

Searching for Explanations: How the Internet Inflates Estimates of Internal Knowledge

Matthew Fisher, Mariel K. Goddu, and Frank C. Keil
Yale University

As the Internet has become a nearly ubiquitous resource for acquiring knowledge about the world, questions have arisen about its potential effects on cognition. Here we show that searching the Internet for explanatory knowledge creates an illusion whereby people mistake access to information for their own personal understanding of the information. Evidence from 9 experiments shows that searching for information online leads to an increase in self-assessed knowledge as people mistakenly think they have more knowledge “in the head,” even seeing their own brains as more active as depicted by functional MRI (fMRI) images.

Keywords: transactive memory, explanation, knowledge

Just as a walking stick or a baseball glove can supplement the functioning of the body, cognitive tools, computational instruments, and external information sources can supplement the functioning of the mind. The mind can often increase efficiency and power by utilizing outside sources; for tasks like memory, it can rely on cognitive prostheses, such as a diary or a photo album. These external archives can become necessary components of an interdependent memory system (Harris, 1978).

The mind can also become dependent on other minds. When others serve as externalized repositories of information, transactive memory systems can emerge (Wegner, 1987). In these systems, information is distributed across a group such that individuals are responsible for knowing a specified area of expertise. For instance, one person could be responsible for knowing where to find food while another knows how to prepare it. Members of the systems must also track where the rest of the knowledge is stored. Thus, these systems consist of two key elements: internal memory (“What do I know?”) and external memory (“Who knows what?”) (Hollingshead, 1998; 2001). By reducing redundancy, transactive memory systems work to encode, store, and retrieve information more effectively than could be done by any individual.

Transactive memory systems explain how intimate couples (Wegner, Giuliano, & Hertel, 1985) and familiar groups (Kozlowski & Ilgen, 2006; Peltokorpi, 2008) divide cognitive labor and perform efficiently. These systems can form even with complete strangers, as stereotypes can serve as “defaults” or proxies for

another person’s expertise (Wegner, Erber, & Raymond, 1991). Better performing memory systems can emerge through communication strategies that allocate domains of knowledge to individuals in the network. Increased group coordination leads to better problem solving than in comparable groups of strangers (Hollingshead, 1998; Littlepage, Robison, & Reddington, 1997). This communication can take place through explicit negotiation (e.g., “you remember the first 3 digits, I will remember the last 4”), but often occurs implicitly. As a relationship develops, members of the system with higher relative self-disclosed expertise will become responsible for knowledge in that domain. Similarly, an individual with access to unique information will become responsible for that information (Wegner, 1987). When groups have not developed these dependencies, decision-making in real-world interactions can be worse than individuals’ decisions (Hill, 1982).

Transactive memory may have origins in children’s early emerging abilities to navigate the social world and access knowledge in others’ minds. External sources of knowledge, especially parents, teachers, and other social partners, play an integral role in children’s conceptual development (Gelman, 2009). Information learned from others also exerts a powerful influence over children’s notions of what to accept as true—for example, the existence of germs or Santa Claus (Harris, Pasquini, Duke, Asscher, & Pons, 2006). From an early age, children show an emerging but initially limited ability to navigate the terrain of distributed knowledge (Keil, Stein, Webb, Billings, & Rozenblit, 2008). With age they become aware of the breadth, depth, and epistemic limitations inherent to particular kinds of expertise (Danovitch & Keil, 2004). These types of early emerging sensitivities to the content and limitations of other minds may underlie adult transactive memory.

A growing body of theoretical and empirical work suggests that transactive memory systems can be technological as well as social. Though these systems are typically thought to be composed of human minds, our reliance on technology, like the Internet, may form a system bearing many similarities to knowledge dependencies in the social world. The Internet is the largest repository of human knowledge and makes vast amounts of interconnected information easily available to human minds. People quickly be-

This article was published Online First March 30, 2015.

Matthew Fisher, Mariel K. Goddu, and Frank C. Keil, Department of Psychology, Yale University.

This project was made possible through the support of a grant from the Fuller Theological Seminary/Thrive Center in concert with the John Templeton Foundation. The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of the Fuller Thrive Center or the John Templeton Foundation.

Correspondence concerning this article should be addressed to Matthew Fisher, Yale University, 2 Hillhouse Avenue, New Haven, CT 06520-8205. E-mail: matthew.fisher@yale.edu

come accustomed to outsourcing cognitive tasks to the Internet. They remember where to find information and rely on the Internet to store the actual information (Sparrow, Liu, & Wegner, 2011). This evidence suggests that the Internet can become a part of transactive memory; people rely on information they know they can find online and thus track external memory (who knows the answer), but do not retain internal memory (the actual answer).

The Internet has been described as a “supernormal stimulus” in that its breadth and immediacy far surpass any naturally occurring transactive partner to which our minds might have adapted (Ward, 2013a). Even if the Internet lacks the agency of human transactive memory partners, it shares many of their features and may thus be easily treated as their cognitive equivalent. Compared with a human transactive memory partner, the Internet is more accessible, has more expertise, and can provide access to more information than an entire human transactive memory network. These features leave Internet users with very little responsibility for internal knowledge and may even reduce the extent to which users rely on social others in traditional, interpersonal transactive memory systems.

In a sense, a transactive memory partnership with the Internet is totally one-sided: the Internet stores all the knowledge, and the human is never queried for knowledge. Furthermore, there is no need to negotiate responsibility because the Internet is the expert in all domains. However, to access knowledge in the transactive memory system, the Internet user must navigate the Internet’s information in much the same way that one transactive memory partner might know about and query the knowledge contained in another’s mind. This interactive aspect of accessing knowledge on the Internet distinguishes it from the way our minds access other information sources. With its unique, supernormal characteristics that allow us to access it much the same way we access human minds, the Internet might be more similar to an ideal memory partner than a mere external storage device. In short, the cognitive systems may well be in place for users to treat the Internet as functionally equivalent to an all-knowing expert in a transactive memory system.

The particular features of the Internet may make it difficult for users to clearly differentiate internally and externally stored information. In most cases of information search, the boundary between information stored “in the head” and information out in the world is quite clear. When we do not know something ourselves, we must take the time and effort to query another source for the answer. If we go to the library to find a fact or call a friend to recall a memory, it is quite clear that the information we seek is not accessible within our own minds. When we go to the Internet in search of an answer, it seems quite clear that we are consciously seeking outside knowledge. In contrast to other external sources, however, the Internet often provides much more immediate and reliable access to a broad array of expert information. Might the Internet’s unique accessibility, speed, and expertise cause us to lose track of our reliance upon it, distorting how we view our own abilities?

One consequence of an inability to monitor one’s reliance on the Internet may be that users become miscalibrated regarding their personal knowledge. Self-assessments can be highly inaccurate, often occurring as inflated self-ratings of competence, with most people seeing themselves as above average (Alicke, Klotz, Breitenbecher, Yurak, & Vredenburg, 1995; Dunning, 2005; Pronin,

2009). For example, people overestimate their own ability to offer a quality explanation even in familiar domains (Alter, Oppenheimer, & Zemla, 2010; Fernbach, Rogers, Fox, & Sloman, 2013; Fisher & Keil, 2014; Rozenblit & Keil, 2002). Similar illusions of competence may emerge as individuals become immersed in transactive memory networks. They may overestimate the amount of information contained in their network, producing a “feeling of knowing,” even when the content is inaccessible (Hart, 1965; Koriat & Levy-Sadot, 2001). In other words, they may conflate the knowledge for which their partner is responsible with the knowledge that they themselves possess (Wegner, 1987). And in the case of the Internet, an especially immediate and ubiquitous memory partner, there may be especially large knowledge overestimations. As people underestimate how much they are relying on the Internet, success at finding information on the Internet may be conflated with personally mastered information, leading Internet users to erroneously include knowledge stored outside their own heads as their own. That is, when participants access outside knowledge sources, they may become systematically miscalibrated regarding the extent to which they rely on their transactive memory partner. It is not that they misattribute the source of their knowledge, they could know full well where it came from, but rather they may inflate the sense of how much of the sum total of knowledge is stored internally.

We present evidence from nine experiments that searching the Internet leads people to conflate information that can be found online with knowledge “in the head.” One’s self-assessed ability to answer questions increased after searching for explanations online in a previous, unrelated task (Experiment 1a and b), an effect that held even after controlling for time, content, and features of the search process (Experiments 1c). The effect derives from a true misattribution of the sources of knowledge, not a change in understanding of what counts as internal knowledge (Experiment 2a and b) and is not driven by a “halo effect” or general overconfidence (Experiment 3). We provide evidence that this effect occurs specifically because information online can so easily be accessed through search (Experiment 4a–c).

Experiment 1a

Experiment 1a used a between-subjects design to test whether searching the Internet for explanations leads to higher subsequent ratings for the ability to answer entirely different questions in unrelated domains. In one condition, participants used the Internet to find the answers to common explanatory knowledge questions; then, in the second phase, they evaluated their ability to explain the answers to unrelated sets of questions in various domains of knowledge. In the other condition, participants were asked *not* to use the Internet in the initial induction portion of the study and then, in the second phase, assessed their ability to explain the same unrelated questions seen by the participants who had used the Internet.

Method

Participants. Two hundred two participants (119 men, 83 women, $M_{Age} = 32.59$, $SD = 12.01$) from the United States completed the study through Amazon’s Mechanical Turk (Rand, 2012). Based on pilot testing, it was determined that a sample size

of 75–100 participants per condition would be required to detect an effect. Once the requested amount of participants completed the experiment, data collection ended. Five participants were eliminated for failing to follow the instructions to look up answers online; failure to follow instructions was assessed via participants' answers to the following question at the end of the survey: "For how many of the trivia questions at the beginning of this survey did you use the Internet to find the answer? Please answer honestly, this will aid us in our research." Participants in the no Internet condition who chose any number greater than zero were eliminated, and participants in the Internet condition who chose any number fewer than four were eliminated. Participants did not complete multiple experiments; each experiment contains a unique naïve sample. Informed consent was obtained from all participants in all experiments.

Procedure and design. Experiment 1a consisted of two conditions, each with two components: induction and self-assessment. During the induction phase, participants were either instructed to use the Internet to find explanations to common questions (Internet condition) or were instructed not to use the Internet to find the answers to those same questions (no Internet condition). During the second, entirely separate self-assessment phase, participants in both conditions were asked to evaluate how well they could explain the answers to groups of questions in a variety of domains. These questions were entirely unrelated to the questions in the induction phase.

In the induction phase, participants in the Internet condition saw a random subset of four out of six questions about explanatory knowledge such as, "How does a zipper work?" and were asked to search the Internet to "confirm the details of the explanation" (see Appendix A for the full set of questions and Appendix B for the exact instructions). The question contained the phrase "confirm the details" because the explanations were common enough that most people could offer some account without looking up the comprehensive answer. The idea was that participants should have some sense of the answers they were searching for, such that they might more readily and consistently inflate their internal knowledge with the knowledge they were accessing. For this reason, all induction questions were selected from a group of Google autocompleted queries beginning with "Why" and "How"; the questions were selected through piloting with both Internet and no Internet condition instructions to avoid possible ceiling and floor effects. After finding a good explanation for each of the induction questions in this first phase of the experiment, participants in the Internet condition reported the URL of the "most helpful website" and rated their ability to explain the answer to the question on a 1 (*very poorly*) to 7 (*very well*) Likert scale. Participants in the no Internet condition viewed the same random subset of questions and were asked to rate their ability to explain the answers to the questions "without using any outside sources." The purpose of asking participants to rate how well they could explain the answers to the induction questions in this first phase was to track likely differences between confidence in no Internet users and Internet users, who might feel more sure of the answers after looking them up.

During the second phase, the self-assessment phase, all participants rated their ability to answer questions about knowledge in six domains unrelated to the questions posed in the induction phase: weather, science, American history, food, health treatments, and the human body. For each set, participants considered four

questions, for example, "Why are there more Atlantic hurricanes in August and September?", "How do tornadoes form?", "Why are cloudy nights warmer?" (see Appendix C for complete list of questions for each domain set). Participants were asked, "How well could you answer detailed questions about [topic] similar to these?" on a 1 (*very poorly*) to 7 (*very well*) Likert scale.

Results

Participants who had looked up explanations on the Internet in the induction phase rated themselves as being able to give significantly better explanations to the questions in the unrelated domains during the self-assessment phase ($M = 3.61$, $SD = 1.27$, 95% confidence interval [CI] = [3.40, 3.91]) than those who had not used the Internet ($M = 3.07$, $SD = 1.06$, 95% CI = [2.88, 3.27]), $t(195) = 3.24$, $p = .001$, Cohen's $d = 0.50$ (Figure 1). The effect was observable across all six domains for which participants were asked to assess their knowledge.

However, participants in the Internet condition spent longer in the induction phase ($M = 214.40$ s, $SD = 129.93$) than participants in the no Internet condition ($M = 26.26$ s, $SD = 26.26$), $t(195) = 14.65$, $p < .001$. Extended reflection on the initial questions may have increased explanatory confidence, accounting for the difference between conditions. Furthermore, participants in the Internet condition rated themselves as having a better ability to explain the items in the induction phase ($M = 5.00$, $SD = 1.42$) than those in the no Internet condition ($M = 3.34$, $SD = 1.19$), $t(195) = 8.90$, $p < .001$.

In addition, the results of this experiment failed to address the possibility that Internet use was not inflating Internet condition participants' confidence in their knowledge, but rather that the No Internet participants' self-assessed knowledge ratings were *deflated* from baseline by lack of Internet use. Experiment 1b was designed to test whether the Internet participants' self-assessed knowledge ratings were in fact rising from a baseline.

Experiment 1b

Experiment 1b used a nearly identical design to Experiment 1a; the key difference was that this experiment added a self-assessment phase *prior* to the induction phase for both the Internet and the no Internet conditions. This additional knowledge self-assessment phase was identical to the second phase of Experiment 1a: It asked participants to evaluate their knowledge about different domains with representative questions through the question, "How well could you answer detailed questions about [topic] similar to these?" Participants provided ratings on a 1 (*very poorly*) to 7 (*very well*) Likert scale. Because this preinduction self-assessment phase was also identical to the third phase of this new Experiment 1b, its addition was intended to allow for direct comparison between pre- and postinduction self-assessed knowledge ratings, testing whether the observed effect (the difference in postinduction self-assessed knowledge ratings between the Internet and no Internet conditions) occurred because Internet use inflated users' confidence from baseline or because a lack of Internet use *deflated* confidence from baseline in the no Internet condition. Because Internet use is quite widespread in the United States, with more than 71.1% of households reporting accessing the Internet in 2011, we wanted to determine the directionality of

the effect (File, 2013). This was perhaps especially important given that our participants were Amazon Mechanical Turk workers who are presumably heavier Internet users than average.

Method

Participants. One hundred fifty-two participants (76 men, 76 women, $M_{\text{Age}} = 32.14$, $SD = 10.24$) from the United States completed the study through Amazon's Mechanical Turk. Ten participants were eliminated for not following instructions in the induction phase to look up answers online, as judged by their responses to the "Internet check" question at the end ("For how many of the trivia questions at the beginning of this survey did you alter the search provided in the link to find the answer? Please answer honestly, this will aid us in our research").

Procedure and design. Experiment 1b used the same design as Experiment 1a but also included an additional, preinduction self-assessment. During the preinduction self-assessment phase, participants in both the Internet and the no Internet conditions were asked to evaluate how well they could explain the answers to questions about specific subject matter (these questions were identical to the self-assessment questions in Experiment 1a and may be viewed in full in Appendix C). Next, during the induction phase, and exactly as in Experiment 1a, participants either used the Internet to find explanations to common questions or were shown the same questions with instructions not to use the Internet to find the answers. Finally, in the postinduction self-assessment phase, participants responded to the same set of questions from the preinduction self-assessment phase.

Results

There was no difference in the preinduction self-assessment baseline ratings between the Internet ($M = 3.21$, $SD = 1.16$, 95% CI = [3.08, 3.34]) and the no Internet condition ($M = 3.21$, $SD = 1.33$, 95% CI = [3.04, 3.37]), $t(140) = -0.04$, $p = .99$. Replicating the results from Experiment 1a, participants who looked for explanations on the Internet during the induction phase rated themselves as being able to give significantly better explanations in the unrelated domains during the postinduction self-assessment phase ($M = 3.63$, $SD = 1.52$, 95% CI = [3.44, 3.81]) than participants who had not used the Internet ($M = 3.15$, $SD = 1.21$, 95% CI = [2.96, 3.29]), $t(140) = -2.1$, $p < .05$, Cohen's $d = 0.35$ (Figure 1). The effect was observable across all six domains for which participants were asked to assess their knowledge, consistent with the account that the baseline of self-assessed knowledge is systematically inflated because of Internet use.

Just as in Experiment 1a, however, participants in the Internet condition spent longer in the induction phase ($M = 196.24$ seconds, $SD = 166.31$) than participants in the no Internet condition ($M = 28.75$ seconds, $SD = 21.67$), $t(140) = 7.17$, $p < .001$. And just as in Experiment 1a, participants in the Internet condition rated themselves as having a better ability to explain the items in the induction phase ($M = 4.18$, $SD = 1.58$) than those in the no Internet condition ($M = 3.56$, $SD = 1.34$), $t(140) = 2.50$, $p < .05$. Experiment 1c addressed these confounds by equating both time spent in the induction task and amount of learning (i.e., self-assessed ability to explain induction questions) during the induction task across conditions.

Experiment 1c

The results of Experiment 1a and b provided initial evidence for an effect whereby searching the Internet for explanations results in an increase from baseline in self-assessed knowledge for unrelated domains. However, in those experiments, both time spent in the induction phase and information learned during the induction phase may have accounted for the difference between the Internet and no Internet conditions. Experiment 1c sought to rule out these alternative explanations.

Experiment 1c was designed to match both the amount of time spent and the content (i.e., the explanations viewed in the induction phase) across the Internet and no Internet conditions. The design of Experiment 1c also addressed a third possibility: that features of autonomous searching behaviors (including source scrutiny, a sense of self-directed learning, etc.) could explain the difference in self-assessed knowledge ratings between the conditions.

Method

Participants. Two hundred four participants (120 men, 84 women, $M_{\text{Age}} = 32.85$, $SD = 10.29$) from the United States completed the study through Amazon's Mechanical Turk. Nine participants were eliminated for failing to look up the answers to the questions in the Internet condition, as assessed by the Internet check question at the end of the survey.

Procedure and design. The design for Experiment 1c was identical to that of Experiment 1a, with three changes. First, participants in the Internet condition were provided with specific instructions for how to find each explanation to the induction questions, thereby constraining their searching to prespecified sources. For example, participants in the Internet condition were asked, "Why are there dimples on a golf ball?" and instructed, "Please search for the scientificamerican.com page with this information." We specifically selected these web sources because they contained primarily textual content and included little or no graphics.

Participants in the no Internet condition were provided with the *exact text* from the same websites for which Internet participants were instructed to search. (In contrast, in Experiments 1a and b, the no Internet participants received no information at all about the questions in the induction phase.) Because the websites from which the explanations were drawn were specially selected for their heavy textual content, this ensured that participants across conditions viewed the same explanatory content, controlling for the amount of new information participants accessed during the induction phase across conditions.

Last, because the results of a separate pilot study showed that participants took an average of 12.6 s to find the webpage listed in the instructions, Experiment 1c also introduced a 12.6-s delay before explanations were displayed for participants in the no Internet condition, thereby controlling for time spent in the induction phase.

Participants in the Internet condition were excluded either if they failed to provide the URL of the intended page or if they provided any URL different from the intended page for any question in the induction phase.

Results

Participants in the Internet condition spent the same number of time on each explanation ($M = 68.00$ s) as participants in the no Internet condition ($M = 73.36$ s), $t(193) = -1.01$, $p = .31$.

Once again, participants in the Internet condition provided higher self-assessments of knowledge postinduction ($M = 3.78$, $SD = 1.19$, 95% CI = [3.54, 4.03]) than participants in the no Internet condition ($M = 3.07$, $SD = 1.12$, 95% CI = [2.85, 3.28]), $t(193) = 4.30$, $p < .001$, Cohen's $d = 0.63$ (Figure 1).

Though the explanatory content viewed in the induction phase was exactly matched across conditions, when asked how well they could explain the questions in the induction phase, participants in the Internet condition gave higher ratings ($M = 5.07$ $SD = 0.99$) than participants in the no Internet condition ($M = 4.33$, $SD = 1.33$), $t(193) = -4.33$, $p < .001$. A linear regression model, controlling for self-rated ability to explain the questions in the induction phase, $B = .039$, $\beta = .41$, $p < .001$, showed experimental condition to still be a significant predictor of knowledge self-assessments, $B = .042$, $\beta = .17$, $p < .01$. Even though participants in the Internet condition rated their ability to explain the induction questions higher than those in the no Internet condition, this difference did not explain the difference in knowledge ratings in the self-assessment phase.

Notably, in addition to equating for time spent and exact content accessed during the induction phase, the results of Experiment 1c suggest that features of autonomous searching such as evaluating, comparing, or choosing between multiple sources of information cannot explain the effect because participants in the Internet condition were told exactly where to go to retrieve their explanatory information.

Discussion

Across three studies, Experiment 1 demonstrated that searching for explanations online increases self-assessed knowledge in unrelated domains. The effect is observed even when time spent in the induction phase is the same for participants in both Internet and no Internet conditions, when the content viewed across conditions is identical, and when Internet condition participants' autonomous search behavior is restricted through the assignment of a particular web source.

Experiment 2a

Participants in the Internet and no Internet conditions in Experiments 1a–c could have interpreted the dependent measure differently. That is, the phrasing of the dependent measure (“How well could you explain the answers to questions similar to these about [topic]?”) may have left the meaning of “you” open to interpretation. If participants in the Internet condition were considering both the knowledge in their heads *and* online information when assessing their ability to answer questions in various domains, then it would be entirely unsurprising that they deemed themselves more knowledgeable. Experiments 2a and b were designed to resolve this ambiguity inherent to the phrasing of the dependent measure in Experiments 1a–c, thus allowing for more accurate interpretations of findings from those experiments.

Method

Participants. Two hundred three participants (99 men, 104 women, $M_{\text{Age}} = 33.24$, $SD = 11.35$) from the United States completed the study through Amazon's Mechanical Turk. Eleven participants were eliminated for either not looking up the answers to the questions in the Internet condition or using the Internet to look up answers to the questions in the no Internet condition as determined by answers to the end-of-survey Internet check question.

Procedure and design. In Experiment 2a, a new dependent measure replaced those used in the self-assessment phase of Experiment 1. Instead of asking participants to rate how well they could answer questions about topics using a Likert scale ranging from 1 (*very poorly*) to 7 (*very well*), participants were shown a scale consisting of seven functional MRI (fMRI) images of varying levels of activation, as illustrated by colored regions of increasing size (Figure 2). Participants were told, “Scientists have shown that increased activity in certain brain regions corresponds with higher quality explanations.” This dependent variable was designed to unambiguously emphasize one's brain as the location of personally held knowledge. Participants were then asked to select the image that would correspond with their brain activity when they answered the self-assessed knowledge questions in each of the six domains. To further ensure that participants accurately interpreted this new dependent measure as pertaining to their own, independently held knowledge, at the end of the experiment participants explained all of the factors they considered when making judgments about their brain activity via free response. The procedure was otherwise identical to Experiment 1a.

Results

Replicating the effect found in Experiment 1, participants in the Internet condition chose the images with more brain activity ($M = 4.66$, $SD = .99$, 95% CI = [4.40, 4.83]) than those in the no Internet condition ($M = 4.12$, $SD = 1.13$, 95% CI = [3.94, 4.39]), $t(190) = 3.52$, $p = .001$, Cohen's $d = 0.43$ (Figure 1).

Two independent raters coded free responses about the factors participants considered when making judgments about their brain activity and found that Internet participants were not considering knowledge online when making their ratings in the self-assessment phase. When asked what they did consider, 41% spontaneously mentioned their current knowledge, 37% the complexity of their explanation, 30% the complexity of the question, 13% the amount of thinking required, 17% other explanations, and only 2% cited access to other sources. Interrater reliability was high ($\kappa = .83$, $p < .001$) with disagreements resolved through discussion. Removing participants who considered accessing outside sources as a relevant factor had no effect on the significance of the results.

Experiment 2b

The findings of Experiment 2a suggested that participants in both conditions had interpreted the dependent measure in the self-assessment phase as intended—that is, as pertaining solely to the knowledge they held in their heads, rather than to their own knowledge plus knowledge accessible from outside sources. Experiment 2b addressed the possibility of a misinterpretation of the

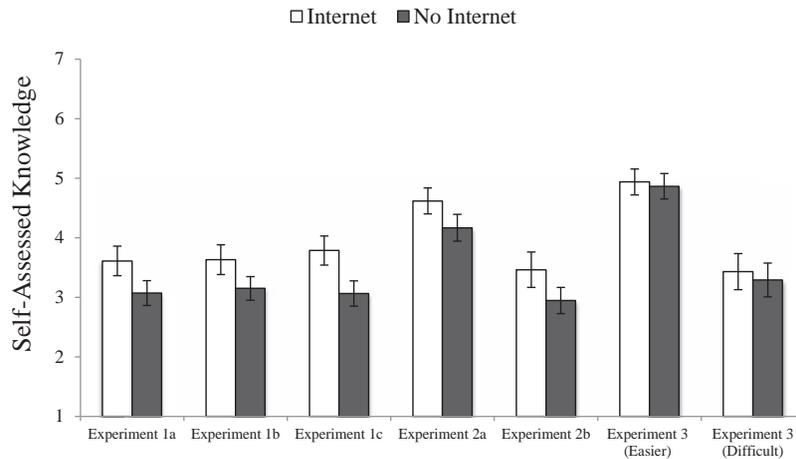


Figure 1. Differences in self-assessed knowledge between the Internet and no Internet conditions for Experiments 1–3. Error bars indicate mean \pm 95% confidence interval (CI).

dependent measures even more directly with instructions clarifying that the ratings in the self-assessment phase should reflect the participant's current knowledge "without any outside sources."

Method

Participants. One hundred ninety-nine participants (127 men, 72 women, $M_{Age} = 31.42$, $SD = 10.92$) from the United States completed the study through Amazon's Mechanical Turk. Twelve participants were eliminated for either not looking up the answers to the questions in the Internet condition or using the Internet to look up answers to the questions in the no Internet condition as determined by the end-of-survey Internet check question.

Procedure and design. The procedure and design of Experiment 2b were identical to those of Experiment 1a with one difference: Each question from the self-assessment phase of the experiment asked participants explicitly how well they could answer questions about this topic "without any outside sources." In previous experiments, this phrase had appeared only in the instructions for the induction phase of the no Internet condition, not in the self-assessment phase questions. The addition of this phrase was intended to explicitly restrict participants' judgments about the boundaries of knowledge to include only their own internal knowledge.

Results

Participants in the Internet condition rated their ability to answer questions without using outside sources higher ($M = 3.41$, $SD =$

1.47, 95% CI = [3.10, 3.72]) than participants in the no Internet condition ($M = 2.94$, $SD = 1.16$, 95% CI = [2.72, 3.16]), $t(193) = 4.30$, $p < .001$, Cohen's $d = 0.36$ (Figure 1).

Discussion

The findings from Experiments 2a and b provide direct evidence that participants interpreted the self-assessed knowledge questions similarly across the Internet and no Internet conditions. In other words, the results of Experiment 2 suggest that participants in the Internet condition of these experiments and Experiment 1 did not consider knowledge available online when rating their knowledge in new domains during the self-assessment phase of the experiment.

Experiment 3

A possible explanation for the observed effect is that using the Internet to access explanations simply increases confidence in one's knowledge or abilities more generally. Experiment 3 was designed to explore the possibility of such a "halo effect" by replacing the topics in the self-assessment phase with explanatory knowledge topics that are similar yet cannot be accessed using the Internet. Detailed autobiographical knowledge is one of the few forms of knowledge that cannot be found online yet closely mirrors the kind of explanatory questions used in the previous experiments. If a difference between the self-assessed knowledge ratings for the Internet and no Internet conditions still exists for these questions, it would suggest that general overconfidence from In-

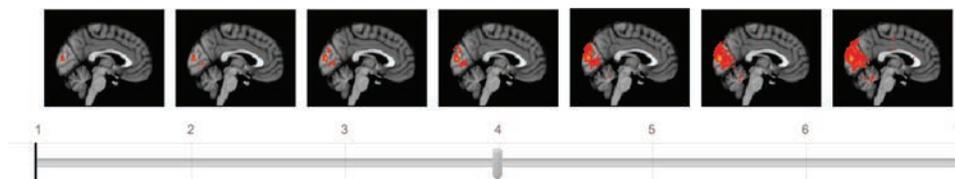


Figure 2. Measure of self-reported brain activity in Experiment 2a. See the online article for the color version of this figure.

ternet use could account for the results of Experiment 1–2. However, no difference between the conditions would be evidence for a boundary condition of the phenomenon, indicating the previous results are not explained by a “halo effect.”

Method

Participants. Three hundred two participants (194 men, 108 women, $M_{\text{Age}} = 31.32$, $SD = 9.86$) from the United States completed the study through Amazon’s Mechanical Turk. Twenty-two participants were eliminated for either not looking up the answers to the questions in the Internet condition or using the Internet to look up answers to the questions in the no Internet condition as determined by responses to the end-of-survey Internet check question.

Procedure and design. In Experiment 3, we changed the type of topics presented during the self-assessment phase of the experiment. Instead of asking about explanatory questions that can be answered using the Internet, we asked participants about autobiographical explanatory knowledge for which the Internet would be of no help. The six knowledge topics were personal history, personal future, relationships, local culture, personal habits, and emotions. For example, questions about relationships were, “Why are you so close with your best friend?”; “How are you similar to your mother?”; and, “How could you become friendlier with your next door neighbor?” (see Appendix D for the full list of autobiographical questions).

In pilot testing, participants viewed these autobiographical knowledge self-assessment questions as significantly easier than the explanatory self-assessment topics from Experiments 1–2 (presumably because they were much more familiar topics); so, we included a second set of questions which were rated to be equally as difficult as the explanatory knowledge questions from Experiments 1–2. The difficult autobiographical questions were grouped into the same categories as the easier autobiographical questions and were chosen from the results of pilot tests in which self-assessed knowledge ratings were similar to the self-assessment questions used in Experiments 1–2 (see Appendix E for the full list of the difficult autobiographical questions).

Experiment 3 thus used a 2 (Internet vs. no Internet) \times 2 (Easier vs. Difficult) between-subjects design. Just as in previous experiments, after the induction phase participants viewed six unrelated knowledge topics (with three representative sample questions each) and were then asked, “How well could you explain the answer to questions about [topic] similar to these?” They provided their responses on a 1 (*very poorly*) to 7 (*very well*) Likert scale.

Results

A one-way analysis of variance (ANOVA) showed that self-assessed knowledge ratings for autobiographical questions were the same after accessing the Internet ($M = 4.24$, $SD = 1.33$) compared with not accessing the Internet ($M = 4.04$, $SD = 1.38$), $F(1, 276) = 1.37$, $p = .30$ (Figure 1). As expected, participants gave higher ratings for the easier autobiographical questions ($M = 4.91$, $SD = 0.96$) compared with the difficult autobiographical questions ($M = 3.35$, $SD = 1.25$), $F(1, 276) = 134.21$, $p < .001$. There was no interaction, indicating that at both levels of difficulty, using the Internet did not boost self-assessed knowledge for questions that could not be found online.

Discussion

Experiment 3 suggests that accessing the Internet does not lead to a general overconfidence, but rather to a more specific illusion of knowledge that occurs only in domains where the Internet would be of use. If induction time and induction content (Experiment 1), imprecise interpretations of personally held knowledge (Experiment 2), and general overconfidence effects (Experiment 3) cannot explain the observed inflation of self-assessed knowledge, what else might? Experiment 4 explores whether the process of querying through Internet search might be the underlying mechanism explaining the effect.

Experiment 4a

The findings from Experiments 1–3 raise important questions about the locus of this effect. Experiment 4a provides evidence that actively posing queries through Internet search engines is the specific mechanism by which Internet usage causes an increase in self-assessed knowledge. Experiment 4a was designed to investigate (a) whether the effect persists when the act of searching is removed from accessing explanations on the Internet and (b) whether the effect persists even when less popular search engines are used for searching. If there is no effect when the act of searching is removed from accessing explanations yet remains even when Internet users couple with unfamiliar “partners,” together these results would be strong evidence that active searching is the element of Internet access that drives the observed effect in which participants inflate their self-assessed knowledge after Internet use.

Method

Participants. One hundred fifty-seven participants (77 men, 80 women; $M_{\text{Age}} = 31.94$, $SD = 11.34$) from the United States completed the study through Amazon’s Mechanical Turk. Nine participants were eliminated for failing to follow instructions by using a search engine they were not instructed to use; this was assessed using an end-of-survey Internet check question that explicitly asked participants whether they had used a search engine other than the one they had been assigned (“For how many of the trivia questions at the beginning of this survey did you alter the search provided in the link to find the answer? Please answer honestly, this will aid us in our research”).

Procedure and design. Experiment 4a contained two conditions, each a variation on the design of Experiment 1c. The first condition, the no search condition, was designed to test whether there is an increase in knowledge self-assessments when the search component is entirely removed from the induction task for participants in the Internet condition. In this condition, participants saw an experimental setup that was identical to that of the Internet condition of Experiment 1c; however, instead of searching for the answers to induction questions online themselves, participants in this condition were provided with a link that took them directly to the website with the explanation.

Second, the other search engines condition was designed to test whether the effect might only occur using a search engine with which one has successfully queried for knowledge in the past. If an association between a particular search engine and successfully

accessing knowledge drives the increase in self-assessed knowledge in the Internet conditions of the previous experiments, then using less popular search engines should yield a weaker effect (or perhaps none at all). However, if actively querying an information-rich source via search engine drives the effect, then participants should increase their knowledge self-assessments after using *any* search engine. The other search engine condition was identical to the Internet condition of Experiment 1c, except that participants were instructed to look up the answers to the induction questions using one of the following 5 search engines (varying in popularity): duckduckgo.com, blekko.com, ixquick.com, search.yahoo.com, ask.com.

Results

Participants who used links to access information instead of searching provided lower self-knowledge ratings ($M = 3.20$, $SD = 0.99$, 95% CI = [3.09, 3.31]) than participants who used other search engines ($M = 3.63$, $SD = 1.27$, 95% CI = [3.49, 3.77]), $t(146) = -2.28$, $p = .03$, Cohen's $d = 0.37$. A one-way ANOVA showed no difference in self-assessed knowledge across the different website that participants accessed, $F(4, 77) = .63$, $p = .64$.

Experiment 4b

The findings of Experiment 4a suggest that active search is necessary in order for Internet usage to result in inflated estimates of self-assessed knowledge. To further explore the account that searching is the specific mechanism by which the effect occurs, Experiment 4b investigated whether the effect still holds when searching the Internet yields unhelpful information. Does self-assessed knowledge increase even if the search engine does not provide a satisfactory answer?

Method

Participants. One hundred fifty-one participants (106 men, 45 women, $M_{\text{Age}} = 29.79$, $SD = 7.94$) from the United States completed the study through Amazon's Mechanical Turk. Six participants were eliminated for not following instructions by removing the Google search filters; this was assessed using an end-of-survey Internet check question that explicitly asked participants whether they had removed the Google search filters ("For how many of the trivia questions at the beginning of this survey did you alter the search provided in the link to find the answer? Please answer honestly, this will aid us in our research").

Procedure and design. The design of Experiment 4b's two conditions was identical to that used in the Internet condition induction in Experiment 1c except for one difference: in place of the original explanatory induction questions, each of the two conditions of Experiment 4b used a new set of induction questions. These two new sets of induction questions were matched in structure and content to each other; the only difference between them was that one set contained questions that a top Google search result could answer comprehensively (the answer condition), while the other set consisted of questions with answers that could not be found using Google (no answer condition). For example, participants in the Answer condition would be asked "Why is ancient Egyptian history more peaceful than Mesopotamian history?",

which returns an article that clearly answers the question, and participants in the no answer condition would be asked "Why is ancient Kushite history more peaceful than Greek history?", which is parallel in content and structure yet does not have an answer easily found online (see Appendix F for the full set of questions). Just as in the induction phase of Experiment 1c, participants saw a random subset of 4 induction questions and were instructed to "search the Internet to confirm the details" of their explanations to these questions. After the induction phase, they completed the same general knowledge self-assessment questions used in previous experiments.

Results

Knowledge ratings did not differ between participants in the answer condition ($M = 4.00$, $SD = 1.19$, 95% CI = [3.74, 4.26]) and those in the no answer condition ($M = 4.11$, $SD = 1.22$, 95% CI = [3.81, 4.41]), $t(143) = -0.55$, $p = .58$. To draw meaningful comparisons between these results and those of previous experiments, we combined the results of relevant earlier experiments. Experiments 1a–c and 2b provided four successful demonstrations of the effect (though Experiment 2a also successfully replicated the effect, the change in the dependent measure to fMRI pictures shifted ratings higher compared with the other experiments). Using the data from the no Internet condition of these four previous studies for comparison with the results of Experiment 4b, we can determine whether search activity, even if unsuccessful, leads to increased ratings of knowledge.

Pooling across the no Internet conditions from the previous studies, we found that participants in the answer condition of Experiment 4b increased their knowledge ratings compared with the aggregate no Internet baseline formed by combining Experiments 1a–c and Experiment 2b ($M = 3.05$, $SD = 1.13$, 95% CI = [2.94, 3.16]), $t(476) = -6.81$, $p < .001$. The results of this comparison hold if the answer condition ratings are compared individually to the no Internet condition from Study 1a, $t(153) = -4.43$, $p < .001$, Study 1b, $t(188) = -5.66$, $p < .001$, Study 1c, $t(184) = -5.46$, $p < .001$, and Study 2b, $t(185) = -6.15$, $p < .001$. Surprisingly, even despite unsuccessful search efforts, participants in the no answer condition *also* increased their self-assessed knowledge compared with the aggregated no Internet ratings, $t(461) = -6.94$, $p < .001$. This result also holds when the no answer condition is compared with each of the previous no Internet conditions individually, Study 1a, $t(138) = -4.69$, $p < .001$, Study 1b, $t(173) = -5.91$, $p < .001$, Study 1c, $t(169) = -5.70$, $p < .001$, and Study 2b, $t(170) = -6.32$, $p < .001$. This is strong support for searching as the mechanism that gives rise to illusions of knowledge from Internet use.

Experiment 4c

The findings from Experiment 4b provide initial evidence that search activity, even when unsuccessful because of hard-to-find relevant results, drives the observed effect. Experiment 4c further explores the extent to which the retrieval of search results causes Internet users to inflate their self-assessed knowledge in that it uses an even stronger test. Experiment 4c asks whether the illusion persists even when searching returns only irrelevant results or no results at all. If search *success* is causing participants to inflate their self-assessed knowledge, then participants who access irrelevant results or zero

results will have lower ratings of self-assessed knowledge than participants in the analogous Experiments 1a–c and 2b who successfully access relevant search results. However, if search activity alone, regardless of search success, is driving the inflation of self-assessed knowledge, then participants who search unsuccessfully will rate themselves higher than participants in the no Internet conditions from previous experiments.

Method

Participants. One hundred thirty-eight participants (men = 86, women = 52; $M_{\text{Age}} = 31.26$, $SD = 10.39$) from the United States completed the study through Amazon's Mechanical Turk. Seven participants were eliminated for removing the filters placed on the search; this was assessed using an end-of-survey Internet check question that explicitly asked participants whether they had removed the search filters they had been assigned.

Procedure and design. Experiment 4c consisted of two conditions that were each variations on the design of the Internet condition of Experiment 1a. These conditions were together designed to investigate whether impeding the effectiveness of participants' search activity by filtering search results affected participants' subsequent self-assessed knowledge ratings.

In the filtered results condition, participants were instructed to search for the explanations to the induction questions using a filtered Google search that provided only the most recently posted results (i.e., within the past week). These recent results from the filtered Google search did not provide direct answers to the induction questions. In the no results condition, participants were instructed to search for the answers to the induction questions using a Google filter that blocked *all* results, with the Google results page displaying a message to participants that their search "did not match any documents."

Results

Participants in the filtered results did not differ in their self-assessed knowledge ($M = 3.57$, $SD = 1.27$, 95% CI = [3.27, 3.87]) compared with those in the no results condition ($M = 3.75$, $SD = 1.17$, 95% CI = [3.46, 4.04]), $t(129) = -.82$, $p = .42$. Again pooling together ratings from participants in the no Internet conditions of Experiments 1a–c and 2b to form a baseline for comparison, we found that participants in the filtered results condition rated their knowledge higher ($M = 3.57$, $SD = 1.27$) than participants who had not searched online for answers to the induction questions ($M = 3.05$, $SD = 1.13$, 95% CI = [2.94, 3.16]), $t(465) = -3.49$, $p = .001$. The results holds if the filtered results ratings are compared individually to the no Internet condition from Study 1a, $t(142) = -2.06$, $p < .05$, Study 1b, $t(177) = -2.84$, $p < .01$, Study 1c, $t(173) = -2.74$, $p < .01$, and Study 2b, $t(174) = -3.43$, $p = .001$.

Strikingly, the ratings from the no results condition, in which participants' searching activities returned *zero* search results at all, were also higher ($M = 3.75$, $SD = 1.17$) than the aggregate no Internet condition ($M = 3.05$, $SD = 1.13$, 95% CI = [2.94, 3.16]), $t(458) = -4.50$, $p < .001$. When compared individually, the no results ratings were also higher than the no Internet condition from Study 1a, $t(135) = -2.93$, $p < .01$, Study 1b, $t(170) = -3.86$, $p < .001$, Study 1c, $t(166) = -3.71$, $p < .001$, and Study 2b, $t(167) = -4.37$, $p < .001$.

Discussion

The illusion of knowledge from Internet use appears to be driven by the act of searching. The effect does not depend on previous success on a specific search engine, but rather generalizes to less popular search engines as well (Experiment 4a). It persists when the queries posed to the search engine are not answered (Experiment 4b) and remains even in cases where the search query fails to provide relevant answers or even any results at all (Experiment 4c). Even when stripped of such potentially integral features, Internet searching still results in increases in self-assessed knowledge. This suggests that the illusion is driven by the act of searching itself.

General Discussion

Searching for answers online leads to an illusion such that externally accessible information is conflated with knowledge "in the head" (Experiment 1a and b). This holds true even when controlling for time, content, and search autonomy during the task (Experiment 1c). Furthermore, participants who used the Internet to access explanations expected to have increased brain activity, corresponding to higher quality explanations, while answering unrelated questions (Experiment 2a). This effect is not driven by a misinterpretation of the dependent measure (Experiment 2b) or general overconfidence (Experiment 3) and is driven by querying Internet search engines (Experiment 4a–c).¹

In many ways, our minds treat the Internet as a transactive memory partner, broadening the scope of knowledge to which we have access. The results of these experiments suggest that searching the Internet may cause a systematic failure to recognize the extent to which we rely on outsourced knowledge. Searching for explanations on the Internet inflates self-assessed knowledge in unrelated domains. Our results provide further evidence for the growing body of research suggesting that the Internet may function as a transactive memory partner (Sparrow, Liu, & Wegner, 2011).

People tend to inaccurately recall the original source of their internal memories (Johnson, 1997; Johnson, Hashtroudi, & Lindsay, 1993). In this regard, our findings might be initially unsurprising: When searching online, people misattribute the source of the specific answers they find because they think the answer was stored in their own mind instead of on the Internet. However, in the current set of studies, we first asked people one set of questions in the induction phase and then asked them an entirely separate set of questions in different domains of knowledge. This design rules out the explanation that participants merely failed to monitor the fact that the Internet was the source of their knowledge. Rather, our results suggest that what participants failed to accurately monitor was the proportion of internal and external memory comprising the sum total of accessible knowledge, mistaking outsourced knowledge for internal knowledge. People neglect the extent to which they would rely on their partner in the transactive memory system to access explanatory knowledge.

This illusion of knowledge might well be found for sources other than the Internet: for example, an expert librarian may experience a similar illusion when accessing a reference Rolodex.

¹ See Table 1 for summary of experimental findings.

Table 1
Summary of Experimental Results

| Experiment | Method | Results (self-assessed knowledge ratings) | Conclusions |
|------------|--|--|---|
| 1a | Internet condition uses Internet to look up explanations to common questions; no Internet condition does not. | Internet > no Internet | Internet condition gives higher self-knowledge ratings than no Internet condition |
| 1b | Same as Experiment 1a, but all participants make self-assessed knowledge ratings both before and after induction phase. | Preinduction, no difference between Internet and no Internet | Searching the Internet increases self-assessed knowledge from baseline |
| 1c | Internet condition searches constrained to specific sources; no Internet condition sees identical explanations | Controlling for induction phase ratings, Internet > no Internet | Time, content, and autonomous search activity do not explain the effect |
| 2a | Same as 1a, but DV for self-assessed knowledge questions = fMRI "brain activity" | Internet > no Internet | Participants are not misinterpreting the dependent measure |
| 2b | Same as 1a, but DV for self-assessed knowledge questions = "on your own, with no outside sources" | Internet > no Internet | Participants are not misinterpreting the dependent measure |
| 3 | Same as 1a, but questions for self-assessed knowledge phase are autobiographical explanatory questions | No difference between Internet and No Internet conditions | Effect not explained by general overconfidence |
| 4a | Link condition clicked on a link to explanation instead of searching; Other search engines condition used 5 alternative engines for searching | Other search engines > Link | Effect driven by active search independent of search engine |
| 4b | Answer condition searched for induction questions easily found on Internet; no Answer condition searched for matched-content questions not answered in any search result | Both answer and no answer conditions > no Internet baseline | Effect holds even without direct answers to search query |
| 4c | Recent results condition searched for induction explanations in Google search returning irrelevant recent results only; zero results condition returned zero search results. | Both recent results and no results conditions > no Internet baseline | Effect holds even without any results for search query |

An individual in a highly integrated social environment (Hutchins, 1995) may conflate knowledge "in the head" with knowledge stored in other human sources, such as fellow members of a cockpit crew. While such effects may be possible, the rise of the Internet has surely broadened the scope of this effect. Before the Internet, there was no similarly massive, external knowledge database. People relied on less immediate and accessible inanimate stores of external knowledge, such as books—or, they relied on other minds in transactive memory systems. In contrast with other sources and cognitive tools for informational access, the Internet is nearly always accessible, can be searched efficiently, and provides immediate feedback. For these reasons, the Internet might become even more easily integrated with the human mind than other external sources of knowledge and perhaps even more so than human transactive memory partners, promoting much stronger illusions of knowledge.

Recent evidence suggests similar illusions occur when users search for fact-based information online (Ward, 2013b). After using Google to retrieve answers to questions, people seem to believe they came up with these answers on their own; they show an increase in "cognitive self-esteem," a measure of confidence in one's own ability to think about and remember information, and predict higher performance on a subsequent trivia quiz to be taken without access to the Internet. These fact-based effects are dependent on the reliability and the familiarity of the search engine, suggesting the processes by which the Internet affects cognition function differently across types of knowledge. These differences across informational contexts highlight the need for further research on the effects of different forms of online information search.

Confusion of accessible knowledge with one's personal knowledge may not be a gradual result of cultural immersion. Instead, it may be an early emerging tendency that remains even as children learn to access to the division of cognitive labor in the world around them. Tasks like learning the meanings of new words may be facilitated by a tendency for children to believe that they "knew it all along" (Kominsky & Keil, 2014; Mills & Keil, 2004; Taylor, Esbensen, & Bennett, 1994). Such misattributions may endow children with an adaptive confidence that their understandings are well grounded. The Internet may exaggerate this bias even in adults, leading to failures in recognizing the limits of internal explanatory knowledge.

The participants in our studies completed the experiments online and presumably use Internet search engines frequently. Why might we still observe an effect if the participants are already closely connected with the Internet as a transactive memory partner? It may be that chronic (overall frequency of use) and recent (experimental induction) search both influence knowledge self-assessments. In the area of social priming, similar effects have been found. Chronic and recent influences combine additively, so experimental exposures can be an effective way of mimicking chronic states (Bargh, Bond, Lombardi, & Tota, 1986; Higgins & Bargh, 1987). In political psychology, for example, where self-interest is assumed to drive political and economic choices, when participants are primed with self-interest its influence is even stronger (Young, Thomsen, Borgida, Sullivan, & Aldrich, 1991). In the case of the Internet, frequent use may boost ratings of self-assessed knowledge, but in-the-moment online search independently increases ratings as well.

There are clearly benefits to the freely accessible information on the Internet; however, there may be costs inherent to the strategy of accessing that information. The boundary between personal and interpersonal knowledge is becoming increasingly blurred (Clark & Chalmers, 1998). As technology makes information ever more easily available and accessible through searching, the ability to assess one's internal "unplugged" knowledge will only become more difficult. Erroneously situating external knowledge within their own heads, people may unwittingly exaggerate how much intellectual work they can do in situations where they are truly on their own.

References

- Alicke, M. D., Klotz, M. L., Breitenbecher, D. L., Yurak, T. J., & Vredenburg, D. S. (1995). Personal contact, individuation, and the better-than-average effect. *Journal of Personality and Social Psychology, 68*, 804–825. <http://dx.doi.org/10.1037/0022-3514.68.5.804>
- Alter, A. L., Oppenheimer, D. M., & Zemla, J. C. (2010). Missing the trees for the forest: A construal level account of the illusion of explanatory depth. *Journal of Personality and Social Psychology, 99*, 436–451. <http://dx.doi.org/10.1037/a0020218>
- Bargh, J. A., Bond, R. N., Lombardi, W. J., & Tota, M. E. (1986). The additive nature of chronic and temporary sources of construct accessibility. *Journal of Personality and Social Psychology, 50*, 869–878. <http://dx.doi.org/10.1037/0022-3514.50.5.869>
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis, 58*, 7–19. <http://dx.doi.org/10.1093/analys/58.1.7>
- Danovitch, J. H., & Keil, F. C. (2004). Should you ask a fisherman or a biologist? Developmental shifts in ways of clustering knowledge. *Child Development, 75*, 918–931. <http://dx.doi.org/10.1111/j.1467-8624.2004.00714.x>
- Dunning, D. (2005). *Self-insight: Roadblocks and detours on the path to knowing thyself*. New York, NY: Psychology Press. <http://dx.doi.org/10.4324/9780203337998>
- Fernbach, P. M., Rogers, T., Fox, C. R., & Sloman, S. A. (2013). Political extremism is supported by an illusion of understanding. *Psychological Science, 24*, 939–946. <http://dx.doi.org/10.1177/0956797612464058>
- File, T. (2013). Computer and Internet use in the United States. *Population Characteristics, 1*–14.
- Fisher, M., & Keil, F. C. (2014). The illusion of argument justification. *Journal of Experimental Psychology: General, 143*, 425–433. <http://dx.doi.org/10.1037/a0032234>
- Gelman, S. A. (2009). Learning from others: Children's construction of concepts. *Annual Review of Psychology, 60*, 115–140. <http://dx.doi.org/10.1146/annurev.psych.59.103006.093659>
- Harris, J. E. (1978). External memory aids. In M. M. Gruneberg, P. E. Morris, & R. N. Sykes (Eds.), *Practical aspects of memory* (pp. 172–180). London, England: Academic Press.
- Harris, P. L., Pasquini, E. S., Duke, S., Asscher, J. J., & Pons, F. (2006). Germs and angels: The role of testimony in young children's ontology. *Developmental Science, 9*, 76–96. <http://dx.doi.org/10.1111/j.1467-7687.2005.00465.x>
- Hart, J. T. (1965). Memory and the feeling-of-knowing experience. *Journal of Educational Psychology, 56*, 208.
- Higgins, E. T., & Bargh, J. A. (1987). Social cognition and social perception. *Annual Review of Psychology, 38*, 369–425. <http://dx.doi.org/10.1146/annurev.ps.38.020187.002101>
- Hill, G. W. (1982). Group versus individual performance: Are $N + 1$ heads better than one? *Psychological Bulletin, 91*, 517–539. <http://dx.doi.org/10.1037/0033-2909.91.3.517>
- Hollingshead, A. B. (1998). Communication, learning, and retrieval in transactive memory systems. *Journal of Experimental Social Psychology, 34*, 423–442. <http://dx.doi.org/10.1006/jesp.1998.1358>
- Hollingshead, A. B. (2001). Cognitive interdependence and convergent expectations in transactive memory. *Journal of Personality and Social Psychology, 81*, 1080–1089. <http://dx.doi.org/10.1037/0022-3514.81.6.1080>
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Johnson, M. K. (1997). Source monitoring and memory distortion. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences, 352*, 1733–1745. <http://dx.doi.org/10.1098/rstb.1997.0156>
- Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114*, 3–28. <http://dx.doi.org/10.1037/0033-2909.114.1.3>
- Keil, F. C., Stein, C., Webb, L., Billings, V. D., & Rozenblit, L. (2008). Discerning the division of cognitive labor: An emerging understanding of how knowledge is clustered in other minds. *Cognitive Science, 32*, 259–300. <http://dx.doi.org/10.1080/03640210701863339>
- Kominsky, J. F., & Keil, F. C. (2014). Overestimation of knowledge about word meanings: The "misplaced meaning" effect. *Cognitive Science, 38*, 1604–1633. <http://dx.doi.org/10.1111/cogs.12122>
- Koriat, A., & Levy-Sadot, R. (2001). The combined contributions of the cue-familiarity and accessibility heuristics to feelings of knowing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 34–53. <http://dx.doi.org/10.1037/0278-7393.27.1.34>
- Kozlowski, S. W. J., & Ilgen, D. R. (2006). Enhancing the effectiveness of work groups and teams. *Psychological Science in the Public Interest, 7*, 77–124.
- Littlepage, G., Robison, W., & Reddington, K. (1997). Effects of task experience and group experience on group performance, member ability, and recognition of expertise. *Organizational Behavior and Human Decision Processes, 69*, 133–147. <http://dx.doi.org/10.1006/obhd.1997.2677>
- Mills, C. M., & Keil, F. C. (2004). Knowing the limits of one's understanding: The development of an awareness of an illusion of explanatory depth. *Journal of Experimental Child Psychology, 87*, 1–32. <http://dx.doi.org/10.1016/j.jecp.2003.09.003>
- Peltokorpi, V. (2008). Transactive memory systems. *Review of General Psychology, 12*, 378–394. <http://dx.doi.org/10.1037/1089-2680.12.4.378>
- Pronin, E. (2009). The introspection illusion. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (pp. 1–67). Burlington, VT: Academic Press.
- Rand, D. G. (2012). The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. *Journal of Theoretical Biology, 299*, 172–179. <http://dx.doi.org/10.1016/j.jtbi.2011.03.004>
- Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science, 26*, 521–562. http://dx.doi.org/10.1207/s15516709cog2605_1
- Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *Science, 333*, 776–778. <http://dx.doi.org/10.1126/science.1207745>
- Taylor, M., Esbensen, B. M., & Bennett, R. T. (1994). Children's understanding of knowledge acquisition: The tendency for children to report that they have always known what they have just learned. *Child Development, 65*, 1581–1604. <http://dx.doi.org/10.2307/1131282>
- Ward, A. F. (2013a). Supernormal: How the Internet is changing our memories and our minds. *Psychological Inquiry, 24*, 341–348. <http://dx.doi.org/10.1080/1047840X.2013.850148>
- Ward, A. F. (2013b). *One with the Cloud: Why people mistake the Internet's knowledge for their own* (Unpublished doctoral dissertation). Cambridge, MA: Harvard University.
- Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of group behavior* (pp. 185–208). New York, NY: Springer-Verlag. http://dx.doi.org/10.1007/978-1-4612-4634-3_9
- Wegner, D. M., Erber, R., & Raymond, P. (1991). Transactive memory in close relationships. *Journal of Personality and Social Psychology, 61*, 923–929. <http://dx.doi.org/10.1037/0022-3514.61.6.923>

Wegner, D. M., Giuliano, T., & Hertel, P. T. (1985). Cognitive interdependence in close relationships. In W. J. Ickes (Ed.), *Compatible and incompatible relationships* (pp. 253–276). New York: Springer-Verlag. http://dx.doi.org/10.1007/978-1-4612-5044-9_12

Young, J., Thomsen, C. J., Borgida, E., Sullivan, J. L., & Aldrich, J. H. (1991). When self-interest makes a difference: The role of construct accessibility in political reasoning. *Journal of Experimental Social Psychology*, 27, 271–296. [http://dx.doi.org/10.1016/0022-1031\(91\)90016-Y](http://dx.doi.org/10.1016/0022-1031(91)90016-Y)

Appendix A

Questions Used in the Induction Phase

Why are there leap years?

Why are there more women than men?

Why are there phases of the moon?

Why are there time zones?

How does a zipper work?

Why are there dimples on a golf ball?

Why are there jokers in a deck of cards?

How is glass made?

Appendix B

Induction Phase Instructions From Experiment 1a

Internet condition:

We are interested in how well people can explain the answers to common questions. Please search the Internet to confirm the details of the explanation, and then evaluate. Please copy and paste the URL of the most helpful website in the space provided.

No Internet condition:

In this task, you will be asked a series of questions. We are interested in how well people can explain the answers to common questions. Please evaluate your understanding, using no outside sources.

Appendix C

Topics and Questions Used as the Dependent Measure in Experiments 1–2, 4–5a-b

Weather

Consider the following questions about weather:

1. Why are there more Atlantic hurricanes in August and September?
2. How do tornadoes form?
3. Why are cloudy nights warmer?

Science

Consider the following questions about science:

1. How do scientists determine the dates of fossils?
2. How do scientists know that the universe is expanding?
3. Why can't x-rays penetrate lead?

American History

Consider the following questions about American history:

1. Why did the Civil War begin?
2. How were the first labor unions formed?
3. Why did Nixon resign?

Food

Consider the following questions about food:

1. What is gluten?
2. Why does Swiss cheese have holes?
3. How is vinegar made?

(Appendices continue)

Anatomy and Physiology

Consider the following questions about anatomy and physiology:

1. Why do people laugh?
2. How does the heart pump blood?
3. Why do men go bald?

Health Issues

Consider the following questions about health issues:

1. Why are so many people allergic to peanuts?
2. Why can't HIV be transmitted through saliva?
3. Why can't you drink on antibiotics?

Appendix D**Topics and Easier Questions Used as the Dependent Measure in Experiment 3****Personal History**

Consider the following questions about your personal history:

1. How did you choose your current career?
2. Why did you choose to live where you currently live?
3. How did you decide what to study during high school?

Relationships

Consider the following questions about relationships:

1. Why are you close with your best friend?
2. How are you similar to your mother?
3. How could you become friendlier with your next door neighbor?

Local Culture

Consider the following questions about the local culture where you live:

1. How does the way people dress in your town reflect their socioeconomic status?
2. How is your town different from other parts of the country?

3. How do the restaurants near where you live reflect your state's culture?

Personal Habits

Consider the following questions about personal habits:

1. How do you choose what music to listen to?
2. How do you decide what to wear on important days?
3. How do you decide what to do on the weekend?

Future

Consider the following questions about the future:

1. How will you feel when you become elderly?
2. How will you try to succeed next week?
3. How will your life satisfaction be one year from now?

Emotions

Consider the following questions about emotions:

1. Why do you become annoyed by some things that don't seem to bother others?
2. Why do you become frustrated?
3. What causes you to feel most alive?

(Appendices continue)

Appendix E

Topics and Difficult Questions Used as the Dependent Measure in Experiment 3

Personal History

1. What is the relationship between the classes you chose during freshman year of high school and your current career?
2. How did the number of windows in your current living space influence your feelings of social connectedness after you moved in?
3. How did your learning style in your high school freshman year math class affect your later interest in miniature golf?

Relationships

1. How does your best friend influence your protein intake?
2. What are the origins of the difference in the degree to which you and your mother enjoy the genre of 'Mystery'?
3. How could you discover enough about your next door neighbor's sense of humor enough to reliably predict when he or she will laugh?

Local Culture

1. How does the menu organization at individually owned restaurants in your town compare with the menu organization at chain restaurants in your town?
2. How is your town's or county's governing body different from where your relatives live?

3. In your area, how are people's hairstyles correlated with their religious beliefs?

Personal Habits

1. How do songs in the key of D affect your mood the next day?
2. How do car advertisements affect the clothes you wear on formal occasions?
3. How does the way you make weekend plans reflect the way your father made weekend plans as a child?

Future

1. In what ways will being elderly be similar to the time times of physical discomfort you have already experienced?
2. How will the number of phone calls you make at your job affect the ways in which you try to succeed next week?
3. One year from now, how will your attention to detail affect your life satisfaction?

Emotions

1. How does being annoyed affect how likely you are to attend a sporting event next year?
2. How are your current feelings of frustration related to your first memory?
3. How do lunar patterns affect your emotional well-being?

Appendix F

Questions Used in the Induction Phase of Experiment 4b

Questions with answers online:

Why is ancient Egyptian history more peaceful than Mesopotamian history?

How does the location of Cameroon affect the health of its inhabitants? How do mountains affect the weather? How did the Erie Canal affect New York City?

Questions without answers online:

Why is ancient Kushite history more peaceful than Greek history?

How does the location of Pierre, South Dakota, affect the health of its inhabitants?

How do wheat fields affect the weather?

How did the Erie Canal affect Tioga County?

Received October 10, 2014

Revision received January 12, 2015

Accepted February 20, 2015 ■