Published in final edited form as:

Memory. 2012 July ; 20(5): 487–498. doi:10.1080/09658211.2012.682067.

Inferring Facts From Fiction: Reading Correct and Incorrect Information Affects Memory for Related Information

Andrew C. Butler,
Duke University

Nancy A. Dennis, and
Pennsylvania State University

Elizabeth J. Marsh
Duke University

Abstract

People can acquire both true and false knowledge about the world from fictional stories (Marsh & Fazio, 2007). The present study explored whether the benefits and costs of learning about the world from fictional stories extend beyond memory for directly stated pieces of information. Of interest was whether readers would use correct and incorrect story references to make deductive inferences about related information in the story, and then integrate those inferences into their knowledge bases. Subjects read stories containing correct, neutral, and misleading references to facts about the world; each reference could be combined with another reference that occurred in a later sentence to make a deductive inference. Later, they answered general knowledge questions that tested for these deductive inferences. The results showed that subjects generated and retained the deductive inferences regardless of whether the inferences were consistent or inconsistent with world knowledge, and irrespective of whether the references were placed consecutively in the text or separated by many sentences. Readers learn more than what is directly stated in stories; they use references to the real world to make both correct and incorrect inferences that are integrated into their knowledge bases.

Fictional stories sometimes take place in fantasy worlds that are completely divorced from reality, but more often they include information about real-world people, places, and events. For example, the best-selling novel by Dan Brown, *The Da Vinci Code*, contains references to European cities, history, art, and architecture. When fictional stories are based in reality, they represent a source from which people can learn information about the world. Indeed, fictional stories are commonly incorporated into the curriculum to complement nonfictional sources in classes like social studies (McGowan & Sutton, 1988) because they help students to learn while also increasing their enjoyment and participation (e.g., Palmer & Burroughs, 2002). However, while fictional stories often contain accurate representations of the real world, they also can contain inaccuracies because authors sometimes make mistakes or take liberties with the facts in order to tell a more entertaining story. *The Da Vinci Code* is largely based on fact, but it contains false statements as well, such as the contention that the Roman Emperor Constantine was baptized against his will (historical records indicate otherwise) or that a character in the book could be the sole descendent of Jesus and Mary (which would be genetically impossible).

The presence of both correct and incorrect information makes learning from fictional stories a double-edged sword: people can gain true knowledge about the world, but they can also acquire false knowledge. Previous research has shown that reading a fictional story that contains references to facts about the world improves people's ability to answer questions about those facts on a subsequent general knowledge test; however, if the story contains references to false information, people often acquire this misinformation and then reproduce it on the later test (e.g., Marsh, Meade, & Roediger, 2003; for review see Marsh & Fazio, 2007) even when they know the correct information (Fazio et al., 2011). For example, reading a reference to "*Patrick Henry riding a horse to warn that the British were coming*" increases the likelihood that students taking a general knowledge test (which does not refer back to the stories and is paired with a warning against guessing) will answer the question "*Who rode horseback through Boston to warn that the British were coming?*" with "*Patrick Henry*" instead of the correct answer of "*Paul Revere*". The goal of the present study was to explore whether the benefits and costs of learning about the world from fictional stories extend beyond memory for directly stated pieces of information, such as explicit references to places, things, and people. That is, do people use correct and incorrect information in the story to make deductive inferences about the real world and then integrate those inferences into their knowledge bases?

When authors of fiction make references to the real world, these pieces of information often relate to other information in the story. For example, the sentence "*Next she tried to pay for it with her Indian currency and gave him all the rupees she had in her purse!*" might be followed by a sentence that contains related information: "*Heck, I was just so happy that she hadn't spent all of the rupees bargaining in the markets of Mumbai.*" The first and second sentences (correctly) refer to the rupee as a form of currency used in India and the city of Mumbai, respectively. Of interest is whether readers use the transitive relationship (i.e., if $A = B$ and $B = C$, then $A = C$) between these sentences to make the deductive inference that Mumbai is a city in India, and then retain that inference, thereby acquiring additional true knowledge. Alternatively, if the word "*Indian*" is replaced with "*Pakistani*" in the first sentence, people might make an incorrect inference and acquire the false knowledge that Mumbai is located in Pakistan.

In the present study, we conceptualize the process of generating and retaining deductive inferences as consisting of two steps (for more fine-grained ideas see Alba & Hasher, 1983; Lea et al., 2005). First, when a reader encounters a new proposition that is related to previously read elements of the story and / or prior knowledge, these elements are retrieved from memory and used to make the deductive inference. Second, the reader integrates the inference into the mental representation of the text, and possibly into long-term memory as well. Normally these processes are studied in the domain of episodic memory (for an exception with children see Bauer & San Souci, 2010) with a focus on how people make inferences that they later believe that they read in a text (e.g., Johnson, Bransford, & Solomon, 1973). Instead, our goal is to examine whether these inferences can be integrated into people's general world knowledge, with the end result that people may gain both true and false knowledge about the world without it being directly stated in the story. In order for this result to happen, people must both make the deductive inference and integrate it into the knowledge base. We turn now to describing evidence that bears on whether or not each of these two steps will occur.

Evidence for the Generation and Retention of Deductive Inferences

Do readers of fictional stories generate deductive inferences from correct and incorrect information in the text and then retain those inferences? The literature offers some support for the two steps that would be necessary for this outcome to occur. First, when readers

encounter a new proposition, they must activate or retrieve related information from the text and then combine this information with the new proposition to make a deductive inference. Previous research has shown that readers immersed in fiction are less likely to retrieve related prior knowledge, relying instead on information from the text (e.g., Gerrig, 1989; Gerrig & Prentice, 1991). For example, Rapp (2008) found that when people read a scenario in which suspense was created about the outcome of well-known historical event (e.g., George Washington becoming the first president of the United States), they spent less time reading an inaccurate statement (e.g., “*he wrote that he would be unable to accept the nomination*”) relative to people who read a control scenario with no suspense. A few studies also suggest that readers sometimes make deductive inferences even when these inferences are not critical to comprehension (e.g., Champion, 2004; Lea, 1995; Lea et al., 1990). For example, Lea (1995) found that people were faster at making lexical decisions about words that were semantically related to an elaborative inference (i.e. an inference that adds extraneous detail but does not facilitate understanding of the story) when they were probed after reading a text from which that inference could be drawn relative to a non-inference control version of the text.

Second, once a deductive inference is made during text processing, the reader must integrate it into the mental model of the text and long-term memory. In the memory literature, there is a wealth of evidence to show that people remember the inferences that they make over long periods of time, regardless of whether those inferences are based on correct or incorrect information (e.g., Bransford & Franks, 1971; Johnson et al., 1973; Loftus, 1979; Sulin & Dooling, 1974). For example, Brewer (1977) had people read sentences (e.g., “*The flimsy shelf weakened under the weight of the books*”) that led to an inference that was neither objectively stated nor necessarily implied (“*the flimsy shelf collapsed under the weight of the books*”). On a subsequent cued recall test, people often remembered the inference, indicating that it had been generated during study and integrated into memory. Inferences can also be made during the process of remembering when people use their prior knowledge to help reconstruct a previously experienced event (e.g., Bartlett, 1932). Dooling and Christiaansen (1977) had people read the short biographical story about a fictional dictator named Gerald Martin and then take a delayed recognition test for sentences in the story. When people were told that the story was about Adolf Hitler, they often falsely recognized lure sentences that were related to Hitler, such as one that referred to Gerald Martin’s persecution of the Jews. Similarly, other studies have found that reading incorrect information slows down both the rejection of that incorrect information (e.g., Gerrig, 1989) and the retrieval of related correct information (e.g., Lewis & Anderson, 1976) on a later task.

Limits on the Generation and Retention of Deductive Inferences

Although there is some evidence to indicate that people generate and retain deductive inferences while reading fictional stories, it is not a forgone conclusion that both steps in the process will occur. First, most models of text processing agree that readers typically generate inferences that are critical to comprehension of the text (e.g., Graesser, Singer, & Trabasso, 1994; Kintsch, 1988; McKoon & Ratcliff, 1992), but many studies have found that people do not make elaborative inferences (e.g., Keefe & McDaniel, 1993; McKoon & Ratcliff, 1989; Potts, Keenan, & Golding, 1988). Other research suggests that prior knowledge plays a large role in text comprehension, especially with respect to making inferences (Bransford & Johnson, 1972; Bower, Black, & Turner 1979; Schank & Abelson, 1977); if readers rely on prior knowledge, they would be less likely to make false inferences based on incorrect information from the text because it may contradict prior knowledge. In addition, even if readers do rely on information from the text, they may fail to connect the new proposition to previously read information under certain conditions, such as when the

context of the story has changed (e.g., Albrecht & Myers 1995; Lea et al. 2005). Furthermore, there is also evidence to indicate that readers of fictional stories may not always use information from the text to make deductive inferences, even if that information is retrieved or re-activated. As described above, Rapp (2008) showed that reading a fictional story can decrease access to prior knowledge as measured by the time spent processing an inaccurate statement in the text; however, he also found that people slowed down considerably when reading a related statement that was accurate immediately after the inaccurate statement, which suggests that they noticed the discrepancy between the two statements. When such discrepancies are detected, readers may choose to rely on prior knowledge rather than information from a fictional story.

Second, other studies that focus on deductive inferences have found that people do not always retain them in long-term memory. For example, Lea et al. (1990) showed that people remember deductive inferences being in the story immediately after reading the story, but only when it was critical to the comprehension of the story (i.e. not when it was an elaborative inference). In addition, while studies have found that exposure to incorrect information slows down responding on a later task (e.g., Gerrig, 1989; Lewis & Anderson, 1976), subjects in these experiments rarely made judgment errors, such as endorsing an incorrect statement as correct or rejecting a correct statement. This finding suggests that the correct information was still accessible in prior knowledge even if incorrect information was also retained. Finally, it is possible that even if people integrate inferences into long-term memory, later they will notice incorrect inferences that conflict with their general knowledge and then reject them (e.g., the discrepancy detection principle; Loftus, 1992). In summary, the literature makes competing predictions about whether readers generate deductive inferences from correct and incorrect information in fictional stories, and whether they retain those inferences in long-term memory.

Present Research

In two experiments, we investigated whether people use correct and incorrect information to make deductive inferences about related information and then later use those inferences to answer general knowledge questions. Both experiments utilized the same general methodology. First, subjects read three fictional stories that contained both correct and incorrect information. Each piece of correct or incorrect information was presented in a critical sentence, which was followed by a related sentence that contained related information (see Methods section for an example). The critical sentence presented one of three types of information: correct information (correct frame condition), incorrect information (misleading frame condition), or a general reference to the information (neutral frame condition). Critically, the purpose of the neutral frame condition was to obtain a baseline measure of subjects' general knowledge (i.e. no exposure to correct or incorrect information in the stories). After reading the stories and completing a brief filler task, subjects took a general knowledge test that contained many filler questions as well as questions that could be answered with the information from the stories. To the extent that readers update their knowledge bases with inferences formed while reading, they should answer the general knowledge questions with these inferences; thus, exposure to the correct information should increase the number of correct responses on the general knowledge test relative to the neutral frame baseline (the measure of prior knowledge), while exposure to incorrect information should reduce the number of correct responses below baseline as well as increasing production of the target errors.

Experiment 1

The main goal of Experiment 1 was to ascertain whether people would make the inference that was suggested by the related sentence and then reproduce that inference on the subsequent general knowledge test. In the stories, the related sentence always followed immediately after the critical sentence. The general knowledge test included questions that tested the information that was directly stated in the critical sentence (no inference condition) as well as questions that tested the information that had to be inferred from the transitive relationship between the critical and related sentences (near inference condition). Replicating prior research, we predicted that readers would reproduce story information (both correct and erroneous) on a later general knowledge test. The novel question involved whether readers would go beyond the facts directly stated in the stories by making inferences and using them to answer the test questions.

We predicted that presenting the critical sentence in the correct frame would lead to a greater proportion of correct responses on the no inference general knowledge questions relative to the neutral (baseline) and incorrect frame conditions, which would replicate previous research. Of interest was whether the correct frame would also produce an increase in the proportion of correct responses to the near inference general knowledge questions. We also predicted that presenting the critical sentence in the misleading frame would increase the production of misinformation responses and decrease the production of correct responses on the no inference general knowledge questions (i.e. relative to the neutral frame condition, meaning that correct answers would be reduced below the level expected based on prior knowledge), which would also replicate previous research. However, it was unclear whether the misleading frame would have the same effect on the near inference general knowledge questions.

Methods

Subjects—Forty-eight undergraduates at Duke University participated for course credit or pay. All subjects were treated in accordance with the “Ethical Principles of Psychologists and Code of Conduct” put forth by the APA (2002).

Design—The experiment used a 3 (sentence frame: correct, neutral, misleading) x 2 (inference: no inference, near inference) within-subjects design. The dependent measures were the proportions of correct and misinformation responses on the general knowledge test.

Materials—The materials consisted of three fictional stories and a set of general knowledge questions. The three fictional stories with different themes (“Art Thief”, “Career Fair”, and “Travel”) were adapted from Marsh (2004). On average, each story was 1375 words long, consisted of 116 sentences, and contained references to six facts from the Nelson and Narens (1980) norms. The stories were modified to include six additional facts, each of which was related to one of the original Nelson and Narens facts. Thus, there were six pairs of facts in total. Every fact provided the answer to a question on the general knowledge test. Each pair of facts was incorporated into the story in the form of a critical sentence and a related sentence. Within each pair, the assignment of a fact to the critical or related sentence was counterbalanced across subjects (i.e. which determined whether the fact would be tested in the no inference or near inference condition).

In the story, the critical sentence and the related sentence were presented consecutively, and the critical sentence always preceded the related sentence. For example:

“‘I must see that famous ceiling artwork of Michelangelo in the Sistine Chapel,’ the man said. ‘To paint the ceiling of the Sistine Chapel and create ‘David,’ one of the most famous and beautiful sculptures known to man – what an artist!’ ”

The critical sentence presents the correct information that Michelangelo painted the ceiling of Sistine Chapel and the related sentence allows the reader to make the inference that Michelangelo also created the sculpture of “David”. Importantly, the related sentence never directly stated the to-be-inferred piece of information (e.g., Michelangelo created “David”). In the correct sentence frame condition, the critical sentence presented the correct information (e.g., *Michelangelo*). In the neutral frame condition, neither the correct nor the incorrect information was presented – instead the sentence only made a general reference to the critical fact (e.g., *that artist*), to later allow an estimate of what subjects knew prior to the experiment. In the misleading frame condition, the critical sentence presented the target incorrect information (e.g., *Leonardo da Vinci*).

Within each story, two pairs of facts were presented in each sentence frame condition (correct, neutral, misleading). Thus, there were 36 facts in total – 12 presented in the correct frame (six no inference, six near inference), 12 in the neutral frame (six no inference, six near inference), and 12 in the misleading frame (six no inference, six near inference). Six versions of each story were created to counterbalance sentence frame and inference conditions across subjects. The general knowledge test consisted of 100 short answer questions: 36 questions about the facts from the stories (18 in the no inference condition and 18 in the near inference condition) as well as 64 filler questions. The questions on the general knowledge test were blocked into two sets in order to counterbalance the order in which the two facts from each pair were tested. Although the test appeared as a single set of questions to subjects, the blocking ensured that the related questions (corresponding to the no and near inference conditions) were not tested immediately in succession. The first block included nine no inference questions and nine near inference questions, and the second block contained the corresponding question from each pair. Both blocks contained the same number of filler questions and the presentation of questions within a block was randomized.

Procedure—The entire experiment was conducted on a computer. First, subjects read the three fictional stories. They were instructed to read as quickly as possible while also making sure that they understood the story. Each story was presented on the screen one sentence at a time and subjects advanced to the next sentence by pressing the spacebar. After reading the stories, they engaged in a 7-minute filler task (solving visual spatial puzzles). Finally, subjects took a general knowledge test that consisted of 100 short answer questions. They were given instructions to answer as many questions as possible, but not to guess. If they could not answer the question, they were told to enter “I don’t know”. Test questions were presented one at a time and responding was self-paced. Subjects typed in their responses using the keyboard. Once they finished the test, subjects were debriefed and dismissed.

Results

All results, unless otherwise stated, were significant at the .05 level. Pair-wise comparisons were Bonferroni-corrected to the .05 level. Eta-squared and Cohen’s *d* are the measures of effect size reported for all significant effects in the ANOVA and *t*-test analyses, respectively. A Geisser-Greenhouse correction was used for violations of the sphericity assumption of ANOVA (Geisser & Greenhouse, 1958).

Scoring—One of the authors (NAD) and a research assistant independently scored all of the short answer responses. Both scorers were blind to condition and coded all the responses for a given question together in order to increase consistency in scoring. Cohen’s kappa

(Cohen, 1960) was calculated to assess inter-rater reliability. Reliability was perfect ($\kappa = 1.00$) – there were no disagreements in scoring.

General Knowledge Test: Correct Responses—The left side of Figure 1 shows the proportion of correct responses on the general knowledge test as a function of sentence frame and inference condition. The results were analyzed with a 3 (sentence frame) x 2 (inference) repeated measures ANOVA. Replicating prior research, there was a significant main effect of sentence frame [$F(2, 94) = 24.13$, $MSE = .06$, $\eta^2 = .29$]. Subjects answered more questions correctly after reading the critical sentence in the correct frame than reading it in the neutral frame [.78 vs. .63: $t(47) = 3.80$, $SEM = .04$, $d = .59$]; correspondingly, when they read the critical sentences in the misleading frame, they produced fewer correct responses relative to when they read it in the neutral frame [.63 vs. .53: $t(47) = 2.98$, $SEM = .04$, $d = .41$]. In other words, relative to what subjects knew prior to the experiment (as indicated by the neutral baseline), the expected benefits and costs of story-reading were observed. There was no significant main effect of inference ($F < 1$), but the interaction was marginally significant [$F(2, 82) = 2.65$, $MSE = .01$, $p = .08$]. The interaction is likely driven by the fact that the difference between the neutral and misleading frame conditions in the no inference condition (.64 vs. .51) was greater than in the near inference condition (.62 vs. .55). The latter difference is particularly important because it indicates that the deductive inferences that subjects made and retained while reading the story interfered with their prior knowledge (as measured by the neutral condition). Unfortunately, this comparison was only marginally significant [.62 vs. .55; $t(47) = 1.82$, $SEM = .04$, $p = .07$].

Since the information directly stated in the critical sentence (no inference condition) and the inferred information from the related sentence (near inference condition) were both tested on the general knowledge test, a follow-up analysis was conducted to assess whether the order in which the two questions were presented affected the results (i.e. responses to the near inference questions might be affected by prior responses to the no inference questions and vice versa). When presentation order (1st vs. 2nd) was included as a factor in a 3 x 2 x 2 ANOVA, there was no main effect of order ($F < 1$) and none of the interactions were significant (all F s < 1).

General Knowledge Test: Misinformation Responses—The left side of Figure 2 shows the proportion of misinformation responses on the general knowledge test as a function of sentence frame and inference condition. The results were analyzed with a 3 (sentence frame) x 2 (inference) repeated measures ANOVA. Again replicating prior research, there was a significant main effect of sentence frame [$F(2, 62) = 43.20$, $MSE = .02$, $\eta^2 = .38$]. When subjects read the critical sentence in the misleading frame, they were more likely to produce a misinformation response on the later general knowledge test relative to when they read it in either the correct frame [.19 vs. .02: $t(47) = 6.62$, $SEM = .03$, $d = 1.18$] or the neutral frame [.19 vs. .03: $t(47) = 7.30$, $SEM = .02$, $d = 1.16$]. Misinformation production in the neutral condition (the measure of prior knowledge) was low, suggesting that participants did not enter the experiment believing the target errors. The correct frame and neutral frame conditions did not differ ($t < 1$). Neither the main effect of inference ($F < 1$) nor the interaction [$F(2, 75) = 2.03$, $MSE = .01$, $p = .15$] reached significance.

A follow-up analysis was conducted to assess whether the order in which the two questions were presented affected the results. When presentation order (1st vs. 2nd) was included as a factor in a 3 x 2 x 2 ANOVA, there was a marginally significant main effect of order [$F(1, 47) = 4.01$, $MSE = .01$, $p = .051$, $\eta^2 = .003$] in which subjects produced a slightly higher proportion of misinformation on the first question asked ($M = .09$) than on the second question ($M = .06$). However, critically, there were no interactions between presentation

order and the other factors (all $F_s < 1$) and all of the other results were the same as in the 3 x 2 ANOVA that collapsed across presentation order.

General Knowledge Test: Other Incorrect and “I Don’t Know” Responses—The proportions of other incorrect ($M = .07$) and “I don’t know” ($M = .20$) responses on the general knowledge test were low and essentially uniform across all the conditions, and thus the analyses for these data will not be reported in the interest of brevity.

Discussion

Experiment 1 replicated previous research by showing that subjects acquired both correct and incorrect information that was directly stated in the stories and produced that information on the subsequent general knowledge test. However, the results also revealed a novel finding: subjects made deductive inferences using both the correct and incorrect information and then produced those inferences on the subsequent general knowledge test. Critically, performance on the general knowledge test in the correct and incorrect frame conditions was compared to the neutral frame condition that provided a baseline measure of subjects’ general knowledge (i.e. no exposure to correct or incorrect information in the stories). The finding that exposure to correct or incorrect information significantly improved and worsened performance, respectively, relative to baseline general knowledge effectively rules out the alternative explanation that these differences were due to knowledge that subjects had before reading the stories.

Interestingly, the benefits and costs of presenting correct and incorrect information, respectively, were roughly equivalent in the no inference and near inference conditions. That is, subjects were just as likely to produce the correct information later on when it was directly stated in critical sentence as when they had to make the deductive inference from the related sentence, and likewise for the incorrect information. However, one issue is that the difference between the neutral and misleading frame conditions in the near inference condition, which is a measure of whether reading the story interfered with subjects’ prior knowledge, was only marginally significant.

Experiment 2

The main goal of Experiment 2 was to replicate and extend the novel finding from Experiment 1. The proximity of the critical and related sentences in the near inference condition may have facilitated subjects’ ability to make the deductive inference because both pieces of information were in working memory. Thus, it is unclear whether subjects would still make the inference if the related sentence appeared later in the text. Several studies have shown that people reactivate previously read information when they encounter a new proposition with related information after reading many other unrelated sentences (e.g., Albrecht & Myers, 1995) and then use that information to make deductive inferences (e.g., Lea et al., 2005). However, they only make such inferences when the new proposition successfully reinstates the context in which the original information was presented. No prior study has looked at whether readers integrate deductive inferences into long-term memory when the inferences involve combining information across large portions of text. In order to investigate this question, the no inference condition from Experiment 1 was replaced by a “far inference” condition in which the critical sentence was separated from the related sentence by an average of 15 sentences. An additional goal of Experiment 2 was to re-examine whether there was any significant difference between the neutral and misleading frame conditions in the inference conditions, to allow stronger conclusions about whether reading the story interfered with subjects’ prior knowledge.

Methods

Subjects—Thirty-two undergraduates at Duke University participated in the study for pay.

Design—The experiment used a 3 (sentence frame: correct, neutral, misleading) x 2 (inference: near inference, far inference) mixed design. Sentence frame was manipulated within-subjects as in Experiment 1, but inference was manipulated between-subjects to assess whether the results of Experiment 1 generalize across design. The dependent measures were the proportions of correct and misinformation responses on the general knowledge test.

Materials—The materials were same except for the following changes. First, the piece of incorrect information for two of the items was changed to be more plausible because no subject produced either of the original errors in Experiment 1. Second, the stories were modified to create new versions for the far inference condition. In the near inference condition, the critical and related sentences occurred consecutively (as in Experiment 1). In the far inference condition, the critical sentence and the related sentence were separated by an average of 15 sentences ($SD = 4$). In order to ensure that the story remained coherent regardless of where the related sentences were placed (i.e. next to the critical sentence or later in the text), minor changes were made to the text surrounding the critical sentences and the related sentences. Overall, the stories were exactly the same in the near and far inference conditions, except for the location of the related sentence relative to the critical sentence. Third, all of the questions that tested the direct statement of information in the story (i.e. the no inference condition in Experiment 1) were dropped from the general knowledge test, in order to focus on the novel inference conditions. Thus, the test consisted of 82 questions: 18 questions about information in the stories and 64 filler questions.

Procedure—The procedure was the same as in Experiment 1, except that subjects were randomly assigned to one of the two inference conditions at the start of the experiment.

Results

Scoring—The data were scored in the same manner as for Experiment 1 by one of the authors (ACB) and a research assistant. Reliability was essentially perfect ($\kappa = 1.00$) and the author resolved the one disagreement in scoring.

General Knowledge Test: Correct Responses—The right side of Figure 1 shows the proportion of correct responses on the general knowledge test as a function of frame and inference condition. A 3 (sentence frame) x 2 (inference) repeated measures ANOVA was used to analyze the data. There was a main effect of sentence frame [$F(2, 80) = 23.73$, $MSE = .04$, $\eta^2 = .38$], but neither the main effect of inference ($F < 1$) nor the interaction [$F(2, 80) = 1.23$, $MSE = .04$, $p = .30$] were significant. Follow-up pair-wise comparisons indicated that subjects produced more correct responses when they had read the critical sentence in the correct frame than in the neutral frame [.68 vs. .53; $t(42) = 3.59$, $SEM = .04$, $d = .53$]; likewise, they produced fewer correct responses on the general knowledge test when they had read the critical sentence in the misleading frame than in the neutral frame [.53 vs. .38; $t(31) = 3.33$, $SEM = .04$, $d = .56$], indicating that reading the story interfered with subjects' prior knowledge.

General Knowledge Test: Misinformation Responses—The right side of Figure 2 illustrates the proportion of misinformation responses on the general knowledge test as a function of sentence frame and inference conditions. A 3 (sentence frame) x 2 (inference) repeated measures ANOVA revealed a main effect of sentence frame [$F(1, 56) = 38.01$, $MSE = .02$, $\eta^2 = .37$]. Again, misinformation production was low in the neutral condition,

indicating that subjects rarely believed the target errors without reading references to them in the stories. Subjects produced more misinformation responses when they had read the critical sentence in the misleading frame relative to reading it in either the correct frame [.20 vs. .02: $t(41) = 6.08$, $SEM = .03$, $d = 1.12$] or the neutral frame [.20 vs. .03: $t(41) = 6.42$, $SEM = .03$, $d = 1.08$]. There was also a significant main effect of inference [$F(1, 40) = 6.89$, $MSE = .01$, $\eta^2 = .03$], which was driven by the differences between the two inference conditions after having read misinformation, as reflected by a significant interaction between frame and inference condition [$F(1, 56) = 4.99$, $MSE = .02$, $\eta^2 = .05$]. Subjects answered more general knowledge questions with misinformation when the critical and related sentences were presented consecutively in the near inference condition than when the two sentences were separated in the far inference condition.

General Knowledge Test: Other Incorrect and “I Don’t Know” Responses—

Again, the proportions of other incorrect ($M = .12$) and “I don’t know” ($M = .26$) responses were low and varied only slightly as a function of condition.

Discussion

Experiment 2 replicated the novel finding in Experiment 1: subjects in the near inference condition made deductive inferences using both the correct and incorrect information and then produced those inferences on the subsequent general knowledge test. Although a higher proportion of misinformation responses were produced in the near inference condition in Experiment 2 relative to Experiment 1, this difference is likely due to the change in design (within-subjects in E1 vs. between-subjects in E2) and the fact that the incorrect information for two items was changed to be more plausible, among other possible factors. In addition, subjects made deductive inferences even when the critical and related sentences were separated by a substantial amount of text, and then produced those inferences on the final test. Interestingly, the near and far inference conditions showed roughly the same benefit when the text contained correct information, but the near inference condition produced a much greater cost than the far inference condition when the text contained incorrect information.

General Discussion

The present research shows that the benefits and costs of learning about the world from fictional stories extend beyond memory for directly stated pieces of information. Readers used the transitive relationship between the critical and related sentences to make deductive inferences, regardless of whether the sentences were placed consecutively in the text or separated by many sentences. The deductive inferences were retained and produced on the subsequent general knowledge test, indicating that subjects had acquired knowledge from reading the fictional stories. Importantly, this effect occurred irrespective of whether the critical sentence contained correct or incorrect information, which means that subjects acquired both true and false knowledge. In addition, exposing people to incorrect information interfered with their prior knowledge, as indicated by the reduction in the proportion of correct responses in the misleading frame condition relative to the neutral frame condition. We turn now to discussing these findings in more detail and placing them in the context of the broader literature.

When correct information about the world was presented in the stories, readers both learned the correct information (Experiment 1) and combined it with related information to make deductive inferences (Experiments 1 and 2). If interpreted as a two-step process, the results suggest that readers often generated deductive inferences by combining previously read correct information with related information in the text, and then integrated the inference into long-term memory. However, it is possible that subjects did not generate the inference

while reading the related information, but rather retained the information in the critical and related sentences and then generated the inference during the final test. One piece of evidence against this idea is that there was a significant difference in the production of misinformation between the near and far inference conditions in Experiment 2. If subjects were generating the inference when prompted by the question at test rather than while they were reading, then the distance between the two pieces of information should not matter because they are retained separately. Nevertheless, this possibility should be investigated in follow up studies.

Regardless of when the inferences are generated, these findings replicate and extend previous research that shows people can acquire true knowledge from reading fictional sources (e.g., Marsh et al., 2003; for review see Marsh & Fazio, 2007). More generally, these findings are consistent with many studies showing that people often activate related knowledge and generate inferences during encoding of events, and these inferences are incorporated into memory as if they had been part of the objective event (e.g., Bransford & Franks, 1971; Brewer, 1977; Johnson, Bransford, & Solomon, 1973; Roediger & McDermott, 1995). This finding also adds to the text processing literature in that it builds upon work by Lea and colleagues (Lea, 1995; Lea et al., 2005; Lea et al., 1990) by showing that deductive inferences can be integrated into long-term memory even if they are not critical to the comprehension of the story.

A similar pattern of results occurred when the story contained incorrect information about the world: readers learned the incorrect information (Experiment 1) and used it to generate deductive inferences (Experiments 1 and 2). These findings replicate and extend previous research showing that people learn incorrect information from fictional stories and integrate that information into their general knowledge about the world (e.g., Marsh et al., 2003; see Marsh & Fazio, 2007). Critically, these effects were measured relative to a neutral baseline that showed that people did not come into the experiment believing these errors. The findings fit well within the memory literature in which there is ample evidence to show that people generate and remember the inferences that they make regardless of whether those inferences are based on correct or incorrect information (e.g., Bransford & Johnson, 1972; Sulin & Dooling, 1974; Loftus, 1979). In addition, the results of our study add to the text processing literature by showing that readers make deductive inferences from incorrect information in the text as well, and then integrate the false inferences into long-term memory. The finding that readers made and remembered false inferences suggests that they had decreased access to related prior knowledge while processing the fictional story, as previous studies have indicated (e.g., Gerrig, 1989; Gerrig & Prentice, 1991). Importantly, much like inferences made from correct information, the inferences based on incorrect information were not central to comprehending the story, which shows that elaborative inferences can be retained under certain circumstances.

The present findings also have implications for education and, more broadly, learning about the world from fictional sources. Fictional stories contain a mix of accurate and inaccurate representations of the world. As such, people can acquire both true and false knowledge about the world from reading fictional stories. The results of our two experiments suggest that the effects of reading correct and incorrect information presented in fictional stories are much broader than previously believed (see Marsh & Fazio, 2007). Although the statement of correct and incorrect information can have a direct effect on what is learned, it can also affect memory for other related information, thereby producing additional benefits and costs to readers.

Acknowledgments

This research was supported by a Collaborative Activity Award from the James S. McDonnell Foundation's 21st Century Science Initiative in Bridging Brain, Mind and Behavior (EJM). Partial support was also provided by NIA grant T32 AG00029 (NAD). We thank Suedabeh Walker, Holli Sink, Barbie Huelser, and Rachel Kerns for their help with programming, data collection, and coding. We also thank the Marsh Lab Group for helpful comments on an earlier version of this manuscript.

References

- Alba JW, Hasher L. Is memory schematic? *Psychological Bulletin*. 1983; 93:203–231.
- Albrecht JE, Myers JL. The role of context in accessing distant information during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1995; 21:1459–1468.
- American Psychological Association. Ethical principles of psychologists and code of conduct. *American Psychologist*. 2002; 57:1060–1074. [PubMed: 12613157]
- Bartlett, FC. *Remembering: A study in experimental and social psychology*. Cambridge: Cambridge University Press; 1932.
- Bauer PJ, San Souci P. Going beyond the facts: Young children extend knowledge by integrating episodes. *Journal of Experimental Child Psychology*. 2010; 107:452–465. [PubMed: 20663513]
- Bower GH, Black JB, Turner TJ. Scripts in memory for text. *Cognitive Psychology*. 1979; 11:177–220.
- Bransford JD, Franks JJ. The abstraction of linguistic ideas. *Cognitive Psychology*. 1971; 2:331–350.
- Bransford JD, Johnson MK. Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior*. 1972; 11:717–726.
- Brewer WF. Memory for the pragmatic implication of sentences. *Memory and Cognition*. 1977; 5 (6): 673–678.
- Campion N. Predictive inferences are represented as hypothetical facts. *Journal of Memory and Language*. 2004; 50:149–164.
- Cohen J. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*. 1960; 20:37–46.
- Dooling JD, Christiaansen RE. Episodic and semantic aspects of memory for prose. *Journal of Experimental Psychology: Learning and Memory*. 1977; 3:428–436.
- Fazio LK, Barber SJ, Rajaram S, Ornstein PA, Marsh EJ. Creating illusions of knowledge: Learning errors that contradict prior knowledge. 2011 Manuscript submitted for publication.
- Geisser S, Greenhouse SW. An extension of Box's results on the use of F distribution in multivariate analysis. *Annals of Mathematical Statistics*. 1958; 29:885–891.
- Gerrig RJ. Suspense in the absence of uncertainty. *Journal of Memory and Language*. 1989; 28:633–648.
- Gerrig RJ, Prentice DA. The representation of ctional information. *Psychological Science*. 1991; 2:336–340.
- Graesser AC, Singer M, Trabasso T. Constructing inferences during narrative text comprehension. *Psychological Review*. 1994; 101:371–395. [PubMed: 7938337]
- Johnson MK, Bransford JD, Solomon DK. Memory for tacit implications of sentences. *Journal of Experimental Psychology*. 1973; 98:203–205.
- Keefe DE, McDaniel MA. The time course of durability of predictive inferences. *Journal of Memory and Language*. 1993; 32:446–463.
- Kintsch W. The role of knowledge in discourse comprehension: A construction- integration model. *Psychological Review*. 1988; 95:163–182. [PubMed: 3375398]
- Lea RB. On-line evidence for elaborative logical inferences in text. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1995; 21:1469–1482.
- Lea RB, Mulligan EJ, Walton J. Accessing distant premise information: How memory feeds reasoning. *Journal of Experimental Psychology: Learning, Memory & Cognition*. 2005; 28:303–317.
- Lea RB, O'Brien DP, Fisch SM, Noveck IA. Predicting propositional logic inferences in text comprehension. *Journal of Memory and Language*. 1990; 29:361–387.

- Lewis CH, Anderson JR. Interference with real world knowledge. *Cognitive Psychology*. 1976; 8:311–335.
- Loftus EF. The Malleability of Human Memory. *American Scientist*. 1979; 67:312–320. [PubMed: 475150]
- Loftus EF. When a lie becomes memory's truth: Memory distortion after exposure to misinformation. *Current Directions in Psychological Science*. 1992; 1:121–123.
- Marsh EJ. Story stimuli for creating false beliefs about the world. *Behavior Research Methods, Instruments, and Computers*. 2004; 36:650–655.
- Marsh, EJ.; Fazio, LK. Learning from fictional sources. In: Nairne, J., editor. *The foundations of remembering: Essays in honor of Henry L Roediger III*. New York: Psychology Press; 2007. p. 397–413.
- Marsh EJ, Meade ML, Roediger HL III. Learning facts from fiction. *Journal of Memory and Language*. 2003; 49:519–536.
- McGowan TM, Sutton AM. Exploring a persistent association: Trade books and social studies teaching. *Journal of Social Studies Research*. 1988; 12:8–16.
- McKoon G, Ratcliff R. Semantic associations and elaborative inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1989; 15:326–338.
- McKoon G, Ratcliff R. Inference during reading. *Psychological Review*. 1992; 99:440–446. [PubMed: 1502273]
- Nelson TO, Narens L. Norms of 300 general-information questions: Accuracy of recall, latency of recall, and feeling-of-knowledge ratings. *Journal of Verbal Learning and Verbal Behavior*. 1980; 19:338–368.
- Palmer J, Burroughs S. Integrating children's literature and song into the social studies. *The Social Studies*. 2002; 93:73–78.
- Potts GR, Keenan JM, Golding JM. Assessing the occurrence of elaborative inferences: Lexical decision versus naming. *Journal of Memory and Language*. 1988; 27:399–415.
- Rapp DN. How do readers handle incorrect information during reading? *Memory & Cognition*. 2008; 36:688–701.
- Roediger HL III, McDermott KB. Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, & Cognition*. 1995; 21:803–814.
- Schank, RC.; Abelson, RP. *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*. Oxford, England: Lawrence Erlbaum; 1977.
- Sulin RA, Dooling DJ. Intrusion of a thematic idea in retention of prose. *Journal of Experimental Psychology*. 1974; 103:255–262.

Figure 1. Proportion of correct responses on the general knowledge test as a function of sentence frame and inference condition in Experiment 1 (left side) and Experiment 2 (right side). Error bars represent 95% confidence intervals.

Figure 2.

Proportion of misinformation responses on the general knowledge test as a function of sentence frame and inference condition in Experiment 1 (left side) and Experiment 2 (right side). Error bars represent 95% confidence intervals.