CHAPTER THREE

Believing that Humans Swallow Spiders in Their Sleep: False Beliefs as Side Effects of the Processes that Support Accurate Knowledge

Elizabeth J. Marsh[1], Allison D. Cantor and Nadia M. Brashier
Department of Psychology & Neuroscience, Duke University, Durham, NC, USA
[1]Corresponding author: E-mail: emarsh@psych.duke.edu

Contents

Psychology of Learning and Motivation, Volume 64
ISSN 0079-7421
http://dx.doi.org/10.1016/bs.plm.2015.09.003

Abstract

Humans can store, maintain, and retrieve an impressive amount of information—but the processes that support accurate knowledge can also lead to errors, such as the false belief that humans swallow eight spiders in their sleep each year. In this chapter, we review characteristics of the knowledge base and explore how five adaptive properties that support accurate knowledge can also lead to the learning, storage, and retrieval of falsehoods. First, people exhibit a bias to believe information is true since, most of the time, incoming information is indeed true. Second, we utilize a fluency-based heuristic for judging truth since—again, most of the time—easy processing reliably signals that something is true. Third, the knowledge base is productive: people use existing knowledge to make new inferences, which are typically accurate but occasionally are inappropriate and result in errors. Fourth, existing knowledge supports new learning, so our ingrained misconceptions can foster new errors and interfere with learning the truth. Fifth, because it would be too taxing to carefully compare all incoming information to stored knowledge, we do not require a perfect match and often accept information as "good enough." As a result, errors that are similar to the truth often slip by undetected, and sometimes are later reproduced. Finally, we discuss methods for correcting errors and potential barriers to the correction of misconceptions. In particular, it is essential to refute the error as well as provide a simple alternative to replace it. Overall, the processes that support accurate knowledge and false beliefs are the same, and can lead to competing representations in memory.

1. INTRODUCTION

1.1 The Issue at Hand

The average person swallows eight spiders in her sleep every year.

Many of us have encountered some version of this claim. Is it true or false? Many people are unsure, but become concerned about the possibility;

a quick Internet search reveals many posts (e.g., on Quora, Reddit, Yahoo questions), where people explicitly ask others to verify this claim. Fortunately, science has provided absolutely no evidence to support this claim, and instead offers many reasons to doubt it (e.g., most humans move around a lot in their sleep; spiders avoid predators). So where did this idea come from? The original source is allegedly a journalist, who was poking fun at the ridiculous "facts" people learn on the Internet and unwittingly loosed the spider statistic on the world. In an ironic twist, this origin story may itself be an urban legend, as fact-checkers failed to locate the infamous article or even demonstrate that the author worked for the publication that supposedly published the piece.

The spider example may be laughable, but it demonstrates similarities between misconceptions and accurate knowledge. The cognitive processes used to encode, store, and retrieve veridical concepts (e.g., spider, sleep) form the basis for critical misunderstandings. We examine errors as "by-products" of an otherwise efficient memory system and discuss ways to correct misconceptions after the fact.

1.2 Defining Knowledge

Stating the capital of Peru, solving a differential equation, and translating a text from Russian to English are all examples of successfully using knowledge. Knowing how to traverse airport security, deciding whether a joke is appropriate in a particular context, and calling one's sister by the correct name are also examples of successful use of knowledge. Defining "knowledge" is a tricky business, as the label applies to so many different things. Knowledge includes facts and concepts and an understanding of their relationships; knowledge also includes language, schemas, and inferences.

Most psychologists agree that knowledge is a form of memory extracted from past experience. However, knowledge is often defined by what it is *not*, rather than by what it is. That is, researchers contrast it to event memories (i.e., episodic memories; memories of specific events from particular places and times) with the emphasis on knowledge as information stored in memory that is decontextualized and that does not elicit a feeling of reliving. Depending on one's theoretical orientation, the term *semantic memory* may be considered synonymous with knowledge.

What is uncontroversial is the large role knowledge plays in many different cognitive processes. For example, knowledge drives how we interpret what we see. Imagine a photo of a girl standing on a balcony, looking at buildings. If we took a ruler and measured the image, the girl and the buildings might be the same height, but we do not interpret her as a giant or the buildings as miniatures; instead, we interpret their similar sizes as evidence that the building is further away than the girl, as we have learned from past experience. More generally, knowledge allows inferences, affects decision-making, guides the reconstruction of event memories, and supports communication and emotional responses.

2. GENERAL PROPERTIES OF THE KNOWLEDGE BASE

2.1 The Knowledge Base Has No Known Capacity Limit

There is no known limit to the amount of knowledge that can be stored in memory. Storage space does not become "full" over the years; to the contrary, older adults generally outperform younger adults on measures of vocabulary and general knowledge (Botwinick & Storandt, 1980; Mitchell, 1989; Perlmutter, 1978). As will be further described in Section 2.2, research on domain experts highlights just how much knowledge can be stored. Chess experts, for example, store an estimated 50,000 "game boards" in memory, allowing them to move quickly and automatically upon recognizing a particular board layout during a game (see Bedard & Chi, 1992; for a review). Such impressive memory feats are not specific to chess; expert musicians, bridge players, and computer programmers possess similar amounts of domain knowledge (see Ross, 2006; for a review). Computer simulations support these demonstrations of impressive knowledge; attempts to estimate the storage capacity of human memory by examining the rate at which people acquire new information suggest that we can store virtually limitless amounts of information (e.g., Landauer, 1986).

2.2 Knowledge Is Interconnected and Organized

Of course, knowledge does not consist of infinite separate pieces of information; that would imply that "the more one knows about a concept, the longer it would take to retrieve any particular fact about the concept" (i.e., *the paradox of the expert*, Anderson, 1983, p. 28). Instead, newly acquired information becomes integrated with existing information, creating an interconnected web of knowledge (e.g., Bartlett, 1932; Chi, Glaser, & Farr, 1988). This idea can be visually represented as a collection of "nodes," each of which represents a concept, with links to other related nodes (Collins & Loftus, 1975). The result is nonindependence among concepts; activating any one concept (e.g., by reading it, hearing it, etc.) "spreads" activation to other related concepts. Behaviorally, spreading activation manifests itself in *semantic priming*: people are faster to decide if a target (*nurse*) is a word after reading a related word (*doctor*) than after reading an unrelated one (*butter*) (e.g., Meyer & Schvaneveldt, 1971). This facilitation in reaction time, or *priming*, occurs because the concept of *doctor* was already partially activated after reading *nurse*. Activation continues to spread to concepts that are further away in semantic space, although the amount of activation decreases with semantic distance from the original concept. For example, *lion* can prime *stripes,* even though there is not a direct relationship between *lions* and *stripes*; *lion* primes *tiger*, and *tiger* primes *stripes*, meaning that exposure to *lion* yields observable priming of *stripes*, albeit less than for *tiger* (Balota & Lorch, 1986).

Of course, knowledge can be higher level than individual concepts, representing generalizations and extractions from past experience. A classic example comes from Thorndyke and Hayes-Roth (1979), where students read multiple texts describing different exemplars of a category. For example, a student might read about the constellations, Pisces, Aries, and Scorpio, before reading a target text about a new constellation. Memory for the target text depended on how many related passages preceded it; reading more passages boosted memory for the commonalities across passages, but this occurred at the expense of passage-specific details. In other words, participants extracted a *schema*, or generalized representation of constellation texts, which supported new learning at the expense of details. The more specific term *scripts* refers to action schemas, such as the steps involved in getting a haircut, shopping, and eating at restaurants. Supporting the existence of scripts, people are remarkably consistent when asked to generate the steps of common events like "eating at restaurant" or "getting a haircut" (Bower, Black, & Turner, 1979). Even though two strangers have never shared a restaurant meal together, they both know that a prototypical event begins with the hostess seating you, followed by the delivery of menus and ordering of food, and that at the end of the meal there is a bill and an expectation to tip. Both schemas and scripts are extracted from past experiences, and can be powerful tools for predicting outcomes in new experiences.

Expertise illustrates exactly how well-organized knowledge can be. That is, experts differ from novices in more than just the *amount* or *strength* of information stored in memory; expert knowledge differs qualitatively from that of novices in its *structure*. Chi and Koeske (1983), for example, compared a child's mappings of two sets of dinosaurs, one familiar to the child and the other unfamiliar. Unsurprisingly, the child's map for well-known dinosaurs boasted more structure than the one for unknown dinosaurs. That is, even though the novice map contained a similar number of property nodes (e.g., *eats plants*), the expert map contained many more linkages among dinosaurs, yielding a more interconnected and cohesive network (see also Gobbo & Chi, 1986). In addition, an expert's knowledge yields concepts that are more clearly differentiated from one another. For example, bird experts differentiate between warblers and finches more rapidly than do novices (Johnson & Mervis, 1997).

Experts also represent knowledge at a deeper level, whereas novices focus on surface similarities. Physics experts, for example, sort physics problems into groups based on principles of mechanics (e.g., problems pertaining to the Work—Energy Theorem) while novices group by literal features (e.g., problems with "blocks on an inclined plane"; Chi, Feltovich, & Glaser, 1981). A similar pattern emerges when people categorize fish; novices group by physical similarities across fish (e.g., groups for "all long and skinny fish"), whereas experts take functional information into consideration and form categories based on common habitats and for "fish that I eat" (Boster & Johnson, 1989). Furthermore, different types of expertise lead to distinct, but focused, organizations; for example, different tree experts sort trees differently depending on their particular expertise: maintaining city trees, landscape design, or science education/research. Both landscape designers and maintenance workers formed functional categories, whereas the scientists sorted the trees according to their actual scientific classifications (Medin, Lynch, Coley, & Atran, 1997).

2.3 Knowledge Is Surprisingly Durable

We forget many of our daily experiences, but knowledge proves surprisingly resilient, persisting across time and changes in context. In contrast to other kinds of memory, knowledge does not decline steadily with age; as mentioned earlier, older adults often outperform young adults on tests of general knowledge and vocabulary. Even after the onset of dementia, knowledge can sometimes remain intact and accessible. Hodges and Patterson (1995), for example, showed remarkable heterogeneity in the performance of patients diagnosed with minimal and mild Alzheimer's disease. All patients demonstrated event memory deficits, but some performed perfectly on measures of knowledge (e.g., category fluency: generating exemplars of a category).

Of course, much knowledge is encountered, accessed, and applied repeatedly over the years, effectively rehearsing the information. However, even knowledge that is *not* rehearsed appears to be remarkably durable over time. Knowledge of the Spanish language, the layout of a city one previously lived in, and names and faces of high school classmates all remain fairly stable over time, following an initial period of forgetting, even though participants do not report using or rehearsing the material since the time of original learning (see Bahrick, Hall, & Baker, 2013 for a review).

For example, Bahrick (1984) measured participants' retention of high school and college Spanish up to 50 years after initial learning. To estimate typical levels of acquisition, a subset of participants were currently enrolled in or had recently finished a Spanish course. The remaining participants had completed their last Spanish course between one and 50 years earlier. All participants took a number of recall and recognition tests for Spanish vocabulary, grammar, and reading comprehension. Performance on these tests was impressive: after an initial drop, the functions were quite consistent across the remaining interval. Based on these data, Bahrick argued that a portion of the knowledge base is so long-lasting that it is essentially permanent, or a *permastore* (see Figure 1 for a schematic based on hypothetical data).

Psychology instructors may be interested to know that a similar pattern occurs when examining students' retention of cognitive psychology course material (Conway, Cohen, & Stanhope, 1991). Since the majority of Bahrick's work examined the long-term retention of procedural knowledge (e.g., knowing *how* to speak a language), Conway and colleagues examined whether the same principles apply to declarative knowledge (e.g., knowing *that* cones allow color vision). To estimate students' retention of the

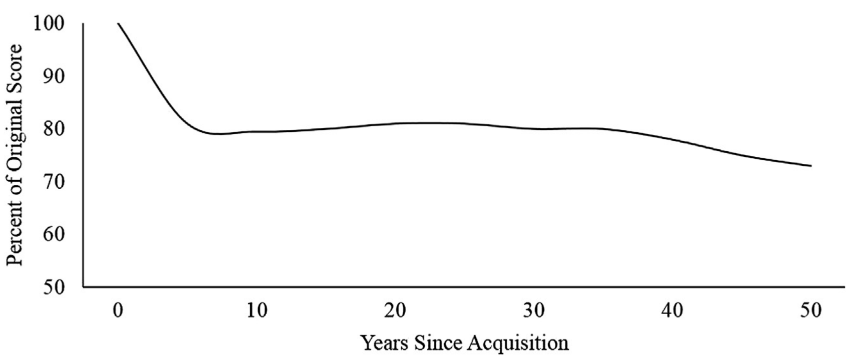

Figure 1 Hypothetical long-term retention of knowledge depicting Bahrick's (1984) "permastore."

material, the researchers administered a range of knowledge tests (tapping knowledge of both specific details and broader concepts) to students who had finished the course between 3 and 125 months prior. Consistent with Bahrick's findings, student performance on retention tests declined across the first 39 months of the interval, but then stabilized across the remaining portion.

2.4 Much, but Not All, Knowledge Is "Sourceless"

Retrieving knowledge and remembering an event "feel" different. While people "just know" that Washington, D.C. is the capital of the US, remembering a recent trip to the nation's capital evokes many associated details. Retrieving an event memory typically involves the feeling of traveling back in time and reliving the episode (Tulving, 1985). For example, remembering the events that occurred at a recent party might involve reexperiencing the music that was playing, the party decorations, the people in attendance, etc. In contrast, people often (but not always) fail to remember the source of their knowledge; when remembering that Washington was the first president of the United States, people do not normally think back to the original time and place of learning. In other words, people often report just "knowing" facts rather than "remembering" them. While knowledge may be linked to its source at first, this information is often lost over time, probably due to lack of rehearsal and to repeated encounters with the information that were associated with different sources. Supporting this claim, students initially judge course material as "remembered" but shift to "knowing" over the course of a single semester (Conway, Gardiner, Perfect, Anderson, & Cohen, 1997).

Consistent with these ideas, context appears to exert little or no influence on knowledge retrieval. For example, students do not always benefit from taking an exam in the same room where they attended class. While a few studies show small benefits of a contextual match (e.g., Abernathy, 1940; Metzger, Boschee, Haugen, & Schnobrich, 1979; Van Der Wege & Barry, 2008), others found no differences (e.g., Farnsworth, 1934; Saufley, Otaka, & Bavaresco, 1985). In contrast, there are many studies showing a benefit of contextual match for event memories, with people remembering more of a studied word list if they are tested in the same physical context. Overall, this benefit for event memories is "modest ($d = .28$) but reliable" (Smith & Vela, 2001, p. 213).

However, knowledge is occasionally associated with a source; these cases are often ones where there is a *reason* to remember the source—either because one will want to find the information again, or because the source might cast doubt on the veracity of the information. For example, there is some evidence that people are better at remembering *where* they found information on the Internet, as opposed to the information itself (Sparrow, Liu, & Wegner, 2011). Another example involves fictional sources, with the logic that readers/viewers should be hesitant to integrate everything from a fictional source into their knowledge base. In the study supporting this argument, subjects read a passage about the takahe bird that was labeled as *fiction* for some subjects but *factual* for others (Potts & Peterson, 1985). After reading the passage, participants made true—false decisions about the takahe either in blocks of other questions related to the takahe (passage context) or in blocks of unrelated questions (new context). Readers who believed the passage was fictional were slower to access their knowledge about the takahe when in a new context; no such effect occurred with readers of nonfiction. In other words, only information from the fictional story still retained some links to its original context.

2.5 Access to Specific Knowledge Can Shift

Storing and maintaining information in the knowledge base is not sufficient; just as important is the ability to retrieve that information when needed. Tulving and Pearlstone (1966) made the classic distinction between *available* and *accessible* memories: information stored in memory is available, but only information that can be retrieved is accessible. Although the availability—accessibility distinction comes out of the literature on event memory, the same idea applies to knowledge. People do not produce exactly the same knowledge at different points in time, reflecting the shifting accessibility of knowledge. Brown (1923) demonstrated this for the first time when participants attempted to recall the United States twice, 30 min apart in time. Even though participants' knowledge of the US states could not have changed over the course of the experiment, Brown found that participants forgot, or "lost," some states between the first and second tests and "gained" others. To get a better estimate of how much the accessibility of knowledge shifts across retrieval attempts, Bellezza (1984a, 1984b, 1984c, 1987, 1988) quantified the within-individual reliability of knowledge retrieval. His participants recalled as many category exemplars, noun meanings, facts about friends and family, scripts, and pieces of self-information as they could; critically, they did this twice, with the two attempts separated by 1 week. Across the various types of knowledge, reliability between two retrievals was modest but not very high; common-element correlations (McNemar, 1969) between two retrievals ranged from .38 to .69.

Bahrick coined the term *marginal knowledge* to describe knowledge that is stored in memory, but is currently inaccessible. Perhaps the most famous example is the nearly universal *tip-of-the-tongue* (TOT) experience, where one feels very confident that one knows a word or name or other piece of information but cannot produce it (e.g., Brown & McNeil, 1966; see Brown, 1991 for a review). TOT states likely reflect knowledge stored in memory that is available, albeit not accessible; people can frequently report the first letter or number of syllables of the target (e.g., Brown & McNeil, 1966; Yarmey, 1973) and such states often resolve with time (e.g., Choi & Smith, 2005; Read & Bruce, 1982). Presumably, TOT states resolve upon encountering different cues, something that could not happen if the targets were not actually stored in memory.

It can be difficult to distinguish the recovery of marginal knowledge from new learning. Berger, Hall, and Bahrick (1999) tackled this problem with a clever methodology, creating a set of fictitious questions that paralleled real ones. The fictitious questions matched the real questions in structure and sentence length, but had no factual basis (i.e., the researchers made them up). For example, for the real question *What is the name of the constellation that looks like a flying horse*? the parallel fictitious version asked, *What is the name of the constellation that looks like a sitting bull*? Critically, improvement on fictitious questions after an intervention (e.g., study phase) must reflect new learning, rather than reactivation of marginal knowledge. Berger and colleagues tested the ability of a 5-s exposure to stabilize access to answers that participants failed to produce on an earlier test. This intervention benefited real questions much more than fictitious ones, suggesting the existence of marginal knowledge for the real items. The benefits decreased with time; performance on an immediate test (90%) dropped continuously over 9 days (to 49%, see Table 1). In our own work, we showed that answering a multiple-choice question effectively reactivates marginal

Table 1 Proportion of final test questions answered with targets for real and fictitious questions (given initial test failure)

	.83 min	1.68 min	5 min	20 min	1 day	9 days
Real	.90 (.01)	.89 (.01)	.81 (.01)	.72 (.02)	.68 (.02)	.49 (.02)
Fictitious	.66 (.01)	.60 (.01)	.44 (.01)	.31 (.01)	.14 (.01)	.03 (.00)

Note. Standard errors are presented in parentheses.
Table adapted from Berger et al. (1999)Table 1.

knowledge (Cantor, Eslick, Marsh, Bjork, & Bjork, 2015), and such questions need not be paired with feedback.

2.6 People Are Good, but Not Perfect, at Judging What They Know

People generally have a good sense of whether or not they know something (see Nelson, 1988 for a review). In a classic example, Hart (1965) examined the accuracy of *feeling-of-knowing* (FOK) judgments. Participants answered a series of general knowledge questions; when they could not answer a question within 10 s, they judged whether or not they could successfully recognize the correct answer among several wrong answers. Participants' FOK judgments accurately predicted their stored knowledge: when participants claimed to hold knowledge, they correctly recognized the target 66% of the time.

Conversely, people demonstrate awareness of what they do *not* know. Hart (1965) found that when participants judged that they did not know the answer, they subsequently failed to select the target 62% of the time. Furthermore, Glucksberg and McCloskey (1981) proposed two different types of *don't know* judgments. The first is the slow, low-confidence *don't know* made in response to questions like *Is Kiev in the Ukraine?* The other is the fast, high-confidence decision made when asked questions like *What is Jimmy Carter's favorite color?* These two *don't know* judgments differ because of the amount of related knowledge stored in memory. When people have some knowledge about the topic (e.g., Ukraine), they search memory for the target answer. However, when people know little to no related information, no search occurs, resulting in a quick *don't know* response. In line with this dichotomy, Glucksberg and McCloskey found that people responded *don't know* more quickly to questions like *Does Bert Parks have a degree in journalism?* than to questions where the participants presumably drew on some relevant knowledge (e.g., *Does Anne Landers have a degree in journalism?*).

Of course, people's judgments of what they know are not perfect; the classic example is the *hindsight bias*, whereby people claim to have "known it all along" when told the answer to a question they could not answer (see Roese & Vohs, 2012 for a review). Similarly, people sometimes claim to know about nonexistent topics, for example, reporting use of fictitious products (i.e., *overclaiming;* e.g., Phillips & Clancy, 1972). Why do people overclaim? While multiple factors are likely involved, recent work highlights the role of self-perceived domain knowledge (Atir, Rosenzweig, & Dunning, 2015). In one experiment, participants rated their general knowledge about finance compared to the average American, then rated their knowledge about specific financial concepts. While some of these topics were real (e.g., tax bracket), the researchers fabricated others (e.g., pre-rated stocks). Finally, participants took a financial literacy quiz to estimate their actual domain knowledge. Higher self-perceived knowledge predicted overclaiming, independent of participants' actual domain knowledge. Furthermore, this effect appeared to be domain-specific, in that self-perceived expertise in a given domain related specifically to over-claiming in that domain and not others.

3. EXAMPLES OF ERRORS

3.1 Overview

Knowledge is impressive but not perfect; it is virtually unlimited in capacity and lasts for years, but it is not always available when one needs it. Gaps in knowledge are not surprising; what is more interesting than simple errors of omission are errors of commission. That is, people also believe things that are not actually true: errors can be stored in the knowledge base. For example, many people believe that we only use 10% of our brains, that the crime rate increases during the full moon phase, that seasons reflect differences in the physical distance between the Earth and the Sun, and that raindrops are tear-drop shaped. Misconceptions arise across domains; people hold false beliefs about science (e.g., McCloskey, 1983; Munson, 1994; Nakhleh, 1992), health (e.g., Lee, Friedman, Ross-Degnan, Hibberd, & Goldmann, 2003; Wynn, Foster, & Trussell, 2009), and probability and chance (e.g., Ferland, Ladouceur, & Vitaro, 2002), among many others. We will now describe a few examples that have both real-world parallels and laboratory analogs, before turning to some general principles that help explain *why* these errors occur.

3.2 The Grading Problem

Most educators have experienced the unfortunate feeling of becoming "dumber" after grading error-ridden exams and papers. Brown (1988) captured this phenomenon experimentally, examining how exposure to spelling errors hurts one's ability to spell correctly. After checking that participants knew how to spell the target words, they read or generated misspellings; later, they spelled the words again. People tended to switch from a correct to an incorrect spelling on the final test after seeing errors in the interim. This finding is striking, given that participants had likely spelled and read the correct version hundreds of times in the past. Why does a professor's brief exposure to a spelling mistake matter so much?

3.3 Side Effects of Reading Novels and Watching Movies

Movies and novels often are set in real places, refer to actual objects, and occur in familiar time periods, so they constitute a source of information about the world. For this reason, educators sometimes incorporate fiction into their course materials to better engage the students in learning (e.g., Dubeck, Bruce, Schmuckler, Moshier, & Boss, 1990). However, by definition, works of fiction contain inaccuracies, and as such have the potential to serve as sources of misinformation about the world. For example, viewers pick up errors from historically inaccurate portrayals in films and furthermore often misattribute that information to a historically accurate text (Butler, Zaromb, Lyle, & Roediger, 2009; Umanath, Butler, & Marsh, 2012). In one study, participants first read a historical text about the Satsuma Rebellion and then watched a clip from the popular film, *The Last Samurai*. While the text accurately stated that the Emperor Meiji hired a French military advisor to help quell the rebellion, the film inaccurately portrayed the advisor's nationality as American. Even after instructions to rely only on their memory of the text, participants answered questions (e.g., *From what country did Emperor Meiji hire military advisors*?) with inaccuracies depicted in the films (Butler et al., 2009).

We captured the fiction reader's experience in the lab by giving participants short stories with characters, dialogue, and plot. Critically, each story contained references to facts (see Marsh, 2004 for materials). The references were correct (e.g., *paddling around the largest ocean, the Pacific*), neutral (*paddling around the largest ocean*), or misleading (e.g., *paddling around the largest ocean, the Atlantic*). Participants later took a general knowledge test that probed the critical facts from the stories (e.g., *What is the largest ocean on Earth*?). In multiple experiments, reading misinformation (e.g., *Atlantic*) dramatically increased students' production of that error on the final general knowledge test. This effect occurred even after explicit warnings that authors of fiction often take liberties with the truth (Marsh & Fazio, 2006) and drawing attention to the errors with text signals (Eslick, Fazio, & Marsh, 2011). Why do students continue to rely on such low-credibility sources?

3.4 Repeated Claims Feel True

Used car salesmen, politicians, advertisers, and carnival barkers all capitalize on one truism: repeating something makes it seem truer. In the laboratory, falsehoods like *Zachary Taylor was the first President to die in office* appear truer if they were also seen earlier in the experiment. This phenomenon, coined *illusory truth* (Hasher, Goldstein, & Toppino, 1977), generalizes to political opinions (Arkes, Hackett, & Boehm, 1989) and claims about consumer products (Hawkins & Hoch, 1992). The influence of repetition persists minutes (e.g., Begg & Armour, 1991), weeks (e.g., Bacon, 1979), and even months (e.g., Brown & Nix, 1996) after initial exposure to the claim. Reminding participants that claims come from untrustworthy sources reduces the effect (e.g., Begg, Anas, & Farinacci, 1992); as mentioned earlier, though, people often fail to monitor source. Henkel and Mattson (2011), for example, found an illusory truth effect for statements that participants later identified (correctly or incorrectly) as coming from an unreliable source. Why do the repeated statements made by politicians or salesmen influence us, even when voters and consumers realize that they may not be credible sources?

3.5 Tests Can Teach Errors

Considerable controversy surrounds the educational value of tests; concerns include teaching to the test and decreasing motivation in students. Another possible side effect involves the potential of some tests, particularly multiple-choice tests, to introduce errors into the knowledge base. Multiple-choice tests by definition pair a correct answer with multiple plausible, but incorrect, answers. In other words, such tests expose students to more errors than correct answers.

Across multiple studies, people reproduced some of the multiple-choice lures from an initial test on later tests (e.g., Bishara & Lanzo, 2015; Fazio, Agarwal, Marsh, & Roediger, 2010; Odegard & Koen, 2007; Roediger & Marsh, 2005). In one study, students took a multiple-choice test consisting of retired SAT II questions about biology, chemistry, U.S. history, and world history (Marsh, Agarwal, & Roediger, 2009). The correct answer and four lures accompanied each question; participants also had the option to skip questions. After a short delay, participants took a final general knowledge test consisting of short-answer questions; some of these corresponded to the earlier multiple-choice items, whereas others were new. Prior testing helped overall: participants were more likely to answer correctly if a question had appeared on the earlier multiple-choice test. However, recent exposure to multiple-choice questions also increased the probability of incorrectly answering the questions with lures. This effect diminishes over a delay and with feedback (Butler & Roediger, 2008; Marsh, Fazio, &

Goswick, 2012), but increases if students rehearse the multiple-choice lure on a later test. Why do students pick up errors from assessment tools that otherwise boost learning?

4. ADAPTIVE PROCESSES THAT ALSO SUPPORT ERRORS

4.1 Overview

This section focuses on the mechanisms by which errors enter the knowledge base with implications for how to correct them (see Section 6). It is easy to identify potential sources of errors: we could point to textbooks that unfortunately contain errors (e.g., Cho, Kahle, & Nordland, 1985), the learner's own illogical reasoning (e.g., Clement, Narode, & Rosnick, 1981), and other people such as family and friends (e.g., Landau & Bavaria, 2003). Knowing that errors exist in the real world, however, does not tell us why people accept them as facts, reproduce them later, and let them influence their beliefs. For example, how do incorrect portrayals of amnesia in films (e.g., Uncle Fester in *The Addams Family*, Dory in *Finding Nemo*) contribute to people's misunderstanding of amnesia and traumatic brain injury? When viewers simply know nothing about neuroscience, it is unsurprising that they rely on such depictions. However, as discussed below, the problem extends beyond mere naiveté. Below we consider five interrelated properties of how knowledge is encoded, stored, and retrieved; all are properties that normally support accurate knowledge, but sometimes backfire and allow the introduction of errors into the knowledge base.

4.2 Property #1: Bias to Believe Information Is True

Daily life barrages people with new information, some true and some false. How do people decide the truthfulness of claims in the environment? Do they automatically know the truthfulness of statements like *The Pacific is the largest ocean on Earth* and *Lexington is the capital of Kentucky*? One argument is that comprehending a statement requires automatically accepting it as true; "unbelieving" involves a second, resource-demanding step. Gilbert reintroduced this idea to psychology in the 1990s, borrowing from the philosopher Baruch Spinoza. He illustrates the automatic acceptance of new information with a library analogy: the librarian assumes all books to be nonfiction *unless* they are marked by a special "fiction" tag. The reader may wonder why a librarian would ever take this approach to sorting books; put simply, it saves time and resources. It would take a librarian a lot more time to tag every book as fiction or nonfiction, as opposed to simply tagging the fictional ones. This strategy makes sense in the real world as well, where truths occur more frequently than falsehoods.

In most situations, automatically accepting claims conserves time and energy. However, strains on cognitive resources, like competing demands on one's attention, may prevent readers from actually reappraising and "unbelieving" false claims. Gilbert, Krull, and Malone (1990) demonstrated this phenomenon experimentally by blocking the evaluative, unbelieving step. Participants first learned fictional statements (e.g., *A twyrin is a doctor*) that were explicitly labeled as "true" or "false." While reading some items, participants performed a second task, which presumably interrupted the unbelieving step. In a second phase, participants judged the truth of the claims seen earlier. Compared to reading statements alone, distraction led participants to make more "true" judgments later. In other words, participants never reached the stage of evaluating and tagging false statements.

Critically, Gilbert, Tafarodi, and Malone (1993) replicated this effect with real-world judgments. Participants read two crime reports, each containing both true (black font) and false (red font) information about robberies, with instructions to evaluate the information carefully; they knew that they would later play the role of a judge. In one report, the false information exacerbated the crime (e.g., *The robber had a gun*); in the other, the false information extenuated the crime. Half of the participants proceeded through the reports uninterrupted while the other half completed a second task while reading. Later, they decided how long the prison term of each perpetrator should be (0–20 years). Participants who read the reports under full attention assigned similar sentences to the crimes exacerbated and extenuated by false information. Interrupted participants, however, incorporated false information into their judgments: the prison terms for exacerbated crimes were twice as long as those for extenuated crimes. Disturbingly, this pattern emerged even after a strong warning to pay attention to detail and with false information explicitly marked. Even when preparing to make consequential decisions, people initially accept claims as true. This "economical" and "adaptive" bias (Gilbert, 1991) ultimately leaves us vulnerable to errors.

4.3 Property #2: Fluency-Based Heuristic for Judging Truth

As explained in Section 2.4, people often experience knowledge as information they "just know" without necessarily "remembering" where they first learned it (e.g., Tulving, 1985). Instead of thinking back to a particular time and place, people judge their knowledge based on how easy or hard it is to retrieve information. Assuming one knows the capitals of France and Turkey, Paris likely comes to mind more easily than Ankara. This experience of *retrieval fluency* is in turn interpreted as confidence in one's answer.

Kelley and Lindsay (1993) experimentally demonstrated that retrieval fluency causes "illusions of knowledge." Participants read a list of words before completing a general knowledge test. Critically, some of the studied words were semantically related to the answers to the subsequent test questions. For example, participants studied *Hickock,* a reference to the American folk character "Wild Bill" Hickock, and then later answered a question about a different folk character with a similar name: *What is Buffalo Bill's last name? Hickock* came to participants' minds easily due to the recent exposure; people answered the general knowledge questions with related, but wrong, answers because they were easy to retrieve. People made these errors with high confidence, demonstrating that their responses were not just guesses. In other words, participants mistook retrieval fluency for actually knowing an answer.

In our own work, we showed that this illusion can occur even if people remember the source of the misinformation. That is, people read stories containing errors (e.g., *St. Petersburg is the capital of Russia*), which increased their likelihood of answering later general knowledge questions with misinformation (e.g., answering *What is the capital of Russia?* with *St. Petersburg*). Critical for present purposes is that readers made two judgments about each answer; first, they indicated whether or not they had read each answer in one of the experiment's stories and second, they indicated whether or not they had known the answer prior to coming to the experiment. The bottom line is that readers were good at remembering the story sources, but they *also* misattributed these answers to pre-experimental knowledge (Marsh, Meade, & Roediger, 2003). This finding likely reflects how information in the real world is often encountered in multiple contexts—meaning that remembering a lower credibility source does not preclude information from also being associated with a reliable pre-experimental source. This illusion of prior knowledge even occurs in young children (Goswick, Mullet, & Marsh, 2013).

More generally, the ease with which we process information (i.e., fluency) serves as an extraneous cue for many judgments, including truth; perceptions of truthfulness increase when information pops to mind or even when statements are easy to read. Because **Antananarivo is the**

capital of Madagascar is much easier to read than Antananarivo is the capital of Madagascar, the former seems more truthful (Reber & Schwartz, 1999). As mentioned earlier in Section 3.3, illusory truth also occurs with repetition; repeated statements are easier to process, and thus receive higher truth ratings (see Dechêne, Stahl, Hansen, & Wänke, 2010, for a meta-analysis).

This heuristic proves to be both cognitively inexpensive and effective, as fluency naturally correlates with truth (Unkelbach, 2007). On average, the single true version of a statement (e.g., *The capital of Argentina is Buenos Aires*) occurs more frequently in the environment than any one of its many possible falsifications (e.g., *The capital of Argentina is La Paz, The capital of Argentina is Lima, The capital of Argentina is Montevideo,* etc.). People learn this relationship between truth and fluency with experience, as relying on fluency typically leads to the correct judgment (Unkelbach, 2007). In the absence of timely and accurate feedback, as is often the case in real life, people accept fluent errors and incorporate them into their knowledge bases. This tendency renders people vulnerable to misinformation in repeated advertisements, political propaganda, and rumors.

As many advertisers seem to understand, repetition is not the only means of creating a fluent experience. Pairing a statement like *The first windmills were built in Persia* with a photograph of a windmill in an unidentifiable field increases truth ratings. This effect, coined *truthiness*, occurs despite the fact that the picture provides no further evidence for the specific claim (Newman, Garry, Bernstein, Kantner, & Lindsay, 2012). Several mechanisms likely contribute to the power of pictures; for example, the photograph of a windmill may encourage people to generate pseudo-evidence for the claim (e.g., "this field looks arid, so maybe it was taken in Persia"). Among these candidates, fluency is bolstered by the most empirical evidence. Critically, truthiness only occurs when people view a mixture of statements, some paired with pictures and others appearing alone (Newman et al., 2015); in other words, statement–picture pairs seem truer only when contrasted with statements appearing alone, which are presumably less fluent. This result parallels the finding that illusory truth only emerges when people rate a mixture of repeated and new statements.

4.4 Property #3: The Knowledge Base Is Productive

Not all information incorporated into the knowledge base needs to be directly encountered in the outside world. People go beyond the information stored in memory to generate new knowledge. Consider a simple example, whereby 6-year-old children successfully integrated facts to arrive at novel inferences (Bauer & San Souci, 2010). Children learned that *Dolphins live in groups called pods* and *Dolphins communicate by clicking and squeaking*, and then later demonstrated knowledge that *Pods communicate by clicking and squeaking* (which was never explicitly stated). Electroencephalography in adults suggests that newly inferred facts possess a phenomenology similar to that of facts learned long ago: well-known and integrated facts resulted in similar P600 responses, which reflect the ease with which information is processed (Bauer & Jackson, 2015). This positive activity peaking at 600 ms also reflects whether information is stored in long-term memory. In other words, people readily generate inferences that "feel like" facts they learned years ago.

This remarkable ability allows us to bridge gaps but also has the potential to introduce errors into the knowledge base. For example, consider the reader faced with the following passage:

> *That's why we had to go to St. Petersburg, but at least I got to see the Kremlin while there. Her family came too—even though they lived in Russia's capital city, they had never visited the Kremlin!*

What happens when the reader is later asked *What is the capital of Russia?* The passage incorrectly implies, but never explicitly states, that the capital is St. Petersburg. Butler, Dennis, and Marsh (2012) demonstrated that inferences formed while reading persist, leading participants to reproduce errors on a later general knowledge test. In addition to forming incorrect inferences following misinformation, people may self-generate errors by false analogy or other misapplications of logical processes.

People also generate inappropriate inferences when they can retrieve knowledge relevant to one, but not both, objects of a comparison. When presented with the question *Which city is larger, Milan or Modena?* most people respond *Milan* simply because they recognize its name. After successfully retrieving knowledge about Milan and failing to retrieve knowledge about Modena, people arrive at the inference that Milan is larger. This *recognition heuristic* extends beyond judgments of city size: people infer that a recognizable option scores higher than an unknown option on any criterion.

The recognition heuristic provides a cognitive shortcut, allowing people to make a judgment where they otherwise could not. As is the case for fluency, relying on recognition leads to accurate judgments in many cases, even when pitted against more complex approaches. When asked to judge the relative sizes of German cities, for example, American students outperformed their German counterparts, while the reverse pattern emerged for American cities (Katsikopoulos, 2010). Here, using additional cues actually harmed performance (i.e., *less-is-more effect*). However, recognition can be misleading in plenty of situations; for example, many people would incorrectly answer *Which city is larger, Hong Kong or Chengdu?* with *Hong Kong* on the basis of recognition alone. A remarkable proportion of people (nearly 50%) consistently base their judgments on recognition even when faced with three contradicting cues (e.g., learning that a recognized city does *not* have an international airport; Pachur, Bröder, & Marewski, 2008). In addition to leading people astray at the time, the errors generated during these comparisons may persist over time.

4.5 Property #4: Existing Knowledge Supports New Learning

Learners do not enter new situations as blank slates; they bring along a wealth of knowledge, from specific concepts to generalized structures (see Section 2.2). Most of the time, this knowledge supports understanding of, and later memory for, novel concepts. Experts generally learn new information within their expert domain more readily than novices, because they can scaffold and integrate this new learning with an impressive foundation. For example, Van Overschelde and Healy (2001) looked at the ability of baseball experts and nonexperts to learn fictitious statements about real baseball players (e.g., *Sammy Sosa likes to read books by John Steinbeck*); the fictitious materials ensured that participants could not have these details stored in memory. Even though the facts were not relevant to the game, baseball experts outperformed novices on a memory test.

However, knowledge can also interfere with new learning, particularly when used inflexibly. For example, children typically learn that a *particle* refers to a specific entity (e.g., a particle of sand), which then interferes with their ability to grasp a different conceptualization of particle, as used in physics. This problem is an example of *linguistic transfer*, whereby an established use of a word interferes with learning a new usage (Clancy, 2004). Another example involves the simple terms *while* and *then*, which take on different meanings for computer programming than they do in everyday English. Specifically, the everyday usage of the word *then* has temporal implications, which interfere with thinking about *then* as a conditional (e.g., Bonar & Soloway, 1985).

Similarly, knowledge interferes with problem solving when the rules change. For example, *functional fixedness* occurs when a learner cannot think of a new way to use a familiar object. In the classic candle problem, Duncker

(1945) instructed participants to attach a candle to a wall, after giving them a box of tacks, candles, and matches. Most people attempted to complete the task by attaching the candle directly to the wall, either with a tack or by melting the candle. Few participants realized that the most effective solution utilized the box by attaching it to the wall with a tack, then resting the candle in the box. People "fixated" on the well-learned function of the box as a container, rendering them unable to conceptualize its alternative function as a shelf.

Problems with linguistic transfer and functional fixedness reflect a functional system; people adaptively default to the typical meaning of a word or use of an object rather than regularly reevaluate the range of a meaning or function, as this typically suffices. Reliance on a schema is similar; it works like heuristics in that it is cognitively efficient and generally leads one to the correct answer, but can also occasionally lead one astray. For example, when learning about complex phenomena, teachers and students often make comparisons to well-known objects and processes; sometimes these examples lead people astray when they overgeneralize the example. When computer science instructors discuss the concept of a *variable*, they often compare them to boxes—boxes hold things, just like a variable does. People possess a detailed schema for boxes (e.g., rectangular, often made of cardboard, holds things), and to some extent this helps. However, students may be misled by this analogy, as multiple objects fit in a box, but a variable possesses only one definition.

Larger problems arise when a learner brings along misconceptions or other erroneous beliefs. For example, many introductory psychology students possess preconceived notions about a wide range of course topics. Specifically, many believe that suicide rates peak in adolescence, that hypnosis can recover memories, and that polygraphs can accurately detect lies (Taylor & Kowalski, 2004). These misguided ideas set the stage for *proactive interference* (see Anderson & Neely, 1996 for a review), where students' misconceptions interfere with learning and remembering new, updated information.

4.6 Property #5: Partial Matches Are Often "Good Enough"

As described above, people often hold false beliefs or incorrectly apply otherwise correct knowledge; there are also situations in which people *neglect* their knowledge. That is, despite "knowing better," people miss referents that contradict their knowledge (i.e., *knowledge neglect;* see Marsh and Umanath 2013 for a review). Consider a passage about a plane that crashes on the border of France and Spain, which ends with *The authorities were trying to decide where to bury the survivors*. Readers generated a solution to this problem, although the question does not make sense, as survivors are alive and should not be buried. Despite knowing the definition of *survivors*, people often responded with answers like, "let the relatives decide" instead of noting that survivors should not be buried (Barton & Sanford, 1993). Similarly, in what is called the Moses Illusion, people willingly answer questions like *How many animals of each kind did Moses take on the ark?* despite later demonstrating knowing that the correct referent is *Noah* (Erickson & Mattson, 1981). The bottom line is that people often do not take advantage of the information stored in memory.

In line with Gilbert's ideas, it is adaptive to accept new information given that it is "close enough" to the correct answer, as opposed to carefully comparing incoming information to knowledge (*partial match theory;* Reder & Kusbit, 1991). This strategy makes sense when one considers that speech is full of disfluencies, including breaks, nonlexical utterances (e.g., "uh," "erm"), mispronunciations, and inappropriate word choices. Anyone who has ever sat through a recording of his own speech or presentation can attest to this; we have all probably thought, *do I really sound like that?* Comprehension presents an enormous challenge if people always algorithmically process language. Instead, we form shallow representations and use knowledge to fill in the gaps, allowing us to interpret garbled input streams like *in the grand scream of things* (*good enough theory;* Ferreira, Bailey, & Ferraro, 2002). Accepting speech that is "good enough" offers necessary flexibility, given that the same information often presents itself in slightly different forms.

Unfortunately, people accept close matches even when errors are not "slipped in" and unexpected; warnings are generally ineffective (e.g., Marsh & Fazio, 2006). Even trial-by-trial detection does not eliminate knowledge neglect. In one study, participants read stories sentence-by-sentence and (in one condition) made a keypress every time they detected an error. Those in the control condition received no special instructions other than a general warning that stories could contain errors. All participants read stories that included misleading references (e.g., *paddling around the largest ocean, the Atlantic*). Following the reading phase, participants answered general knowledge questions like *What is the largest ocean on Earth*?

Readers demonstrated some ability to detect errors: they were more likely to press the error key when sentences contained misinformation than when they contained correct or neutral references. However, this performance was hardly impressive; readers inappropriately flagged 20% of error-free statements, while also missing approximately two-thirds of the errors. Critically, Table 2 demonstrates that this pattern occurred regardless of participants' knowledge. Based on Nelson and Narens' (1980) norms, half of the facts were well-known (easy; answered on average by 70% of norming participants) and half were unknown (hard; answered by only 15% of norming participants). In both cases, participants caught approximately one-third of the errors (Marsh & Fazio, 2006). Additional experiments yielded similar results and suggested no advantage for detecting errors that contradicted known facts (e.g., Fazio & Marsh, 2008; Umanath & Marsh, 2012). Detection instructions do reduce the later reproduction of errors on a general knowledge test, but the reduction in suggestibility is relatively small as most errors still slip by the reader.

The assumption in the literature has been that knowledge neglect only occurs for weakly-held knowledge and that individuals would not fail to notice errors that contradict strong or expert knowledge. That is, "Biblically trained people [would] not false alarm to the Moses question" (Reder & Cleeremans, 1990, p. 249). However, there are reasons to question this prediction, given that experts rely on heuristics in many different tasks (e.g., Tversky & Kahneman, 1971). To test these ideas, we recruited biology and history graduate students to look for errors embedded in biology (e.g., *Water contains two atoms of helium and how many atoms of oxygen?*) and history questions (e.g., *In which state were the forty-niners searching for oil?*) (Cantor & Marsh, 2015). Overall, expertise helped: participants were more likely to detect errors that contradicted their expert knowledge. However, even experts missed approximately one-third of total errors, and they reproduced a small portion of these on a later general knowledge test. The experts did not appear to be "reading over" the errors, as **bolding and underlining** the key concepts did not help the experts any more than the non-experts. Instead, experts seemed to be relying on the same "good enough" strategy as the non-experts.

Table 2 *Error Detection During Story Reading*. Proportion of critical sentences labeled as containing an error as a function of question ease (easy or hard) and fact framing (misleading, correct, or neutral)

	Misleading	Correct	Neutral
Easy	.35	.17	.26
Hard	.31	.26	.17

Data are from Marsh and Fazio (2006) Experiment 3.

Of course, all of these examples include a referent that is "good enough," or semantically close enough to the truth for the partial match to be accepted. The Moses Illusion decreases when the referent is semantically or phonologically further from the target, such as when *Moses* is replaced with *Adam* (van Oostendorp & de Mul, 1990; see also Hinze, Slaten, Horton, Jenkins, & Rapp, 2014 for a related effect of plausibility). Monitoring takes effort, and accepting "good-enough" representations is a shortcut that normally works. Accordingly, use of this strategy increases in generally valid contexts, where monitoring is not worth the effort. For example, readers are less likely to catch a problem with *How many animals of each kind did Moses bring on the ark?* when tricky questions are rare (Bottoms, Eslick, & Marsh, 2010). Use of a "good-enough" approach also increases in attention-demanding contexts, given that few resources are left over for monitoring. For example, consumers of entertaining media, like novels and movies, who are busy building mental models to track multiple characters, storylines, and goals, can devote relatively few resources to catching factual inaccuracies.

5. LINGERING QUESTIONS ABOUT ERROR REPRESENTATION AND RETRIEVAL

5.1 Co-existence versus Overwriting

Importantly, we do not believe that errors overwrite or otherwise erase truths already stored in memory; instead, the literature suggests that both the error and correct knowledge can coexist in memory. The most telling evidence is that people can regain access to their correct knowledge, even after producing and using the misinformation. First, many knowledge checks occur *after* the main part of the experiment, and people access correct knowledge that they failed to leverage earlier (e.g., Bottoms et al., 2010). Similarly, the effects of misinformation dissipate over time; as the misinformation is forgotten, people recover access to their correct knowledge (Fazio et al., 2010). On the flip side, one can correct a misconception (see Section 6), but the misconception may re-emerge over time—again meaning that it must still be stored in memory (Butler, Fazio, & Marsh, 2011).

This claim is consistent with event memory, as argued within the eyewitness testimony literature, where the general conclusion was that the eyewitness' original memory was blocked rather than overwritten (McCloskey & Zaragoza, 1985). The effect of misinformation, whether about events or knowledge, is a simple example of retroactive interference, with a recently encountered error-blocking access to correct information stored in memory. The effects at retrieval are different, in that the eyewitness' problem is more likely to be one of source misattribution (Lindsay & Johnson, 1989), whereas the student's problem is that she interprets the ease with which misinformation comes to mind as evidence of truth. But in both cases, the representation is the same: two competing memories.

5.2 Direct Retrieval versus Construction

People often construct, rather than directly retrieve, truth. When one agrees with the statement that Lima is the capital of Portugal or that San Diego is larger than San Antonio, one is not necessarily retrieving that information directly from memory. Direct retrieval is often unnecessary given the shortcuts in the system for judging truth. As already reviewed, people have a bias to claim information is true, and evaluations may be based on how fluently something comes to mind rather than direct retrieval per se. It is clear that direct retrieval is not always attempted; people are very quick to say that they don't know the capital of Jupiter, for example—clearly no attempt at exhaustive search was made (Glucksberg & McCloskey, 1981).

The point here is that there is often an assumption in the literature that heuristics and constructions of knowledge only occur in cases of ignorance (e.g.,Unkelbach & Stahl, 2009)—but that is not the case. Having knowledge does *not* always preclude the use of these heuristics. We further unpack how this might happen in the next section.

5.3 A Fluency-Conditional Model of Illusory Truth

Intuitively, it seems unlikely that repeating *A date is a dried plum* makes people believe it. After all, most people know that drying plums produces *prunes*, not dates. Indeed, researchers assumed that illusory truth only emerges when statements are "ambiguous, that is, participants have to be uncertain about their truth status because otherwise the statements' truthfulness will be judged on the basis of their knowledge" (Dechêne et al., 2010, p. 239). In other words, people presumably only rely on fluency when they lack knowledge. Along these lines, Unkelbach and Stahl's (2009) model of illusory truth includes a knowledge parameter that is intentionally low; they used obscure materials, assuming that knowledge eliminates illusory truth.

Recent work in our lab, however, demonstrates the opposite: repetition can influence people's belief that *A date is a dried plum*. Participants read statements that contradicted well-known and unknown facts, as estimated by norms. After rating these and new statements' truthfulness, they completed a knowledge check to determine which specific facts they knew. In contrast to dominant assumptions, repetition swayed judgments of both known and unknown facts. This pattern emerged regardless of whether knowledge was estimated via norms or defined on the basis of individuals' knowledge check performance. In other words, fluency does sometimes "trump" knowledge (Fazio, Brashier, Payne, & Marsh, 2015).

The behavioral data are clear on their own, but we also tested these ideas using multinomial models. Figure 2 shows the two different models: the

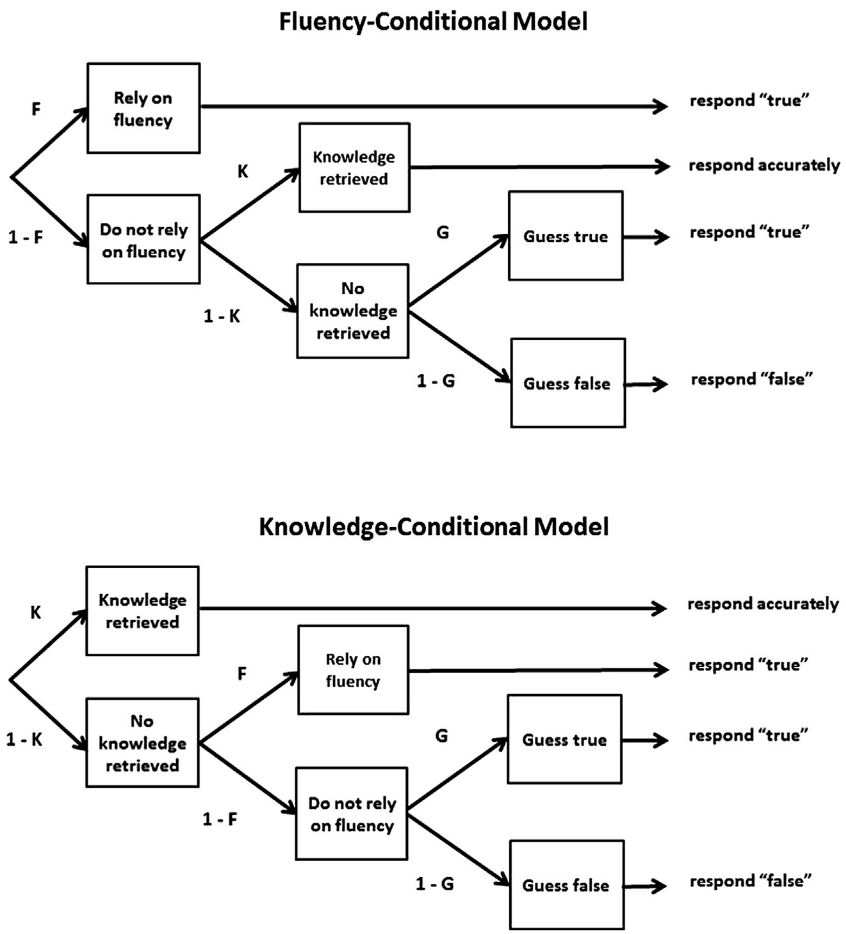

Figure 2 Fluency-conditional and knowledge-conditional models of illusory truth. K = knowledge; F = fluency; G = guess true. *Figure adapted from Fazio, et al. (2015).*

fluency-conditional model, wherein fluency can supersede knowledge, and the *knowledge-conditional model* assumed in the literature, where fluency only influences behavior in the absence of knowledge. In a new experiment requiring binary judgments, the behavioral result replicated and model testing revealed that the fluency-conditional model fit the data better than the knowledge-conditional model. In other words, people *sometimes* rely on fluency despite the fact that contradictory knowledge is stored in memory. Of course, people do not always rely on fluency over knowledge; fluency may be discounted (i.e., explained away) or absent due to inattention on a given trial. However, the success of the knowledge-conditional model demonstrates that it is possible to rely on this fluency signal, even when knowledge can be retrieved from memory. Extraneous factors, like repetition, can encourage people to accept contradictions of even well-learned facts.

6. CORRECTING ERRORS IN THE KNOWLEDGE BASE

6.1 Overview

Mistakes are not necessarily a bad thing. Much research shows the advantage of *desirable difficulties* (Schmidt & Bjork, 1992): challenges should be introduced into the learning phase, even if they cause frustration, slow the learning process, and lead to errors. More recently, Kornell and colleagues explicitly argued that trying to answer a question, and responding with an error, benefits memory, so long as feedback is provided. To avoid item-selection problems, they used fictitious stimuli that were *impossible* for participants to answer (borrowed from Berger et al., 1999, see Section 2.5) and weakly related word pairs (pond—frog). They compared the consequences of guessing and erring (followed by correct answer feedback) to the outcomes of simply studying the correct answers. Across experiments, forced guessing yielded performance that matched, and sometimes exceeded, studying (Kornell, Hays, & Bjork, 2009; see also Grimaldi & Karpicke, 2012; Kang et al., 2011). Why did error generation benefit learners? One possibility is that as people attempt an answer, they elaborate and think of related information (e.g., Carpenter, 2009), which facilitates the eventual encoding of the truth. Another possible explanation is that the error serves as a mediator (e.g., Pyc & Rawson, 2010) that connects the question and the correct answer (i.e., the error *green* could serve as an intermediary cue between *pond* and *frog*). Of course, learners needed to receive feedback to correct their answers; guesses unaccompanied by feedback can cause the same problems as the undetected errors discussed earlier.

6.2 Basic Advice for Correction

Correcting misinformation cannot take the form of simply pointing out the error (see Ecker, Swire, & Lewandowsky, 2013; Lewandowsky, Ecker, Seifert, Schwarz, & Cook, 2012 for reviews). Merely telling a learner that she is wrong leaves her with a gap in her knowledge, but nothing to fill it in with. Even if she knows that the information is wrong, she may continue to rely on it (i.e., *continued influence effect*). Studies illustrating this point involve participants reading stories (e.g., about a warehouse fire) and later learning that pieces of information from the story were false (e.g., the negligent storage of volatile materials) and should be disregarded (e.g., Johnson & Seifert, 1994; Wilkes & Leatherbarrow, 1988). Even though participants acknowledge that the retracted information is false, they still make inferences based on the misinformation (e.g., when asked the cause of the fire).

Similarly, research on feedback in educational contexts finds very limited benefits of telling students *whether* their test answers are correct or incorrect (*verification feedback*) (Pashler, Cepeda, Wixted, & Rohrer, 2005). Verification feedback works only when this feedback is delivered on a multiple-choice test (Marsh, Lozito, Umanath, Bjork, & Bjork, 2012). Knowing that one has chosen the wrong multiple-choice alternative allows students to narrow down the remaining possible answers. But when the test is composed of short-answer questions, rather than multiple-choice questions, verification feedback leaves students no closer to the correct answer.

Instead, the key to correcting errors is to replace the error with the correct answer (if there is one) or at least an alternative response. For example, in the story retraction example above, participants relied less on the false information if provided with an alternative account (e.g., that the fire was caused by petrol-soaked rags). Additionally, telling the learner the correct answer to a test question as feedback (e.g., via text, presentation, etc.) effectively decreases error reliance and increases correct performance on a follow-up test (e.g., Pashler et al., 2005). Surprisingly, this advice proves even more effective for correcting erroneous beliefs held with high confidence; when errors take the form of *Sydney is the capital of Australia* or *Water is stored in a camel's hump*, high confidence in the error predicts correction (*hypercorrection effect;* Butterfield & Metcalfe, 2001). This unexpected result has been replicated numerous times (e.g., Butler, Karpicke, & Roediger, 2008; Kulhavy, Yekovich, & Dyer, 1976), although the errors do start to re-emerge over time (Butler, et al., 2011). These results appear to reflect an attentional mechanism, whereby learners pay more attention to feedback when it surprises them (see Butterfield & Metcalfe, 2006 and Fazio & Marsh, 2009).

Unfortunately, sometimes there is no ideal alternative explanation. For example, while the general consensus is that vaccines do not cause autism (e.g., DeStefano & Thompson, 2004), the actual cause is still unknown. Without an alternative response to fill the gap left, such misconceptions are particularly challenging to combat (Ecker et al., 2013). Furthermore, sometimes the truth is known, but it is more complicated than the misconception. Because people tend to prefer simple explanations (Lombrozo, 2007), they may continue to fall back on a simple misconception if the truth is too complex.

A note of caution: when providing feedback, it is important to present the truth without actually repeating the misinformation (Ecker et al., 2013). While the intent may be to repeat the error to then provide the correction, repetition of errors also makes them more fluent and familiar. As a result, this attempt to correct misinformation can actually backfire and increase error reliance (e.g., Skurnik, Yoon, Park, & Schwarz, 2005). For example, correctly stating *The Big Dipper is not a constellation* repeats the association between the *Big Dipper* and *constellation*, potentially strengthening the misconception.

6.3 When Should Feedback Do More than Provide the Answer?

Sometimes exposure to the correct answer is not enough. For example, work from our lab finds that explanation feedback can be even better than correct answer feedback in certain situations (Butler, Godbole, & Marsh, 2013). To examine the benefits of explanation feedback, participants read passages about various topics (e.g., respiratory system, tropical cyclones) and took an initial short-answer test on concepts from the passages. After answering a question, participants either received correct answer feedback, explanation feedback, or no feedback. For example, for the question, *What is the process by which gas exchange occurs in the part of the human respiratory system called the alveoli*? correct answer feedback contained only the answer (*Gas exchange occurs within the alveoli through diffusion*.), while explanation feedback expanded on the answer by including two additional sentences that elaborated on the answer. (*Diffusion is the movement of particles from a region of high concentration to a region of low concentration. The oxygen concentration is high* *in the alveoli and the carbon dioxide concentration is high in the pulmonary capillaries, so the two gases diffuse across the alveolar membrane in opposite directions toward lower concentrations*.)

Two days later, participants returned to take a final test consisting of both repeated questions from the initial test as well as new inference questions that required participants to transfer their knowledge to a new context. Performance on repeated questions benefited from either kind of feedback relative to no feedback, with no difference between explanation and correct answer feedback. However, when the final test required transfer of learning, receiving an explanation improved performance above and beyond receiving the correct answer alone (Butler et al., 2013).

Similarly, research investigating strategies to invoke conceptual change finds that telling the learner the truth (e.g., via expository texts that only explain the truth) does not eliminate science misconceptions (see Guzzetti, Snyder, Glass, & Gamas, 1993 for a meta-analysis). Instead, strategies that induce "cognitive conflict" by directly refuting the misconception effectively promote conceptual change. For example, Hynd and Alvermann (1986) compared the effectiveness of standard texts that presented the correct information to ones that explicitly refuted the misconception while also presenting the correct information. The target was misconceptions about motion theory (i.e., about the path an object would take when launched off a cliff); a pretest confirmed that the majority of participants did hold a misconception about motion theory. Most students endorsed a false "impetus theory," where the projectile curves downward and then falls straight down. Next, students either read a standard text that explained Newtonian mechanics or a refutational text that directly contrasted impetus theory with Newtonian mechanics. On a final test administered 2 days later, students who read the refutational text changed their prior misconceptions more often than students who read the nonrefutational text.

6.4 Problem: Teacher and Student Preferences

One additional challenge is that the most effective feedback does not always match people's intuitions of what will help them the most. For example, many laboratory studies show that delaying feedback is a good thing, likely because it spreads out exposure to the information over time (spaced study; e.g., Butler & Roediger, 2008; Metcalfe, Kornell, & Finn, 2009). However, students do not appreciate this benefit. In our work, we delivered feedback either immediately after a homework deadline or 1 week later in a real engineering course (Signals and Systems) at the University of Texas El

Paso. Students thought they learned more with immediate feedback and disliked the delayed feedback—but the delayed feedback led to better performance on the final exam (Mullet, Butler, Verdin, von Borries, & Marsh, 2014).

A similar discrepancy emerges in students' preferences for the *source* of the feedback they receive. Students strongly prefer instructor feedback over peer feedback, even though peer feedback has the potential to match, or even exceed, the benefit of instructor feedback (e.g., Ertmer et al., 2007; Topping, 1998). We should note, of course, that not all peer feedback is high quality; successful peer feedback requires training and a careful scoring rubric (e.g., Falchikov & Goldfinch, 2000)—the point for present purposes is that it *can* be high quality. Regardless, students feel that instructors are unbiased, motivated to provide quality feedback, and more knowledgeable than other students (e.g., Ertmer et al., 2007). The teacher, on the other hand, may want to assign exercises for which there would be a large amount of time-consuming grading, or the instructor may view the act of providing feedback to be a valuable educational activity (e.g., Ertmer et al., 2007). This issue will continue to assert itself as coursework moves online and increasing numbers of students forces the instructor to choose between assignments that can be automatically graded (multiple-choice questions) versus using peer feedback to evaluate writing and solutions to problems. The Massive Open Online Courses (MOOCs) hosted by Coursera still draw thousands of students, and consequently depend upon peer feedback.

6.5 Problem: Misconceptions May Be Motivated

Misconceptions will be particularly sticky and hard to correct when they support someone's belief system and worldview (Lewandowsky et al., 2012). For example, one study showed that following a retraction, Republicans were less likely than Democrats to correct the misconception that weapons of mass destruction were found in Iraq (Nyhan & Reifler, 2010). Another study demonstrated that Americans were more likely than German participants to continue to rely on retracted information about the Iraq War (e.g., that weapons of mass destruction were found; Lewandowsky, Stritzke, Oberauer, & Morales, 2005). Furthermore, providing corrections that contradict an individual's worldview can even backfire and strengthen his or her belief in the misconception (e.g., Ecker, Lewandowsky, Fenton, & Martin, 2014; Nyhan & Reifler, 2010).

So how can these worldview-consistent misconceptions be corrected? At the moment, there is no perfect solution. However, Lewandowsky and colleagues (2012) offer two tactics that may help. First, to the extent that the correction can be framed within the individual's worldview, the correction will be more successful (e.g., Kahan, 2010). Second, people seem to be more receptive to correcting misconceptions after an opportunity for self-affirmation. For example, when participants first wrote about a value that was important to them and made them feel good about themselves, they were subsequently less likely to report misconceptions (Nyhan & Reifler, 2011).

7. CONCLUSIONS

7.1 The Science of Learning Is Not Different for Errors

Cognitive psychologists have uncovered many learning strategies that promote long-term retention of knowledge (see Dunlosky, Rawson, Marsh, Nathan, & Willingham, 2013 for a review). For example, actively retrieving concepts leads to better memory than passively restudying them (e.g., Roediger & Karpicke, 2006), as does spacing out learning opportunities over time relative to massing them (e.g., Cepeda, Pashler, Vul, Wixted, & Rohrer, 2006). The literature focuses heavily on the learning of true information, but the same principles apply to errors. For example, retrieving errors increases their likelihood of being reproduced on later tests (Barber, Rajaram, & Marsh, 2008). We predict that other learning strategies that promote the learning of truths can promote the learning of falsehoods, such as spaced practice.

7.2 Comparing Errors of Event Memory to Illusions of Knowledge

We know a lot about errors of event memory, and how to encourage the misremembering of words, pictures, and even events from people's own lives. We know less, however, about how errors enter the knowledge base. In some ways, the two types of errors appear to be similar: both show interference effects, for example. Other properties of knowledge suggest that the two kinds of errors differ qualitatively from one another; for example, source-monitoring errors play a large role in errors of event memory (Johnson, Hashtroudi, & Lindsay, 1993), but are unlikely to produce "slips" of knowledge, which people often retrieve without reference to source.

7.3 Open Questions

Obviously many open questions remain, but we end by highlighting three with particularly important practical and theoretical implications. First, we still do not know enough about how to revise strongly held misconceptions such as *vaccines cause autism*. Such misconceptions pose practical challenges as they obviously cannot be recreated in the laboratory, and because the believers of such misconceptions do not want to be corrected. Second, little research speaks to individual differences and whether particular types of people may be especially prone to misconceptions. Third, from a theoretical perspective, we need to better understand the differences between illusions of knowledge and errors of event memories; knowledge differs from events in countless ways, including but not limited to the number of past occurrences, distance in the past, and richness of the original encounter. Future research should examine whether it is these factors, rather than a distinction between knowledge and events, that matter most in the learning and correction of errors.

REFERENCES

Abernathy, E. M. (1940). The effect of changed environmental conditions upon the results of college examinations. *The Journal of Psychology, 10*, 293–301.
Anderson, J. R. (1983). Retrieval of information from long-term memory. *Science, 220*, 25–30.
Anderson, M. C., & Neely, J. H. (1996). Interference and inhibition in memory retrieval. In E. L. Bjork, & R. A. Bjork (Eds.), *Memory (Handbook of perception and cognition)* (2nd ed., pp. 237–313). San Diego, CA: Academic Press.
Arkes, H. R., Hackett, C., & Boehm, L. (1989). The generality of the relation between familiarity and judged validity. *Journal of Behavioral Decision Making, 2*, 81–94.
Atir, S., Rosenzweig, E., & Dunning, D. (2015). When knowledge knows no bounds: self-perceived expertise predicts claims of impossible knowledge. *Psychological Science, 26*(8), 1295–1303. Advanced online publication.
Bacon, F. T. (1979). Credibility of repeated statements: memory for trivia. *Journal of Experimental Psychology: Human Learning and Memory, 5*, 241–252.
Bahrick, H. P. (1984). Semantic memory content in permastore: fifty years of memory for Spanish learning in school. *Journal of Experimental Psychology: General, 113*, 1–29.
Bahrick, H. P., Hall, L. K., & Baker, M. K. (2013). *Life-span maintenance of knowledge*. New York: Psychology Press.
Balota, D. A., & Lorch, R. F. (1986). Depth of automatic spreading activation: mediated priming effects in pronunciation but not in lexical decision. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12*, 336–345.
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.
Bartlett, F. C. (1932). *Remembering: An experimental and social study*. Cambridge: Cambridge University.
Barton, S. B., & Sanford, A. J. (1993). A case study of anomaly detection: shallow semantic processing and cohesion establishment. *Memory and Cognition, 21*, 477–487.
Bauer, P. J., & Jackson, F. L. (2015). Semantic elaboration: ERPs reveal rapid transition from novel to known. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 41*, 271–282.

Bauer, P. J., & San Souci, P. (2010). Going beyond the facts: young children extend knowledge by integrating episodes. *Journal of Experimental Child Psychology, 107*, 452–465.

Bedard, J., & Chi, M. T. H. (1992). Expertise. *Current Directions in Psychological Science, 1*, 135–139.

Begg, I. M., Anas, A., & Farinacci, S. (1992). Dissociation of processes in belief: source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General, 121*, 446–458.

Begg, I., & Armour, V. (1991). Repetition and the ring of truth: biasing comments. *Canadian Journal of Behavioural Science, 23*, 195–213.

Bellezza, F. S. (1984a). Reliability of retrieval from semantic memory: common categories. *Bulletin of the Psychonomic Society, 22*, 324–326.

Bellezza, F. S. (1984b). Reliability of retrieval from semantic memory: information about people. *Bulletin of the Psychonomic Society, 22*, 511–513.

Bellezza, F. S. (1984c). Reliability of retrieval from semantic memory: noun meanings. *Bulletin of the Psychonomic Society, 22*, 377–380.

Bellezza, F. S. (1987). Reliability of retrieving information from knowledge structures in memory: self information. *Bulletin of the Psychonomic Society, 25*, 407–410.

Bellezza, F. S. (1988). Reliability of retrieving information from knowledge structures in memory: scripts. *Bulletin of the Psychonomic Society, 26*, 11–14.

Berger, S. A., Hall, L. K., & Bahrick, H. P. (1999). Stabilizing access to marginal and submarginal knowledge. *Journal of Experimental Psychology: Applied, 5*, 438–447.

Bishara, A. J., & Lanzo, L. A. (2015). All of the above: when multiple correct response options enhance the testing effect. *Memory, 23*, 1013–1028.

Bonar, J., & Soloway, E. (1985). Preprogramming knowledge: a major source of misconceptions in novice programmers. *Human-Computer Interaction, 1*, 133–161.

Boster, J. S., & Johnson, J. C. (1989). Form or function: a comparison of expert and novice judgments of similarity among fish. *American Anthropologist, 91*, 866–889.

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.

Botwinick, J., & Storandt, M. (1980). Recall and recognition of old information in relation to age and sex. *Journal of Gerontology, 35*, 70–76.

Bower, G. H., Black, J. B., & Turner, T. J. (1979). Scripts in memory for text. *Cognitive Psychology, 11*, 177–220.

Brown, W. (1923). To what extent is memory measured by a single recall trial? *Journal of Experimental Psychology, 6*, 377–382.

Brown, A. S. (1988). Encountering misspellings and spelling performance: why wrong isn't right. *Journal of Educational Psychology, 80*, 488–494.

Brown, A. S. (1991). A review of the tip-of-the-tongue experience. *Psychological Bulletin, 109*, 204–223.

Brown, R., & McNeil, D. (1966). The "tip of the tongue" phenomenon. *Journal of Verbal Learning and Verbal Behavior, 5*, 325–337.

Brown, A. S., & Nix, L. A. (1996). Turning lies into truths: referential validation of falsehoods. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*, 1088–1100.

Butler, A. C., Dennis, N. A., & Marsh, E. J. (2012). Inferring facts from fiction: reading correct and incorrect information affects memory for related information. *Memory, 20*, 487–498.

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

Butler, A. C., Godbole, N., & Marsh, E. J. (2013). Explanation feedback is better than correct answer feedback for promoting transfer of learning. *Journal of Educational Psychology, 105*, 290–298.

Butler, A. C., Karpicke, J. D., & Roediger, H. L. (2008). Correcting a metacognitive error: feedback increases retention of low-confidence correct responses. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*, 918–928.
Butler, A. C., & Roediger, H. L., III (2008). Feedback enhances the positive effects and reduces the negative effects of multiple-choice testing. *Memory and Cognition, 36*, 604–616.
Butler, A. C., Zaromb, F. M., Lyle, K. B., & Roediger, H. L., III (2009). Using popular films to enhance classroom learning: the good, the bad, and the interesting. *Psychological Science, 20*, 1161–1168.
Butterfield, B., & Metcalfe, J. (2001). Errors committed with high confidence are hypercorrected. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 1491–1494.
Butterfield, B., & Metcalfe, J. (2006). The correction of errors committed with high confidence. *Metacognition and Learning, 1*, 69–84.
Cantor, A. D., Eslick, A. N., Marsh, E. J., Bjork, R. A., & Bjork, E. L. (2015). Multiple-choice tests stabilize access to marginal knowledge. *Memory and Cognition, 43*, 193–205.
Cantor, A. D., & Marsh, E. J. (2015). *In what US state were the forty-niners searching for oil: Sometimes history experts miss the error* (Manuscript in preparation).
Carpenter, S. K. (2009). Cue strength as a moderator of the testing effect: the benefits of elaborative retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35*, 1563–1569.
Cepeda, N. J., Pashler, H., Vul, E., Wixted, J. T., & Rohrer, D. (2006). Distributed practice in verbal recall tasks: a review and quantitative synthesis. *Psychological Bulletin, 132*, 354–380.
Chi, M. T., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science, 5*, 121–152.
Chi, M. T., Glaser, R., & Farr, M. J. (Eds.). (1988). *The nature of expertise*. Hillsdale, NJ: Erlbaum.
Chi, M. T. H., & Koeske, D. R. (1983). Network representation of a child's dinosaur knowledge. *Developmental Psychology, 19*, 29–39.
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.
Choi, H., & Smith, S. M. (2005). Incubation and the resolution of tip-of-the-tongue states. *The Journal of General Psychology, 132*, 365–376.
Clancy, M. (2004). Misconceptions and attitudes that interfere with learning to program. In S. Fincher, & M. Petre (Eds.), *Computer science education research* (pp. 85–100).
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.
Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review, 82*, 407–428.
Conway, M. A., Cohen, G., & Stanhope, N. (1991). On the very long-term retention of knowledge acquired through formal education: twelve years of cognitive psychology. *Journal of Experimental Psychology: General, 120*, 395–409.
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.
Dechêne, A., Stahl, C., Hansen, J., & Wänke, M. (2010). The truth about the truth: a meta-analytic review of the truth effect. *Personality and Social Psychology Review, 14*, 238–257.
DeStefano, F., & Thompson, W. W. (2004). MMR vaccine and autism: an update of the scientific evidence. *Expert Review of Vaccines, 3*, 19–22.
Dubeck, L. W., Bruce, M. H., Schmuckler, J. S., Moshier, S. E., & Boss, J. E. (1990). Science fiction aids science teaching. *The Physics Teacher, 28*, 316–318.

Duncker, K. (1945). On problem solving. *Psychological Monographs, 58*(5) (Serial No. 270).
Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J., & Willingham, D. T. (2013). Improving students' learning with effective learning techniques: promising directions from cognitive and educational psychology. *Psychological Science in the Public Interest, 14*, 4—58.
Ecker, U. K. H., Lewandowsky, S., Fenton, O., & Martin, K. (2014). Do people keep believing because they want to? Preexisting attitudes and the continued influence of misinformation. *Memory and Cognition, 42*, 292—304.
Ecker, U. K. H., Swire, B., & Lewandowsky, S. (2013). Correcting misinformation — a challenge for education and cognitive science. In D. N. Rapp, & J. Braasch (Eds.), *Processing inaccurate information: Theoretical and applied perspectives from cognitive science and the educational sciences*. MIT Press.
Erickson, T. D., & Mattson, M. E. (1981). From words to meaning: a semantic illusion. *Journal of Verbal Learning and Verbal Behavior, 20*, 540—551.
Ertmer, P. A., Richardson, J. C., Belland, B., Camin, D., Connolly, P., Coulthard, G., et al. (2007). Using peer feedback to enhance the quality of student online postings: an exploratory study. *Journal of Computer-Mediated Communication, 12*, 412—433.
Eslick, A. N., Fazio, L. K., & Marsh, E. J. (2011). Ironic effects of drawing attention to story errors. *Memory, 19*, 184—191.
Falchikov, N., & Goldfinch, J. (2000). Student peer assessment in higher education: a meta-analysis comparing peer and teacher marks. *Review of Educational Research, 70*, 287—322.
Farnsworth, P. R. (1934). Examinations in familiar and unfamiliar surroundings. *Journal of Social Psychology, 5*, 128—129.
Fazio, L. K., Agarwal, P. K., Marsh, E. J., & Roediger, H. L., III (2010). Memorial consequences of multiple-choice testing on immediate and delayed tests. *Memory and Cognition, 38*, 407—418.
Fazio, L. K., Brashier, N. M., Payne, B. K., & Marsh, E. J. (2015). Knowledge does not protect against illusory truth. *Journal of Experimental Psychology: General, 144*(5), 993—1002.
Fazio, L. K., & Marsh, E. J. (2008). Slowing presentation speed increases illusions of knowledge. *Psychonomic Bulletin and Review, 15*, 180—185.
Fazio, L. K., & Marsh, E. J. (2009). Surprising feedback improves later memory. *Psychonomic Bulletin and Review, 16*, 88—92.
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.
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.
Gilbert, D. T., Tafarodi, R. W., & Malone, P. S. (1993). You can't not believe everything you read. *Journal of Personality and Social Psychology, 65*, 221—233.
Glucksberg, S., & McCloskey, M. (1981). Decisions about ignorance: knowing that you don't know. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 311—325.
Gobbo, C., & Chi, M. (1986). How knowledge is structured and used by expert and novice children. *Cognitive Development, 1*, 221—237.
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.
Grimaldi, P. J., & Karpicke, J. D. (2012). When and why do retrieval attempts enhance subsequent encoding? *Memory and Cognition, 40*, 505—513.

Guzzetti, B. J., Snyder, T. E., Glass, G. V., & Gamas, W. S. (1993). Promoting conceptual change in science: a comparative meta-analysis of instructional interventions from reading education and science education. *Reading Research Quarterly, 28*, 116—159.

Hart, J. T. (1965). Memory and the feeling-of-knowing experience. *Journal of Educational Psychology, 56*, 208—216.

Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior, 16*, 107—112.

Hawkins, S. A., & Hoch, S. J. (1992). Low-involvement learning: memory without evaluation. *Journal of Consume Research, 19*, 212—225.

Henkel, L. A., & Mattson, M. E. (2011). Reading is believing: the truth effect and source credibility. *Consciousness and Cognition, 20*, 1705—1721.

Hinze, S. R., Slaten, D. G., Horton, W. S., Jenkins, R., & Rapp, D. N. (2014). Pilgrims sailing the Titanic: plausibility effects on memory for misinformation. *Memory and Cognition, 42*, 305—324.

Hodges, J. R., & Patterson, K. (1995). Is semantic memory consistently impaired early in the course of Alzheimer's disease? Neuroanatomical and diagnostic implications. *Neuropsychologia, 33*, 441—459.

Hynd, C., & Alvermann, D. E. (1986). The role of refutation text in overcoming difficulty with science concepts. *Journal of Reading, 29*, 440—446.

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

Johnson, K. E., & Mervis, C. B. (1997). Effects of varying levels of expertise on the basic level of categorization. *Journal of Experimental Psychology: General, 126*, 248—277.

Johnson, H. M., & Seifert, C. M. (1994). Sources of the continued influence effect: when misinformation in memory affects later inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 1420—1436.

Kahan, D. (2010). Fixing the communications failure. *Nature, 463*, 296—297.

Kang, S. H. K., Pashler, H., Cepeda, N. J., Rohrer, D., Carpenter, S. K., & Mozer, M. C. (2011). Does incorrect guessing impair fact learning? *Journal of Educational Psychology, 103*, 48—59.

Katsikopoulos, K. V. (2010). The less-is-more effect: predictions and tests. *Judgment and Decision Making, 5*, 244—257.

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.

Kornell, N., Hays, M. J., & Bjork, R. A. (2009). Unsuccessful retrieval attempts enhance subsequent learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35*, 989—998.

Kulhavy, R. W., Yekovich, F. R., & Dyer, J. W. (1976). Feedback and response confidence. *Journal of Educational Psychology, 68*, 522—528.

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

Landauer, T. K. (1986). How much do people remember? Some estimates of the quantity of learned information in long-term memory. *Cognitive Science, 10*, 477—493.

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.

Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: continued influence and successful debiasing. *Psychological Science in the Public Interest, 13*, 106—131.

Lewandowsky, S., Stritzke, W. G. K., Oberauer, K., & Morales, M. (2005). Memory for fact, fiction, and misinformation: the Iraq War 2003. *Psychological Science, 16*, 190—195.

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

Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology, 55*, 232–257.

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

Marsh, E. J., Agarwal, P. K., & Roediger, H. L., III (2009). Memorial consequences of answering SAT II Questions. *Journal of Experimental Psychology: Applied, 15*, 1–11.

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

Marsh, E. J., Fazio, L. K., & Goswick, A. E. (2012). Memorial consequences of testing school-aged children. *Memory, 20*, 899–906.

Marsh, E. J., Lozito, J. P., Umanath, S., Bjork, E. L., & Bjork, R. A. (2012). Using verification feedback to correct errors on a multiple-choice test. *Memory, 20*, 645–653.

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

Marsh, E. J., & Umanath, S. (2013). Knowledge Neglect: failures to notice contradictions with stored knowledge. In D. N. Rapp, & J. Braasch (Eds.), *Processing inaccurate information: Theoretical and applied perspectives from cognitive science and the educational sciences*. MIT Press.

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

McCloskey, M., & Zaragoza, M. (1985). Misleading postevent information and memory for events: arguments and evidence against memory impairment hypotheses. *Journal of Experimental Psychology: General, 114*, 1–16.

McNemar, A. (1969). *Psychological statistics* (4th ed.). New York: Wiley.

Medin, D. L., Lynch, E. B., Coley, J. D., & Atran, S. (1997). Categorization and reasoning among tree experts: do all roads lead to Rome? *Cognitive Psychology, 32*, 49–96.

Metcalfe, J., Kornell, N., & Finn, B. (2009). Delayed versus immediate feedback in children's and adults' vocabulary learning. *Memory and Cognition, 37*, 1077–1087.

Metzger, R. L., Boschee, P. F., Haugen, T., & Schnobrich, B. L. (1979). The classroom as learning context: changing rooms affects performance. *Journal of Educational Psychology, 71*, 440–442.

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: evidence of a dependence between retrieval operations. *Journal of Experimental Psychology, 90*, 227–234.

Mitchell, D. B. (1989). How many memory systems? Evidence from aging. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 31–49.

Mullet, H. G., Butler, A. C., Verdin, B., von Borries, R., & Marsh, E. J. (2014). Delaying feedback promotes transfer of knowledge despite student preferences to receive feedback immediately. *Journal of Applied Research in Memory and Cognition, 3*, 222–229.

Munson, B. H. (1994). Ecological misconceptions. *The Journal of Environmental Education, 25*, 30–34.

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

Nelson, T. O. (1988). Predictive accuracy of the feeling of knowing across different criterion tasks and across different subject populations and individuals. In M. M. Gruneberg, P. Morris, & R. N. Sykes (Eds.), *Practical aspects of memory* (Vol. 2, pp. 190–196). New York: Wiley.

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.

Newman, E. J., Garry, M., Bernstein, D. M., Kantner, J., & Lindsay, D. S. (2012). Nonprobative photographs (or words) inflate truthiness. *Psychonomic Bulletin and Review, 19*, 969–974.

Newman, E. J., Garry, M., Unkelbach, C., Bernstein, D., Lindsay, D. S., & Nash, R. A. (2015). Truthiness and falsiness of trivia claims depend on judgmental contexts. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 41*(5), 1337—1348.

Nyhan, B., & Reifler, J. (2010). When corrections fail: the persistence of political misperceptions. *Political Behavior, 32*, 303—330.

Nyhan, B., & Reifler, J. (2011). *Opening the political mind? The effects of self-affirmation and graphical information on factual misperceptions* (Unpublished manuscript). Hanover, NH: Dartmouth College.

Odegard, T. N., & Koen, J. D. (2007). "None of the above" as a correct and incorrect alternative on a multiple-choice test: Implications for the testing effect. *Memory, 15*, 873—885.

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

Pachur, T., Bröder, A., & Marewski, J. N. (2008). The recognition heuristic in memory-based inference: is recognition a non-compensatory cue? *Journal of Behavioral Decision-Making, 21*, 183—210.

Pashler, H., Cepeda, N. J., Wixted, J. T., & Rohrer, D. (2005). When does feedback facilitate learning of words? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 3—8.

Perlmutter, M. (1978). What is memory aging the aging of? *Developmental Psychology, 14*, 330—345.

Phillips, D. L., & Clancy, K. J. (1972). Some effects of "social desirability" in survey studies. *American Journal of Sociology, 77*, 921—940.

Potts, G. R., & Peterson, S. B. (1985). Incorporation versus compartmentalization in memory for discourse. *Journal of Memory and Language, 24*, 107—118.

Pyc, M. A., & Rawson, K. A. (2010). Why testing improves memory: mediator effectiveness hypothesis. *Science, 330*, 335.

Read, J. D., & Bruce, D. (1982). Longitudinal tracking of difficult memory retrievals. *Cognitive Psychology, 14*, 280—300.

Reber, R., & Schwartz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness and Cognition, 8*, 338—342.

Reder, L. M., & Cleeremans, A. (1990). The role of partial matches in comprehension: the Moses illusion revisited. In A. Graesser, & G. Bower (Eds.), *The psychology of learning and motivation* (Vol. 25, pp. 233—258). New York: Academic Press.

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., & Karpicke, J. D. (2006). The power of testing memory: basic research and implications of educational practice. *Perspectives on Psychological Science, 1*, 181—210.

Roediger, H. L., III, & Marsh, E. J. (2005). The positive and negative consequences of multiple-choice testing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 1155—1159.

Roese, N. J., & Vohs, K. D. (2012). Hindsight bias. *Perspectives on Psychological Science, 7*, 411—426.

Ross, P. E. (August 2006). The expert mind. *Scientific American*, 64—71. Retrieved from http://www.arts.uwaterloo.ca/~dkoehler/ACC784/Ross2006.pdf.

Saufley, W. H., Jr., Otaka, S. R., & Bavaresco, J. L. (1985). Context effects: classroom tests and context independence. *Memory and Cognition, 13*, 522—528.

Schmidt, R. A., & Bjork, R. A. (1992). New conceptualizations of practice: common principles in three paradigms suggest new concepts for training. *Psychological Science, 3*, 207—217.

Skurnik, I., Yoon, C., Park, D. C., & Schwarz, N. (2005). How warnings about false claims become recommendations. *Journal of Consumer Research, 31*, 713—724.

Smith, S. M., & Vela, E. (2001). Environmental context-dependent memory: a review and meta-analysis. *Psychonomic Bulletin and Review, 8*, 203–220.

Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: cognitive consequences of having information at our fingertips. *Science, 333*, 776–778.

Taylor, A. K., & Kowalski, P. (2004). Naïve psychological science: the prevalence, strength, and sources of misconceptions. *The Psychological Record, 54*, 15–25.

Thorndyke, P. W., & Hayes-Roth, B. (1979). The use of schemata in the acquisition and transfer of knowledge. *Cognitive Psychology, 11*, 82–106.

Topping, K. (1998). Peer assessment between students in colleges and universities. *Review of Educational Research, 68*, 249–276.

Tulving, E. (1985). Memory and consciousness. *Canadian Psychology, 26*, 1–12.

Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior, 5*, 381–391.

Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin, 76*, 105–110.

Umanath, S., Butler, A. C., & Marsh, E. J. (2012). Positive and negative effects of monitoring popular films for historical inaccuracies. *Applied Cognitive Psychology, 26*, 556–567.

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

Unkelbach, C. (2007). Reversing the truth effect: learning the interpretation of processing fluency in judgments of truth. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33*, 219–230.

Unkelbach, C., & Stahl, C. (2009). A multinomial modeling approach to dissociate different components of the truth effect. *Consciousness and Cognition, 18*, 22–38.

Van Der Wege, M., & Barry, L. A. (2008). Potential perils of changing environmental context on examination scores. *College Teaching, 56*, 173–176.

Van Overschelde, J. P., & Healy, A. F. (2001). Learning of nondomain facts in high- and low-knowledge domains. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 1160–1171.

Wilkes, A. L., & Leatherbarrow, M. (1988). Editing episodic memory following the identification of error. *The Quarterly Journal of Experimental Psychology, 40A*, 361–387.

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.

Yarmey, A. D. (1973). I recognize your face but I can't remember your name: further evidence on the tip-of-the-tongue phenomenon. *Memory and Cognition, 1*, 287–290.