

How Warnings about False Claims Become Recommendations

IAN SKURNIK
CAROLYN YOON
DENISE C. PARK
NORBERT SCHWARZ*

Telling people that a consumer claim is false can make them misremember it as true. In two experiments, older adults were especially susceptible to this "illusion of truth" effect. Repeatedly identifying a claim as false helped older adults remember it as false in the short term but paradoxically made them more likely to remember it as true after a 3 day delay. This unintended effect of repetition comes from increased familiarity with the claim itself but decreased recollection of the claim's original context. Findings provide insight into susceptibility over time to memory distortions and exploitation via repetition of claims in media and advertising.

In everyday life, people are bombarded with consumer information, from a wide variety of sources that differ in credibility (such as news reports, advertisements, Web sites, and so on) and in diverse consumption domains (from packaged goods to medical care). A crucial task for people is to determine whether the information they see is true or false. And because people often act on consumer information long after encoding it (e.g., Alba, Hutchinson, and Lynch 1991), their memory for the truth of information should be as accurate as possible.

We argue that when people try to determine the truth of a remembered consumer claim, memory for the original context of the claim can be as important as memory for the claim itself. For example, suppose the claim "shark cartilage will help your arthritis" feels familiar to consumers because they have encountered it recently. They might trust it less if they remember reading it in a tabloid headline than if they remember hearing it as advice from their physician. A weakness of this strategy is that memory for prior exposure to a claim is often much better than, and can be wholly independent of, memory for the context in which the claim

appeared (for reviews, see Johnson, Hashtroudi, and Lindsay 1993; Mandler 1980). And, when people find a claim familiar because of prior exposure but do not recall the original context or source of the claim, they tend to think that the claim is true (e.g., Hasher, Goldstein, and Toppino 1977; Hawkins and Hoch 1992).

Against this background, we document in two experiments some paradoxical effects of warnings about false information for older adults. In the first study, when older adults were repeatedly told that a claim was false, repetition helped them remember the claim immediately thereafter as false. But paradoxically, after 3 days had passed, the more times older adults had been warned that a claim was false, the more likely they were to misremember the claim as true. In the second study, trying to discredit claims after making them familiar to older adults backfired and increased their tendency to call those claims true.

We assume that repeated warnings about a false claim strengthen a feeling of familiarity for the claim and improve memory for the truth-specifying context of its presentation after a short delay. However, detailed memory for the warning fades more quickly for older than for younger adults. Once the needed contextual details are lost, the remaining feeling of familiarity fosters the paradoxical acceptance of false claims as true, rendering older adults particularly susceptible to this bias. Similarly, repeating a claim without identifying its truth strengthens familiarity without enhancing memory for any truth-specifying context; hence eventual attempts to discredit the familiar claim can backfire. Taken together, these findings favor a theoretical explanation of the truth bias in terms of constructive memory inferences

*Ian Skurnik is assistant professor of marketing, University of Toronto, Toronto, ON, Canada M5S3E6 (iskurnik@rotman.utoronto.ca). Carolyn Yoon is assistant professor of marketing, University of Michigan, Ann Arbor, MI 48109-1234 (yoonc@umich.edu). Denise C. Park is professor of psychology, University of Illinois, Champaign, IL 61801-9381 (denisep@uiuc.edu). Norbert Schwarz is professor of psychology, University of Michigan, Ann Arbor, MI 48106-1248 (nschwarz@umich.edu). The authors wish to acknowledge the helpful input from the editor, associate editor, and reviewers. The authors also thank Fred Feinberg and Scott Hawkins for comments on an earlier version of the article. This work was supported by National Institute on Aging (NIA) grant F32 AG05829-01 and Seed grant AG11715.

at the time of retrieval rather than truth-biased encoding at the time of exposure.

MEMORY FOR TRUTH

Numerous studies have demonstrated that exposing people to claims increases the perceived truth of the claim when it is seen again later (e.g., Hasher et al. 1977; Hawkins and Hoch 1992). This effect occurs even for statements that are explicitly identified as false on initial presentation (Begg, Anas, and Farinacci 1992; Gilbert, Krull, and Malone 1990). One interpretation of this effect is that it comes from a constructive inference that people make when they have little information with which to judge the truth of a claim other than the realization that they have seen it before. Without supporting information to help determine truth (along the lines of memory for truth-specifying contextual details such as the source of communication or expertise and background knowledge related to the claim), people often judge a claim to be true on the basis of some partial information in memory, such as the claim's subjective familiarity. As a general strategy, this type of constructive inference is appropriate if most information that people encounter or remember is actually true (Grice 1989; Hasher et al. 1977; Skurnik, Schwarz, and Winkielman 2000), but it creates a memory-based illusion of truth for familiar information that is actually false.

In accordance with this interpretation of the illusion of truth effects, a large body of research in recognition memory shows that familiarity from prior exposure remains intact while memory for the context of presentation is selectively impaired by a number of variables (for reviews, see Johnson et al. 1993; Mandler 1980). Generally speaking, context memory declines with conditions and tasks that prevent controlled thinking and deliberation, while familiarity is relatively unaffected by such conditions. For example, dividing people's attention while they encode or retrieve information can severely weaken their recollection of contextual details but spares their later sense of familiarity for the information.

AGE, REPETITION, AND DELAY

This reasoning suggests that we may become increasingly susceptible to the illusion of truth as we age. Normal human aging is associated with declines in memory for context or source, whereas experienced familiarity is largely unaffected by age (for a review, see Spencer and Raz 1995). Various factors ranging from neurobiological (Schacter, Norman, and Koutstaal 1998) to social/developmental (Rahhal, May, and Hasher 2002) seem to underlie these age differences. We expect that older adults are less likely than younger adults to recall whether consumer information was acquired from a trustworthy or untrustworthy source. Law, Hawkins, and Craik (1998) and Mutter, Lindsey, and Pliske (1995) found that prior presentation can enhance truth judgments more for older than for younger adults, suggesting that older adults may rely on experienced familiarity more than their

younger counterparts do. In both of these studies, claims were not explicitly identified as true or false during encoding, so the role of context memory in creating the illusion of truth could not be directly assessed.

Helping older adults to encode the truth-specifying context of claims should help them avoid remembering false information as true. One of the simplest ways to enhance memory is through repetition. Repeatedly informing people that a claim is false should enhance accurate memory of the information's context as well as its subjective familiarity and help older adults remember it as false.

Our emphasis on recognition memory, however, suggests a paradoxical effect of repetition. Specifically, one cause of the dissociation of context memory and familiarity is the passage of time, with context memory fading from memory much more quickly than subjective familiarity. This relatively rapid loss of contextual cues from memory is accelerated in older adults; after a long delay, older adults are more likely than younger adults to have experienced a decline in context memory but not in familiarity. Therefore we suggest a different pattern of results among older and younger adults after a short versus long delay in memory accuracy for repeatedly presented false claims. We predict that, after a short delay, repetition will help both younger and older adults remember false claims as false. However, after a longer delay, older adults may still find the repeated claims highly familiar, but details of the context may have faded from memory. After a longer delay, repeated identification of false claims should backfire and increase the chances that older adults will misremember those claims as true.

This predicted pattern is consistent with some past findings in recognition-based research where repetition decreases false alarms for younger adults but increases the same false alarms for older adults. Jacoby (1999) studied memory for the presentation modality of individual words. He found that increasing the number of presentations of words improved younger adults' memory for whether they originally read or heard each word but decreased older adults' memory for this information. Bartlett, Strater, and Fulton (1991) studied a "false fame" effect (Jacoby, Woloshyn, and Kelley 1989) in face recognition. Repeated exposure to nonfamous faces improved younger adults' ability to distinguish these faces from a set of famous faces but led to a pronounced false fame effect for older adults. Our research will add to these studies by demonstrating cognate effects in memory for truth, therein strengthening our explanation of illusion of truth effects as constructive memory inferences.

Our research will also help answer a question raised by the Bartlett et al. (1991) and Jacoby (1999) studies: Does repetition have opposite effects on younger and older adults' memory because older adults fail to encode the contextual information or because they lose the information more rapidly once it has been encoded? Our predicted results suggest that the deficit lies in loss of information after it has been encoded. If repeated presentation of false claims aids older

adults' memory for truth in the short term, then they have encoded the repeated claims and their context more thoroughly. If the effect of repetition reverses after a long delay, then the overall pattern can be explained only by loss of contextual information after it was encoded. On the other hand, if repetition strengthens familiarity for older adults without helping them encode contextual information, then repeated false claims should be misremembered as true more after both short and long delays.

The effects of age, repetition, and delay also bear on a different encoding-based alternative explanation for illusion of truth findings. According to the "Spinozan processor" explanation (Gilbert et al. 1990), people automatically encode all new information as true, if only for a moment. People can change this default representation to false only by engaging a subsequent effort-dependent mechanism through which they attach a false "tag" to the representation. According to this explanation, the illusion of truth arises when people fail to attach false tags to information, and the information is mistakenly stored in memory in its default true form. Manipulations such as dividing attention during encoding create the illusion of truth by disabling the false-tagging mechanism.

The Spinozan account and the constructive memory explanation make many of the same predictions. But the two explanations make different predictions about the interaction of age, repetition, and delay. The Spinozan account predicts a main effect of repetition over time. Repeatedly identifying a claim as false should increase the chances that a false tag will be attached to the claim's representation, which will decrease the chances of misremembering the claim as true (Gilbert et al. 1990). Hence false claims that are presented several times should be remembered more accurately after both short and long delays, compared to false claims presented once. The Spinozan model could account for age differences by adding an assumption that older adults are less able to tag false information as false, but this addition would not change the predicted effect of repetition across delay. In contrast, the constructive memory account predicts an interaction of delay and repetition for older adults. Repetition should improve both context memory and familiarity for false claims, but context memory should fade over a longer delay, making these claims more likely to seem true because of their high familiarity.

EXPERIMENT 1

The primary goal of this experiment is to test the predictions about how age and delay interact with repetition to lead to a greater likelihood of misremembering false statements as true. We exposed younger and older adults once or three times to claims that were explicitly labeled "false" or "true." We then asked participants to remember the truth value of the claims after a 30 min. filler task or after a delay of 3 days.

Method

Older and younger adults participated by studying individual statements (e.g., "Aspirin destroys tooth enamel," "Corn chips contain twice as much fat as potato chips") that were immediately identified as "true" or "false." Half of these statements and their truth values were presented once, and half were presented three times. After either half an hour or 3 days, participants saw the list of statements again with new statements mixed in and indicated whether each statement was "true," "false," or "new." The experiment had a 2 (age: older or younger) \times 2 (delay: short or long) \times 2 (number of presentations at study: one or three) mixed factorial design; age and delay were between-subjects variables.

Participants. Thirty-two younger adults at a large university (ages 18–25, $M = 21.3$) and 32 community-dwelling older adults (ages 71–86, $M = 77.7$) participated in exchange for \$30. Participants were screened individually for serious medical disorders or impairment (e.g., cerebrovascular conditions, past head injuries, recent hospitalization). Consistent with prior research, older adults scored higher ($M = 36.1$) than did younger adults ($M = 31.6$) on the Shipley Vocabulary Test ($F(1, 60) = 16.16$, $p < .01$).

Materials. Past research has demonstrated enhanced truth ratings for a wide range of types of claims, including brand and product claims (Hawkins and Hoch 1992; Law et al. 1998), trivia statements (Begg et al. 1992; Hasher et al. 1977), political opinions (Arkes, Hackett, and Boehm 1989), and invented vocabulary (Gilbert et al. 1990). For our experiment, we chose to use claims about health and medicine. In a separate pretest of 15 younger and 15 older adults, the general topic was deemed of equivalent interest and relevance to the two age groups. In addition, beliefs about health care are a pragmatic concern for older adults, who comprise the largest consumer segment of medical services and products.

In another pretest, 31 younger and 15 older adults rated 150 health claims for perceived truth and familiarity. Because of our ethical concern about presenting objectively false health and medical information to participants, we used claims that were objectively true according to the Web sites of the National Institutes of Health and a large university hospital. We selected 54 statements that were not obviously true or false to both younger and older people. These claims were divided into six sublists that balanced truth and familiarity ratings across ages. Sublists were rotated and counterbalanced so that each claim had an equal chance of appearing as true or false at each level of repetition across participants.

Procedure. The procedure was divided into a study phase where participants encoded the claims and their truth value and a test phase where their memories for truth were tested. Depending on condition, the test phase took place 30 min. after the study phase or 3 days later during the

second session. During each phase, participants read instructions on the computer screen, while the experimenter simultaneously read them aloud. Participants were told that they would see statements about medicine and health collected from a variety of sources and that each statement would immediately be identified as true or false. Participants were also told that (1) some of the statements would appear once and some three times, (2) overall there would be an equal number of true and false statements, and (3) later they would see other statements and answer some health-related questions. Participants were then guided through practice trials.

During the study phase, each claim appeared on the screen for 5 sec., followed by a blank screen for 750 milliseconds (ms.), then the word "true" or "false" for 1,500 ms., and finally a blank screen for 1,500 ms. before the next claim appeared. Presentation order of claims was random with the exception that repetitions of claims were separated by at least six other claims. At the test phase, all 36 study phase and 18 new claims were presented individually. Participants were asked to indicate for each claim whether they thought it was "true," "false," or "new" (i.e., not on the study phase list) by pressing one of three marked keys on the keyboard. Participants were guided through several practice trials. There was no response deadline, and claims remained on the screen until participants pressed a valid response key. Following the test phase, participants completed a questionnaire, were debriefed, and were given a printed list of all the health/medical claims used in the experiment.

Results

Manipulation Check. There were no effects of the study phase list counterbalances on test phase responses ($F < 1$). The positions of the "true," "false," and "new" response keys were also varied, and there were no effects of response key placement on test phase responses ($F < 1$).

Response Type. Participants' responses to each claim were coded according to their answers ("true," "false," or "new") and the claims' real status (true, false, or new). Of primary interest for the illusion of truth effect are the likelihoods of saying "true" to an originally false claim and "false" to an originally true claim. When participants make these two types of responses, they correctly think that they have seen these claims before (otherwise they would call them "new") but incorrectly remember the original truth value. If participants guessed at truth randomly, these two responses would be equally likely. But if they tend to think that previously seen false claims are true, then participants should be more likely to call false statements "true" than vice versa.

Memory for Truth and Falsity. A 2 (age) \times 2 (delay) \times 2 (number of presentations) \times 2 (response type: "true" to false or "false" to true) mixed analysis of variance on the response proportions revealed a number of significant effects. All ANOVAs were conducted on arcsine transfor-

mations of the response proportions. All statistical tests (F -tests) were based on (1, 60) degrees of freedom. For easier comprehension, however, we report untransformed means in subsequent discussion of results.

For all four factors, significant main effects on response proportions were found. First, a main effect of age emerged, with older adults ($M = .21$) making more errors in remembering truth than younger adults ($M = .12$, $F = 13.66$, $p < .01$). In addition, repetition of claims and their truth values improved memory for truth status (mean error rates of .15 for claims presented three times and .18 for claims presented once; $F = 5.20$, $p < .05$). There was also a main effect of delay, with participants making more errors in remembering truth after a longer delay ($M = .21$ and .13 for long and short delays, respectively; $F = 11.19$, $p < .01$). Finally, participants were more likely to misremember false information as "true" ($M = .21$) than true information as "false" ($M = .12$, $F = 20.03$, $p < .01$).


The main effects were further qualified by several significant interactions: a four-way interaction ($F = 4.95$, $p < .05$); a three-way interaction of age, delay, and number of presentations ($F = 7.55$, $p < .01$); and two-way interactions of age and number of presentations ($F = 6.32$, $p < .05$) and age and response type ($F = 7.41$, $p < .01$). To interpret the highest order interaction, we present separate analyses for the condition in which responses were elicited after the short delay and after the long delay.

Short Delay. A 2 (age) \times 2 (number of presentations) \times 2 (response type) mixed ANOVA produced a main effect of age, showing that older adults made more errors than younger adults ($M = .17$ and .08, respectively; $F = 5.71$, $p < .05$). Reduced error rates were found with claims seen three times ($M = .09$) versus once ($M = .16$; $F = 9.87$, $p < .01$). In addition, a main effect of response type reflected an overall illusion of truth effect (i.e., more false claims were misremembered as "true" ($M = .16$) than true claims misremembered as "false" ($M = .09$; $F = 4.92$, $p < .05$). These effects were qualified by a significant three-way interaction ($F = 4.01$, $p < .05$). Figure 1 shows the mean responses after a short delay.

Decomposition of the interaction by age showed that repetition led to reduced error rates for both age groups ($F = 4.43$, $p < .05$ for younger adults; $F = 5.47$, $p < .05$ for older adults). Whereas no illusion of truth effect emerged for younger adults across the two repetition conditions ($F < 1$), older adults were more likely to respond "true" to originally false information than to respond "false" to originally true information, when claims were presented only once ($F = 8.25$, $p < .01$) but not when they were presented three times ($F = 1.69$, $p = .19$). Further, older adults' tendency to misremember false statements as "true" decreased with repetition ($M = .28$ and .17 for one and three presentations, respectively; $F = 7.00$, $p < .01$). Multiple presentations thus improved older adults' context memory in the short run.

Long Delay. Figure 2 shows the mean responses made

FIGURE 1
MEMORY FOR TRUTH AND FALSITY AFTER SHORT DELAY, BY AGE AND REPETITION



after a long delay. Analysis of these data yielded a main effect of age, with older adults ($M = .26$) more likely than younger adults ($M = .16$) to misremember the truth value of claims ($F = 8.05, p < .01$); a main effect of response type, showing an overall illusion of truth effect ($M = .26$ for false information as "true" and $M = .15$ for true information as "false"; $F = 16.90, p < .01$); and an interaction of age and number of presentations ($F = 13.84, p < .01$). These effects were qualified by a significant three-way interaction ($F = 4.29, p < .05$).

To clarify this interaction, we performed further analyses by age. Whereas younger adults had reduced error rates with repetition ($M = .20$ and $.12$ for one and three presentations, respectively; $F = 7.23, p < .01$), the opposite was true for older adults ($M = .22$ and $.30$ for one and three presentations, respectively; $F = 6.61, p < .05$). Further, although both younger and older adults exhibited an illusion of truth effect, the pattern of results for claims presented once or three times was quite different across the two age groups. For younger adults, once-presented statements yielded a marginally significant illusion of truth effect ($F = 2.80, p < .10$), and thrice-presented statements produced no illusion of truth effect ($F < 1$). Moreover, the tendency to misremember false statements as "true" decreased with repetition ($M = .24$ and $.14$ for one and three presentations, respectively; $F = 5.36, p < .05$). Thus repetition improved younger adults' context memory, decreasing the likelihood

that they would misremember false claims as "true" even after a delay of 3 days.


Older adults, however, exhibited an illusion of truth effect for both once-presented and thrice-presented statements. Older adults were more likely to respond "true" to originally false information than to respond "false" to originally true information when claims were presented only once ($F = 4.90, p < .05$), and the bias was particularly pronounced for claims presented three times ($F = 14.30, p < .01$).

Thus in contrast to the reduced tendency with repetition to misremember false statements as "true" in the short run, we found a "backfire" effect among older adults after a 3 day delay. Far from improving their context memory for the original information, after 3 days repeated exposure left older adults only with increased feelings of familiarity, resulting in the erroneous endorsement of 40% of the originally false information as "true." Moreover, the illusion of truth bias by the elderly was accounted for solely by an increase in the likelihood of responding "true" to false claims presented three times ($M = .40$) compared to those presented once ($M = .28; F = 4.90, p < .05$). There was no difference across number of presentations among older adults to call true claims "false" ($M = .19$ and $.16$ for three and one presentation, respectively; $F < 1$).

Old-New Memory. These biases in remembering truth were not mirrored in responses to new claims. After a short

FIGURE 2

MEMORY FOR TRUTH AND FALSITY AFTER LONG DELAY, BY AGE AND REPETITION



delay, older and younger adults were equally able to identify new claims as "new" ($M = .97$ and $.96$, respectively) and were equally likely to call new claims either "true" or "false" by mistake ($M = .02$ and $.01$, respectively, for older adults and $.01$ and $.03$, respectively, for younger adults). Older adults were also just as likely as younger adults to mistakenly call previously presented true and false claims "new" ($M = .04$ and $.04$, respectively, for older adults and $.03$ and $.04$, respectively, for younger adults; all F 's < 1). Repetition led only to an overall decrease in mistakenly calling true and false claims "new" under the short delay condition ($M = .06$ and $.02$ for one and three presentations, respectively; $F = 4.96$, $p < .05$).

After a delay of 3 days, older adults were less able than younger adults to distinguish repeated from new claims. Older adults ($M = .72$) were less likely than younger adults ($M = .93$) to identify new claims as "new" ($F = 7.34$, $p < .01$) and were more prone to call new claims "true" or "false" by mistake ($M = .14$ for older adults, and $M = .03$ for younger adults; $F = 11.88$, $p < .01$). Additionally, older adults ($M = .21$) were more apt than younger adults ($M = .11$) to mistakenly call previously presented true or false claims "new" ($F = 7.53$, $p < .01$). Despite the higher rates of recognition errors for older adults, there was no truth bias to these responses for either younger or older adults (so, for example, older adults were equally likely to call a claim "new" by mistake whether it was originally identified as true or false). Repetition decreased these mis-

takes for both age groups ($M = .22$ and $.10$ for one and three presentations, respectively; $F = 42.27$, $p < .01$).

Because older adults' old-new recognition was reduced in the long delay condition relative to the other between-subjects conditions, we repeated all analyses reported above using participants' responses as a proportion of items that they correctly identified as old (see Murname and Bayen 1996 for an extended discussion). Results of these analyses matched those reported above in direction and significance. Figures 1 and 2 present the means of raw response proportions, rather than the conditionalized measures, as a clear indication of the total amount of false information that participants thought was true.

Discussion

This experiment produced findings consistent with the constructive memory predictions. The key result highlights an extremely undesirable, and previously unidentified, side effect of warnings: the more often older adults were told that a claim was false, the more likely they were to remember it erroneously as true after a 3 day delay. The size of this effect is far from negligible. After 3 days, older adults misremembered 28% of false statements as true when they were told once that the statement was false but 40% when told three times that the statement was false. There was no parallel tendency to misremember true information as false. Older adults also showed the illusion of truth effect after a

short delay, but this effect was limited to statements presented once. It thus appears that repeated warnings about false information can help older adults in the short run to remember false information as false.

The interaction of repetition and delay for older adults strongly suggests that older adults did not fail to encode the truth-specifying context of false claims, because repetition improved their accuracy with false information after 30 min. Instead, older adults seem to have failed to retain contextual information in memory over the course of 3 days. Their experienced familiarity for repeated claims remained relatively intact over the longer delay. We analyze evidence for this dissociation further in the general discussion.

These findings are also difficult for the Spinozan processing model to explain. Specifically, making it more likely that people attach a false "tag" to their representation of false claims during encoding (by repeating the claims) cannot also make people less likely to have the tag attached after 3 days.

In contrast to the performance of older adults, younger adults' memories for truth benefited from repeated warnings after both short and long delays. Younger adults did not exhibit a bias toward truth except when there was a long delay following exposure to once-presented statements. After a 3 day delay, the younger adults' responses were similar to those of older adults after a short delay.

Moreover, for both age groups the bias to call claims true occurred only for statements whose subjective familiarity was presumably enhanced by prior presentation and not for new, unrepeated statements. Finally, repetition helped older adults distinguish old (familiar) statements from new statements even over the long run, a task for which they could rely on the familiarity of the repeated old statements. Hence, both findings—that old-new discrimination and erroneous acceptance of a false statement as true each increased with repetition—are consistent with older adults' increased reliance on feelings of familiarity in constructive judgment.

EXPERIMENT 2

Experiment 1 provided evidence that people misremember a false claim as true when they recognize the claim from prior exposure but have no contextual information in memory that would specify the claim as false. Increasing perceived familiarity through repetition of the claim and its falsifying context helps only when accompanied by increased context memory at the time of retrieval. Although older adults clearly were able to encode context in experiment 1, as indicated by their results after 30 min., their overall performance showed a deficit in comparison to younger adults. This deficit is consistent with prior research on age differences in memory showing that older adults have somewhat more trouble encoding contextual information than do younger adults (e.g., Law et al. 1998; Spencer and Raz 1995). If so, then minimizing the opportunity to encode a claim's context, while simultaneously maximizing the claim's familiarity, should produce the same

"backfire" effect of repetition for older adults even after a short delay.

For example, presenting a claim several times without providing definitive information about its truth would enhance the claim's familiarity. Presenting the claim again in order to warn people explicitly that it is actually false may be inadequate to overcome the effects of heightened familiarity for older adults. Hence for older adults it may be especially difficult to discredit claims that are already familiar. Consumers face this situation if they do not initially evaluate the truth of every product claim they see and do so only later when more complete information becomes available.

Consumers' possible exposure to health claims and rumors whose truth they cannot initially assess is a grave public policy concern. The U.S. Federal Trade Commission (FTC) recently found that more than half of 300 ads for products related to weight loss (a \$35 billion consumer market in the United States in 2000) made at least one claim that was blatantly false or unsubstantiated (Federal Trade Commission 2002). Similarly, Morris and Avorn (2003) found that over 65% of Web sites selling popular herbal supplements contained unsubstantiated claims about the benefits of their products. To combat this problem, the FTC can order disclosures or disclaimers (Jacoby, Nelson, and Hoyer 1982; Johar and Simmons 2000) or even corrective advertising (Mizerski, Allison, and Calvert 1980; cf. Johar 1996). Such measures can be effective, but research so far has been carried out on younger adults.

In experiment 2 we make claims familiar to older and younger adults through repetition and manipulate the point at which we disclose the actual truth of each claim. Some participants see claims twice without learning if the claims are true or false; they learn this information only on a third and final presentation. This single opportunity to encode truth-specifying context may not be adequate for older adults and may lead them to think that discredited but repeated claims are truer than discredited claims they saw only once. Other participants see the truth-specifying context every time they have seen the claims. We expect their results to mirror those of the short-delay condition in experiment 1.

Method

Experiment 2 had a 2 (age: older or younger adults) \times 2 (number of presentations at study: one or three) \times 2 (truth disclosure of claim at study: every presentation or last presentation only) mixed factorial design, with the second factor within subjects. The same stimuli, dependent variables, and basic procedure as in experiment 1 were used. One procedural difference was that all participants made truth judgments after a 20 min. filler task (i.e., there was no long delay, and this delay was shorter than the short delay in experiment 1). The other difference was the between-subjects manipulation of the timing of disclosure of claims' truth values during the study phase: on the last presentation of each claim or on every presentation. Participants in the former condition were instructed that some of the claims

would first appear with a set of question marks rather than with a true or false designation but that by the end of the list they would have seen the truth value of all the claims. Thus for claims presented three times this group viewed the claims twice before finding out if they were true or false, whereas for claims presented once, the truth value was always provided (hence this manipulation did not change the presentation of claims presented once). The test phase was the same as in experiment 1; participants saw the original claims, along with new claims, and indicated whether each claim was "true," "false," or "new."

Participants. Fifty-one younger adults (ages 18–27, $M = 22.6$) participated for course credit at a large university, and 31 community-dwelling older adults (ages 66–81, $M = 73.9$) participated in exchange for \$30. As in experiment 1, older adults did not have any serious medical disorders or indications of abnormal cognitive impairment. Older adults ($M = 35.6$) scored higher than younger adults ($M = 29.1$) on the Shipley Vocabulary Test ($F(1, 78) = 37.72, p < .01$).

Results

Manipulation Check. There were no effects of the study phase list counterbalances on test phase responses ($F < 1$). The placement of the response keys at the test phase was also varied, and there were no effects of key position on test phase responses ($F < 1$).

Memory for Truth and Falsity. A 2 (age) \times 2 (number of presentations) \times 2 (truth disclosure) \times 2 (response type: "true" to false or "false" to true) mixed ANOVA on the response proportions revealed a number of significant effects. As in experiment 1, all ANOVAs were conducted on arcsine transformations of the response proportions. All F -tests were based on (1, 78) degrees of freedom. In subsequent discussion of results, we report untransformed means.

First, a main effect of age was found, with older adults ($M = .21$) making more errors than younger adults ($M = .12$) in remembering truth ($F = 18.79, p < .01$). There was also a main effect of truth disclosure, indicating reduced error rates in remembering truth when truth value was identified on each presentation ($M = .13$) compared to when truth was disclosed only on the last presentation ($M = .19; F = 9.05, p < .01$).

These main effects were qualified by several significant interactions: a four-way interaction ($F = 4.50, p < .05$); a three-way interaction of age, truth disclosure, and response type ($F = 8.10, p < .01$); and two-way interactions of truth disclosure and response type ($F = 10.33, p < .01$) and number of presentations and truth disclosure ($F = 27.21, p < .01$). To interpret the four-way interaction, we conducted separate analyses for the condition in which truth value was identified every time and when truth value was revealed only on the last presentation (i.e., the condition in which

statements of some claims were familiar before their truth value was disclosed).

Truth Disclosed on Each Presentation. Figure 3 shows the mean proportions of error responses when the truth of claims was identified on every presentation. A 2 (age) \times 2 (number of presentations) \times 2 (response type) mixed ANOVA on these data showed that older adults ($M = .18$) made more errors than did younger adults ($M = .09; F = 7.83, p < .01$). Repetition of claims and their truth values, however, improved overall memory for their truth status (overall $F = 16.12, p < .01$; error rates of $M = .18$ and $.09$ for claims presented once and thrice, respectively). There were no other significant effects for these data.

As might be expected, these results are generally consistent with those found under the short delay condition in experiment 1. Older adults in these conditions, in both studies, were aided by repetition and did not exhibit an illusion of truth bias when claims were initially presented three times. For once-presented claims, however, an illusion of truth effect obtained among the elderly after a short delay in experiment 1 but not in experiment 2. A possible explanation for this difference may be the shorter delay interval of 20 min. between the study and test phases in experiment 2 compared to the 30 min. delay in experiment 1. Although older adults made roughly the same number of errors in these two conditions across the experiments, these errors did not demonstrate an illusion of truth bias until about 30 min. had passed.


Truth Disclosed on Last Presentation Only. A different pattern of results emerged when truth had not been identified until the last presentation. Although older adults in this condition also exhibited poorer memory for truth status compared to younger adults ($M = .24$ and $.15; F = 11.44, p < .01$), repetition of claims led to poorer, not better, memory for truth ($M = .16$ for claims presented once and $.23$ for claims presented three times; $F = 11.12, p < .01$). In addition, there was an overall illusion of truth effect: participants were more likely to misremember false information as "true" ($M = .23$) than true information as "false" ($M = .16; F = 12.23, p < .01$).

These main effects were qualified by a significant three-way interaction ($F = 7.53, p < .01$; see fig. 4 for all means). More detailed analyses of the illusion of truth effect by age yielded no significant effects for younger adults (F 's < 1). For older adults, however, there was a strong illusion of truth effect for claims presented three times ($M = .35$ for "true" to false, and $M = .18$ for "false" to true responses; $F = 21.30, p < .01$) and a marginally significant effect for claims presented once ($M = .24$ for "true" to false, and $M = .18$ for "false" to true responses; $F = 2.92, p < .10$).

Further, the illusion of truth effect found for repeated claims among the elderly was due solely to an increase in the likelihood of calling false statements "true" ($M = .35$ and $.24$ for three and one presentation, respectively; $F = 8.44, p < .01$). There was no difference across number of

FIGURE 3

MEMORY FOR TRUTH AND FALSITY WHEN TRUTH VALUE IS DISCLOSED ON EACH PRESENTATION, BY AGE AND REPETITION



presentations to call true statements "false" ($M = .18$ and $.18$; $F < 1$).

Old-New Memory. These biases in remembering truth by older adults were not reflected in responses to new statements. Participants were generally accurate during the test phase in indicating that new claims were not presented at the study phase ($M = .97$ and $.91$ for younger and older adults, respectively). Once in a while, however, participants mistakenly thought that a new claim had been presented at the study phase and reported it as "true" or "false." Younger adults ($M = .06$) made an equivalent number of these errors as older adults ($M = .09$; $F = 1.77$, $p = .18$), and there was no systematic pattern in these responses to account for the truth bias.

In some instances, participants mistakenly responded "new" to true and false claims that had been presented at the study phase. An ANOVA on these types of errors revealed a main effect of number of presentations, with claims presented three times ($M = .03$) less likely to be misidentified as "new" than those presented once ($M = .13$; $F = 65.49$, $p < .01$). This effect was qualified by an interaction of truth disclosure and number of presentations ($F = 10.25$, $p < .01$) reflecting a greater effect of repetition when truth was revealed only on last presentation. Specifically, multiple presentations drove these mistakes toward zero in both truth disclosure conditions. Errors decreased, from a mean of $.09$ for once-presented to $.03$ for thrice-presented

claims when truth was always disclosed ($F = 10.61$, $p < .01$) and from $.16$ to $.02$ when truth was disclosed only on last presentation ($F = 73.09$, $p < .01$). We found no other significant effects with respect to old-new discrimination.


Discussion

When truth value of claims was identified on every presentation, repetition helped both younger and older adults remember whether the claims were true or false. These results match the short-delay results in experiment 1. However, when the truth value of repeated claims was not disclosed until the last presentation, repeating the claim in order to call it "false" led to poorer memory for truth for both age groups. In addition, only older adults exhibited an illusion of truth bias.

These findings are consistent with the view that repetition enhances the subjective familiarity of a claim but without also improving recollection of the truth-specifying context for older adults. Discrediting a familiar claim made it more likely to seem true, compared to discrediting a claim that was not already familiar. Thus loss of context memory but intact familiarity for claims among older adults appeared to have made it more likely that they misremember discredited familiar information as true. These findings raise the concern that attempts to correct false and questionable claims may paradoxically lead to negative consequences for older adults.

FIGURE 4

MEMORY FOR TRUTH AND FALSITY WHEN TRUTH VALUE IS DISCLOSED ON LAST PRESENTATION ONLY, BY AGE AND REPETITION



GENERAL DISCUSSION

Our experiments document a paradoxical effect of warnings: the more often older adults were told that a given claim was false, the more likely they were to accept it as true after several days have passed (experiment 1). Similarly, warning older adults that a previously seen claim of unknown validity is actually false increases acceptance of this claim as true (compared to warning them about false claims that they have not read before; experiment 2). Both of these effects have the potential to put older adults at considerable risk, rendering it important that we understand the underlying processes and their applied implications.

Together, findings from the two studies suggest that people have multiple bases for making constructive judgments about the truth of remembered claims. One basis can be remembered contextual details, such as explicit "true" and "false" designations. A second basis can be partial information about prior exposure, such as experienced familiarity. When incomplete information such as familiarity is the only available cue for judging truth, people generally tend to infer that the information is true. As in prior research (Bartlett et al. 1991; Jacoby 1999), we suggest that repeating claims and truth-specifying contextual cues enhances recall of the information's context as well as its subjective familiarity. Hence repeated warnings about falseness have the desired effect on memory immediately (experiment 1). But

after contextual information has faded, leaving only enhanced familiarity, more warnings about the falsity of information are paradoxically more likely to make the information seem true. The same enhanced familiarity still helped our participants distinguish between new and old information, even as it biased them to think of the old information as true.

Older adults' results after a short delay show that they do encode claims and their context better with repetition. Therefore, the increased illusion of truth effects over time reflect a loss of context information from memory rather than from a failure to encode context in the first place. In addition, the pattern of findings obtained in experiment 1 is difficult for the Spinozan processing account to explain. The results after a short delay are consistent with predictions of the Spinozan model: increased exposure to false claims improves memory, by raising the probability that the claims will be "tagged" false as they are encoded. But the Spinozan model cannot explain the results at both delays: if repetition has increased the probability of tagging of false claims, as suggested by the short delay results, then those claims cannot also be less likely to have a tag, as suggested by the long delay results.

In experiment 2, some claims were repeated without their context, strengthening familiarity but not improving context memory. In this condition, older adults' memory for truth

after 20 min. was similar to older adults' performance in experiment 1 after 3 days. Thus, seeing a claim several times before one learns that it is false does not improve memory for its truth value. To the contrary, the repetition merely increases familiarity and hence reported memory of the false claim as true, thereby interfering with older adults' ability to make good decisions and choices. One interesting question for future research about the effects of repetition is whether a single warning about the falseness of familiar information is effective in counteracting the illusion of truth effect that comes from sheer repetition. Because of the categorical memory task at the test phase in our studies, we could not present people with claims at the study phase whose truth was never identified, so our data do not address this question.

Our findings have important public policy implications for protecting older consumers, who constitute the most vulnerable targets of scams. The FTC reports that a disproportionate number of consumer fraud victims, some 80% or more, are 65 or older (Federal Trade Commission 1999–2001). Our findings highlight the risks of merely identifying a given claim as unsubstantiated or false, a communication strategy commonly employed by regulators. Attempts to update people's beliefs about outdated information, fight rumors, and discredit misleading advertisements may have the unintended effect of increasing the familiarity of a false claim, rendering it more likely that older adults accept it as true later on, after the details of the phrasing have faded from memory. Statements such as "It is not true that X is good for your arthritis" always involve repetition of the core claim, thus increasing the familiarity of the claim when it is considered again in a different context. Repetition increases overall memory accuracy in the short term as long as the truth value accompanies a claim each time, but repetition can backfire over time.

It may seem reasonable to counsel against relying on familiarity when judging truth, in order to avoid the illusion of truth effect. But the reason people infer the truth of claims from partial information such as subjective familiarity is probably that familiarity is generally a valid cue to truth. Ignoring this cue altogether, even if it were possible to do so, could deprive people of a broadly useful heuristic. This inference is consistent with the tacit assumptions underlying the conduct of conversations in daily life (Grice 1989), where communicated information comes with a tacit "guarantee of relevance" (Sperber and Wilson 1986). In support of this reasoning, Skurnik, Moskowitz, and Johnson (2005) found that the illusion of truth effect reversed altogether when people thought that past claims were likely to be false.

A more useful strategy may be to provide people, and older adults in particular, with environmental support, such as written materials (e.g., Park and Shaw 1992) or visual imagery (e.g., Law et al. 1998), to supplement or improve memory. For older consumers, additional research should explore situations in which they are less susceptible to the bias, especially when recollection of context is needed to establish truth. One possible direction for future research is

to generate insights about helping older adults encode and retain the original context. For example, syntactically embedding a warning about falseness in a claim or providing meaningful contextual cues may increase the likelihood among the elderly that truth-specifying cues are retrieved as part of the core claims themselves.

The illusion of truth bias might be attenuated for claims that were worded "X is bad for your arthritis." Unfortunately, such statements would expose many information campaigns to high legal risks. In addition, scientific research that would support such claims is likely to couch its conclusions in the language of significance testing and avoid asserting that a null effect has been demonstrated conclusively. In many cases, it is merely documented that no positive effects of X have yet been found, giving laypeople the impression that definitive research has not been done. Because trying to discredit false information often involves increasing recipients' familiarity with its core components, our results suggest that information campaigns should focus more on what is true than on reiterating what is false.

[Dawn Iacobucci served as editor and Gita V. Johar served as associate editor for this article.]

REFERENCES

- Alba, Joseph W., J. Wesley Hutchinson, and John G. Lynch Jr. (1991), "Memory and Decision Making," in *Handbook of Consumer Behavior*, ed. Thomas S. Robertson and Harold K.assarjian, Englewood Cliffs, NJ: Prentice Hall, 1–49.
- Arkes, Hal R., Catherine Hackett, and Larry Boehm (1989), "The Generality of the Relation between Familiarity and Judged Validity," *Journal of Behavioral Decision Making*, 2 (April–June), 81–94.
- Bartlett, James C., Laura Strater, and Annette Fulton (1991), "False Recency and False Fame of Faces in Young Adulthood and Old Age," *Memory and Cognition*, 19 (March), 177–88.
- Begg, Ian Maynard, Ann Anas, and Suzanne Farinacci (1992), "Dissociation of Processes in Belief: Source Recollection, Statement Familiarity, and the Illusion of Truth," *Journal of Experimental Psychology: General*, 121 (December), 446–58.
- Federal Trade Commission (1999–2001), "Activities Affecting Older Americans," staff summary, <http://www3.ftc.gov/os/2001/10/agingrpt.htm>.
- (2002), "Weight Loss Advertising: An Analysis of Current Trends," staff report, <http://www.ftc.gov/opa/2002/09/weightlossrpt.htm>.
- Gilbert, Daniel T., Douglas S. Krull, and Patrick S. Malone (1990), "Unbelieving the Unbelievable: Some Problems in the Rejection of False Information," *Journal of Personality and Social Psychology*, 59 (April), 601–13.
- Grice, H. Paul (1989), *Studies in the Way of Words*, Cambridge, MA: Harvard University Press.
- Hasher, Lynn, David Goldstein, and Thomas Toppino (1977), "Frequency and the Conference of Referential Validity," *Journal of Verbal Learning and Verbal Behavior*, 16 (February), 107–12.
- Hawkins, Scott A. and Stephen J. Hoch (1992), "Low-Involvement Learning: Memory without Evaluation," *Journal of Consumer Research*, 19 (September), 212–25.
- Jacoby, Jacob, Margaret C. Nelson, and Wayne D. Hoyer (1982),

- "Corrective Advertising and Affirmative Disclosure Statements: Their Potential for Confusing and Misleading the Consumer," *Journal of Marketing*, 19 (Winter), 61-72.
- Jacoby, Larry L. (1999), "Ironic Effects of Repetition: Measuring Age-Related Differences in Memory," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25 (January), 3-22.
- Jacoby, Larry L., Vera Woloshyn, and Colleen Kelley (1989), "Becoming Famous without Being Recognized: Unconscious Influences of Memory Produced by Dividing Attention," *Journal of Experimental Psychology: General*, 118 (June), 115-25.
- Johar, Gita Venkataramani (1996), "Intended and Unintended Effects of Corrective Advertising on Beliefs and Evaluations: An Exploratory Analysis," *Journal of Consumer Psychology*, 5 (3), 209-30.
- Johar, Gita Venkataramani and Carolyn J. Simmons (2000), "The Use of Concurrent Disclosures to Correct Invalid Inferences," *Journal of Consumer Research*, 26 (March), 307-22.
- Johnson, Marcia K., Shahin Hashtroudi, and Stephen D. Lindsay (1993), "Source Monitoring," *Psychological Bulletin*, 114 (July), 3-28.
- Law, Sharmistha, Scott A. Hawkins, and Fergus I. M. Craik (1998), "Repetition-Induced Belief in the Elderly: Rehabilitating Age-Related Memory Deficits," *Journal of Consumer Research*, 25 (September), 91-107.
- Mandler, George (1980), "Recognizing: The Judgment of Previous Occurrence," *Psychological Review*, 87 (May), 252-71.
- Mizerski, Richard W., Neil K. Allison, and Stephen Calvert (1980), "A Controlled Field Study of Corrective Advertising Using Multiple Exposures and a Commercial Medium," *Journal of Marketing Research*, 17 (August), 341-48.
- Morris, Charles A. and Jerry Avorn (2003), "Internet Marketing of Herbal Products," *Journal of the American Marketing Association*, 290 (September), 1505-9.
- Murnane, Kevin and Ute J. Bayen (1996), "An Evaluation of Empirical Measures of Source Identification," *Memory and Cognition*, 24 (July), 417-28.
- Mutter, Sharon A., Sue E. Lindsey, and Rebecca M. Pliske (1995), "Aging and Credibility Judgment," *Aging and Cognition*, 2 (2), 89-107.
- Park, Denise C. and Raymond J. Shaw (1992), "Effect of Environmental Support on Implicit and Explicit Memory in Younger and Older Adults," *Psychology and Aging*, 7 (December), 632-42.
- Rahhal, Tamara A., Cynthia P. May, and Lynn Hasher (2002), "Truth and Character: Sources That Older Adults Can Remember," *Psychological Science*, 13 (March), 101-5.
- Schacter, Daniel L., Kenneth A. Norman, and Wilma Koutstaal (1998), "The Cognitive Neuroscience of Constructive Memory," *Annual Review of Psychology*, 49 (February), 289-318.
- Skurnik, Ian, Gordon B. Moskowitz, and Marcia K. Johnson (2005), "Biases in Remembering True and False Information: Illusions of Truth and Falseness," unpublished manuscript, University of Toronto.
- Skurnik, Ian, Norbert Schwarz, and Piotr Winkielman (2000), "Drawing Inferences from Feelings: The Role of Naive Beliefs," in *The Message Within: The Role of Subjective Experience in Social Cognition and Behavior*, ed. Herbert Bless and Joseph P. Forgas, Philadelphia: Psychology Press, 162-75.
- Spencer, Wesley D. and Naftali Raz (1995), "Differential Effects of Aging on Memory for Content and Context: A Meta-analysis," *Psychology and Aging*, 10 (December), 527-39.
- Sperber, Dan and Deirdre Wilson (1986), *Relevance: Communication and Cognition*, Cambridge, MA: Harvard University Press.