

21

Populist Beliefs

Byproducts of an Adaptive System?

Elizabeth J. Marsh, Matthew Stanley, and Morgan K. Taylor

Gun violence is increasing in America. U.S. taxes are the highest in the world. Teenage pregnancy rates are climbing. The solar industry has more jobs than oil. These claims have all been made by politicians (some liberal, others conservative) and are objectively false but believed to be true by many Americans (for additional examples, see Wood & Porter, 2019). We are interested in understanding how people come to believe such claims, with an approach grounded in the basic science of memory. This approach differs from that of social psychologists, political scientists, media scientists, and other researchers who focus on *motivational* factors that underlie belief. That is, much research shows that people are often committed to their beliefs and will reject new information that contradicts their beliefs (e.g., Ecker & Ang, 2019). We build on this important research by highlighting *cognitive* factors that also affect the acceptance and updating of false beliefs. We see the problem filtered through the lens of what we know about the construction, representation, and updating of knowledge more generally. Humans are impressive learners, in part because we rely on shortcuts during learning to increase cognitive efficiency. Such shortcuts evolved because they normally support accurate cognition and correct judgments about truth, but sometimes they lead the learner astray and result in the acceptance of unsupported information, misinformation, and even disinformation. Critically, these shortcuts are fast, and they often take the place of direct retrieval of information—meaning that people may have the knowledge to reject the misinformation and still fall prey to misinformation.

Before reviewing our guiding principles, we provide an example that highlights two important themes: (1) A motivational account is not mutually exclusive with a cognitive account and (2) demonstrated knowledge does not necessarily protect one from susceptibility to misinformation. This

example comes from a study about misinformation spread during the Gulf War (Lewandowsky et al., 2005). During the war, as with any fast-changing situation, information was reported, updated, debated, and sometimes retracted. Of interest are specific instances in which the press covered claims that later turned out to be false—for example, about the supposed executions of American prisoners of war by Iraqi forces and the existence of weapons of mass destruction (WMDs). Critical for present purposes is that the experiment focused on claims that were officially retracted by newspapers—meaning that people had a chance to know better. Participants were Americans, Germans, and Australians, who did three tasks. First, they rated their memory for each of a series of events from the war (some true, some false but later retracted, and some fictional ones made up for the experiment). Second, they rated their belief about the veridicality of each claim. Finally, later in the experiment, each individual indicated whether they remembered each claim as retracted by the press. Critically, all three groups of subjects showed similar levels of awareness that the retracted stories were actually retracted—but the Americans reported a higher level of belief in the false retracted claims compared to the Germans and Australians. Motivation likely played a role in updating beliefs; on average, Germans (and to a lesser extent, Australians) did not believe that the war's original purpose was to “destroy WMDs”—meaning that the retractions were a better fit to their justifications for the war than was true of Americans. However, Americans demonstrated as much knowledge about the retractions as did the other participants—an example of how knowing better did not prevent them from holding false beliefs. One key procedural choice in the experiment (which we return to later in this chapter) was to ask about memory for retractions at the *end* of the experiment—meaning that Americans may not have retrieved that information when prompted about their beliefs (previewing one of the principles covered in this chapter: that knowledge can be stored in memory but not retrieved).

In this chapter, we review how cognitive processes can contribute to the acceptance of misinformation. Although our focus is on misinformation, it is important to note that these processes are not *rouge* ones but, rather, are processes that normally support veridical learning and inferences. Our cognitive system is efficient, relying on heuristics to reduce the processing load. Heuristics are shortcuts that are *normally* correct—for example, a feeling of familiarity upon encountering a company's name might lead you to infer that it is a good company or at least that you have heard of it before, because in

the past you have experienced this feeling with companies you know. But the shortcut does not have to be correct; for example, perhaps you just saw an advertisement for the company and that is why the name seems familiar. To the extent that the heuristic is normally correct, it is adaptive, even if it occasionally leads the learner astray. So while the chapter focuses on misinformation, the reader should keep in mind the glass half full version of the chapter: that the cognitive system is designed to promote quick and robust learning.

Here, we focus on the heuristics people use to judge truth (for a review, see Brashier & Marsh, 2020) and how truth judgments are made in the larger context of the noisy and complicated world (for a complementary point, see Chapter 22 in this volume). We highlight how having stored knowledge does not preclude the use of heuristics because it may be stored in memory but inaccessible at a particular point in time. We take a deep dive into this last example, showing how activation of one's national identity can change what comes to mind, even when all participants are similarly knowledgeable.

The Role of Source Information

One frequently used cue to truth is the *source of the information*—not surprisingly, people are more likely to trust and be persuaded by information coming from higher credibility sources than lower credibility ones (e.g., Priester & Petty, 1995; for a review, see Pornpitakpan, 2004). Such behavior is generally rational because higher credibility sources are normally more likely to share correct information. Most interventions for misinformation revolve around the idea that critically evaluating sources is the answer to the problem; for example, the first recommendation in an infographic distributed via the webpage of the International Federation of Library Associations and Institutions (IFLA, 2017) is to “consider the source,” followed by “check the author” and to examine supporting sources. Even IFLA’s recommendation to ask oneself “Is it a joke?” revolves around source, telling readers “if it is too outlandish, it might be satire. Research the site and the author to be sure.” However, people often judge sources heuristically rather than doing research to judge the quality of a source. For example, people are more likely to believe an article with additional citations (Putnam & Phelps, 2017) or trust a native speaker over an accented one (Lev-Ari & Keysar, 2010). And it is not difficult for a malicious actor to pose as a credible source. For example, the website PeaceData.net appeared to be operated by a left-wing U.S. news site but was

in fact attributed to Russia (the interested reader will find that the site has been removed, with only warnings by the Federal Bureau of Investigation, Facebook, and *The Washington Post* remaining).

The solution is not as simple as removing the burden of assessing source credibility, although this solution is also a common suggestion. Twitter, for example, uses warning tags to alert users to the possibility of fake news. Facebook has experimented with “disputed by 3rd party fact-checker” tags on posts that algorithms suggest are problematic. One problem is that sources are not mutually exclusive; even if people remember a low-credibility source, they may be unconcerned if they believe the information came from other sources as well (e.g., in our work, people are unbothered when information comes from fictional stories because they believe they encountered it other places as well; E. Marsh et al., 2003). Furthermore, source information is forgotten more quickly than the information itself. This forgetting of source information is actually adaptive given our general goal to have knowledge that can be used flexibly—it is only problematic for the rarer cases when sources are not credible. Social psychologists first documented this “sleeper effect” during World War II (Hovland & Weiss, 1951), and much modern research suggests that misinformation from a lower credibility source is more potent after time has passed (e.g., Underwood & Pezdek, 1998). The problem is that there is another cue to truth that is much less affected by the passage of time: how easily a given statement is processed (*fluency*), which we discuss in the next section.

Fluency as a Cue to Truth

Easy-to-read statements are more likely to be judged as true, as shown in numerous experiments. This interpretation of *processing fluency* is a learned one; based on experience, we have learned that true statements are on average more easily processed than false ones. This correlation occurs because the true statement is on average more likely to be repeated than any one of the infinite possible negations of it (Unkelbach, 2007). Again, it is a heuristic that normally leads us to the correct answer, but the problem is that there are numerous ways to make a false statement fluent.

How does one make a falsehood fluent and easy to read? One way is to simply repeat the statement, a tactic that many used cars salesmen (and politicians) are aware of. For example, consider President Donald Trump’s

arguments for why a wall is needed on the U.S.–Mexico border. He has repeatedly made claims such as

In El Paso they have close to 2,000 murders right on the other side of the wall. And they had 23 murders. It's a lot of murders. But it's not close to 2,000 murders right on the other side of the wall in Mexico.

In his 2019 State of the Union Address, he claimed that

The border city of El Paso, Texas, used to have extremely high rates of violent crime—one of the highest in the country, and considered one of our nation's most dangerous cities. Now, with a powerful barrier in place, El Paso is one of our safest cities.

The truth is much more complicated, but it is clear that El Paso follows national trends of a decrease in violent crime since the mid-1990s, whereas the fence referenced by President Trump was constructed in 2007. However, each time Trump makes a claim about the wall, it becomes more familiar and easier to understand. It does not matter if he uses slightly different language each time (Begg et al., 1985); the claim becomes more *fluent* with repetition.

This effect of repetition on truth is known as the *illusory truth effect*. It has been demonstrated in dozens of studies (see the meta-analysis by Dechêne et al., 2010) and, more important, is not limited to obscure trivia statements. Pennycook and colleagues (2018) used actual headlines in a standard illusory truth experiment, exposing people to one-half of a set of real but fake headlines, such as “Election Night: Hillary Was Drunk, Got Physical With Mook and Podesta” and “Trump on Revamping the Military: We're Bringing Back the Draft.” During the exposure phase, all headlines were presented with a photograph to mimic how a news story appears on social media sites such as Facebook. Half the headlines were fake and half were real. Participants then completed a filler task before rating the entire set of headlines for their truth value. As expected, participants rated the fake headlines as truer when the headlines appeared earlier in the exposure phase compared to headlines that only appeared in the test phase.

Recent evidence suggests that repetition can influence truth judgments even in the face of partisan motivations. Murray and colleagues (n.d.) collected more than 100 actual tweets from Donald Trump posted between November 2016 and October 2019. Sample tweets include “Originally, almost

all models predicted Dorian would hit Alabama,” “California admitted there were a million illegal votes in the 2016 presidential election,” and “Chicago is the city with the strongest gun laws in our nation.” In an initial phase of the study, participants saw half of the tweets along with the relevant context—namely that they came from Donald Trump. Later, they rated a series of Trump tweets for their truth value (on a scale of 1 to 6); half were old (in that they had been seen earlier in the experiment) and half were new (to the experiment). Old tweets were rated as truer than the new ones—an illusory truth effect—even among Democrats who, as a group, are presumably disposed to disbelieve Trump and even when the tweets were patently false (as determined by Politifact). This study is particularly compelling because it is unlikely participants forgot the source of the statements, given that they were all real tweets from the President of the United States. Furthermore, many of these tweets were likely experienced outside of the experiment—and yet a single additional exposure in the study was enough to boost truth ratings, highlighting the power of the fluency cue.

Repetition is not the only way to make information easier to process. High-contrast font is easier to read than low contrast font and accordingly is rated as truer on average (Reber & Schwarz, 1999). High-quality audios of scientific talks are deemed better quality science than lower quality audios of the same content (Newman & Schwarz, 2018). Similarly, accented speech, for example, is more difficult to process and thus is rated as less true than native speech (Lev-Ari & Keysar, 2010). Simple language (e.g., language that rhymes) leads to higher truth ratings (McGlone & Tofiqbakhsh, 2000). And simply adding a photo to a claim increases belief, even if that photo adds no information to support the claim (Newman et al., 2012). Thus, Pennycook et al.'s (2018) study with the actual headlines used two techniques to instill belief—repetition and the presence of a photo—making the claim even easier to process.

Unfortunately, knowing better does not confer immunity against illusory truth. In Pennycook et al.'s (2018) study on fake news, half of the participants had the fake news stories flagged for them during the exposure phase. That is, as has been used by some social media platforms, the fake news headlines were tagged as “disputed by third party checkers” for half of the participants in the study. The illusory truth effect was similar in the two conditions; an explicit warning/negative tag did nothing to help people avoid the fluency cue. Of course, one possibility is that people simply forgot the falsehood tags by the time they did their final set of truth ratings (akin to forgetting

a source)—but there are at least two other possibilities. One is a motivational account: Not all people will believe “falsehood” tags provided by “fact checkers” from the mainstream media (which some people assume is biased). For example, many people believe that the 2020 U.S. presidential election was stolen, regardless of the lack of evidence to support such claims (and the almost immediate dismissals of dozens of lawsuits related to this claim). These believers have not forgotten that the media and the courts disputed these claims (effectively labeling them as falsehoods); rather, they simply do not believe that the media and the courts are correct.

A second possibility is cognitive, namely that people know the truth and yet sometimes fail to retrieve it. In our own studies, we have shown that people who demonstrate knowledge of basic facts such as “A kilt is the short pleated skirt worn by Scotsmen” are nevertheless affected by exposure to the false statement “A sari is the short pleated skirt worn by Scotsmen” (Fazio et al., 2015). The illusory truth effect persists regardless of how plausible (or implausible) the repeated statements are (once floor and ceiling effects are taken into account; Fazio et al., 2019). That is, repetition increases belief in statements such as “The United States was founded in 1979” and “The earth is a perfect square”—even though it seems almost 100% certain that people have corrective information stored in memory for these statements. In such cases, people’s focus on the false claim (and its fluency) is at the expense of the corrective information coming to mind—an example of Tulving and Pearlstone’s (1966) classic distinction between *availability* and *accessibility*. That is, correct information can be stored in memory (available) but not accessible (retrievable), a distinction we return to throughout this chapter.

Tools for Dealing with the Messy World

Consider the challenges facing listeners when parsing President Trump’s response to a question about why fewer African Americans than Whites trust the police:

They have to get better than what they’ve been doing. I mean obviously that was a terrible thing. And I’ve spoken about it numerous times in various speeches. And what’s interesting is I spoke about it when we launched a very successful rocket—a tremendous program that culminated on that day and obviously it goes on from there.

As another example, here is his reference to the Obama administration’s response to the H1N1 pandemic: “And they got very bad marks on the job they did on the swine flu. H1N1, he calls it N1H1. H1N1 got very poor marks from Gallup on the job they did on swine flu.” Such statements highlight how spoken language can be a challenge for the listener, given that utterances are often made rapidly, disfluently, and ungrammatically—and sometimes are garbled, muffled, or accented.

Humans have developed multiple strategies for dealing with the complexity of language, two of which we focus on here. First, we simplify the incoming information stream, creating rough representations that only need to be “good enough” for the task at hand—not perfect (Ferreira et al., 2002; Ferreira & Lowder, 2016; Karimi & Ferreira, 2016). Returning to our examples, the reader likely extracted that Trump thinks Obama did a bad job with swine flu—and we doubt the reader remembers that quote verbatim as that level of detail is unnecessary to understand Trump’s point. Second, humans are predictors, anticipating what a speaker will say based on constraining factors such as semantic context; correct predictions prime the speaker’s upcoming utterance (making it easier to process), and incorrect predictions provide the listener with chances to learn (Bar, 2009). Of course, a rambling speaker may make it almost impossible to make correct predictions about upcoming utterances, putting the listener at a disadvantage. In the case of the second Trump quote in the previous paragraph, the reader may have missed the misinformation in it—namely the reference to Obama polling worse than Trump over the pandemic (Obama’s numbers never reached as low as Trump’s ratings during the pandemic).

Experimentally, psychologists study this situation with a nonpolitical experimental analog in which participants answer a series of general knowledge questions. Critically, they are explicitly warned that some of the questions will be trick questions that contain errors (and that they should respond “wrong” to those). In some ways, this instruction is akin to the suspicion with which one might receive incoming information from a source one expected to spew misinformation—you know there will be errors and you are trying to catch them while following along. Returning to the experimental analog, of key interest are the responses to general knowledge questions that contain incorrect presuppositions such as “How many animals of each kind did Moses take on the Ark?” Many people fail to notice anything wrong with this question and answer “two,” even though later in the study they demonstrate knowledge that the referent should be to Noah, not Moses

(the Moses illusion; Erickson & Mattson, 1981). The Moses illusion is robust and persists even when people see sample erroneous questions (Erickson & Mattson, 1981), when the errors are phonologically related to the correct answer (as opposed to semantically related; Shafto & MacKay, 2000), and across younger and older adults (Umanath et al., 2014). The Moses illusion is reduced when a greater proportion of questions contain errors (Bottoms et al., 2010) and when the questions are printed in a more difficult-to-read font (Song & Schwarz, 2008)—both conditions that likely cause vigilance.

The reader may be wondering about the relevance of the Moses illusion to the real world. Two points are important to note. First, failing to notice an error matters because people repeat some subset of the errors at a later time (Bottoms et al., 2010). That is, when people miss contradictions with stored knowledge, they later produce a subset of those contradictions as fact. Second, the effect occurs even if “knowledge” is defined more stringently, by looking at experts’ ability to detect contradictions of information in their expert domain. Not only do experts know more than novices (Chi et al., 1988) but also their concepts are more differentiated (Johnson & Mervis, 1997) and organized differently from those of novices (Chi et al., 1981). Experimentally, we defined expertise as being students in a PhD program; although these students may not meet all definitions of expertise, they clearly knew more about their PhD discipline than the control domain. Students were recruited from biology and history departments to answer a series of questions about biology and history, with the explicit warning that they were not to answer any questions with incorrect presuppositions. Embedded in the test were distorted questions from both domains; a history question asked “In what state were the forty-niners searching for oil?” and a biology question asked “Water contains two atoms of helium and how many atoms of oxygen?” (the reader should note that the references should be to *gold* and to *hydrogen*). PhD students were more sensitive in their expert domain, answering fewer trick questions—but they still missed approximately 30% of the errors! A similar pattern occurred when the erroneous presuppositions were bolded and underlined in the questions, meaning that the effect occurred even when the critical terms were “flagged” for readers (Cantor & Marsh, 2017).

Returning to the larger topic of misinformation, the point is that people may be relatively unaware of errors made by politicians or others, so long as the errors are “close enough” to the truth. A listener or reader sometimes fails to notice a mismatch between their simplified representation (e.g., a major

figure from the Bible) and what they actually hear or see—but the good news is that they typically do notice blatant errors. That is, a reference to Noah may pass as good enough, but a reference to Nixon would not. Certain linguistic structures provide more camouflage for errors than others; for example, people are less likely to notice an error in the question “How many animals of each kind did Moses take on the ark?” than in the sentence “Moses took two animals of each kind on the ark” (Büttner, 2007). The question presupposes Moses as the sailor on the ark and directs the reader’s attention toward retrieving the number “two” from memory. Burying the error later in the question (as opposed to at the beginning) also reduces the chance that people will notice it (Bredart & Modolo, 1988). In other words, it would be easier to slip misinformation by an informed audience if it occurred in complicated speech that camouflaged errors in questions or near the end of a ramble.

Shifting Access to Knowledge

We are all familiar with the dreaded tip-of-the-tongue state, where one knows one knows something but simply cannot retrieve it. During the 2016 U.S. presidential election, this state was nicknamed an Aleppo moment, referring to the time when presidential candidate Gary Johnson asked a newscaster “what’s Aleppo?” even though the war-torn Syrian city was prominent in news coverage at the time. Later, he failed to name any world leader when asked during a town hall—eventually referring to a former president of Mexico whose name he could not remember (Vicente Fox). Most people do not think that Johnson did not know any world leaders—rather, for a variety of reasons, he could not at that moment generate the names of leaders he knew about. These incidents are additional examples of Tulving and Pearlstone’s (1966) classic distinction between availability and accessibility, discussed previously in this chapter. Information (i.e., the names of world leaders) is stored in memory (available) but not accessible (retrievable). This distinction is particularly well researched in the domain of episodic memory, with an emphasis on understanding the types of retrieval cues that might help people remember events that they are struggling to remember. In contrast, here we focus on shifting accessibility of knowledge.

There are many simple paradigms to study knowledge inaccessibility in the lab; we begin with Brown’s (1923) classic demonstration of the phenomena. He asked undergraduates to recall the U.S. states, reasoning that

these should be well known to his students. Then, 30 minutes later, he asked them to recall the states again. There was no reason to think that students permanently forgot any states during the 30-minute interval, nor did they have the opportunity to learn any new states—but the two recalls were not identical, highlighting that some knowledge was accessible at one point in time but not the other. Subjects remembered an average of 36 states on the first recall attempt; on average, 2 of these were absent and 5 new states were recalled during the second recall attempt.

What are the cues that drive the fluctuation of access to knowledge? In Brown's (1923) case, he did not provide any specific cues; he simply told students to retrieve the U.S. states. Presumably there was some shift in retrieval cues with time, but because students internally generated their own cues, there is no direct evidence for this explanation. Other studies have directly manipulated properties of retrieval cues to better understand which cues affect the accessibility of knowledge. Well-learned information and information used recently are both (not surprisingly) more likely to be accessible. Multiple cues help (Solso & Biersdorff, 1975). Finally, the target itself can serve as a retrieval cue, albeit an exact one. For instance, when knowledge is marginal (stored but not accessible), individuals are equally likely to regain access to it following exposure (e.g., Berger et al., 1999) or recognizing the target among a list of alternatives (Cantor et al., 2015).

Returning to the problem of misinformation, why should we care if knowledge is temporarily inaccessible? We remind the reader of the example provided at the beginning of the chapter, namely that Americans (when prompted) remembered that claims about WMDs had been retracted—but immediately beforehand indicated that they still believed the claims. To the extent that someone knows something but does not retrieve it, it will not affect belief. In the remainder of this chapter, we take a deep dive into this property to examine how one's national identity can affect what knowledge is accessible at a given time.

An Empirical Example Related to Nationalism

As previously described, not all information stored in memory is always accessible. We sometimes struggle to remember well-known names or the meaning of words we have encountered repeatedly in the past, as well as other types of information. Sometimes, knowledge is not accessible because it has

not been used in a long time; other times, something in the environment cues us toward the wrong information or makes it more difficult to access the correct information. In a series of studies recently published in the *Journal of Applied Research in Memory and Cognition*, we investigated whether the activation of one's national identity influences the accessibility of information relevant to that identity (Stanley et al., 2021). Some theorists have argued that our personal identities are complex and multifaceted, composed of many different, contextually activated constructions (H. Marsh & Hattie, 1995; McConnell, 2011). On this view, different identities can be active at different times, and individuals, environments, and institutions all have the power to activate a particular identity across varied contexts. For example, a person's occupational identity as a mechanic might be active at work, but that same person's identity as a father might be active when playing with his children after work. At yet another time, that person's partisan identity might be active when exposed to a political commercial during campaign season, when attending a political rally for a preferred candidate, or when commenting on a political news story on social media. Beyond partisan identities, a person's identity as an American might be active when listening to the Pledge of Allegiance, after exposure to particular symbols (e.g., the American flag or a bald eagle), or when taking a walk through the National Mall in Washington, DC. However, American symbols are not required; one's national identity can also be activated when confronted with an out-group member of another national identity (Wertsch, 2021).

In our studies, we asked whether the activation of one's national identity affects the accessibility of information relevant to that identity. We created a simple laboratory paradigm to examine whether thinking about one's national identity affects the accessibility of knowledge about one's country. Namely, we had participants write about one of two identities (one's national identity, American, or one's self identity), and then we simply asked participants to remember the 50 U.S. states. We followed Brown's (1923) logic that the states are well known to Americans (through school, travel, the media, state-themed postage stamps and quarters, etc.) and are also related to the American identity. No information was taught during the study; it was assumed to exist. At a minimum, there was no reason to believe that knowledge differed as a function of experimental condition, given random assignment to conditions.

Both groups of subjects began by writing about an identity. Subjects in the American identity condition read that

One of the things we are interested in is the American (United States) identity. Please take some time to think about the way your American identity is important to you. Why and how is your American identity important to you? Please imagine that you are describing to a stranger why and how your American identity is important to you. In 3–7 sentences, please type what you would say to that stranger about the importance of your American identity.

In contrast, participants in the self identity condition were prompted

One of the things we are interested in is your personal self identity. Please take some time to think about the way your personal self identity is important to you. Why and how is your personal self identity important to you? Please imagine that you are describing to a stranger why and how your personal self identity is important to you. In 3–7 sentences, please type what you would say to that stranger about the importance of your personal self identity.

Subjects had no problems responding to these writing prompts. For example, sample responses in the American identity condition included that

American identity to me is the ability to walk freely outside without fear of danger. I can go where I please, do what I want to do, and essentially face no restrictions in doing so. Its [*sic*] a feeling of safety in America knowing that your future is pretty much secured as long as you take care of yourself.

Another subject wrote that

I am proud to be an American not because I am a nationalist but because I am proud of my countries [*sic*] history. America has been a beacon of freedom to the world pretty much since it's [*sic*] inception. It has also been a scientific leader because capitalism allows for very efficient research. America is looked up to by pretty much every country in the world.

Participants in the identity condition wrote responses such as

My self-identity is important to me because it always keeps me grounded. If I know who I am it saves me from ever acting outside of myself. I think it gives me hope and peace to know who I am, especially in Christ.

Another wrote,

It's important to me that I have personal integrity. I can't stand liars and value truth. No matter how hard truth may be, it's important to me that it not be obscured. I'm also careful to make sure that, in the pursuit of truth, if I give my word, make a promise, I always keep that promise, stick to my word, no matter how hard it may become. I stand by my friends unconditionally and that's important, too.

After writing about one of their two identities, all participants were asked to recall the 50 U.S. states. To make the task challenging and avoid ceiling effects, they were given only 2 minutes to do so. On average, participants in the American identity condition remembered 22.7 U.S. states, whereas those in the self identity condition remembered 19.9 states. Although the difference between conditions amounted to approximately 3 states, it was still impressive given the overlearned nature of the information. That is, participants had no chance to learn additional states before being asked to recall them, and knowledge of the U.S. states was probably similar between the two groups (given random assignment to conditions).

The reader may be wondering about the claim that the states represent overlearned knowledge, given that people recalled fewer than half of the states. However, it must be noted that the measure here is free recall, not recognition. We predict that if we gave participants a list of states plus foils, they would be at ceiling at identifying the states. Recall is not a perfect measure of what people know, but it is an ideal way of measuring changes in accessibility.

As an aside, we note that preliminary content analyses of recall did not suggest interesting differences in *which* states were recalled or in *when* they were recalled (the order). For example, one might have hypothesized that subjects in the American condition would be more likely to recall the original 13 colonies or other states with particular historical value—but that was not the case. The most recalled states were the most populous ones (California, New York, and Texas). Similarly, recall order did not suggest strategic differences across conditions; it was not, for example, the case that subjects in the American condition were more likely to retrieve the original 13 colonies first. Instead, the state most likely to be recalled first was "Alabama," suggesting that many people used an alphabetical strategy to retrieve the states. Given these results, we did not score the remaining studies for specific state names or recall order.

In a second study, we replicated the effect with two new control conditions. First, we added a no-writing control group, which was not instructed to think of any particular identity (in case the self identity was somehow interfering with recall of the states). Second, a group of participants was asked to write two to five sentences explaining how and why their family identity matters to them. Sample responses from the family identity condition included "My family identity is an important part of who I am. Being married and having children is my life. This creates the person I am. My life revolves around my family" and

My family is a vital part of who I am. They raised me with the values and morals that have shaped my life and I appreciate the way they approach life. Similarly, my family and I share many important characteristics, so I feel very similar to them. At the same time, my family and I help each other frequently, so this has built very strong reciprocal relationships.

Similar to the first study, participants in the American identity condition recalled approximately three more states than those in the no identity condition and two more states than those in the family identity condition. There was no difference in the number of states recalled between the two control conditions.

Originally, we allowed 2 minutes for recall to make the task challenging; however, later studies showed that the identity priming effect disappeared when people received more time to recall the states. When given 7 minutes or unlimited time to recall states, participants remembered the same number of states regardless of whether they had written about their American or self identities. Our explanation is that activating one's cultural identity primes (speeds retrieval of) related information. The mechanism is essentially what is called "semantic priming" in cognitive psychology, which is demonstrated when, for example, people read or make a word-nonword decision faster after reading a related word. Participants are faster to respond to "doctor" if they have just responded to "nurse" rather than "bread" (Meyer & Schvaneveldt, 1971). Our effects are more about speeded access than, for example, inhibition of some states; thus, with sufficient time, the difference between conditions disappears.

We tested these ideas more directly in a final study using a variation of a cumulative recall paradigm (e.g., Roediger & Thorpe, 1978) to obtain a

more fine-grained analysis of recall. After writing about their American or self identities, participants recalled as many states as they could for 1 minute. At the end of each minute, the screen refreshed, the states already recalled were listed, and participants were given an additional minute to recall as many *new* states as they could, taking care to avoid repeating any items. Participants recalled states over six blocks (i.e., 6 minutes total); Figure 21.1 shows the number of additional states recalled each minute. Not surprisingly, participants in both conditions recalled the most states in the first minute and recalled fewer and fewer (new) states as time progressed. Critically, the difference between the two identity conditions was driven by early recall: In the first minute of recall, participants in the American condition recalled more states than did participants in the self condition. Over time, the two conditions recalled similar amounts.

Returning to the issue of nationalism and misinformation, our results highlight how a brief activation of one's American identity changed the accessibility of overlearned information. Our sample was random and not selected for being particularly nationalistic—suggesting that these effects could be much stronger in populations in which the identity was crucial.




Figure 21.1 Number of unique states recalled in each of six 1-minute blocks, as a function of whether subjects wrote about their American or self identities prior to recalling the states.

Conclusion

Humans take a series of shortcuts when evaluating incoming information: We start biased toward truth (Gilbert, 1991) and use cues such as source and ease of processing as markers for truth. Direct retrieval of knowledge is not a given; it may be inaccessible at a given moment, or our fast but cursory processing of complex information may mean that we do notice discrepancies with what is stored in memory. These shortcuts are adaptive because they promote cognitive efficiency and normally lead people to the correct answers—but they can also lead to misinformation acceptance.

Of course, any one false belief cannot be tied definitively to any one action or mechanism, but we can draw on basic science to identify behaviors or phrasing that likely increase belief in populist ideas. Understanding these properties provides a starting point for solutions; for example, knowing that repetition increases belief highlights why corrections should avoid referring to the targeted myth or error (Lewandowsky et al., 2012). Similarly, interventions aimed at getting people to evaluate sources are unlikely to be successful, unless it means that people avoid a communication altogether (E. Marsh & Yang, 2017). Source information is not always accessed and is quickly forgotten, meaning that interventions focused on source are likely short-term interventions. Finally, we are not optimistic about purely informational campaigns. Aside from people's biases to reject information inconsistent with their beliefs, it is not necessarily protective even for people motivated to reject misinformation.

References

- Bar, M. (2009). Predictions: A universal principle in the operation of the human brain. *Philosophical Transactions of the Royal Society: B, Biological Sciences*, 364, 1181–1182.
- Begg, I., Armour, V., & Kerr, T. (1985). On believing what we remember. *Canadian Journal of Behavioural Science*, 17(3), 199–214.
- Berger, S. A., Hall, L. K., & Bahrack, H. P. (1999). Stabilizing access to marginal and sub-marginal knowledge. *Journal of Experimental Psychology: Applied*, 5, 438–447.
- 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(6), 670–678.
- Brashier, N. M., & Marsh, E. J. (2020). Judging truth. *Annual Review of Psychology*, 71, 499–515.
- Bredart, S., & Modolo, K. (1988). Moses strikes again: Focalization effect on a semantic illusion. *Acta Psychologica*, 67(2), 135–144.
- Brown, W. (1923). To what extent is memory measured by a single recall trial? *Journal of Experimental Psychology*, 6, 377–382.
- Büttner, A. C. (2007). Questions versus statements: Challenging an assumption about semantic illusions. *Quarterly Journal of Experimental Psychology*, 60, 779–789.
- 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 & Cognition*, 43, 193–205.
- Cantor, A. D., & Marsh, E. J. (2017). Expertise effects in the Moses illusion: Detecting contradictions with stored knowledge. *Memory*, 25, 220–230.
- 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*. Erlbaum.
- 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.
- Ecker, U. K., & Ang, L. C. (2019). Political attitudes and the processing of misinformation corrections. *Political Psychology*, 40(2), 241–260.
- Erickson, T. D., & Mattson, M. E. (1981). From words to meaning: A semantic illusion. *Journal of Verbal Learning & Verbal Behavior*, 20, 540–551.
- 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, 993–1002.
- Fazio, L. K., Rand, D. G., & Pennycook, G. (2019). Repetition increases perceived truth equally for plausible and implausible statements. *Psychonomic Bulletin & Review*, 26, 1705–1710.
- Ferreira, F., Bailey, K. G., & Ferraro, V. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science*, 11, 11–15.
- Ferreira, F., & Lowder, M. W. (2016). Prediction, information structure, and good-enough language processing. *Psychology of Learning and Motivation*, 65, 217–247.
- Gilbert, D. T. (1991). How mental systems believe. *American Psychologist*, 46, 107–119.
- Hovland, C. I., & Weiss, W. (1951). The influence of source credibility on communication effectiveness. *Public Opinion Quarterly*, 15(4), 635–650.
- International Federation of Library Associations and Institutions. (2017, February 1). *How to spot fake news—IFLA in the post-truth society*. Retrieved January 2021, from <https://www.ifla.org/node/11175>
- 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.
- Karimi, H., & Ferreira, F. (2016). Good-enough linguistic representations and online cognitive equilibrium in language processing. *Quarterly Journal of Experimental Psychology*, 69(5), 1013–1040.
- Lev-Ari, S., & Keysar, B. (2010). Why don't we believe non-native speakers? The influence of accent on credibility. *Journal of Experimental Social Psychology*, 46(6), 1093–1096.
- Lewandowsky, S., Ecker, U. K., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest*, 13(3), 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(3), 190–195.

- Marsh, E. J., Meade, M. L., & Roediger, H. L., III. (2003). Learning facts from fiction. *Journal of Memory and Language*, 49(4), 519-536.
- Marsh, E. J., & Yang, B. W. (2017). A call to think broadly about information literacy. *Journal of Applied Research in Memory and Cognition*, 6, 401-404.
- Marsh, H. W., & Hattie, J. (1995). Theoretical perspectives on the structure of self-concept. In B. A. Bracken (Ed.), *Handbook of self-concept: Developmental, social, and clinical considerations* (pp. 38-90). Wiley.
- McConnell, A. R. (2011). The multiple self-aspects framework: Self-concept representation and its implications. *Personality and Social Psychology Review*, 15(1), 3-27.
- McGlone, M. S., & Tofiqbakhsh, J. (2000). Birds of a feather flock conjointly (?): Rhyme as reason in aphorisms. *Psychological Science*, 11(5), 424-428.
- 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.
- Murray, S., Stanley, M., McPhetres, J., Pennycook, G., & Seli, P. (n.d.). "I've said it before and I will say it again": Repeating statements made by Donald Trump increases perceived truthfulness for individuals across the political spectrum. Unpublished manuscript. <https://doi.org/10.31234/osf.io/9evzc>
- Newman, E. J., Garry, M., Bernstein, D. M., Kantner, J., & Lindsay, D. S. (2012). Nonprobative photographs (or words) inflate truthiness. *Psychonomic Bulletin & Review*, 19, 969-974.
- Newman, E. J., & Schwarz, N. (2018). Good sound, good research: How audio quality influences perceptions of the research and researcher. *Science Communication*, 40, 246-257.
- Pennycook, G., Cannon, T. D., & Rand, D. G. (2018). Prior exposure increases perceived accuracy of fake news. *Journal of Experimental Psychology: General*, 147(12), 1865-1880.
- Pornpitakpan, C. (2004). The persuasiveness of source credibility: A critical review of five decades' evidence. *Journal of Applied Social Psychology*, 34(2), 243-281.
- Priester, J. R., & Petty, R. E. (1995). Source attributions and persuasion: Perceived honesty as a determinant of message scrutiny. *Personality and Social Psychology Bulletin*, 21(6), 637-654.
- Putnam, A. L., & Phelps, R. J. (2017). The citation effect: In-text citations moderately increase belief in trivia claims. *Acta Psychologica*, 179, 114-123.
- Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness and Cognition*, 8, 338-342.
- Roediger, H. L., III., & Thorpe, L. A. (1978). The role of recall time in producing hypernesia. *Memory & Cognition*, 6(3), 296-305.
- Shafto, M., & MacKay, D. G. (2000). The Moses, mega-Moses, and Armstrong illusions: Integrating language comprehension and semantic memory. *Psychological Science*, 11(5), 372-378.
- Solso, R. L., & Biersdorff, K. K. (1975). Recall under conditions of cumulative cues. *Journal of General Psychology*, 93, 233-246.
- Song, H., & Schwarz, N. (2008). Fluency and the detection of misleading questions: Low processing fluency attenuates the Moses illusion. *Social Cognition*, 26(6), 791-799.
- Stanley, M. L., Taylor, M., & Marsh, E. J. (2021). Cultural identity changes the accessibility of knowledge. *Journal of Applied Research in Memory and Cognition*, 10(1), 44-54.
- Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, 5(4), 381-391.

- Umanath, S., Dolan, P. O., & Marsh, E. J. (2014). Aging and the Moses illusion: Older adults fall for Moses but if asked directly, stick with Noah. *Memory*, 22, 481-492.
- Underwood, J., & Pezdek, K. (1998). Memory suggestibility as an example of the sleeper effect. *Psychonomic Bulletin & Review*, 5(3), 449-453.
- 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(1), 219.
- Wertsch, J. (2021). *How nations remember: A narrative approach*. Oxford University Press.
- Wood, T., & Porter, E. (2019). The elusive backfire effect: Mass attitudes' steadfast factual adherence. *Political Behavior*, 41(1), 135-163.