

Truthiness and Falsiness of Trivia Claims Depend on Judgmental Contexts

Eryn J. Newman¹, Maryanne Garry¹, Christian Unkelbach², Daniel M. Bernstein³, D. Stephen Lindsay⁴,
& Robert A. Nash⁵

Victoria University of Wellington, New Zealand¹

University of Cologne, Germany²

Kwantlen Polytechnic University, Canada³

University of Victoria, Canada⁴

University of Surrey, United Kingdom⁵

Author Note

We are grateful for the support of the New Zealand Government through the Marsden Fund, administered by the Royal Society of New Zealand on behalf of the Marsden Fund Council.

Correspondence regarding this manuscript should be sent to Eryn Newman, who can be contacted by email at eryn.newman@vuw.ac.nz or at Victoria University of Wellington, Psychology, Box 600, Wellington, New Zealand.

Word Count: 6866 (Abstract, 211; References, 1975).

Abstract

When people rapidly judge the truth of claims presented with or without related but nonprobative photos, the photos tend to inflate the subjective truth of those claims—a "truthiness" effect (Newman et al., 2012). For example, people more often judged the claim "Macadamia nuts are in the same evolutionary family as peaches" to be true when the claim appeared with a photo of a bowl of macadamia nuts than when it appeared alone. We report several replications of that effect and three qualitatively new findings: (a) in a within-subjects design, when people judged claims paired with a mix of related, unrelated, or no photos, related photos produced truthiness but unrelated photos had no significant effect relative to no photos; (b) in a mixed design, when people judged claims paired with related (or unrelated) and no photos, related photos produced truthiness and unrelated photos produced "falsiness;" and (c) in a fully between design, when people judged claims paired with either related, unrelated, or no photos, neither truthiness nor falsiness occurred. Our results suggest that photos influence people's judgments when a discrepancy arises in the expected ease of processing, and also support a mechanism in which—against a backdrop of an expected standard—related photos help people generate pseudoevidence to support claims.

Keywords: Truth judgments, photographs, cognitive fluency, truthiness

Truthiness and Falsiness of Trivia Claims Depend on Judgmental Contexts

True or false: The liquid metal inside a thermometer is magnesium. If you are like the subjects in our studies, you will find this judgment difficult. The cognitive psychology literature suggests that you will try to arrive at a judgment about the claim by retrieving information from memory—related thoughts (“Which metals are liquid?”) and images (your high school chemistry teacher)—to help you decide (Graesser & Hemphill, 1991). The literature also suggests that you will probably search for information confirming the hypothesis that the claim is true (Gilbert, 1991; Nickerson, 1998). If your retrieved thoughts and images are not diagnostic for a true/false judgment, you may make a feeling-based judgment instead (Reber & Unkelbach, 2010; see Schwarz, 2010, for a review). A large body of work shows that these feeling-based judgments are influenced by beliefs, expectations, and other aspects of previous experience that shape processing in the moment, which in turn influence true/false judgments (Begg, Anas, & Farinacci, 1992; Bransford & Johnson, 1972; Jacoby, Kelley, & Dywan, 1989; Mandler, 1980; Ozubko & Fugelsang, 2011; Unkelbach, 2007).

Now suppose you evaluate the claim with the photo in the left panel of Figure 1. Would this photo influence your answer? It seems obvious that the photo should have no influence, since it is nonprobative: It does not provide any evidence about the nature of the liquid inside the thermometer. But recent research suggests that the photo probably would influence your answer (Newman, Garry, Bernstein, Kantner, & Lindsay, 2012). When we asked people to judge the truth of difficult claims that appeared with or without related nonprobative photos, photos biased people toward believing the claims.

Thus related nonprobative photos can produce a “truthiness¹” effect: When making rapid

¹ From comedian Stephen Colbert, who defined truthiness as “truth that comes from the gut, not books.” Incidentally, when Colbert discussed our research on his show, he noted that, “Of course, everyone knows that the metal in a thermometer is not magnesium: One taste and you know it is made of dancing light and liquid time” (see tinyurl.com/truthiness2012).

judgments about the truth of a claim, nonprobative photos nudge people towards believing that claim (Newman et al., 2012). But why? If a picture of a thermometer does not tell you whether the metal inside is magnesium, why would it bias you to conclude the claim is true? Several lines of research fit with the idea that photos help people generate pseudoevidence² about the claim, and suggest two broad ways photos might have this effect.

First, nonprobative photos might promote truthiness in an automatic, bottom-up way by providing rich semantic contexts for target claims. There is evidence that such contexts can boost conceptual processing and produce illusions of familiarity and truth (see Ozubko & Fugelsang 2011; Whittlesea, 1993). For example, people tend to claim that they saw a target word (“test”) earlier when the word appears in a semantically related sentence (“The anxious student wrote a test”) rather than in a more neutral sentence (“Later that afternoon she took a test;” Whittlesea, 1993). Semantically related contexts produce easier conceptual processing relative to neutral contexts, an experience that people may interpret as evidence of familiarity. Likewise, semantically priming, repeating, or retrieving related information can produce illusions of frequency, familiarity, and truth—presumably due to increased ease of retrieval or cognitive availability (Begg et al., 1992; Jacoby et al., 1989; Kelley & Lindsay, 1993; Ozubko & Fugelsang, 2011; Tversky & Kahneman, 1973; Whittlesea, 2011). People draw on the characteristics of a mental event when deciding its accuracy, judging easily imagined propositions as more probable and evaluating the vividness and detail of imagery as a cue to reality (Johnson, 2006; Lindsay, 2008; Sherman, Cialdini, Schwartzman, & Reynolds, 1985). Photos should be particularly good at bootstrapping the generation of images related to a claim and thereby fostering belief in that

² We use “pseudoevidence” as an umbrella term for the kinds of nonprobative, non-sensical evidence people could extract from the photo rather than as a process per se. That evidence could be semantic activation and an experience of conceptual fluency or it could be a more active process where people trawl through the photo interpreting details in the photo as evidence the claim is true.

claim.

Second, nonprobative photos might promote truthiness because people may "trawl" through the photo, deliberately interpreting information they find as support for a default bias to see the claim as true. Such a process might be described as a confirmation bias (Gilbert, Tafarodi, & Malone, 1993; Nickerson, 1998). For instance, someone might examine the photo of a thermometer and think, "I can see a liquid metal inside, and I think magnesium is a metal." By supporting selective searches for hypothesis-consistent evidence, photos might produce a feeling of knowing the correct answer and steer people away from spending extra effort considering reasons why the claim might be false (see Hart, 1965; Thomas, Bulevich, & Dobois, 2012).

Both the conceptual fluency and trawling/confirmation bias mechanisms (which are not mutually exclusive) hinge on the relationship between the photo and the claim, and fit with our earlier findings that semantically related photos can increase truthiness (Newman et al., 2012). But consider now the right panel of Figure 1. There is, of course, no obvious semantic relationship between a thermometer and a lizard. What might be the effect of such a discordant pairing on your answer? This is the question we address in the eight experiments we present here, as part of our ongoing exploration of the mechanisms underlying the effects of photos on truth judgments.

There are reasons to predict that relative to related photos, unrelated photos will cause conceptual disfluency. That is, unrelated photos should make it more difficult for people to recall related thoughts and images. Consistent with that idea, Lee and Labroo (2004) reported that when people saw a target word that was semantically unrelated to a preceding sentence ("The librarian reached for the top shelf and pulled down a book; napkin"), they rated that target word as less pleasant than when they saw a target that was semantically related ("read"). The semantically related sentence might help subjects more quickly identify and comprehend the subsequent target word, and/or the

semantically unrelated sentence might have prepared subjects for a meaning that mismatched the target and thereby slowed processing. Thus it could be that the semantically unrelated words were relatively disfluent compared to their semantically related counterparts, an experience that people interpreted as a signal that the semantically unrelated words were relatively less pleasant.

Parallel effects can occur with perceptual fluency and manipulations such as degraded images, difficult fonts, and low contrast colors that require people to invest more cognitive effort to make sense of stimuli, which in turn can bias people toward evaluating these stimuli more negatively (Diemand-Yauman, Oppenheimer, & Vaughan, 2011; Hernandez & Preston, 2013; Petrova & Cialdini, 2005; Reber & Schwarz, 1999; see Yue, Castel, & Bjork, 2013). In the context of a truth judgment, additional cognitive effort might be taken as a signal that the information being evaluated is false. Considered together research on conceptual and perceptual fluency suggests that, if unrelated photos lead to disfluent processing relative to related photos, then they should produce “falsiness.” That is, nonprobative, unrelated information should bias people toward disbelieving a claim.

The fluency literature also tells us that the effect of photos should depend not only on the semantic relationship between the photo and the claim, but also the experimental context in which they appear. That is, truthiness or falsiness may come about because the photos make the claims feel relatively easy (or difficult) to process compared to the no photo items. Put simply, the effect of photos may depend on the comparison of processing the no photo items in the experiment. Indeed research suggests that a feeling of easy retrieval or easy imagery is driven by a comparison against a standard (Jacoby & Dallas, 1981; Unkelbach & Greifeneder, 2013). When it is easier than expected to retrieve something, people interpret that processing discrepancy as a cue to truth; conversely, when it is more difficult than expected to retrieve something, people interpret this discrepancy the opposite way. But when processing matches expectations, there is no discrepancy to interpret. In other words, the no

photo items may produce a discrepancy in processing, making the claims with photos easier (for related photos) or more difficult (for unrelated photos) to process than expected. Moreover, processing standards do not have to arise from sustained prior experiences; experimental manipulations can forge them in the moment (Jacoby & Dallas, 1981; Westerman, 2008; Whittlesea & Williams, 2001a, 2001b). For example, repetition is thought to produce truth because repeated statements are more easily processed, and people interpret this processing fluency as a sign that statements are true (Bacon, 1979; Begg et al., 1992; Bernstein, 2005; Dechêne, Stahl, Hansen, & Wänke, 2010; Kelley & Lindsay, 1993; Unkelbach, 2006). But as is the case with some of psychological science's well-known effects—such as the mere-exposure effect—the effect of repetition on truth disappears when repetition is manipulated between-subjects (Dechêne, Stahl, Hansen, & Wänke, 2009; Roediger, 2008). In the experiments reported here, we examine the possibility that no photo items serve as a benchmark for processing by manipulating the presence of photos within and between-subjects. If the effect of photos depends on a comparison against a (no photo) standard, we should expect to see the most pronounced effects of photos in our within-subject designs.

But there are also reasons to predict that unrelated photos would not produce falsiness in any of our experimental designs. As we suggested earlier, when people are making a true/false decision they might trawl through and selectively interpret information from a related photo as evidence that the claim is true. It seems unlikely that people would engage in this deliberate “trawling” strategy with unrelated photos. It makes little sense to search a photo of a lizard for evidence about magnesium's putative role in thermometers. Moreover, we know from research on confirmation bias that if people encounter information inconsistent with the hypothesis at hand, they often ignore it, allocate less weight to that evidence, or even distort it to fit with their hypothesis (Darley & Gross, 1983; Kuhn, 1989; Snyder & Cantor, 1979; see Nickerson, 1998, for a review; see also Fischhoff, 1975). Some

research indicates that adding related nonprobative information can sway people's judgments about others, but that adding unrelated information produces the same effect as giving people no information at all (Gill, Swann, & Silvera, 1998). Thus, a trawling mechanism would predict that unrelated photos would have little if any effect on people's decisions.

We examined the effects of semantically related and unrelated photos in eight experiments. In each, we showed people a series of trivia claims that appeared with or without a photo that was or was not semantically related to the claim.

Experiments 1-3

In Experiments 1-3 we examine the effects of semantically unrelated photos on judgments of truth in a fully within-subjects design. If the effects of photos depend on an expected standard of processing, then they should have a particularly potent effect when people experience related and unrelated photos within the same experimental context.

Method

Subjects. In Experiment 1-3, we used Amazon's Mechanical Turk (MTurk; www.mturk.com/mturk) to recruit subjects in the US. MTurk is an online subject pool, in which "workers" complete experiments and surveys and receive small amounts of Amazon credit (they received \$0.60 for this experiment) that they can use to purchase things on amazon.com. Studies run online using MTurk attract diverse subjects and tend to produce similar results to those run in a laboratory (Buhrmester, Kwang, & Gosling, 2011; Germine, Nakayama, Duchaine, Chabris, Chatterjee, & Wilmer, 2012; Mason & Suri, 2011; see also Lindsay, Allen, Chan, & Dahl, 2004). We predetermined a sample size of 200 subjects based on pilot testing. Because of a quirk in the way MTurk assigns subjects, there were 208 subjects in Experiment 1, 204 in Experiment 1, and 216 in Experiment 3.

Design. For Experiments 1-3, we used a single-factor (photo: related, unrelated, no photo) -

within-subjects design (a table summarizing all of our experimental designs appears in supplementary materials).

Procedure. We used trivia claims from previous research to assemble sets of difficult true/false trivia claims sampling general knowledge (Newman et al., 2012; see also Nelson & Narens, 1980; Unkelbach, 2007). People typically answer these critical claims correctly 40-60% of the time.

We used Qualtrics software to present 40 trivia claims to subjects. We told subjects that sometimes they would see a photo with these claims, and sometimes they would not. We did not provide any instructions about how they should use the photo. We asked subjects to decide the truth of the claim “...as quickly as possible, but not so quickly that you start making errors.”

The claims appeared individually, in large black font against a white background. In Experiments 1 and 2, to orient people to the task, for the first 16 trials they saw easy trivia claims (which tend to be answered correctly 80-100% of the time; for example, “The player who guards the net in soccer is called the goalie”). Half these easy claims appeared with a photograph, half with no photograph. To ensure the practice phase did not teach subjects a rule about the relationship between truth and the presence of photos, we paired photos equally often with true and false claims. Experiments 1 through 3 differed only in terms of this set of 16 practice items seamlessly preceding the critical items: In Experiment 1, a random half of the practice items appeared with a related photo and the remainder appeared with no photo; in Experiment 2, a random half of the 16 practice items appeared with an unrelated photo and the remainder appeared with no photo; in Experiment 3 there were no practice items. To anticipate, the pattern of means for the critical items was the same in all three experiments.

Immediately after these easy practice trivia claims, the experimental phase began. Subjects saw 24 difficult trivia claims, half true and half false. For one third of trials, a related nonprobative photo depicted the grammatical subject of the claim. A few examples will convey the flavor of this phase:

“Macadamia nuts are in the same evolutionary family as peaches” sometimes appeared with a photo of macadamia nuts; “Cactuses can reproduce by parthenogenesis” with a photo of a cactus, and “The plastic things on the ends of shoelaces are called aglets” with a photo of a shoe with a shoelace. For another third of trials, a semantically unrelated nonprobative photo appeared with the claim: The claim about macadamia nuts sometimes appeared with a photo of a trash can; the claim about cactuses with a photo of a bicycle, and the claim about shoelaces with a photo of a pig. For the other third of the trials, people saw trivia claims presented without a photo. We used the set of related nonprobative photos from Newman et al. (2012) and created a new set of semantically unrelated nonprobative photos for the experiments reported here. A semantically unrelated photo was selected for each trivia claim. As a set, the unrelated photos represented a similar range of living and non-living objects as the original set of related photos from Newman et al. None of the photos revealed the accuracy of the trivia claims. We randomized the order of claims for each subject, and counterbalanced so that claims appeared equally often with a related photo, unrelated photo, or no photo. We used an online script to assign subjects to conditions randomly.

Results and Discussion

Our primary aim was to examine the effects of semantic relatedness on truthiness and falsiness. To address this question we calculated people’s bias (C) to say a claim was true (Stanislaw & Todorov, 1999).³ As the left panel of Figure 2 shows, pairing a claim with a related nonprobative photo produced truthiness (as shown by the negative value of C in Experiment 1). But unrelated photos did not produce falsiness; instead, trials with unrelated photos behaved more like trials with no photos.

Consistent with the pattern displayed in the figure, a one-way ANOVA of C showed a main effect

³ We corrected floor and ceiling responses by converting values of 1 to .99 and values of 0 to .01 (Wickens, 2002). Note that for all experiments, when we ran parallel analyses using the total proportion of true responses for each condition we found the same pattern of results (see Stanislaw & Todorov, 1999; Wickens, 2002).

for photo $F(2, 206) = 4.10, p = .02, \eta_p^2 = .04$. Follow-up tests showed that related photos produced more bias to say true than unrelated photos or no photos, $t_{\text{related-unrelated}}(207) = 2.30, p = .02$, Cohen's $d = .16$, $t_{\text{related-no photo}}(207) = 2.77, p = .01, d = .20$. In contrast, unrelated photos did not differ from no photos, $t_{\text{unrelated-no photo}}(207) = 0.47, p = .64, d = .04$. We found no consistent effects for photographs on d' but we report d' analyses for each experiment in supplemental materials.

On the one hand, it is surprising that unrelated photos did not produce falsiness given that they should have felt especially difficult to process compared to the related photos. On the other hand, a critic might argue that unrelated photos failed to produce falsiness because the combination of the practice task (comprised of related and no photos) and experiment proper meant that unrelated photos occurred rarely. That is, 16 practice items appeared with a related photo or no photo, followed by a random series of 24 trivia claims in the experiment proper of 8 related, 8 unrelated, and 8 no photo. Therefore, people saw trivia claims with no photos 40% of the time; trivia claims with related photos 40% of the time, and trivia claims with unrelated photos only 20% of the time—in other words, when a photo appeared with a claim, two-thirds of the time it related to the claim and one-third of the time it did not. Such a mix might have led subjects to “oversee” relationships between some unrelated photos and their associated claims, diluting the intended effect of the unrelated condition. To address this issue, in Experiment 2 we reran Experiment 1 except that the 16 practice items appeared with an unrelated photo or no photo and in Experiment 3 we ran another version of Experiment 1 in which there was no practice task.

Did the practice phase bear on the effects we observed in Experiment 1? The answer is no. Regardless of whether people saw unrelated photos during the practice phase, or even when there was no practice phase, Table 1 shows that we replicated the primary findings in Experiment 1: that is, related photos produced truthiness, and unrelated photos acted just like no photo trials.

A one-way ANOVA of *C* showed a main effect for photo (Experiment 2 $F(2, 202) = 4.29, p = .02, \eta_p^2 = .04$; Experiment 3 $F(2, 214) = 4.25, p = .02, \eta_p^2 = .04$). Related photos produced more bias to say true than unrelated photos or no photos (Experiment 2 $t_{\text{related-unrelated}}(203) = 2.18, p = .03, d = .15$; $t_{\text{related-no photo}}(203) = 2.78, p = .01, d = .20$; Experiment 3 $t_{\text{related-unrelated}}(215) = 2.85, p < .01, d = .19$; $t_{\text{related-no photo}}(215) = 1.90, p = .06, d = .13$), bias for unrelated did not differ from bias for no photos (Experiment 2, $t_{\text{unrelated-no photo}}(203) = .77, p = .44, d = .06$; Experiment 3 $t_{\text{unrelated-no photo}}(215) = 1.03, p = .30, d = .06$).

In addition to these analyses we used data across Experiments 1-3 to arrive at a more precise estimate of the size of both the truthiness and falsiness effects by subjecting them to random effects model mini meta-analyses (Cumming, 2012). Those data appear in Figure 3. In the first analysis we focused on truthiness, comparing the related photo and no photo conditions to calculate an estimated raw effect size of 0.16, 95% CI [0.09, 0.23], $z = 4.27, p < .01$. We ran a parallel analysis focusing on falsiness, this time comparing unrelated photo and no photo conditions, to calculate an estimated raw effect size of 0.00, 95% CI [-0.07, 0.08], $z = .13, p = .90$, a range of values that plausibly include falsiness (negative values), truthiness (positive values), or zero.

Experiment 4

In Experiments 1-3 (and in our prior work on truthiness) subjects judged trivia claims as true or false. The primary aim of Experiment 4 was to replicate the pattern in 1-3 using a procedure in which subjects rated each trivia claim on a scale from 1 (definitely false) to 9 (definitely true). In Experiment 4, we also examined generalizability by using a new set of trivia claims—all of which were false—and on a sample of students at a UK university who participated in one-on-one lab sessions. Finally, in Experiment 4 we also added a response latency measure, which allowed us to examine the possibility that people were simply ignoring the unrelated photos.

Method

Subjects. Sixty-five University of Surrey undergraduates participated in exchange for course credit.

Design. The design was identical to Experiment 1-3.

Procedure. We developed a new set of 87 difficult claims (with mean ratings that fell between 3-5 on a 7 point scale: 1 (definitely false) to 7 (definitely true), range = 3.13 to 4.96 and $M = 3.91$; $SD = 0.48$). All the claims were false. We used E-Prime software to present the 87 claims to subjects, tested individually in the lab for course credit. One third of the claims appeared with an accompanying related photo, one third with an unrelated photo, and one third with no photo. We did not provide any instructions about how subjects should use the photos, but simply asked them to be as quick and as accurate as possible. Unlike our previous experiment in which we asked for binary true/false judgments, this time we asked participants to rate their confidence in the truth of each statement, using a scale from 1 (definitely false) to 9 (definitely true). Subjects entered their responses using the computer keyboard and we measured subjects' response latency for each claim.

Results and Discussion

Although we changed the response measure and the trivia claims, we replicated the results from Experiment 1-3: Pairing a claim with a related nonprobative photo produced truthiness, but unrelated photos behaved more like trials with no photos. There was a significant effect of photo on subjects' confidence ratings, $F(2, 63) = 4.88$, $p = .01$, $\eta^2 = .13$. Related photos ($M = 5.25$, $SD = .66$) produced greater confidence ratings than did unrelated photos ($M = 5.11$, $SD = .73$) or no photos ($M = 5.06$, $SD = .68$), $t_{\text{related-unrelated}}(64) = 2.24$, $p = .03$, Cohen's $d = .28$, $t_{\text{related-no photo}}(64) = 3.00$, $p < .01$, $d = .37$, whereas confidence for unrelated and no photos was similar $t_{\text{unrelated-no photo}}(64) = 0.88$, $p = .38$, $d = .11$. We next coded subjects' responses as "true" whenever their confidence ratings were 6 or above (i.e., greater than the scale's midpoint), and we calculated the proportion of statements rated as true. This

analysis revealed the same main effect of photo, $F(2, 63) = 7.39, p < .01, \eta_p^2 = .19$; $t_{\text{related-unrelated}}(64) = 3.57, p < .01, d = .44$, $t_{\text{related-no photo}}(64) = 3.14, p < .01, d = .39$, $t_{\text{unrelated-no photo}}(64) = 0.23, p = .82, d = .03$.

That related photos produced truthiness replicates our prior findings and shows that non-probative photos that are related to trivia claims can foster belief in those claims, true or false. But why did unrelated photos have no effect relative to no photos? One possibility is that people adopted a strategy of ignoring unrelated photos, much as people in a Stroop task can adopt a strategy of ignoring the word and focusing on the color (Kane & Engle, 2003; Stroop, 1935; see also Besner, Stolz, & Boutilier, 1997). But there are reasons to doubt that possibility. For one thing, subjects must have processed unrelated photos to some extent to determine that they were unrelated. For another, in Experiment 4 we measured response latency and observed that responses were approximately 300 ms slower when there was a photo (related or unrelated) than when there was no photo, which suggests that subjects did not ignore the unrelated photos.⁴ Another explanation is that presenting unrelated photos mixed with related photos led subjects to look for semantic relationships between the claims and the photos, for both types of photo. For instance, they might have looked at the thermometer claim paired with a lizard and thought, “Well, the lizard is long and thin, like a thermometer, and has a stripe up the center just like a thermometer.” Such a strategy might have mitigated the perceived incongruity between the photos and trivia claims, leading unrelated photos to become more like related photos and thus diluting their falsiness. This idea fits with research showing that people will find or create meaning (Bartlett, 1932), especially when they are faced with pairings that do not have an obvious semantic relationship (for instance, novel metaphors; Grimshaw, Stewart, & Lauwereyns, 2011; Lynott &

⁴ Subjects took longer to respond to claims presented with photos than to those without photos, whereas response times for claims with related and unrelated photos were similar ($Mdn_{\text{related photo}} = 6899.28\text{ms}$, $Mdn_{\text{unrelated photo}} = 6979.28\text{ms}$, $Mdn_{\text{no photo}} = 6604.97\text{ms}$; Friedman’s $\chi^2 = 7.02, p = .03$), Wilcoxon’s $Z_{\text{related-unrelated}} = 1.61, p = .11$; $Z_{\text{related-no photo}} = 2.51, p = .01$; $Z_{\text{unrelated-no photo}} = 2.89, p < .01$.

Connell, 2010; see also Foster & Kokko, 2009).

Taken together, the findings from Experiments 1-4 suggest that in the context of related photos, unrelated photos exert the same influence as no photos. But part of this context likely includes subjects' expectations that each photo will be meaningfully related to a claim. Might that expectation modulate the effect of unrelated photos? If we manipulated the semantic relatedness of photos between subjects, those shown unrelated photos should come to have minimal expectations about meaningful relationships between photos and claims. What then should be the effect of unrelated photos relative to no photos? One possibility is that unrelated photos would make it more difficult for people to bring to mind related ideas, producing disfluent conceptual processing. If so, then we should see increased falsiness relative to no-photo items among subjects for whom photos were always unrelated. A second possibility is that setting people up to expect incongruity between photos and claims might lead them to ignore all of the photos—a strategy that should be much easier to apply when all the photos are unrelated (Kane & Engle, 2003). We addressed these questions in Experiments 5 and 6 (the latter was an exact replication of the former).

Experiments 5 and 6

Method

Subjects. We used MTurk to recruit 196 subjects (Experiment 5) and 185 subjects (Experiment 6) in the US. They received \$0.60 Amazon credit for participating.

Design. We used a 2 (photo: yes, no) x 2 (relatedness: related, unrelated) mixed design, manipulating the presence of a photo within subjects and relatedness of the photo between subjects.

Procedure. We used the same procedure as Experiment 1 with the following changes. Immediately after seeing the easy practice trivia claims, people saw 32 difficult trivia claims (half true, half false). Half of the claims appeared with a photo; for half of the subjects the photo was always

related to the claim whereas for the remaining subjects it was always unrelated to the claim.

Results

As the middle panel of Figure 2 shows, related photos again produced truthiness. But the Figure also shows that, in contrast to Experiments 1-3, unrelated photos produced falsiness.

A 2 x 2 mixed ANOVA showed the pertinent Photo x Relatedness interaction, $F(1, 194) = 9.01, p < .01, \eta_p^2 = .04$; related photos produced truthiness, $t_{\text{related-no photo}}(93) = 2.12, p = .04, d = .23$, but unrelated photos produced falsiness, $t_{\text{unrelated-no photo}}(101) = 2.12, p = .04, d = .19$. As Table 1 shows, these patterns replicated in Experiment 6; there was a Photo x Relatedness interaction, $F(1, 183) = 18.22, p < .01, \eta_p^2 = .09$; related photos produced truthiness, $t_{\text{related-no photo}}(83) = 4.08, p < .01, d = .44$, but unrelated photos produced falsiness, $t_{\text{unrelated-no photo}}(100) = 1.94, p = .06, d = .19$.

As in Experiments 1-3, we used data from both the primary experiment and replication to calculate a more precise estimate of the size of both the truthiness and falsiness effects by subjecting these data to random effects model mini meta-analyses (see Figure 3; Cumming, 2012). The estimated raw effect size for truthiness was 0.17, 95% CI [0.10, 0.25], $z = 4.45, p < .01$ and for falsiness it was -0.10, 95% CI [-0.16, -0.03], $z = 2.87, p < .01$.

Discussion

We found that when the semantic relatedness of photos to claims was manipulated between subjects, unrelated photos produced falsiness—a pattern that fits with a fluency account, suggesting that the unrelated photos may have produced an experience of disfluent processing. Taken together, Experiments 1-4 and 5-6 might lead us to conclude that although the effects of unrelated photos depend on the experimental context in which they appear, related photos produce truthiness regardless of context. But in both of these experiments (and in Newman et al., 2012), claims with related photos always appeared among claims without photos; the same is true of claims with unrelated photos. Thus

we cannot rule out the possibility that truthiness also depends on context. That is, perhaps what drives truthiness is that people evaluate their processing experiences with photos against a benchmark of their experiences without photos (e.g. Dechêne, Stahl, Hansen, & Wänke, 2009; Jacoby & Dallas, 1981; Roediger, 2008; Unkelbach & Greifeneder, 2013). If truthiness (and falsiness) depends on a standard, we should see that the pattern from Experiments 5 and 6 disappears when people have no standard against which to interpret the ease or difficulty of processing that accompanies claims paired with photos. Accordingly, in Experiments 7 and 8 (8 was an exact replication of 7) we manipulated the photo factor entirely between subjects.

Experiment 7 and 8

Method

Subjects. We used MTurk to recruit 301 subjects (Experiment 7) and 301 subjects (Experiment 8) in the US. They received a \$0.60 Amazon credit.

Design. We used a single-factor (photo: related, unrelated, no photo) between-subjects design.

Procedure. Subjects saw the same trivia claims as in Experiments 5-6: 16 practice claims followed by 32 trivia claims. The key difference in Experiments 7-8 is that we manipulated the photo factor between subjects. That is, a third of subjects saw the claims paired with related photos, a third saw the claims paired with unrelated photos, and a final third saw the claims paired with no photo. We gave subjects the same instructions as in the prior experiments except that we removed any reference to the presence or absence of photos.

Results and Discussion

As the right side of Figure 2 shows, compared to when there was no photo, related photos did not produce truthiness and unrelated photos did not produce falsiness. In other words, there was no effect of photo, $F(2, 298) = 0.75$, $p = .47$, $\eta_p^2 = .01$. As Table 1 shows, these patterns replicated in Experiment

8 $F(2, 298) = 1.02, p = .36, \eta_p^2 = .01$. Together these findings fit with the idea that truthiness and falsiness depend on expectations acquired in the experimental context, and occur only when there is a discrepancy in the expected ease of processing (Whittlesea & Williams, 1998).

As in our earlier experiments, we used both sets of data to calculate a more precise estimate of the size of both the truthiness and falsiness effects by subjecting them to random effects model mini meta-analyses (see Figure 3, Cumming, 2012). The estimated raw effect size for truthiness was -0.02 , 95% CI $[-0.11, 0.07]$, $z = .48, p = .63$. The estimated raw effect size for falsiness was -0.06 , 95% CI $[-0.14, 0.02]$, $z = 1.44, p = .15$. These analyses show that the effects of photos on truthiness or falsiness are non-existent or trivial in the context of a between-subjects design.

General Discussion

Across eight experiments, we found that the effects of nonprobative photos on the perceived truth of trivia statements vary with experimental context. We also found that the effects of nonprobative photos depend on the semantic relationship between the photo and the claim. When trivia statements were accompanied by a mix of related, unrelated, and no photos, related photos produced truthiness relative to no photos but unrelated photos had no effect relative to no photos. When people had a no-photo standard against which to evaluate either related or unrelated photos (Experiments 5-6), related photos increased the truth of claims (truthiness), and unrelated photos decreased the truth of claims (falsiness). But when we removed people's ability to compare against a standard, employing a fully between-subjects design, neither related nor unrelated nonprobative photos influenced true/false judgments. Considered as a whole, this pattern of results suggests that photos influence people's judgments when a discrepancy arises in the expected ease of processing—that is, when subjects find claims with photos easier (or more difficult) to evaluate compared to claims without photos (Hansen, Dechêne, & Wänke, 2008; Westerman, 2008; Whittlesea & Williams, 2001a, 2001b).

Of course, one might argue that these effects of photos are produced by some kind of demand characteristic. But there are at least three reasons why we don't think that is the case. First, we told subjects that the study was about visual and verbal learning, so subjects did not know the purpose of the study. Second, if subjects did figure out our hypothesis, we wouldn't necessarily expect them to act in a way that would promote the effect. In fact, the literature on fluency shows that when people detect the source of fluent processing (or they can "see through" the manipulation), they tend to discount that fluency experience and try to counteract any influence that source of fluency might have on their judgements (e.g., Jacoby & Whitehouse, 1989; Oppenheimer, 2004). In other words, this literature suggests that if subjects had some inkling of the idea that photos could influence their decisions, we would expect them to say "false" more when they saw a related photo paired with a claim. But we see the opposite pattern of results. Third, if demand characteristics were driving truthiness, then photos should have a systematic effect not just across studies but across all of our items. In our earlier work, we found instead that photos produced truthiness for unfamiliar claims—that is, in conditions of uncertainty, which are conditions that also make people susceptible to fluency effects (see Newman et al., 2012; Schwarz, 2010; Unkelbach, 2007).

Rather, these findings support a mechanism in which—against a backdrop of an expected standard—related photos help people generate pseudoevidence to support the claim. Related photos might help people generate pseudoevidence by facilitating conceptual processing, perhaps increasing cognitive availability and (or) helping people trawl for evidence (Kelley & Lindsay, 1993; Nickerson, 1998; Tversky & Kahneman, 1973; Whittlesea, 2011). Our anecdotal observations of a small number of lab-tested pilot subjects (who provided think-aloud protocols while performing the procedure of Experiments 5-6) also fit with such a mechanism. For instance, when faced with the claim "Macadamia nuts are in the same evolutionary family as peaches" paired with a photo of macadamia nuts, one

subject said, “I’m going to go with yes because they kind of look like peaches, so that would make sense.”

The finding that unrelated photos produced falsiness in Experiments 5 and 6 fits with the idea that unrelated photos should not help people generate confirming pseudoevidence, and may actually generate a feeling of incongruity or difficult processing when evaluating the claim. Nonetheless, it is puzzling that unrelated photos did not produce falsiness in Experiments 1-4; instead, subjects seemed to treat them as no photo items. One reason why unrelated photos did not produce falsiness is that presenting them along with related photos led subjects to look for semantic relationships between the claims and the unrelated photos. This kind of process might dilute the perceived incongruity between the unrelated photos and trivia claims, leading unrelated photos to become more like related photos and thus reducing their falsiness. Indeed, people have a tendency to find meaning (Bartlett, 1932), especially when they are faced with pairings that do not have an obvious semantic relationship (for instance, novel metaphors; Grimshaw, Stewart, & Lauwereyns, 2011; Lynott & Connell, 2010; see also Foster & Kokko, 2009). We think this is an intriguing result worthy of future research.

Photos did not produce truthiness or falsiness in Experiments 7 and 8. That finding is consistent with research on discrepant fluency, where conceptually fluent items have the most pronounced effects when they are presented against a backdrop of relatively disfluent items (when ease of processing is manipulated within subjects; Hansen et al., 2008; Westerman, 2008; Whittlesea & Williams, 2001a, 2001b). But there are at least two other possibilities. First, there are many cues that people could draw on to decide whether a claim is true (e.g., credibility of the source of the claim, familiarity, logical coherence) and it is possible that the within-subject design produces truthiness because it emphasizes photos as the important cue, because their presence changes across trials (see Schwarz, *in press*). Such an emphasis is not placed on photos in the between-subjects design (the presence or absence of photos

is constant across trials), so people may turn to other cues they notice changing across items such as the familiarity of a name or the ease or difficulty of a statement (see Newman et al., 2012). An interesting avenue for future research might be to disentangle these accounts of the between-subjects data.

A second and less interesting account of the between-subjects data is that people simply ignored the photos and that is why they had little influence on people's judgments. But there are several reasons why the ignoring hypothesis is not a compelling explanation of the results of Experiment 7 and 8. For instance, while people might find the unrelated photos distracting, it is unclear why people would be motivated to ignore the related photos in the between-subjects condition. In addition, eye-tracking studies show that we tend to be drawn to visual information, so it would take cognitive effort and some kind of explicit strategy to ignore the photos (e.g., Sargent, 2007). In fact, we should be able to detect if people are engaging in this kind of strategy by examining people's responses and response times across trials. Although we do not have response times for Experiments 7 and 8, we did collect RTs when we ran a between-subjects replication of some of our earlier photo research (We conducted a conceptual replication of Experiment 1 from Newman et al., 2012, manipulating the presence of photos between-subjects). If the ignoring hypothesis is true, we should have seen a difference in people's responses and RTs across trials. That is, people should initially show a photo effect, or spend more time on photo trials, an effect that should fade across trials as people refine their ignoring strategy. In contrast to the ignoring hypothesis, we found that if we compare the photo and no photo trials, there were no differences in people's response times for the first quarter block and the last quarter block. In fact, in both blocks and across all trials, people took slightly longer to respond to photo trials (a pattern consistent with our within-subject data in Experiment 4). That is, people started the same way they ended—they did not learn to ignore across trials. Moreover, there was no evidence that photos caused diminishing truthiness over trials. Both of these findings are at odds with the proposal that people

learned to use some kind of ignoring strategy in the between-subjects design.

In all of our analyses we examined the effect of photos relative to the no photo condition. So what we have demonstrated is that the difference in response bias between photo and no photo conditions varied across experimental contexts. But we could also ask how response bias in each condition—the absolute bias for related, unrelated and no photo trials—varied across experimental contexts. This question gains piquancy from the fact that the response bias to no-photo items (which served as our baseline control condition) varied across conditions (see Figure 4 for a forest plot of estimates of C against the benchmark of zero).

Lindsay and Kantner (2011) observed that old/new recognition memory response bias tends to be conservative when the stimuli are scans of paintings. Subsequent work (Lindsay, Kantner, & Fallow, in press) found that when the stimuli include a mix of paintings and words, bias was conservative on paintings and liberal on words. But in a between-subjects design, although bias continued to be conservative on paintings it was neutral on words. It may be that the liberal recognition bias on words in the mixed case was an artifact of subjects' reluctance to classify paintings as old; that is, subjects were reluctant to claim that paintings were old, but they also did not want to say "No" to too many probes so they compensated by being liberal on words. We find a similar pattern in the data here: when claims were a mix of related, unrelated, and no photos, bias was neutral for claims presented with an unrelated photo or with no photo and liberal for claims presented with related photos. But in a between-subjects design, bias was neutral for all conditions. Taken together, these findings show that absolute response bias varies according to attributes of the stimuli, as well as the context in which those stimuli appear.

Taken together with our prior work, the eight experiments reported here suggest that nonprobative photographs do more than simply decorate claims: they wield a significant and immediate

influence on people's judgments. As with many effects in the cognitive psychology literature, our photo-truthiness and falsiness effects depend on the way in which people process and interpret photos when evaluating the truth of claims.

References

- Alter, A. L., & Oppenheimer, D. M. (2008). Effects of fluency on psychological distance and mental construal (or why New York is a large city, but New York is a civilized jungle). *Psychological Science*, 19, 161-167. doi: 10.1111/j.1467-9280.2008.02062.x.
- Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review*, 13, 219-235. doi: 10.1177/1088868309341564.
- Bacon, F. T. (1979). Credibility of repeated statements: Memory for trivia. *Journal of Experimental Psychology: Human Learning and Memory*, 5, 241-252. doi: 10.1037/0278-7393.5.3.241.
- Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. London: Cambridge University Press.
- Begg, I. M., Anas, A., & Farinacci, S. (1992). Dissociation of processes in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General*, 121, 446-458. doi: 10.1037/0096-3445.121.4.446.
- Bernstein, D.M. (2005). Making sense of memory. *Canadian Journal of Experimental Psychology*, 59, 199-208. doi:10.1037/h0087475
- Besner, D., Stoltz, J. A., & Boutilier, C. (1997). The Stroop effect and the myth of automaticity. *Psychonomic Bulletin & Review*, 4, 221-225. doi: 10.3758/BF03209396.
- Bransford, J. D., & Johnson, M. K. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior*, 11(6), 717-726. doi: 10.1016/S0022-5371(72)80006-9.
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6, 3-5. doi:

10.1177/1745691610393980

Cumming, G. (2012). *Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis*. New York, NY: Routledge.

Darley, J. M., & Gross, P. H. (1983). A hypothesis- confirming bias in labeling effects. *Journal of Personality and Social Psychology*, 44, 20-33. doi: 10.1037/0022-3514.44.1.20.

Dechêne A, Stahl C, Hansen J, & Wänke M. (2009). Mix me a list: Context moderates the truth effect and the mere-exposure effect. *Journal of Experimental Social Psychology*, 45, 117–1122. doi: 10.1016/j.jesp.2009.06.019.

Dechêne, A., Stahl, C., Hansen, J., & Wänke, M. (2010). The truth about the truth: A meta-analytic review of the truth effect. *Personality and Social Psychology Review*, 14, 238–257. doi: 10.1177/1088868309352251.

Derzon, J. H., & Alford, A. A. (2013). Forest plots in Excel: Moving beyond a clump of trees to a forest of visual information. *Practical Assessment, Research & Evaluation*, 18(7), 1-9.

Diemand-Yauman, C., Oppenheimer, D. M., & Vaughan, E. B. (2011). Fortune favors the bold (and the italicized): Effects of disfluency on educational outcomes. *Cognition*, 118, 111–115. doi: 10.1016/j.cognition.2010.09.012.

Fischhoff, B. (1975). Hindsight does not equal foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1, 288-299.

Foster, K. R., & Kokko, H. (2009). The evolution of superstitious and superstition-like behaviour. *Proceedings of the Royal Society B: Biological Sciences*, 276(1654), 31-37. doi: 10.1098/rspb.2008.0981

Germine, L., Nakayama, K., Duchaine, B. C., Chabris, C. F., Chatterjee, G., & Wilmer, J. B. (2012). Is

the web as good as the lab? Comparable performance from web and lab in cognitive/perceptual experiments. *Psychonomic Bulletin & Review*, 847-857. doi:10.3758/s13423-012-0296-9

Gilbert, D. T. (1991). How mental systems believe. *American Psychologist*, 46, 107–119. doi: 10.1037/0003-066X.46.2.107.

Gilbert, D. T., Tafarodi, R. W., & Malone, P. S. (1993). You can't not believe everything you read. *Journal of Personality and Social Psychology*, 65, 221-233. doi: 10.1037/0022-3514.65.2.221.

Gill, M.J, Swann, W.B., Jr., & Silvera, D.H. (1998). On the genesis of confidence. *Journal of Personality and Social Psychology*, 75, 1101-1114. doi: 10.1037/0022-3514.75.5.1101.

Graesser, A. C., & Hemphill, D. (1991). Question answering in the context of scientific mechanisms. *Journal of Memory and Language*, 30, 186-209. doi: 10.1016/0749-596X(91)90003-3.

Grimshaw, G.M., Stewart, J., & Lauwereyns, J. (2011, April). *The subjective nature of meaning*. Poster presented at the meeting of the Cognitive Neuroscience Society, San Francisco CA.

Hansen, J., Dechêne, A., & Wänke, M. (2008). Discrepant fluency increases subjective truth. *Journal of Experimental Social Psychology*, 44(3), 687-691. doi: 10.1016/j.jesp.2007.04.005

Hansen, J., & Wanke, M. (2010). Truth from language and truth from fit: The impact of linguistic concreteness and level of construal on subjective truth. *Personality and Social Psychology Bulletin*, 36, 1576-1588. doi: 10.1177/0146167210386238.

Hart, J. T. (1965). Memory and the feeling-of-knowing experience. *Journal of Educational Psychology*, 56, 208–216. doi: 10.1037/h0022263

Hernandez, I., & Preston, J. L. (2013). Disfluency disrupts the confirmation bias. *Journal of Experimental Social Psychology*, 49, 178-182. doi: 10.1016/j.jesp.2012.08.010.

Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and

perceptual learning. *Journal of Experimental Psychology: General*, 110, 306-340. doi: 10.1037/0096-3445.110.3.306.

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

Jacoby L. L. & Whitehouse K. (1989) An illusion of memory: False recognition influenced by unconscious perception. *Journal of Experimental Psychology: General*, 118, 126–135. doi: 10.1037/0096-3445.118.2.126.

Johnson, M.K. (2006). Memory and reality. *American Psychologist*, 61, 760-771. doi: 10.1037/0003-066X.61.8.760.

Kane, M. J., & Engle, R. W. (2003). Working memory capacity and the control of attention: The contributions of goal neglect, response competition, and task set to Stroop interference. *Journal of Experimental Psychology: General*, 132, 47-70. doi: 10.1037/0096-3445.132.1.47.

Kelley, C. M., & Lindsay, D. S. (1993). Remembering mistaken for knowing: Ease of retrieval as a basis for confidence in answers to general knowledge questions. *Journal of Memory and Language*, 32, 1-24. doi: 10.1006/jmla.1993.1001.

Kuhn, D. (1989). Children and adults as intuitive scientists. *Psychological Review*, 96, 674-689. doi: 10.1037/0033-295X.96.4.674.

Lee, A. Y., & Labroo, A. (2004). Effects of conceptual and perceptual fluency on affective judgment. *Journal of Marketing Research*, 41, 151-165. doi: 10.1509/jmkr.41.2.151.28665.

Lindsay, D. S. (2008). Source monitoring. In H. L. Roediger, III (Ed.), *Cognitive psychology of memory. Vol. 2 of Learning and memory: A comprehensive reference, 4 vols.* (J. Byrne, Editor) (pp. 325-348). Oxford:

Elsevier.

- Lindsay, D. S., Allen, B. P., Chan, J. C. K., & Dahl, L. C. (2004). Eyewitness suggestibility and source similarity: Intrusions of details from one event into memory reports of another event. *Journal of Memory and Language*, 50, 96-111. doi:10.1016/j.jml.2003.08.007
- Lindsay, D. S., & Kantner, J. (2011). A search for influences of feedback on recognition of music, poetry, and art. In P. Higham & J. Leboe (Eds.), *Constructions of remembering and metacognition: Essays in honor of Bruce Whittlesea* (pp. 137-154). Houndmills, UK: Palgrave Macmillan.
- Lindsay, D. S., Kantner, J., & Fallows, K. M. (in press). Why is recognition memory response bias conservative on paintings? In D. S. Lindsay, C. M. Kelley, A. P. Yonelinas, & H L. Roediger III (Eds.), *Remembering: Attributions, processes, and control in human memory*. Psychology Press.
- Lynott, D., & Connell, L. (2010). Embodied conceptual combination. *Frontiers in Psychology*, 1, 1–14. doi: 10.3389/fpsyg.2010.00212.
- Mandler, G. (1980). Recognizing: The judgment of previous occurrence. *Psychological Review*, 87, 252–271. doi: 10.1037/0033-295X.87.3.252.
- Mason, W., & Suri, S. (2011). Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods*, 44, 1-12. doi:10.3758/s13428-011-0124-6
- Masson, M. E. J., & Loftus, G. R. (2003). Using confidence intervals for graphically based data interpretation. *Canadian Journal of Experimental Psychology*, 57, 203-220. doi: 10.1037/h0087426.
- Nelson, T. O., & Narens, L. (1980). Norms of 300 general- information questions: Accuracy of recall, latency of recall, and feeling-of-knowing ratings. *Journal of Verbal Learning and Verbal Behavior*, 19, 338-368. doi: 10.1016/S0022-5371(80)90266-2.
- Newman, E. J., Garry, M., Bernstein, D. M., Kantner, J., & Lindsay, D. S. (2012). Nonprobative photos

(or words) promote truthiness. *Psychonomic Bulletin & Review*, 19, 969-974. doi: 10.3758/s13423-012-0292-0.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2, 175–220. doi: 10.1037/1089-2680.2.2.175.

Oppenheimer, D. M. (2004). Spontaneous discounting of availability in frequency judgment tasks. *Psychological Science*, 15, 100-105. doi: 10.1111/j.0963-7214.2004.01502005.x.

Ozubko, J. D., & Fugelsang, J. (2011). Remembering makes evidence compelling: Retrieval from memory can give rise to the illusion of truth. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37, 270-276. doi: 10.1037/a0021323.

Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness and Cognition*, 8, 338-342. doi: 10.1006/ccog.1999.0386.

Reber, R., & Unkelbach, C. (2010). The epistemic status of processing fluency as source for judgments of truth. *Review of Philosophy and Psychology*, 1, 563-581. doi: 10.1007/s13164-010-0039-7.

Roediger, H. L. (2008). Relativity of remembering: Why the laws of memory vanished. *Annual Review of Psychology*, 59, 225-254. doi:10.1146/annurev.psych.57.102904.190139

Sargent, S. L. (2007). Image Effects on Selective Exposure to Computer-Mediated News Stories. *Computers in Human Behavior* 23, 705–726. doi:10.1016/j.chb.2004.11.005

Schwarz, N. (2010). Meaning in context: Metacognitive experiences. In B. Mesquita, L. F. Barrett, & E. R. Smith (eds.), *The mind in context* (pp. 105 -125). New York: Guilford.

Schwarz, N. (in press). *APA Handbook of Personality and Social Psychology: Attitudes and Social Cognition*. Washington, DC: APA.

Sherman, S. J., Cialdini, R. B., Schwartzman, D. F., & Reynolds, K. D. (1985). Imagining can heighten

or lower the perceived likelihood of contracting a disease: The mediating effect of ease of imagery. *Personality and Social Psychology Bulletin*, 11, 118-127. doi: 10.1177/0146167285111011.

Snyder, M., & Cantor, N. (1979). Testing hypotheses about other people: The use of historical knowledge. *Journal of Experimental Social Psychology*, 75, 330-342. doi: 10.1016/0022-1031(79)90042-8.

Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments & Computers*, 31, 137-149. doi: 10.3758/BF03207704.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 18, 643– 62. doi: 10.1037/h0054651.

Thomas, A. K., Bulevich, J. B., & Dubois, S. J. (2012). An analysis of the determinants of the feeling of knowing. *Consciousness and Cognition*, 21(4), 1681-1694. doi:10.1016/j.concog.2012.09.005

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5, 207–232. doi: 10.1016/0010-0285(73)90033-9.

Unkelbach, C. (2006). The learned interpretation of cognitive fluency. *Psychological Science*, 17, 339-345. doi: 10.1111/j.1467-9280.2006.01708.x.

Unkelbach, C. (2007). Reversing the truth effect: Learning the interpretation of processing fluency in judgments of truth. *Psychological Science*, 20, 135-138. doi: 10.1037/0278-7393.33.1.219.

Unkelbach, C. & Greifeneder, R. (2013). A general model of fluency effects in judgment and decision making. In: C. Unkelbach & R. Greifeneder (Eds.), *The experience of thinking*. London: Psychology Press.

Westerman, D. (2008). Relative fluency and illusions of recognition memory. *Psychonomic Bulletin & Review*, 15, 1196-1200. doi:10.3758/PBR.15.6.1196

- Whittlesea, B. W. A. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19, 1235–1253. doi: 10.1037/0278-7393.19.6.1235.
- Whittlesea, B. W. A. (2011). Remembering under the influence of unconscious expectations. In P. A. Higham & J. P. Leboe (Eds.). *Constructions of remembering and metacognition: Essays in honor of Bruce Whittlesea* (pp. 225-236). Houndmills, UK: Palgrave Macmillan.
- Whittlesea, B. W. A., & Williams, L. D. (2001A). The discrepancy attribution hypothesis: I. The heuristic basis of feelings and familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 3-13. doi: 10.1037/0278-7393.27.1.3.
- Whittlesea, B. W. A., & Williams, L. D. (2001B). The discrepancy attribution hypothesis: II. Expectation, uncertainty, surprise, and feelings of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 14-33. doi: 10.1037/0278-7393.27.1.14.
- Wickens (2002). *Elementary Signal Detection Theory*. Oxford University Press: New York.
- Yue, C. L., Castel, A. D., & Bjork, R. A. (2013). When disfluency is-and is not-a desirable difficulty: The influence of typeface clarity on metacognitive judgments and memory. *Memory & Cognition*, 41, 229-241. doi: 10.3758/s13421-012-0255-8.

Figure Captions

Figure 1. Example of related and unrelated photos and associated trivia claims. Thermometer photo courtesy of Andres Rueda, Skink photo courtesy of William Cho, Creative Commons licenses.

Figure 2. Bias for trivia claims presented with or without a photograph that was semantically related or unrelated to the trivia claims. A negative value of c is a bias to say true. Error bars show 95% within-subject confidence intervals (see Masson & Loftus, 2003) for the photo/no-photo effect for Experiment 1 (within-subjects design) and 5 (mixed design) and the 95% between-subjects confidence intervals for each cell mean in Experiments 7 (between-subjects design)

Figure 3. Forest plot of effect sizes for related and unrelated photos between studies (Derzon & Alford, 2013). Each row represents one experiment, starting with the original experiment and then subsequent replications. The top panel of the plot displays the effects for related photos. The bottom panel of the plot displays the effects for unrelated photos. The location of each shape on the horizontal axis represents the raw effect size—the difference between people’s response bias when they saw claims paired with a (related/unrelated) photo compared to when they saw claims without photos. The lines extending either side of a shape represent a 95% confidence interval (note that a 95% CI for a mean difference that excludes zero will be significantly different from zero using null hypothesis testing). The black shape shows the result of the meta-analysis for each experiment, with the center of the shape indicating the estimated effect size, and again the black lines extending either side of a shape represent a 95% confidence interval. Finally, within each panel on the plot, data that fall to the right of the zero line show truthiness (nonprobative related photos lead people to believe a claim) and data that fall to the left side of the zero line show falsiness (nonprobative unrelated photos lead people to disbelieve a

claim). Tables presenting additional results from the meta-analyses can be found in supplementary materials.

Figure 4. Forest plot of mean bias (c) scores (an estimate from a meta-analysis for each experiment) for related photo, unrelated photo and no photo conditions, plotted for each experimental design. Tables presenting additional results from the meta-analyses can be found in supplementary materials.

Table 1. *Bias Data for Replications of Experiments 1, 5 and 7*

	Bias Related Photo	Bias Unrelated Photo	Bias No Photo
Experiment 1			
Experiment 2: Practice phase with unrelated photos	-.13 (.76)	.00 (.71)	.06 (.81)
Experiment 3: With no practice phase	-.17 (.67)	.01 (.70)	-.05 (.68)
Experiment 5			
Experiment 6			
Related Condition	-.16 (.45)	*	.04 (.46)
Unrelated Condition	*	-.02 (.39)	-.11 (.42)
Experiment 7			
Experiment 8	.00 (.43)	.01 (.42)	-.06 (.33)

Note. Standard deviations are in parentheses. Negative bias (*C*) values show a bias to say true.

Figure 1.

The liquid metal inside a thermometer
is magnesium



The liquid metal inside a thermometer
is magnesium



Figure 2.

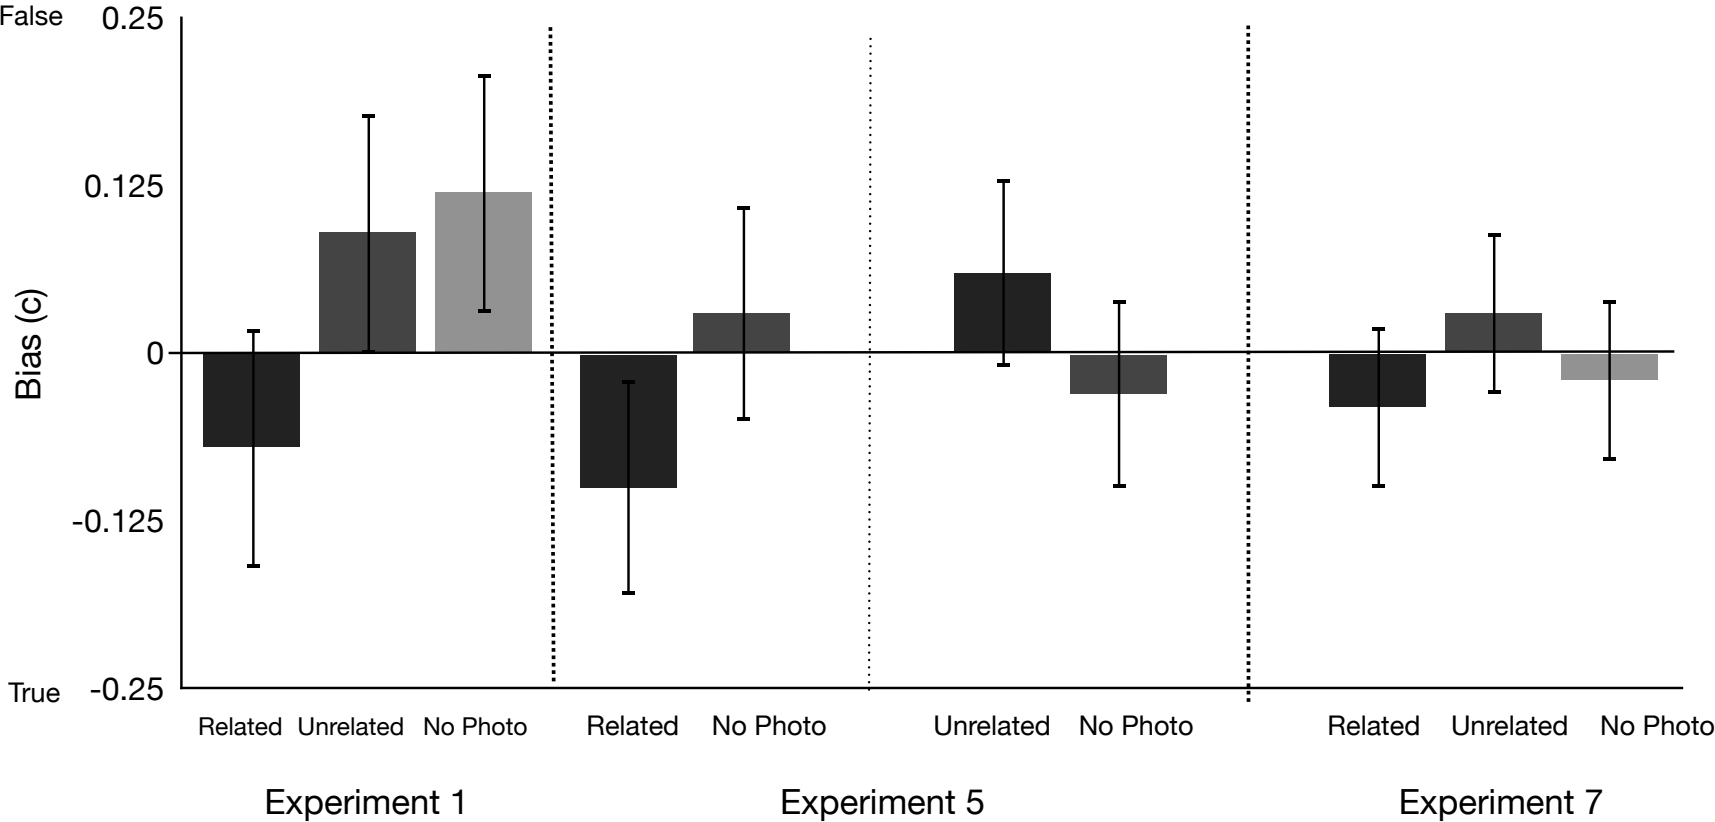


Figure 3.

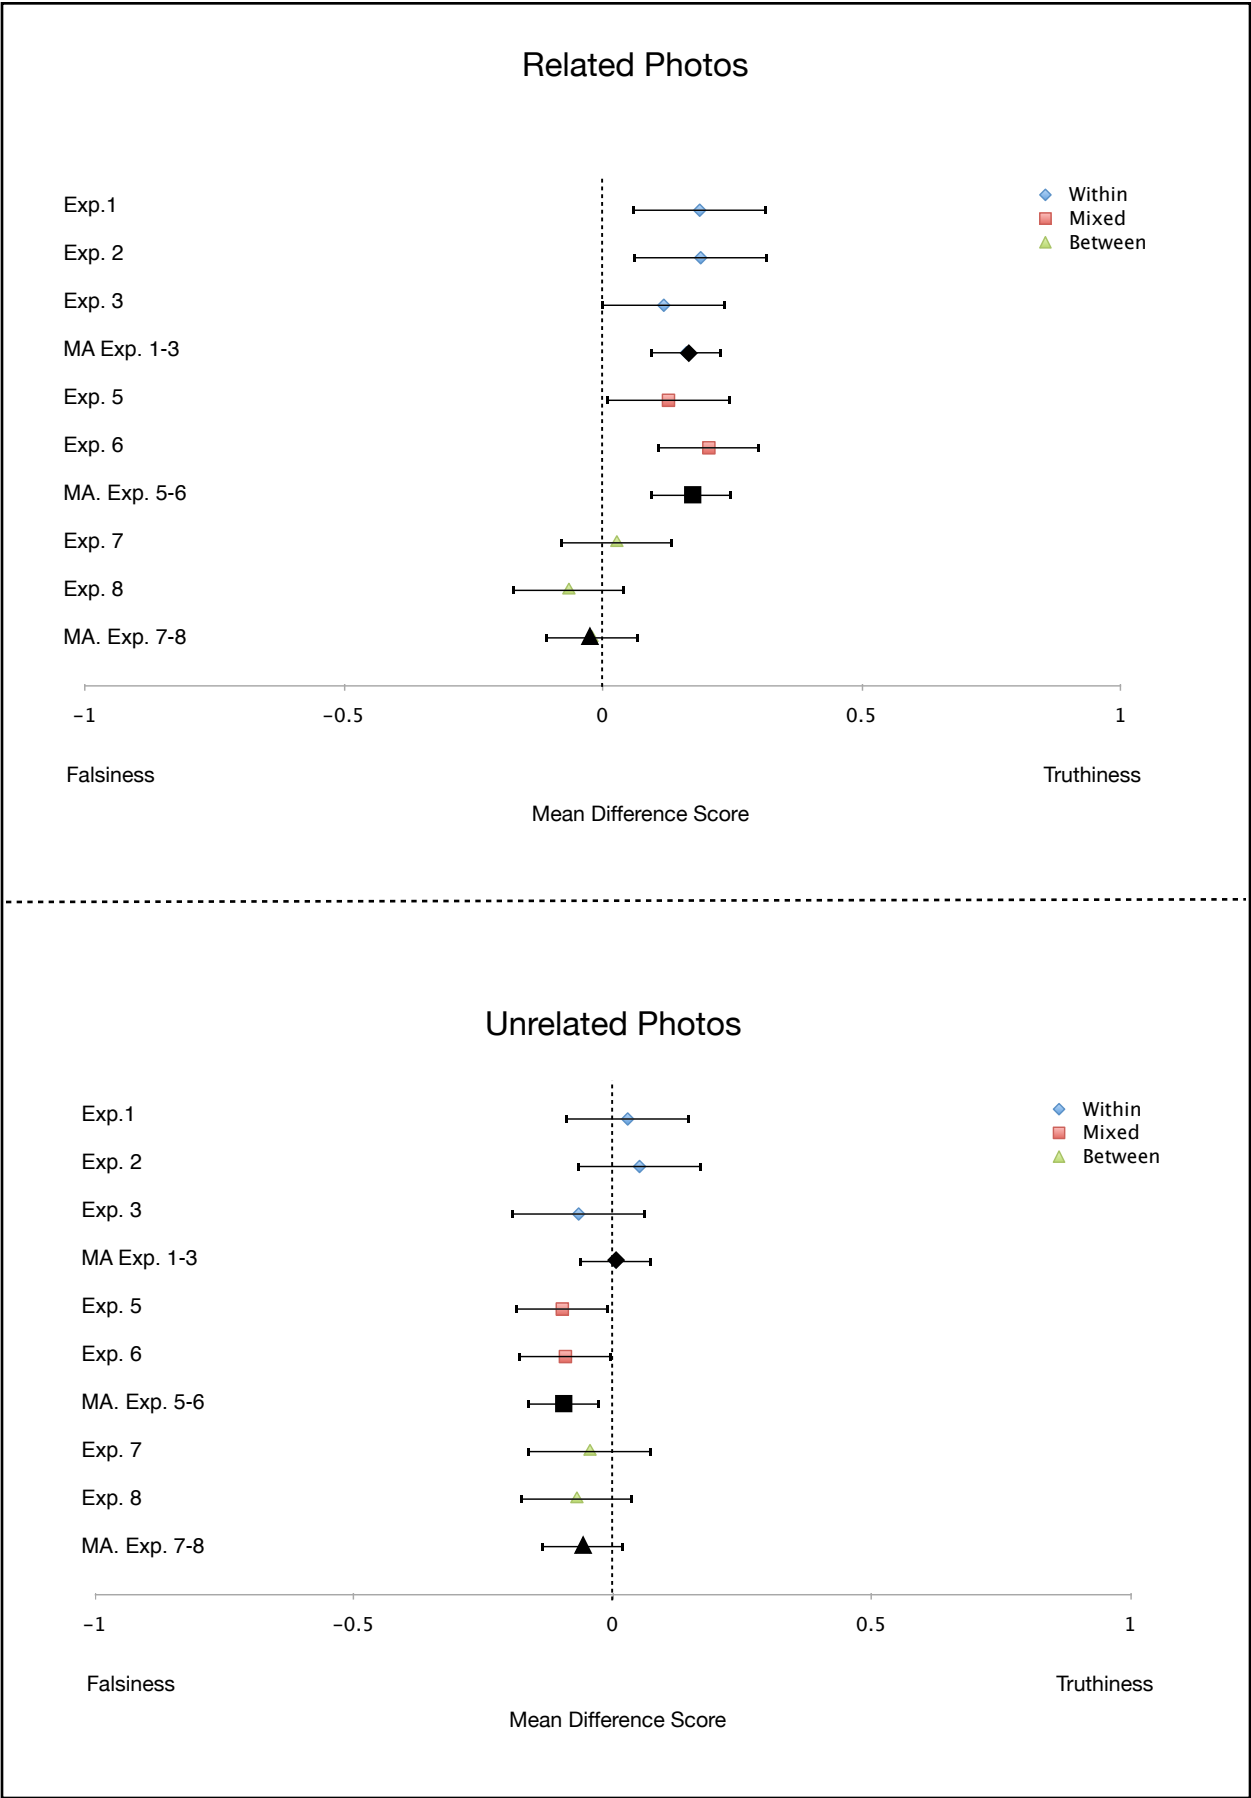



Figure 4.



Supplemental Materials

Table 1.

Discrimination (d') for claims presented with or without a related or unrelated photo.

	d'	Statistical Analyses
Experiment 1		
Related Photos	.22 (1.44)	A one way repeated measures ANOVA showed no main effect for photo $F(2, 206) = .36, p = .70, \eta_p^2 = .003$.
Unrelated Photos	.15 (1.19)	
No Photos	.26 (1.40)	
Experiment 5		
Related Condition		
Related Photos	.34 (0.74)	A 2(photo: photo vs. no photo) x 2(relatedness: related vs. unrelated) mixed ANOVA with repeated measures on the first factor showed no main effect for relatedness, $F(1, 194) = 2.16, p = .14, \eta_p^2 = .01$; no main effect for photos, $F(1, 194) = .55, p = .46, \eta_p^2 = .003$. and no Photo x Relatedness interaction, $F(1, 194) = 2.03, p = .16, \eta_p^2 = .01$.
No Photos	.18 (0.70)	
Unrelated Condition		
Unrelated Photos	.34 (0.72)	
No Photos	.39 (0.65)	
Experiment 7		
Related Photos	.24 (0.44)	A one way ANOVA showed no effect for photo, $F(2, 298) = .171, p = .18, \eta_p^2 = .01$.
Unrelated Photos	.28 (0.47)	
No Photos	.36 (0.53)	

Note. Standard deviations are in parentheses. Higher sensitivity (d') values indicate better discrimination ability.

Table 2. *Discrimination Analyses for Replications of Experiments 1, 5 and 7*

	d' Related Photo	d' Unrelated Photo	d' No Photo	Statistical Analyses
Experiment 1				
Replication 1: Practice phase with unrelated photos	.16 (1.30)	.28 (1.33)	.08 (1.37)	A one way repeated measures ANOVA showed no main effect for photo $F(2, 202) = 1.19, p = .31, \eta_p^2 = .01$.
Replication 2: With no practice phase	.18 (1.30)	.12 (1.29)	.22 (1.26)	A one way repeated measures ANOVA showed no main effect for photo $F(2, 214) = .35, p = .70, \eta_p^2 = .003$.
Experiment 5				
Replication				
Related Condition	.25 (.75)	*	.10 (.66)	A 2(photo: photo vs. no photo) x 2(relatedness: related vs. unrelated) mixed ANOVA with repeated measures on the first factor showed a Photo x Relatedness interaction, $F(1, 183) = 5.95, p = .02, \eta_p^2 = .03$; regardless of condition, when people saw a photograph paired with a statement their discrimination was similar, t related photo-unrelated photo(83) = .45, $p = .65$, but when people evaluated claims without photos, those people in the unrelated condition had better discrimination than people in the related photo conditions, t unrelated no photo-related no photo (183) = 3.73, $p < .01$.
Unrelated Condition	*	.30 (.67)	.47 (.70)	
Experiment 7				
Replication	.20 (.50)	.31 (.45)	.34 (.43)	A one-way ANOVA showed no significant effect for evidence, $F(2, 298) = .255, p = .08, \eta_p^2 = .02$.

Note. Standard deviations are in parentheses. Higher sensitivity (d') values indicate better discrimination ability. Considered together with the experiments reported in the main text, there were no consistent effects of photos on accuracy.

Table 3. *Results from meta-analyses estimating Truthiness (the difference in response bias between related and no photo conditions) for each experimental design.*

Experiment	95% CI			<i>z</i>	<i>p</i>
	ES	<i>LL</i>	<i>UL</i>		
1	0.19	0.05	0.32	2.77	0.01
2	0.19	0.06	0.32	2.78	0.01
3	0.12	-0.00	0.24	1.90	0.06
MA Experiments 1-3	0.16	0.09	0.23	4.27	0.00
5	0.13	0.01	0.24	2.12	0.03
6	0.20	0.11	0.30	4.08	0.00
MA Experiments 5 & 6	0.17	0.10	0.25	4.45	0.00
7	0.03	-0.09	0.14	0.45	0.65
8	-0.07	-0.17	0.04	-1.21	0.23
MA Experiments 7 & 8	-0.02	-0.11	0.07	-0.48	0.63

Note. ES = effect size, calculated as the difference between related photo and no photo means. *LL* and *UL* are the lower and upper limits of the 95% confidence interval of the effect size. Positive effect sizes show truthiness (nonprobative related photos lead people to believe a claim).

Table 4. *Results from meta-analyses estimating Falsiness (the difference in response bias between unrelated and no photo conditions) for each experimental design.*

Experiment	95% CI			<i>z</i>	<i>p</i>
	ES	<i>LL</i>	<i>UL</i>		
1	0.03	-0.09	0.14	0.47	0.64
2	0.05	-0.08	0.18	0.77	0.44
3	-0.07	-0.19	0.06	-1.03	0.30
MA Experiments 1-3	0.00	-0.07	0.08	0.13	0.90
5	-0.10	-0.19	0.01	-2.12	0.03
5	-0.09	-0.19	0.00	-1.94	0.05
MA Experiments 5 & 6	-0.10	-0.16	-0.03	-2.87	0.00
7	-0.04	-0.16	0.07	-0.76	0.45
8	-0.07	-0.18	0.04	-1.26	0.21
MA Experiments 7 & 8	-0.06	-0.14	0.02	-1.44	0.15

Note. ES = effect size, calculated as the difference between unrelated photo and no photo means. *LL* and *UL* are the lower and upper limits of the 95% confidence interval of the effect size. Negative effect sizes show falsiness (nonprobative unrelated photos lead people to disbelieve a claim).

Table 5. Mean bias (*c*) scores and meta-analyses comparing related photo, unrelated photo and no photo conditions to zero, for each experimental design.

Experiment		95% CI			<i>z</i>	<i>p</i>
		<i>M</i>	<i>LL</i>	<i>UL</i>		
No Photo Condition						
	Experiment 1	0.12	0.02	0.22	2.28	0.02
	Experiment 2	0.06	-0.05	0.17	1.06	0.29
	Experiment 3	-0.05	-0.14	0.04	-1.08	0.28
	Result of MA Experiments 1-3	0.04	-0.06	0.14	0.77	0.44
	Experiment 5/unrelated	-0.03	-0.12	0.03	-0.63	0.53
	Experiment 6/unrelated	-0.11	-0.19	-0.03	-2.63	0.01
	Result of MA Experiments 5 & 6/ unrelated	-0.07	-0.15	0.00	-1.84	0.07
	Experiment 5/related	0.03	-0.08	0.14	0.56	0.58
	Experiment 6/related	0.04	-0.06	0.14	0.80	0.43
	Result of MA Experiment 5 & 6/related	0.04	-0.04	0.11	0.97	0.33
	Experiment 7	-0.02	-0.10	0.06	-0.50	0.62
	Experiment 8	-0.06	-0.12	0.00	-1.83	0.07
	Result of MA Experiments 7 & 8	-0.04	-0.09	0.01	-1.73	0.08
Related Photo Condition						
	Experiment 1	-0.07	-0.18	0.04	-1.28	0.20
	Experiment 2	-0.13	-0.23	-0.03	-2.44	0.01
	Experiment 3	-0.17	-0.26	-0.08	-3.73	0.00
	Result of MA Experiments 1-3	-0.13	-0.19	-0.07	-4.42	0.00
	Experiment 5	-0.10	-0.20	0.00	-1.94	0.05
	Experiment 6	-0.16	-0.26	-0.06	-3.26	0.00
	Result of MA Experiments 5 & 6	-0.13	-0.20	-0.06	-3.70	0.00
	Experiment 7	-0.04	-0.11	0.03	-1.10	0.27
	Experiment 8	0.00	-0.08	0.08	0.00	1.00
	Result of MA Experiments 7 & 8	-0.02	-0.08	0.03	-0.84	0.40
Unrelated Photo Condition						
	Experiment 1	0.09	-0.01	0.19	1.75	0.08
	Experiment 2	0.00	-0.10	0.10	0.00	1.00
	Experiment 3	0.01	-0.08	0.10	0.21	0.83
	Result of MA Experiments 1-3	0.03	-0.02	0.09	1.10	0.27
	Experiment 5	0.06	-0.04	0.16	1.19	0.23
	Experiment 6	-0.02	-0.10	0.06	-0.51	0.60
	Result of MA Experiments 5 & 6	0.01	-0.06	0.09	0.34	0.73
	Experiment 7	0.03	-0.06	0.12	0.63	0.53
	Experiment 8	0.01	-0.07	0.09	0.24	0.81
	Result of MA Experiments 7 & 8	0.02	-0.04	0.08	0.60	0.55

Note. M = is a point estimate of response bias (c) in each condition, for each experimental design. LL and UL are the lower and upper limits of the 95% confidence interval of the effect size. Negative values of c (below zero) show a tendency to respond that a claim is true, and positive values (above zero) show a tendency to respond that a claim is false.

Table 6. *Hits and False Alarms for claims presented with or without a photo (Experiments 1-3)*

	Hits	FAs
Experiment 1		
Related Photos	0.55 (0.25)	0.49 (0.29)
Unrelated Photos	0.50 (0.25)	0.46 (0.25)
No Photos	0.51 (0.27)	0.43 (0.26)
Experiment 2		
Related Photos	0.55 (0.25)	0.51 (0.27)
Unrelated Photos	0.54 (0.24)	0.46 (0.27)
No Photos	0.50 (0.27)	0.47 (0.27)
Experiment 3		
Related Photos	0.58 (0.25)	0.52 (0.25)
Unrelated Photos	0.51 (0.23)	0.47 (0.26)
No Photos	0.55 (0.23)	0.48 (0.26)

Running Head: TRUTHINESS AND FALSINESS

Table 7. *Hits and False Alarms for claims presented with or without a photo (Experiments 5-6)*

	Hits	FAs
Experiment 5		
Related Condition		
Related Photos	0.59 (0.18)	0.48 (0.20)
No Photos	0.53 (0.21)	0.46 (0.20)
Unrelated Condition		
Unrelated Photos	0.54 (0.21)	0.42 (.17)
No Photos	0.58 (.17)	0.45 (0.19)
Experiment 6		
Related Condition		
Related Photos	0.60 (0.18)	0.51 (0.20)
No Photos	0.51 (0.19)	0.47 (0.19)
Unrelated Condition		
Unrelated Photos	0.56 (0.18)	0.46 (0.17)
No Photos	0.62 (0.19)	0.46 (0.18)

Table 8. *Hits and False Alarms for claims presented with or without a photo (Experiments 7-8)*

	Hits	FAs
Experiment 7		
Related Photos	0.56 (0.13)	0.47 (0.16)
Unrelated Photos	0.54 (0.16)	0.44 (0.19)
No Photos	0.57 (0.16)	0.44 (0.16)
Experiment 8		
Related Photos	0.54 (0.16)	0.47 (0.16)
Unrelated Photos	0.56 (0.17)	0.45 (0.16)
No Photos	0.59 (0.15)	0.46 (0.13)

Running Head: TRUTHINESS AND FALSINESS

Table 9. *Summary of designs across Experiments 1-8.*

Experiment	Design	Variations
Experiment 1:	Fully within-subjects: 3(photo: related, unrelated, no photo)	Half the practice items appeared with a related photo, and half with no photo
Experiment 2:	Fully within-subjects: 3(photo: related, unrelated, no photo)	Half the practice items appeared with an unrelated photo, and half with no photo
Experiment 3:	Fully within-subjects: 3(photo: related, unrelated, no photo)	There were no practice items.
Experiment 4:	Fully within-subjects: 3(photo: related, unrelated, no photo)	New materials. all statements were false and subjects responded on a likert scale to indicate truth, rather than making a dichotomous true/ false decision.
Experiment 5:	Mixed Design: 2(photo: yes, no) x 2(relatedness: related, unrelated)	
Experiment 6:	Photo was the within-subjects component Mixed Design: 2(photo: yes, no) x 2(relatedness: related, unrelated)	Exact replication of Experiment 5.
Experiment 7:	Photo was the within-subjects component Fully between-subjects: 3(photo: related, unrelated, no photo)	
Experiment 8:	Fully between-subjects: 3(photo: related, unrelated, no photo)	Exact replication of Experiment 7.