of passages were reported by Albrecht and Myers (1995).

However, reactivation cannot be the whole story. Our participants could not have recalled that Nathan did not make his toast, because that information was neither presented in the text nor available from world knowledge. Only after the conditional premise (that if he didn't get back before the meal was served, then he wouldn't be able to make his toast) was combined with the information that people were eating when he got back could a reader conclude, logically, that he had missed his toast. Thus, the second processing step is the integration of two premises to produce a logical inference. A variety of theories have specified a process by which two logical premises are integrated (e.g., mental models, mental logic, and pragmatic reasoning schemas) but by any account, the deduction results from a higher level process than is proposed in the first step. For example, according to one mental-logic approach (Braine & O'Brien, 1998), the simultaneous presence of two premises that match one of the model's inference schemas triggers a reasoning program that applies the relevant schema and produces an inference (see Braine & O'Brien, 1998, for a detailed description of this process). All three of the present experiments, along with other studies (e.g., Lea, 1995; Lea & Mulligan, 2002; Lea et al., 1990), provide evidence for this sort of logical integration of text information. The present work adds to previous research by using texts in which logical inferences are not made during reading. Indeed, an important point of the present work is to help establish the importance of collaboration among cognitive models of inference during reading; alone, these models can only explain behavior in highly constrained circumstances.

The final step is integration of the inference into the reader's mental representation of the text. Previous attempts to establish the fate of pragmatic (i.e., nonlogical) predictive inferences have found only minimal, or feature-based, integration of the inference concept (Cook et al., 2001; McKoon & Ratcliff, 1986). As stated previously, a potentially significant difference between logical deductions and predictive inferences based on world knowledge is that the latter are probabilistic in nature, whereas the former are deterministic (cf. Campion, 2004). Put a different way, the inference drawn in a logical deduction is not only highly constrained; it is precisely and explicitly specified. A reader would have to take the extra measure to "minimalize" or "featurize" an already fully fleshed-out inference to place it at the same level as an inference about the likely result of, for example, a vase being thrown against a wall (Potts et al., 1988). Instead, readers are more likely to encode and integrate the inference as specified in the premise.

The results of Experiment 3 accord well with this view. In Experiment 3, we replaced the inference-contradicting target sentences with ones that were consistent with the inference concept. We found that participants read the target sentences significantly faster both when the contextual cue was presented in inference-evoking passages and when the inference concept had been stated explicitly in the text. Either way, reading times were speeded compared with passages in which the target sentences were thematically consistent with, but not predicted by, the passages. Furthermore, if readers were taking the step beyond logical deduction to integrate the inference into their mental representation of the texts, then we would continue to find a facilitation effect on target sentences that matched the inference concept. Indeed, even three unrelated sentences downstream, target sentence reading times showed a statistically significant reading time advantage for passages in which the inference was drawn (or stated explicitly).

How do these stages fit together theoretically? Our position is that although inferences predicted by psychological models of deduction rely on the sort of low-level memory processes described in the first step above, they cannot be explained fully by them; resonance theory can account for how premises end up in working memory, but it cannot explain how a deduction is derived from them once they are there. Likewise, theories like mental logic are mute with respect to how premises end up in working memory simultaneously when they are not presented simultaneously. Thus, we have argued, understanding the teamwork between low-level, bottom-up processes like resonance and schema-driven, top-down processes like mental logic should be an inevitable goal of complete theories of text comprehension. A decade ago, McKoon and Ratcliff's (1992) minimalist hypothesis, although controversial, forced language theorists to consider just how basic were the comprehension phenomena they studied. It is an exaggeration, though perhaps only a small one, to say that the minimalist–constructionist dichotomy led to a theoretical bifurcation in which text researchers picked sides and worked one end of the processing continuum to the exclusion of the other (see Graesser, Singer, & Trabasso, 1994, for a review of the constructionist position.) A model proposed by Kintsch (1988) that predates the controversy, however, offers an elegant framework within which to view the collaboration between bottom-up and top-down processes in comprehension.

According to Kintsch's (1988) construction-integration (C-I) model, discourse comprehension proceeds in a two-stage fashion. In the first stage, a text base is constructed from the linguistic input that then serves as retrieval cues for associated concepts and propositions in the network of knowledge that the reader brings to the text. This dumb, bottom-up process results in an "initial, enriched, but incoherent and possibly contradictory text base" (Kintsch, 1988, pp. 166). This construction phase shares many fundamental elements of memory-based text processing. The C-I model's second stage is an integration process that strengthens the contextually appropriate elements and inhibits the inappropriate ones by a constraint-satisfaction process. Although the C-I model provides a detailed account of how a proposition from the knowledge base is contacted, it does not address how an abstract schema is activated. However, Kintsch did consider how sets of objects can be represented by a propositional schema. For example, he discussed a part–whole schema in the process of describing how an arithmetic word problem is solved. It follows, therefore, that the three processes under investigation in the present work, reactivation of a premise, schema application to draw an inference, and integration of the inference in a situation model, could, in principle, exist within a discourse comprehension framework such as Kintsch's C-I model. At present, however, no theory of comprehension is adequate to explain the results presented here.

The present work is also consistent with recent work tying memory-based text processing to forward, elaborative inference. Lea, Kayser, Mulligan, and Myers (2002) used passages in which the reunion of two protagonists reactivated backgrounded information about which only one of them knew. When the reunion was then followed by a discussion sentence in which the protagonists "chatted," the reactivation of the backgrounded concept was potentiated. Across four experiments using on- and off-line tasks, Lea et al. (2002) concluded that the discussion sentences were activating schematic knowledge about conversations and that readers were filling the topic slot by drawing inferences about the likely subject of conversation. Thus, a bottom-up process reactivated backgrounded information, and top-down processes used that information to make predictive inferences. That result is in close keeping with the present study.

In conclusion, a substantial amount of data can be explained by memory-based text processing theory and, also, by psychological models of deduction. But unless these two approaches can themselves be integrated, a complete theory of logic-based inferences in discourse is not possible. More generally, any computational or schema-driven comprehension process will require some principled account of how the relevant information is gathered for processing. At the same time, it is clear that humans go beyond the information given while reading by using processes beyond the scope of memory alone. Future theoretical work in this area will require a more precise specification of how these processes collaborate in a way that yields seamless comprehension. The present work constitutes a step in that direction.

References

Albrecht, J. E., & Myers, J. L. (1995). The role of context in accessing distant information during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 1459–1468.

Albrecht, J. E., & Myers, J. L. (1998). Accessing distant text information during reading: Effects of contextual cues. *Discourse Processes, 26,* 87–107.

Almor, A. (1999). Noun-phrase anaphora and focus: The informational load hypothesis. *Psychological Review, 106,* 748–765.

Braine, M. D. S. (1990). The "Natural Logic" approach to reasoning. In W. F. Overton (Ed.), *Reasoning, necessity and logic: Developmental perspectives.* Hillsdale, NJ: Erlbaum.

Braine, M. D. S., & O'Brien, D. P. (1998). *Mental logic.* Hillsdale, NJ: Erlbaum.

Braine, M. D. S., O'Brien, D. P., Noveck, I. A., Samuels, M. C., Lea, R. B., Fisch, S. M., & Yang, Y. (1995). Predicting intermediate and multiple conclusions in propositional logic inference problems: Further evidence for a mental logic. *Journal of Experimental Psychology: General, 124,* 263–292.

Campion, N. (2004). Predictive inferences are represented as hypothetical facts. *Journal of Memory and Language, 50,* 149–164.

Cook, A. E., Limber, J. E., & O'Brien, E. J. (2001). Situation-based context and the availability of predictive inferences. *Journal of Memory and Language, 44,* 220–234.

Corbett, A. T., & Dosher, B. A. (1978). Instrument inferences in sentence encoding. *Journal of Verbal Learning and Verbal Behavior, 17,* 479–491.

Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review, 102,* 211–245.

Gerrig, R. J., & McKoon, G. (1998). The readiness is all: The functionality of memory-based text processing. *Discourse Processes, 26,* 67–86.

Gordon, P. C., Grosz, B. J., & Gilliom, L. A. (1993). Pronouns, names, and the centering of attention in discourse. *Cognitive Science, 17,* 311–347.

Graesser, A. C., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review, 101,* 371–395.

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review, 93,* 411–428.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review, 95,* 163–182.

Klin, C. M. (1995). Causal inferences in reading from immediate activation to long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 1483–1494.

Klin, C. M., Murray, J. D., Levine, W. H., & Guzman, A. E. (1999). Forward inferences: From activation to long-term memory. *Discourse Processes, 27,* 241–260.

Lea, R. B. (1995). On-line evidence for elaborative logical inferences in text. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 1469–1482.

Lea, R. B., Kayser, P. A., Mulligan, E. J., & Myers, J. L. (2002). Do readers make inferences about conversational topics? *Memory & Cognition, 30,* 945–957.

Lea, R. B., Mason, R. A., Albrecht, J. E., Birch, S. L., & Myers, J. L. (1998). Who knows what about whom: What role does common ground play in accessing distant information? *Journal of Memory and Language, 39,* 70–84.

Lea, R. B., & Mulligan, E. J. (2002). The effect of negation on deductive inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28,* 303–317.

Lea, R. B., O'Brien, D. P., Fisch, S. M., Noveck, I. A., & Braine, M. D. S. (1990). Predicting propositional logic inferences in text comprehension. *Journal of Memory and Language, 29,* 361–387.

Long, D. L., & Lea, R. B. (in press). Have we been searching for meaning in all the wrong places? Defining the "search after meaning" principle in comprehension. *Discourse Processes.*

McKoon, G., Gerrig, R. J., & Greene, S. B. (1996). Pronoun resolution without pronouns: Some consequences of memory-based text processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 919–932.

McKoon, G., & Ratcliff, R. (1980). The comprehension processes and memory structures involved in anaphoric reference. *Journal of Verbal Learning and Verbal Behavior, 19,* 668–682.

McKoon, G., & Ratcliff, R. (1986). Inferences about predictable events. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 82–91.

McKoon, G., & Ratcliff, R. (1992). Inference during reading. *Psychological Review, 99,* 440–466.

Murray, J. D., Klin, C. M., & Myers, J. L. (1993). Forward inferences in narrative text. *Journal of Memory and Language, 32,* 464–473.

Myers, J. L., & O'Brien, E. J. (1998). Accessing the discourse representation during reading. *Discourse Processes, 26,* 131–157.

Myers, J. L., O'Brien, E. J., Albrecht, J. E., & Mason, R. A. (1994). Maintaining global coherence during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 876–886.

Noveck, I. A., Lea, R. B., Davidson, G. M., & O'Brien, D. P. (1991). Human reasoning is both logical and pragmatic. *Intellectica, 11,* 81–109.

Potts, G. R., Keenan, J. M., & Golding, J. M. (1988). Assessing the occurrence of elaborative inferences: Lexical versus naming. *Journal of Memory and Language, 27,* 399–415.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85,* 59–108.

Rips, L. J. (1994). *The psychology of proof: Deductive reasoning in human thinking.* Cambridge, MA: MIT Press.

van den Broek, P. (1990). The causal inference maker: Toward a process model of inference generation in text comprehension. In D. Balota, G. Flores d'Arcais, & K. Rayner (Eds.), *Comprehension processes in reading* (pp. 423–455). Hillsdale, NJ: Erlbaum.

Appendix

Example Passage Used in Experiments 1–3

Note that the logical premises are underlined and the ***contextual cue*** is in bold italics (though they were not set off in any way during the experiment). The contextual cue was not present in the ending segment for the no-cue conditions in Experiment 2.

Introduction (All Experiments)

It was early Saturday morning and it looked like it was going to be a beautiful day. Nathan was pleased because he was the best man at his best friend's wedding. The ceremony went smoothly. However, at the reception, they ran out of champagne and beer. Nathan had to go into town and buy some more. He had to hurry, so he ran to his car.

1st Premise

Experiments 1, 2, and 3 Inference condition	As he stared at his **messy** trunk, he realized that if he did not get back before the food was served, then he wouldn't be able to make his toast.
Experiments 2 and 4 No-inference condition	As he stared at his **messy** trunk, he thought about how cute the bride's sister looked in her bridesmaid's gown. "Maybe I'll be able to get her to dance with me," he thought.
Experiment 3 Explicit condition	As he stared at his **messy** trunk, he thought about how glad he was to be getting out of giving the best man's toast. He was sure he would trip over his words and maybe even get choked up.

Intervening Episode for Close Condition (Experiment 1 Only)

On his way to the store, he had a flat tire and had to walk to a gas station.

Intervening Episode for Distant Condition (Experiments 1, 2, and 3)

He quickly started to clear out his trunk to make space for the liquor. He threw away left over food bags from McDonalds and old Coke cans. On his way to the liquor store, Nathan had car trouble. He pulled over to the side of the road. Now he had to try to fix the car. He noticed that the front tire was flat. Unable to fix it himself, he started walking towards town. He passed by many food stores, but no gas stations. On the walk, he saw lots of trash and beer bottles strewn along the side of the road. He worried about getting his tuxedo dirty. Fortunately he was able to hitch a ride to the nearest gas station and found a mechanic who fixed the car.

Ending Segment (All Experiments)

When the mechanic was done, Nathan had the mechanic put the flat tire in the [***messy***] trunk. By the time he got back, people were halfway through the meal.

Ending Segment Filler (Experiment 3 Only)

The large reception hall was filled with the sound of chatter. The band was playing through a set of familiar songs. Everyone seemed to be enjoying themselves.

Target Sentence and Final Sentence (Experiments 1 and 2)

As he put the champagne away, the groom's mother complimented him on his wonderful toast. "I love speaking in public," Nathan replied.

Target Sentence and Final Sentence (Experiment 3)

As he put the champagne away, the groom walked over and accused him of not fulfilling his duties as best-man. "You're not afraid of speaking in public, are you?" he asked.

Received June 17, 2002
Revision received November 8, 2004
Accepted November 12, 2004 ■