8 Knowledge Neglect: Failures to Notice Contradictions with Stored Knowledge

Elizabeth J. Marsh and Sharda Umanath

Why do students often think that Toronto is the capital of Canada and that vitamin C wards off colds? Misconceptions are common across domains, including physics (e.g., Brown, 1992; McCloskey, 1983), health (e.g., Lee, Friedman, Ross-Degnan, Hibberd, & Goldmann, 2003; Wynn, Foster, & Trussell, 2009), chemistry (e.g., Nakhleh, 1992), gambling (e.g., Ferland, Ladouceur, & Vitaro, 2002), and ecology (e.g., Munson, 1994), among many others. Part of the problem is that there are many potential sources of misconceptions, including (but not limited to) misleading content in textbooks (e.g., Cho, Kahle, & Nordland, 1985), logical errors on the part of the learner (e.g., Clement, Narode, & Rosnick, 1981), other people such as family and friends (e.g., Landau & Bavaria, 2003), and fictional sources like movies and novels (e.g., Gouvier, Prestholdt, & Warner, 1988). In many cases, students learn a misconception because they have no knowledge with which to combat that misconception; for example, the average viewer knows little about amnesia and traumatic brain injury and thus finds no fault with movies depicting release from amnesia following a blow to the head. In such situations, the only solutions involve either avoiding the misconceptions in the first place (a nearly impossible task) or correcting the errors after the fact, when they can often prove quite tenacious (e.g., Butler, Fazio, & Marsh, 2011; Chi, 2008; Landau & Bavaria, 2003; Simons & Chabris, 2011).

In contrast, one might expect that having relevant prior knowledge would insulate the learner from picking up misconceptions. For example, it seems plausible that knowing that Neil Armstrong was the first astronaut on the moon would protect one from believing a story or movie claiming that John Glenn was the first astronaut on the moon. Prior knowledge often benefits learning and memory and specifically supports new learning (Anderson, 1981; Ausubel & Blake, 1958; Chiesi, Spilich, & Voss, 1979; Dooling & Mullet, 1973; Kole & Healy, 2007; Rapp, 2008). Thus, it seems intuitive that prior knowledge should make factual inaccuracies so obvious that one would dismiss them immediately and never reproduce them. Yet, surprisingly, the effects of exposure to misconceptions are not limited to cases where people are ignorant of the true state of the world. That is, having accurate knowledge does not guarantee that one

will notice an error (something inconsistent with one's stored knowledge), and furthermore, it does not protect the learner from later repeating that error. It is this particular problem that will be the focus of this chapter. We refer to this problem as *knowledge neglect* and use the term to refer to cases where people fail to retrieve and apply stored knowledge appropriately. For example, consider the reader who, after reading a short passage about a plane crash, answers the question "Where were the survivors buried?" (Barton & Sanford, 1993). This example demonstrates knowledge neglect because the reader has failed to retrieve and apply stored knowledge of the word "survivor" (successfully doing so would make this question an impossible one). Knowledge neglect goes beyond the immediate failure to notice a contradiction, as will be described below, and has consequences for later remembering.

In this chapter, we will describe experimental paradigms that demonstrate knowledge neglect and then discuss the memorial consequences of failing to retrieve and apply stored knowledge. Because error detection is key, we will consider how to promote monitoring, which is more challenging than it seems. We will then discuss how knowledge neglect and the resulting suggestibility compare to what we know about episodic memory, where there is a much larger literature about how people come to misremember events.

Two Examples of Knowledge Neglect

We begin by briefly overviewing two situations that we believe demonstrate knowledge neglect, with learners failing to retrieve and/or apply stored knowledge to a situation where it is highly relevant.

Example 1

Our initial work on knowledge neglect involved a laboratory analogue of learning from fictional sources. Novels, movies, and other fictional works often take place in familiar places and time periods and contain references to real people and events—and as such, can be sources of information about the world. However, fictional works, by definition, are not necessarily accurate, meaning that they can also be sources of misinformation about the world. To explore this situation, we use short stories that contain developed characters, dialogue, and plot and critically refer to facts about the world (Marsh, 2004). Some references are accurate (correct), others are plausible but incorrect (misleading), and the rest refer generally to facts without explicitly stating them (neutral). For example, a correct reference would refer to "*paddling across the largest ocean, the Pacific*," a neutral reference would simply allude to "*paddling across the largest ocean*," and a misleading reference would state "*paddling across the largest ocean, the Atlantic*." Critical for present purposes, we manipulated the preexperimental likelihood that readers would be familiar with the target facts, based on the Nelson and Narens

Table 8.1
Proportion of misleading, correct, and neutral sentences labeled as containing errors

	Misleading (Hits)	Neutral (False Alarms)	Correct (False Alarms)
Well-known facts	.23	.09	.10
Obscure facts	.24	.10	.15
M	.23	.10	.13

Note. Data from the story paradigm (Fazio & Marsh, 2008a).

(1980) norms. For the "well-known" facts, students in the norming study were able to answer about 70% of short-answer questions about these facts (and we would expect recognition to be higher). In contrast, the "obscure" facts were correctly produced by 15% of norming participants on average. Thus, if readers bring their stored knowledge to bear, they should be able to spot more of the errors that contradict well-known facts.

For present purposes, our focus is on whether participants noticed the errors in the stories (successful error detection), or missed them (knowledge neglect). To measure this, participants were explicitly instructed to look for errors while reading the stories and asked to press a key each time they noticed an error. Across experiments, participants did worse than predicted by the norms. For example, consider data from Fazio and Marsh (2008a), as presented in table 8.1. As a reminder, only sentences in the misleading condition actually contained factual inaccuracies. Readers did show some ability to discriminate accurate from erroneous sentences, but the hit rate ($M = 0.23$) was not much higher than the false-alarm rate (which ranged from 0.09 to 0.15, depending upon condition). Two points are crucial for present purposes. First, the detection rate is much lower than would be predicted by the norms; the norms predict that individuals would be able to correctly answer questions about 42% of items, and we would expect recognition rates to be higher. That is, we would expect more subjects to be able to recognize the largest ocean on Earth than would be able to produce the answer to the question "What is the largest ocean?" Second, the detection rate is similar regardless of whether misinformation contradicts well-known or obscure facts; 23% of the errors were successfully detected with no difference as a function of prior knowledge.

In short, readers of fictional stories experienced knowledge neglect because they struggled to identify errors that contradicted information likely to be stored in memory.

Example 2

A second example of knowledge neglect involves the Moses illusion, a semantic illusion in which people fail to notice erroneous presuppositions in questions (Erickson & Mattson, 1981). For example, when asked "How many animals of each kind did Moses take on the ark?" people often answer "two" even though they later demonstrate that they know the biblical reference should be to *Noah*. In a typical Moses illusion experiment, readers are explicitly warned that some of the questions will contain errors, and for each question, the task is to note if a question is unanswerable (because it contains an incorrect presupposition) or to answer it if the question is a valid one (Büttner, 2007; Kamas, Reder, & Ayers, 1996; Reder & Kusbit, 1991). Readers encounter a series of questions, including distorted ones that contain plausible but incorrect references (like the example above) and *un*distorted ones that contain correct references (e.g., "How many animals of each kind did Noah take on the ark?"). The illusion is measured by people's willingness to answer the distorted questions (Erickson & Mattson, 1981). Most critically for present purposes, readers complete a knowledge check at the end of the experiment, so that the experimenter knows exactly what knowledge an individual subject does versus does not have stored in memory (Baker & Wagner, 1987; Bredart & Docquier, 1989; Bredart & Modolo, 1988; Kamas, et al., 1996). This knowledge check typically takes the form of a multiple-choice test (e.g., "Who took two animals of each kind on the Ark": "Moses," "Noah," or "Don't Know"?) and normally includes a "don't know" option as the goal is to be sure of what an individual knows (as opposed to what he or she can guess; e.g., Büttner, 2007; Hannon & Daneman, 2001; van Oostendorp & de Mul, 1990). Strong demonstrations of the Moses illusion conditionalize the results so that the illusion is only examined for questions for which the subject successfully demonstrated knowledge on the final test.

For example, consider data from Bottoms et al. (2010). Readers were able to discriminate between distorted and undistorted questions, correctly saying "wrong" to distorted questions ($M = 0.42$; hits) much more often than mistakenly doing so for undistorted questions ($M = 0.02$; a low false-alarm rate). However, the Moses illusion was robust, with participants answering 35% of distorted questions as if the questions were unproblematic. We want to emphasize that answering distorted questions represents an illusion because we know participants had knowledge stored in memory which could have been used to detect the contradiction (the analysis only includes the 67% of items for which participants demonstrated accurate knowledge on the knowledge check). Overall, the Moses illusion is robust, with readers answering from 14% (van Jaarsveld, Dijkstra, & Hermans, 1997) to 40% (Hannon & Daneman, 2001) to 52% (Erickson & Mattson, 1981) to 77% (Barton & Sanford, 1993) of distorted questions, depending upon the particular experiment.

Why Does Knowledge Neglect Occur?

Why would people be so poor at detecting contradictions with their stored knowledge? Gilbert and colleagues have argued that people automatically believe information when they read it and that a second processing step is required to "unbelieve" information and label it as false (see Gilbert, 1991, for a review). This claim is supported with data from experiments like the following: Subjects learn translations of Hopi words (e.g., *A twyrin is a doctor*), each of which is labeled as correct or incorrect immediately after its presentation. Critically, sometimes there is a dual task during the presentation of the true/false label; a tone sounds, and the subject has to press a button registering that tone. The question is whether completing a secondary task systematically affects later truth judgments (*Is a twyrin a doctor? true, false, no information, never seen*). That is, how do subjects later reconstruct the truth of the statements, given that the secondary task presumably disrupts their ability to process (encode) the true/false label? The key result is that subjects are much more likely to incorrectly say "true" to statements previously labeled as false than they are to say "false" to statements previously labeled as true (Gilbert, Krull, & Malone, 1990). That is, when the secondary task disrupts that second step of "unbelieving," subjects show a bias to call a statement "true" rather than to say that they did not have any knowledge about it.

When reading stories or detecting/answering distorted questions, the participant is doing a lot and may not have the processing resources available to assess the truth value of what is being read. The reader of a story is processing a plot line, keeping track of characters, and more generally, building a mental model of the text (e.g., Bower & Morrow, 1990; Johnson-Laird, 1983); catching contradictions with stored knowledge is not the main focus of the reader. Similarly, the subject in a Moses experiment is focused on answering general knowledge questions, and deciding not to answer a question (because it contains a contradiction with stored knowledge) requires inhibiting one's natural impulse to answer. In both situations, we argue that the monitoring task is not the primary one, if only because errors are relatively infrequent. At least some data suggest that making the task easier (freeing resources for monitoring) may help. For example, Moses subjects are more likely to notice distortions when the task is to verify statements rather than answer questions. That is, people are better at saying "false" to "Snow White slept for a hundred years after she pricked her finger" than at refusing to answer the question "For how many years did Snow White sleep after she pricked her finger?" (Büttner, 2007). Of course, the illusion was not eliminated in the statement condition, but it was smaller when the processing load was reduced.

Even if subjects do engage in evaluation (the assessment stage, using Gilbert's terms), they may not be successful. We will return to this problem later in the chapter. However, briefly, we can draw on the partial match hypothesis (Reder & Kusbit, 1991) to explain why subjects sometimes fail to successfully evaluate the information they encounter. The partial match hypothesis is a theory developed to explain the Moses illusion but is one that can be applied more broadly to other situations that appear to involve knowledge neglect. The idea is that the Moses illusion occurs when the subject fails to notice a mismatch between what is retrieved from memory (e.g., details related

to the biblical story of Noah and the ark) and what one is reading (e.g., a reference to Moses). When what is retrieved is very similar to what one is reading, the match is close enough and is accepted, with the result that highly related (plausible) errors go unnoticed. The reader is much more likely to notice a problem with "How many animals of each kind did Nixon take on the ark?" because there is no partial match, whereas an incorrect reference to Moses overlaps with Noah in multiple ways (e.g., both are men who appeared in the Old Testament of the Bible), and thus, the match is close enough to accept (Erickson & Mattson, 1981; van Oostendorp & de Mul, 1990). A similar idea is that language processing is often shallow and incomplete, leading to distorted but often "good enough" semantic representations (Ferreira, Bailey, & Ferraro, 2002). In the case of knowledge neglect, these incomplete representations are not, in fact, good enough but slip through anyway.

Finally, there are situations that may discourage monitoring or otherwise make noticing factual inaccuracies more difficult, even without a concurrent processing load. Consider the specific case of consuming fiction; readers and moviegoers often enjoy novels and films that contradict the true state of the world (e.g., science fiction) or for which they already know the ending (e.g., that JFK was assassinated). One by-product of this behavior involves what Gerrig (1989) calls "the suspension of disbelief." Gerrig's subjects were slower to verify the well-known fact "*George Washington was elected first president of the United States*" after reading statements that created some doubt about this outcome (e.g., *Washington wanted to retire…*). In such cases, it is almost as if readers/viewers are insulated from their prior knowledge, so that they can enjoy themselves (see also Horton & Rapp, 2003, for a demonstration of the consequences of a reader's taking the point of view of a character). To measure people's involvement in a story (transportation), Green and Brock (2000) developed a scale that includes items such as "*I found myself thinking of ways the narrative could have turned out differently*" and "*I found my mind wandering while reading the narrative*" (reverse scored). Readers who reported greater transportation into a narrative were less likely to indicate that parts of the narrative "rung false" to them (see also Gerrig & Rapp, 2004). In short, some situations may encourage knowledge neglect because participation means ignoring one's general knowledge.

Overall, whether knowledge is inferred from norms (the story paradigm) or directly measured from individual performance on a knowledge check (the Moses illusion), individuals are often poor at noticing errors that contradict information stored in memory. This problem (knowledge neglect) may occur because the situation discourages monitoring (i.e., to allow enjoyment of fictional materials) or because the monitoring is difficult and careful processing is not possible. Monitoring may also fail when the information being evaluated is similar enough to what is stored in memory. The larger issue is that such lapses have consequences for memory, as we will describe in the next section.

Memorial Consequences of Knowledge Neglect

Knowledge neglect is more than a momentary lapse in attention; knowledge neglect has memorial consequences. That is, when the reader fails to notice a contradiction between what he or she is reading and his or her stored knowledge, this failure results in the error's being stored in memory. Because this error was recently encountered, it is accessible in memory and comes to mind fluently on later general knowledge tests. The relative ease with which something comes to mind at test is interpreted as confidence in one's response (Kelley & Lindsay, 1993), meaning that an error that pops to mind quickly is likely to be interpreted as a correct answer. Such memorial consequences are observed in both of the situations described earlier. When a general knowledge test follows the story-reading phase (fiction paradigm) or question-answering phase (Moses illusion), prior exposure to errors affects subjects' responses on the general knowledge test.

For example, figure 8.1 shows the use of the story errors on later memory tests (suggestibility) in the story paradigm. After the story-reading phase, participants answer general knowledge questions, including ones that can be answered with information from the stories (e.g., "*What is the largest ocean on Earth?*"). Readers are very suggestible, meaning that they later use story errors to answer the general knowledge questions

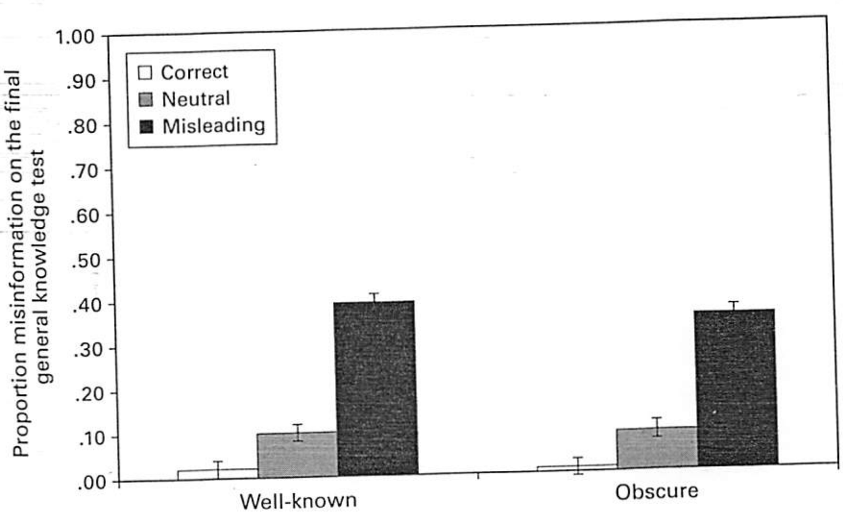

Figure 8.1
The proportion of final general knowledge questions answered with target misinformation as a function of fact framing during story reading (correct, neutral, misleading). The left portion shows responses to questions about well-known facts and the right portion shows responses to questions about lesser-known facts. Data from young adults in Umanath and Marsh (2012).

(e.g., answering *Atlantic*; see Marsh & Fazio, 2007, for a review). As shown in figure 8.1, people answer some questions with the target misinformation even after reading neutral references, likely because the errors are plausible ones and there is some pre-experimental belief in them. However, this baseline is low, and the production of misinformation errors increases dramatically after reading misinformation in the stories. Warning readers that authors of fiction often take liberties with the truth does nothing to reduce this pattern of suggestibility (Marsh & Fazio, 2006, experiment 1), although it does make readers a bit more conservative overall on the final test. Critically, suggestibility is similar regardless of (estimated) prior knowledge. That is, participants are just as likely to later incorrectly claim that the Atlantic is the largest ocean on Earth as they are to state that Jefferson invented dynamite, even though it is much more likely that they know that the Pacific is the largest ocean than that Nobel invented dynamite.

While it is possible that norms (especially old ones like the Nelson and Narens norms, which were published in 1980) may not be the best measure of what a given individual knows, similar results are obtained when individual knowledge is directly measured. To establish knowledge, we simply administered a pretest to measure what each individual participant knew two weeks before reading the stories. Participants were just as likely to later answer general knowledge questions with story errors regardless of whether or not they had demonstrated correct knowledge two weeks earlier (Fazio, Barber, Rajaram, Ornstein, & Marsh, 2013). The results were similar when the analyses were limited to facts produced on the pretest with the highest level of confidence. Of course, knowledge on this pretest was measured two weeks earlier, leaving open the possibility that for some reason knowledge was forgotten before the experimental session where the stories were read (although this possibility seems quite unlikely for the highest-confidence responses). However, in a recent experiment, participants still showed robust suggestibility when the pretest was administered immediately before exposure to the stories, meaning that the relevant knowledge was highly accessible (Mullet, Umanath, & Marsh, 2012). Thus, even when we are sure participants have stored knowledge and are highly confident in those facts, they still reproduce story errors.

Only a handful of studies have examined the memorial consequences of exposure to distorted questions in the Moses illusion, but with these cases a similar pattern occurs (Bottoms et al., 2010; Umanath, Dolan, & Marsh, 2014; see also Kamas et al., 1996). In a typical study, after the initial question-answering phase (where the task is to detect errors and answer valid questions), participants take a general knowledge test containing questions that can be answered with information from the initial questions (e.g., "Who took two animals of each kind on the Ark?"). As in the story paradigm, participants are suggestible, using previously seen errors to answer the later general knowledge questions (e.g., answering *Moses* about 6% of the time; Umanath et al., 2014). Participants rarely answer with the target errors if they have not seen the distorted question during the initial phase (1%). Critically, both analyses are limited to items for

Table 8.2
Probability that an error was reproduced on the final general knowledge test as a function of whether it was detected or missed during the initial phase of the experiment

	Detected	Missed (Knowledge Neglect)
Story paradigm	.09	.36
Moses illusion	.00	.09

Note. The Moses illusion data are conditional on participants having demonstrated knowledge of the correct answer on a later test. Data from Umanath and Marsh (2012) and Bottoms et al. (2010).

which participants are able to select the correct answer on the knowledge check, meaning that suggestibility occurs despite participants' ability to demonstrate the correct knowledge on another test. These results are consistent with evidence of the memorial consequences of knowledge neglect in the story paradigm.

In both paradigms, suggestibility is primarily driven by instances where participants suffered knowledge neglect, meaning they failed to notice the errors initially during the encoding phase (Bottoms et al., 2010; Fazio & Marsh, 2008a; Umanath & Marsh, 2012; Umanath et al., 2014). Table 8.2 shows the probabilities that an error (either in a story or a distorted Moses question) was reproduced as a function of whether or not the reader caught the error during the initial phase of the experiment (successful error detection) or missed it (knowledge neglect). Story readers reproduced about a third of missed errors, but less than 10% of errors they had caught previously. Moses subjects showed lower suggestibility overall but demonstrated the same pattern: They only reproduced errors they had missed, not ones they noticed during the initial phase of the experiment.

Reducing Suggestibility

How do we reduce the likelihood that people will pick up and later reproduce incorrect information that contradicts what they already know? It is clear that warnings alone will not suffice even though pre-encoding warnings often reduce suggestibility in episodic memory paradigms where people are misled about the content of videos or other stimuli (e.g., Loftus, 1979) and can also help readers avoid irrelevant details in texts (Peshkam, Mensink, Putnam, & Rapp, 2011). In both the story paradigm and the Moses illusion, learners are explicitly told that they will be exposed to errors and yet still reproduce some subset of those errors on later tests. One problem may be that readers forget to critically analyze material as they read, especially when they are engrossed

in a narrative or other ongoing task. A general warning imposes a prospective memory load upon the reader (he or she needs to keep that goal in mind while doing something else). Although no study (to our knowledge) has examined the effects of reminding the reader to monitor, increasing the frequency of errors may serve that function, leading to a reduction in the Moses illusion (although this benefit was too small to translate into a reduction in later suggestibility; Bottoms, Eslick, & Marsh, 2010). Similarly, the instruction to explicitly mark errors (via key press) during the story-reading phase reduces later suggestibility by about 5% (Fazio & Marsh, 2008a), likely because it keeps the monitoring task salient. Two points are relevant. First, the drop in suggestibility is modest, because (as alluded to earlier), readers miss the majority of the errors (and eye-tracking data also suggest that readers miss many errors; Bohan & Sanford, 2008). Second, the fact that detection instructions reduce suggestibility means that without this instruction, participants are likely missing even more errors (hence the higher resulting suggestibility).

Given the relationship between error detection and later suggestibility, one obvious solution is to improve the reader's monitoring, as that should reduce later suggestibility. The problem is that manipulations aimed at increasing monitoring sometimes help and sometimes hurt the learner. Consider the perceptual salience of the errors during the initial phase of the experiment. At first glance, the literature suggests that making errors physically salient will make them more noticeable, hence reducing knowledge neglect. For example, when the incorrect presupposition (as opposed to another word in the question) is capitalized and underlined, the Moses illusion is reduced (Bredart & Docquier, 1989; see also Kamas et al., 1996, experiment 1). Participants are also more likely to detect errors in the Moses paradigm when they are syntactically salient in the questions (e.g., placed in the main clause vs. a subordinate clause; Baker & Wagner, 1987; Bredart & Modolo, 1988).

However, drawing attention to errors can have *ironic* effects such that participants become even *more* likely to acquire and reproduce them. Returning to the story paradigm, in an effort to facilitate monitoring and detection of errors, Eslick, Fazio, and Marsh (2011) highlighted some of the general knowledge references in the stories using red font. Thus, readers no longer had to search for the factual information to evaluate it but were just charged with the task of evaluating the references' accuracy. Yet, instead of helping readers notice and later avoid reproducing the errors, the red font actually led readers to reproduce more of those story errors on the final general knowledge test, as compared to ones that were not highlighted. Drawing attention to the specific errors, in this case, increased suggestibility.

How do we resolve such results? The explanation is somewhat unsatisfactory, but the best we can offer is to say that perceptual salience draws attention to the errors, and then the end result depends upon whether readers actually catch the errors or not (which was not directly measured in the Eslick et al. study, unfortunately). If the salient errors are caught, then suggestibility will decrease. However, if the reader suffers from knowledge neglect, then suggestibility will either be constant or even increase (due to the increased attention to the errors). What is lacking from this analysis, of course, is an a priori way of knowing which manipulations will help monitoring and error detection, and in which situations.

In other words, the problem is that many manipulations designed to promote monitoring may not actually do so, and thus their memorial consequences will depend upon how much (indiscriminate) attention they direct toward the errors. Consider the effects of slowing processing; the original idea was that slowed presentation would provide readers with more time and resources to catch inaccuracies embedded in stories (with the idea being that the slower reader is a more careful one; Tousignant, Hall, & Loftus, 1986). To avoid individual differences in reading speed, the materials from the story paradigm were translated to a book-on-tape format and then played in either a faster or slower version, and all subjects later took a general knowledge test. Surprisingly, subjects in the slowed condition were more likely to use story errors to answer the final general knowledge questions even though all participants were warned that the stories might contain errors. This experiment also contained conditions where subjects pressed a key whenever they noticed errors, so there was an explicit measure of the impact of slowing on monitoring and error detection. Critically, even though the slowed version allowed more time for monitoring, readers were equally (un)likely to notice the story errors regardless of presentation speed. Slowing presumably increased attention to the errors, but attention without detection only served to increase suggestibility.

In short, overcoming knowledge neglect is a challenge since it is difficult to predict a priori which manipulations will successfully increase error detection versus simply increasing attention to errors (and, thereby, potentially increasing suggestibility).

Connections to Research on Episodic Memory

Thus far, we have focused on failures to apply knowledge appropriately, both when initially encountering contradictory information and later when one needs to retrieve related knowledge from memory (e.g., when asked general knowledge questions). In this section, we consider the role of episodic memory (remembering past events) in these failures of memory. In addition, we compare the results reviewed here to the large literature examining how people come to misremember events from their lives; in what ways is knowledge neglect different from and similar to failures of episodic memory? Because these two phenomena have not been directly compared in experiments, this section will necessarily be speculative, but we think the comparison raises important directions for future research.

To frame this discussion, we will focus on possible age differences in the memorial consequences of knowledge neglect. First, we describe children's learning from stories. Age-appropriate stories were modified to contain references to facts known to most first graders; some facts were presented in correct frames and others in misleading frames (e.g., referring to autumn as "...another word for fall" versus "...another word for spring"). The stories were presented in a book-on-tape format (with pictures paired with the narrative), so no assumptions about reading ability were made. Of interest was performance on a later general knowledge test. Robust suggestibility was observed, with children using story errors to answer general knowledge questions. However, the key point is that *older* children (age 7) were more suggestible than younger ones (ages 5–6; Fazio & Marsh, 2008b). That is, after hearing an error like "*Mount Washington is the tallest mountain in the world,*" older children were more likely to later repeat that error. Similar results were found when the final test contained multiple-choice questions (e.g., "*What is another word for autumn: fall, spring, summer, winter, or I don't know?*") instead of open-ended questions (Goswick, Mullet, & Marsh, 2013).

At the other end of the life span, consider older adults' learning from stories. College students and older adults (ages 65–85) listened to book-on-tape versions of the stories while also reading them (in large font). Again, the stories contained correct and misleading references, and of interest was performance on a final general knowledge test. The key finding for present purposes is that older adults were *less* suggestible than younger adults (Marsh, Balota, & Roediger, 2005; Umanath & Marsh, 2012). That is, older adults were less likely than younger adults to use story errors like "*Saturn, the largest planet*" when answering later general knowledge questions.

Both of these individual differences contrast with the age differences in suggestibility typically observed in episodic memory experiments. That is, younger children are normally *more* suggestible than older ones, and older adults are typically *more* suggestible than college students. For example, younger children are more likely than older ones to accept experimenter suggestions about mock eyewitness situations (see Ceci & Bruck, 1993, for a review) and to develop false memories for entire events, like riding in a hot air balloon (Ceci, Huffman, Smith, & Loftus, 1994). Similarly, older adults are more likely than college students to be influenced by postevent suggestions and experiences when asked to judge the preexperimental fame of names (Dywan & Jacoby, 1990), remember events portrayed in videos and photographs (Loftus, Levidow, & Duensing, 1992; Schacter, Koutstaal, Johnson, Gross, & Angell, 1997), or recollect experiences from the lab (Mueller-Johnson & Ceci, 2004; Cohen & Faulkner, 1989).

We will use these data on age differences in suggestibility to explore three points about the role of episodic memory. First, these results highlight that the two types of memory errors are not always the same: Age effects may be different for errors of episodic memory than for memorial consequences of knowledge neglect. Other manipulations also appear (at least based on the limited data sets available) to dissociate the two types of memory errors. For example, test instructions that require one to identify the sources of one's memories reduce errors of episodic memory (e.g.,. in mock eyewitness scenarios; Lindsay & Johnson, 1989) but have little impact on use of story errors to answer general knowledge questions (Marsh, Meade, & Roediger, 2003). Similarly, pre-encoding warnings and slower encoding reduce eyewitness suggestibility (Greene, Flynn, & Loftus, 1982; Tousignant, Hall, & Loftus, 1986), but not the learning of story errors (Marsh & Fazio, 2006; Fazio & Marsh, 2008a). On the other hand, there are also variables that have similar effects on both types of errors. For example, repeated exposure to misinformation tends to increase suggestibility across paradigms. Readers who read the same error-filled stories twice are more likely to use story errors to answer later general knowledge questions (Marsh et al., 2003), and mock eyewitnesses who are exposed to misinformation twice are later more likely to rely upon it (Zaragoza & Mitchell, 1996). As we will unpack below, these similarities likely arise because episodic memory is involved when errors enter the knowledge base. However, the fact that dissociations are also observed suggests that these two types of errors are not likely to be identical in mechanism, a point we will also return to later in this section.

Second, as just alluded to, episodic memory is likely to be involved in encoding and later remembering the misinformation. Learning the misinformation can be thought of as encoding a new association (e.g., the association between "the tallest mountain" and "Mount Washington"), a standard episodic memory task. Episodic memory improves across childhood (Ghetti & Lee, 2011) and declines in old age (Craik & Jennings, 1992). Thus, younger children will have poorer memory for the misinformation than will older children, and older adults will have worse memory for the misinformation than will the college students (see Brainerd, Reyna, & Ceci, 2008, and Mitchell, Johnson, & Mather, 2003, for similar arguments in the episodic memory domain). Episodic memory deficits also likely explain why older adults with mild cognitive impairment are even less likely to reproduce story errors on later tests, as compared to healthy older adults (Marsh et al., 2005). In addition, episodic memory mechanisms provide a clear explanation for the effects of manipulations like repeated reading, initial testing, and delaying the final test. Repeated reading increases encoding of the story errors, an initial test provides retrieval practice of the error (something known to benefit memory; Roediger & Karpicke, 2006; Roediger & Butler, 2011), and delaying the final test reduces memory for errors read in the stories due to forgetting. Consistent with this analysis, repeated reading and initial testing increase suggestibility but delaying the final test decreases suggestibility (Marsh et al., 2003; Barber, Rajaram, & Marsh, 2008).

Our final point about episodic memory involves the role of source monitoring in the two memory errors. While there are many differences across false memory paradigms, one commonality is that children and older adults tend to perform worse on memory tasks that depend upon source memory, meaning that they struggle when the task requires participants to identify the origin of their memories (Johnson, Hashtroudi, & Lindsay, 1993). For example, the successful eyewitness would be able to identify which details were actually observed versus only heard about in later media coverage or alluded to in a police interview. When an eyewitness claims to remember seeing a detail that was suggested in a postevent interview, it is a failure of source memory because that detail is misattributed to the original event rather than correctly sourced to the interview. Because older adults often struggle with source memory, they are more likely than college students to become confused between what they thought and what they actually did (Rabinowitz, 1989) or to confuse what they saw presented visually with something someone else said (McIntyre & Craik, 1987). Similarly, the suggestibility of young eyewitnesses has been linked to source problems, such as distinguishing between what was seen in a movie from what was included in an experimenter's subsequent narrative (Giles, Gopnik, & Heyman, 2002). However, with the story paradigm, failure to remember the story source does not drive suggestibility (Marsh et al., 2003). In fact, older children are better able to identify which answers had been encountered in the stories, as compared to younger children (Fazio & Marsh, 2008b). Similarly, college students are better able to identify which answers had been encountered in the stories, as compared to older adults (Marsh et al., 2005). In these studies, the groups with best memory for the story source were the ones who were the most suggestible, instead of the least suggestible, as might be expected based on the large literature showing source memory as protecting against suggestion. This is not to say that subjects had perfect source memory in the story experiments, but that they made a very specific type of source error. Critically, readers experienced an *illusion of prior knowledge*, whereby they believed that they knew the story errors prior to the experiment, even though baseline belief in the errors was low. It does a reader no good to remember that a fact was read in a story if he or she also attributes it to prior experience and, thereby, assumes it is true.

In short, episodic memory is necessarily involved in learning and storing false facts, but illusions of knowledge do not appear to be identical to episodic memory errors. The two types of memory errors can be dissociated.

Conclusions and Future Directions

This chapter makes a simple point: People often have information stored in memory that they fail to apply successfully. We used the term knowledge neglect to refer to people's failures to bring to bear relevant stored knowledge. We believe such failures to be common. They may occur because the reader is not focused on monitoring (e.g., because of processing load) or because monitoring fails when a plausible error is close enough to the correct answer to register as a match. Regardless, a failure to notice an error has repercussions for memory. The problem is that this recently encountered error is more available in memory than the original knowledge, meaning that it is the error that later comes to mind fluently; this fluency is misinterpreted as truth. Thus, people reproduce the errors even though they have stored knowledge that contradicts those errors.

There remain a number of open questions and directions for future research. For instance, we discussed the mixed literature aimed at improving monitoring, with the goal of reducing knowledge neglect and subsequent suggestibility; future research should pinpoint the specific circumstances under which monitoring and detection of errors can be successfully improved, so that we can predict this a priori. It is unsatisfying to claim after the fact that "monitoring was poor" because suggestibility increased, even when a manipulation was designed to increase monitoring.

Theoretically, the most interesting open question involves whether "knowledge neglect" is specific to knowledge or if this behavior is an example of a broader phenomenon that includes episodic memory. That is, one can easily imagine a situation where one has stored personal memories that are relevant to the current situation and yet fails to bring them to bear. To what extent is there likely to be something different about knowledge? Unlike episodic memories, much knowledge is retrieved without a feeling of projecting oneself backward in time to a specific time and place; much knowledge is sourceless and labeled as "known" rather than "remembered" (e.g., you simply "know" that George Washington was the first president of the United States; Conway, Gardiner, Perfect, Anderson, & Cohen, 1997). The result is that learners often judge the veracity of retrieved knowledge by how fluently or easily it comes to mind (Kelley & Lindsay, 1993). Fluency misattributions can also lead to errors of episodic memory, of course (e.g., Jacoby, Kelley, & Dywan, 1989), but we believe people are particularly reliant on fluency when the content involves knowledge. Because of this, readers may be more likely to suffer from knowledge neglect, as opposed to neglecting other types of memories. More generally, future research should directly compare errors of episodic memory to the errors observed following knowledge neglect. We presented cross-experimental evidence for the similarities and differences between these types of errors, but a direct comparison is needed.

Overall, a lot is known about how people come to misremember events, but less is known about how errors enter the knowledge base, especially when the errors contradict stored knowledge. Clearly, there are connections between episodic memory and the knowledge base, but there are differences as well. Future research should better connect these two areas of research, with the goal of better understanding knowledge neglect and the resultant illusions of knowledge.

Acknowledgment

The writing of this chapter was supported in part by a collaborative activity award from the James S. McDonnell Foundation.

References

Anderson, J. R. (1981). Effects of prior knowledge on memory for new information. *Memory & Cognition*, *9*, 237–246.

Ausubel, D. P., & Blake, E. (1958). Proactive inhibition in the forgetting of meaningful school material. *Journal of Educational Research*, *52*, 145–149.

Baker, L., & Wagner, J. L. (1987). Evaluating information for truthfulness: The effects of logical subordination. *Memory & Cognition, 15*, 247–255.

Barber, S. J., Rajaram, S., & Marsh, E. J. (2008). Fact learning: How information accuracy, delay, and repeated testing change retention and retrieval experience. *Memory, 16*, 934–946.

Barton, S. B., & Sanford, A. J. (1993). A case study of anomaly detection: Shallow semantic processing and cohesion establishment. *Memory & Cognition, 21*, 477–487.

Bohan, J., & Sanford, A. (2008). Semantic anomalies at the borderline of consciousness: An eye-tracking investigation. *Quarterly Journal of Experimental Psychology, 61*, 232–239.

Bottoms, H. C., Eslick, A. N., & Marsh, E. J. (2010). Memory and the Moses illusion: Failures to detect contradictions with stored knowledge yield negative memorial consequences. *Memory, 18*, 670–678.

Bower, G. H., & Morrow, D. G. (1990). Mental models in narrative comprehension. *Science, 247*, 44–48.

Brainerd, C. J., Reyna, V. F., & Ceci, S. J. (2008). Developmental reversals in false memory: A review of data and theory. *Psychological Bulletin, 134*, 343–382.

Bredart, S., & Docquier, M. (1989). The Moses illusion: A follow-up on the focalization effect. *European Bulletin of Cognitive Psychology, 9*, 357–362.

Bredart, S., & Modolo, K. (1988). Moses strikes again: Focalization effect on a semantic illusion. *Acta Psychologica, 67*, 135–144.

Brown, D. E. (1992). Using examples and analogies to remediate misconceptions in physics: Factors influencing conceptual change. *Journal of Research in Science Teaching, 29*, 17–34.

Butler, A. C., Fazio, L. F., & Marsh, E. J. (2011). The hypercorrection effect persists over a week, but high confidence errors return. *Psychonomic Bulletin & Review, 18*, 1238–1244.

Büttner, A. C. (2007). Questions versus statements: Challenging an assumption about semantic illusions. *Quarterly Journal of Experimental Psychology, 60*, 779–789.

Ceci, S. J., & Bruck, M. B. (1993). Suggestibility of the child witness: A historical review and synthesis. *Psychological Bulletin, 113*, 403–439.

Ceci, S. J., Huffman, M. L. C., Smith, E., & Loftus, E. F. (1994). Repeatedly thinking about a non-event: Source misattributions among preschoolers. *Consciousness and Cognition, 3*, 388–407.

Chi, M. T. H. (2008). Three types of conceptual change: Belief revision, mental model transformation, and categorical shift. In S. Vosniadou (Ed.), *Handbook of research on conceptual change* (pp. 61–82). Hillsdale, NJ: Lawrence Erlbaum Associates.

Chiesi, H. L., Spilich, G. J., & Voss, J. F. (1979). Acquisition of domain-related information in relation to high and low domain knowledge. *Journal of Verbal Learning and Verbal Behavior, 18*, 257–274.

Cho, H.-H., Kahle, J. B., & Nordland, F. H. (1985). An investigation of high school biology textbooks as sources of misconceptions and difficulties in genetics and some suggestions for teaching genetics. *Science Education, 69*, 707–719.

Clement, J., Narode, R., & Rosnick, P. (1981). Intuitive misconceptions in algebra as a source of math anxiety. *Focus on Learning Problems in Mathematics, 3*, 36–45.

Cohen, G., & Faulkner, D. (1989). Age differences in source forgetting: Effects on reality monitoring and on eyewitness testimony. *Psychology and Aging, 4*, 10–17.

Conway, M. A., Gardiner, J. M., Perfect, T. J., Anderson, S. J., & Cohen, G. M. (1997). Changes in memory awareness during learning: The acquisition of knowledge by psychology undergraduates. *Journal of Experimental Psychology. General, 126*, 393–413.

Craik, F. I. M., & Jennings, J. M. (1992). Human memory. In F. I. M. Craik & T. A. Salthouse (Eds.), *The handbook of aging and cognition* (pp. 51–110). Hillsdale, NJ: Lawrence Erlbaum Associates.

Dooling, D. J., & Mullet, R. L. (1973). Locus of thematic effects in retention of prose. *Journal of Experimental Psychology, 97*, 404–406.

Dywan, J., & Jacoby, L. (1990). Effects of aging on source monitoring: Differences in susceptibility to false fame. *Psychology and Aging, 5*, 379–387.

Erickson, T. D., & Mattson, M. E. (1981). From words to meaning: A semantic illusion. *Journal of Verbal Learning and Verbal Behavior, 20*, 540–551.

Eslick, A. N., Fazio, L. K., & Marsh, E. J. (2011). Ironic effects of drawing attention to story errors. *Memory, 19*, 184–191.

Fazio, L. K., Barber, S. J., Rajaram, S., Ornstein, P. A., & Marsh, E. J. (2013). Creating illusions of knowledge: Learning errors that contradict prior knowledge. *Journal of Experimental Psychology. General, 142*, 1–5

Fazio, L. K., & Marsh, E. J. (2008a). Slowing presentation speed increases illusions of knowledge. *Psychonomic Bulletin & Review, 15*, 180–185.

Fazio, L. K., & Marsh, E. J. (2008b). Older, not younger, children learn more false facts from stories. *Cognition, 106*, 1081–1089.

Ferland, F., Ladouceur, R., & Vitaro, F. (2002). Prevention of problem gambling: Modifying misconceptions and increasing knowledge. *Journal of Gambling Studies, 18*, 19–29.

Ferreira, F., Bailey, K. G. D., & Ferraro, V. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science, 11*, 11–15.

Gerrig, R. J. (1989). Suspense in the absence of uncertainty. *Journal of Memory and Language, 28*, 633–648.

Gerrig, R. J., & Rapp, D. N. (2004). Psychological processes underlying literary impact. *Poetics Today, 25*, 265–281.

Ghetti, S., & Lee, J. (2011). Children's episodic memory. *Wiley Interdisciplinary Reviews: Cognitive Science, 2*, 365–373.

Gilbert, D. T. (1991). How mental systems believe. *American Psychologist, 46*, 107–119.

Gilbert, D. T., Krull, D. S., & Malone, P. S. (1990). Unbelieving the unbelievable: Some problems in the rejection of false information. *Journal of Personality and Social Psychology, 59*, 601–613.

Giles, J. W., Gopnik, A., & Heyman, G. D. (2002). Source monitoring reduces the suggestibility of preschool children. *Psychological Science, 13*, 288–291.

Goswick, A. E., Mullet, H. G., & Marsh, E. J. (2013). Suggestibility from stories: Can production difficulties and source monitoring explain a developmental reversal? *Journal of Cognition and Development, 14*, 607–616.

Gouvier, W. D., Prestholdt, P. H., & Warner, M. S. (1988). A survey of common misconceptions about head injury and recovery. *Archives of Clinical Neuropsychology, 3*, 331–343.

Green, M. C., & Brock, T. C. (2000). The role of transportation in the persuasiveness of public narratives. *Journal of Personality and Social Psychology, 79*, 701–721.

Greene, E., Flynn, M. S., & Loftus, E. F. (1982). Inducing resistance to misleading information. *Journal of Verbal Learning and Verbal Behavior, 21*, 207–219.

Hannon, B., & Daneman, M. (2001). Susceptibility to semantic illusions: An individual-differences perspective. *Memory & Cognition, 29*, 449–461.

Horton, W. S., & Rapp, D. N. (2003). Out of sight, out of mind: Occlusion and the accessibility of information in narrative comprehension. *Psychonomic Bulletin & Review, 10*, 104–110.

Jacoby, L. L., Kelley, C. M., & Dywan, J. (1989). Memory attributions. In H. L. Roediger, III, & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 391–422). Hillsdale, NJ: Lawrence Erlbaum Associates.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114*, 3–28.

Johnson-Laird, P. N. (1983). *Mental models*. Cambridge, MA: Harvard University Press.

Kamas, E. N., Reder, L. M., & Ayers, M. S. (1996). Partial matching in the Moses illusion: Response bias not sensitivity. *Memory & Cognition, 24*, 687–699.

Kelley, C. M., & Lindsay, D. S. (1993). Remembering mistaken for knowing: Ease of retrieval as a basis for confidence in answers to general knowledge questions. *Journal of Memory and Language, 32*, 1–24.

Kole, J. A., & Healy, A. F. (2007). Using prior knowledge to minimize interference when learning large amounts of information. *Memory & Cognition, 35*, 124–137.

Landau, J. D., & Bavaria, A. J. (2003). Does deliberate source monitoring reduce students' misconceptions about psychology? *Teaching of Psychology, 30*, 311–314.

Lee, G. M., Friedman, J. F., Ross-Degnan, D., Hibberd, P. L., & Goldmann, D. A. (2003). Misconceptions about colds and predictors of health service utilization. *Pediatrics, 111*, 231–236.

Lindsay, D. S., & Johnson, M. K. (1989). The eyewitness suggestibility effect and memory for source. *Memory & Cognition, 17*, 349–358.

Loftus, E. F. (1979). Reactions to blatantly contradictory information. *Memory & Cognition, 7*, 368–374.

Loftus, E. F., Levidow, B., & Duensing, S. (1992). Who remembers best? Individual differences in memory for events that occurred in a science museum. *Applied Cognitive Psychology, 6*, 93–107.

Marsh, E. J. (2004). Story stimuli for creating false beliefs about the world. *Behavior Research Methods, Instruments, & Computers, 36*, 650–655.

Marsh, E. J., Balota, D. A., & Roediger, H. L., III. (2005). Learning facts from fiction: Effects of healthy aging and early-stage dementia of the Alzheimer type. *Neuropsychology, 19*, 115–129.

Marsh, E. J., & Fazio, L. K. (2006). Learning errors from fiction: Difficulties in reducing reliance on fictional stories. *Memory & Cognition, 34*, 1140–1149.

Marsh, E. J., & Fazio, L. K. (2007). Learning from fictional sources. In J. Nairne (Ed.), *The foundations of remembering: Essays in honor of Henry L. Roediger III* (pp. 397–413). New York: Psychology Press.

Marsh, E. J., Meade, M. L., & Roediger, H. L., III. (2003). Learning facts from fiction. *Journal of Memory and Language, 49*, 519–536.

McCloskey, M. (1983). Intuitive physics. *Scientific American, 248*, 122–130.

McIntyre, J. S., & Craik, F. I. M. (1987). Age differences in memory for item and source information. *Canadian Journal of Psychology, 41*, 175–192.

Mitchell, K. J., Johnson, M. K., & Mather, M. (2003). Source monitoring and suggestibility to misinformation: Adult age-related differences. *Applied Cognitive Psychology, 17*, 107–119.

Mueller-Johnson, K., & Ceci, S. J. (2004). Memory and suggestibility in older adults: Live event participation and repeated interview. *Applied Cognitive Psychology, 18*, 1109–1127.

Mullet, H. G., Umanath, S., & Marsh, E. J. (November 2012). Accessibility of knowledge and susceptibility to semantic illusions. Poster presented at the 53rd Annual Meeting of the Psychonomic Society, Minneapolis, MI.

Munson, B. H. (1994). Ecological misconceptions. *Journal of Environmental Education, 25*, 30–35.

Nakhleh, M. B. (1992). Why some students don't learn chemistry. *Journal of Chemical Education, 69*, 191–196.

Nelson, T. O., & Narens, L. (1980). Norms of 300 general-information questions: Accuracy of recall, latency of recall, and feeling-of-knowing ratings. *Journal of Verbal Learning and Verbal Behavior, 19*, 338–368.

Peshkam, A., Mensink, M. C., Putnam, A. L., & Rapp, D. N. (2011). Warning readers to avoid irrelevant information: When being vague might be valuable. *Contemporary Educational Psychology, 36*, 219–231.

Rabinowitz, J. C. (1989). Judgments of origin and generation effects: Comparisons between young and elderly adults. *Psychology and Aging, 4*, 259–268.

Rapp, D. N. (2008). How do readers handle incorrect information during reading? *Memory & Cognition, 36*, 688–701.

Reder, L. M., & Kusbit, G. W. (1991). Locus of the Moses illusion: Imperfect encoding, retrieval, or match? *Journal of Memory and Language, 30*, 385–406.

Roediger, H. L., III, & Butler, A. C. (2011). The critical role of retrieval practice in long-term retention. *Trends in Cognitive Sciences, 15*, 20–27.

Roediger, H. L., III, & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science, 17*, 249–255.

Schacter, D. L., Koutstaal, W., Johnson, M. K., Gross, M. S., & Angell, K. E. (1997). False recollection induced by photographs: A comparison of older and younger adults. *Psychology and Aging, 12*, 203–215.

Simons, D. J., & Chabris, C. F. (2011). What people believe about how memory works: A representative survey of the U.S. population. *PLoS ONE, 6*, e22757.

Tousignant, J. P., Hall, D., & Loftus, E. F. (1986). Discrepancy detection and vulnerability to misleading postevent information. *Memory & Cognition, 14*, 329–338.

Umanath, S., Dolan, P. O., & Marsh, E. J. (2014). Aging and the Moses illusion: Older adults fall for Moses but if asked directly, stick with Noah. *Memory, 22*, 481–492.

Umanath, S., & Marsh, E. J. (2012). Aging and the memorial consequences of catching contradictions with prior knowledge. *Psychology and Aging, 27*, 1033–1038.

van Jaarsveld, H. J., Dijkstra, T., & Hermans, D. (1997). The detection of semantic illusions: Task-specific effects for similarity and position of distorted terms. *Psychological Research, 59*, 219–230.

van Oostendorp, H., & de Mul, S. (1990). Moses beats Adam: A semantic relatedness effect on a semantic illusion. *Acta Psychologica, 74*, 35–46.

Wynn, L. L., Foster, A. M., & Trussell, J. (2009). Can I get pregnant from oral sex? Sexual health misconceptions in e-mails to a reproductive health website. *Contraception, 79*, 91–97.

Zaragoza, M. S., & Mitchell, K. J. (1996). Repeated exposure to suggestion and the creation of false memories. *Psychological Science, 7*, 294–300.