

## The “Saw-It-All-Along” Effect: Demonstrations of Visual Hindsight Bias

Erin M. Harley  
University of California, Los Angeles

Keri A. Carlsen and Geoffrey R. Loftus  
University of Washington

The authors address whether a hindsight bias exists for visual perception tasks. In 3 experiments, participants identified degraded celebrity faces as they resolved to full clarity (Phase 1). Following Phase 1, participants either recalled the level of blur present at the time of Phase 1 identification or predicted the level of blur at which a peer would make an accurate identification. In all experiments, participants overestimated identification performance of naive observers. Visual hindsight bias was greater for more familiar faces—those shown in both phases of the experiment—and was not reduced following instructions to participants to avoid the bias. The authors propose a fluency-misattribution theory to account for the bias and discuss implications for medical malpractice litigation and eyewitness testimony.

Imagine the following scenario: A patient at a local hospital undergoes routine chest radiography at Time 1, and Radiologist 1 interprets the films as normal. Three years later, the patient experiences discomfort and undergoes chest radiography a second time, whereupon Radiologist 2 discovers a large tumor. Despite treatment, the patient dies. The patient’s family files a medical malpractice lawsuit against Radiologist 1, claiming the tumor should have been detected at Time 1. Radiologist 2, testifying on the family’s behalf, views the original Time 1 radiographs and, seeing the tumor missed by Radiologist 1, claims that it was visible at Time 1.

Such a scenario is typical; the vast majority of medical malpractice litigation in radiology over the past 20 years has involved diagnostic, perceptual, or decision-making errors (Berlin, 1996a, 1996b; Berlin & Hendrix, 1998). In such cases, the defendant radiologist who is accused of negligence for a decision he or she made prospectively (without the benefit of outcome knowledge) is often judged by radiology experts who have full knowledge of what future radiographs revealed (Berlin, 2000). Are such experts biased by this knowledge? In other words, can a radiologist who views a tumor at Time 2 make an appropriate prediction about whether another radiologist should have detected the tumor when it was smaller and less visible at Time 1? This is a question of *visual hindsight*, which is akin to the verbal hindsight reported in the past by many investigators.

First reported by Fischhoff (1975), *hindsight bias*, or the *knew-it-all-along* effect (Wood, 1978), is the tendency for individuals

with outcome knowledge to claim that they would have estimated a higher probability of occurrence for the reported outcome than was estimated in foresight. In other words, it is the after-the-fact feeling that some outcome was very likely to happen, or was predictable, even though it was not predicted to happen beforehand. Hindsight bias is thought to result from cognitive reconstruction processes that occur after outcome information is received (for a review, see Hawkins & Hastie, 1990). Judges reanalyze the event so that the beginning and the middle connect causally to the end. During this process evidence consistent with the reported outcome is elaborated, and evidence inconsistent with the outcome is minimized or discounted. The result of this rejudgment process is that the given outcome seems inevitable or, at least, more plausible than alternative outcomes.

The verbal hindsight bias effect has been shown to be quite robust, occurring in both within-subject and between-subjects designs, in spite of explicit instructions to subjects to avoid the bias, and across a range of time intervals between initial judgments, outcome feedback, and second judgments (see Christensen-Szalanski & Willham, 1991; Fischhoff, 1982; Hawkins & Hastie, 1990, for reviews). The bias is sensitive to task difficulty—it is greater for difficult than it is for easy items; it is greater for events initially judged to be least plausible (Arkes, Wortmann, Saville, & Harkness, 1981; Fischhoff, 1977; Wood, 1978)—and is greater when evidence supporting the given outcome is more easily brought to mind (Sanna, Schwarz, & Small, 2002). Hindsight bias has been studied extensively in the cognitive domain, with broadly ranging types of events, for example, outcomes of historical events, psychiatric cases, scientific experiments, consumer purchases, sporting events, economic decisions, election outcomes, medical and legal cases, and answers to almanac trivia questions.

In contrast, there has been remarkably little research on the kind of visual hindsight bias exemplified by the radiology example sketched earlier in this article. Our main purpose in this article is to begin to correct this deficit. We first describe how hindsight bias might apply to visual perception, and in this context, introduce the broader topic of metaperception—people’s insights into their perceptual abilities. Next, we present two experiments that show evidence of visual hindsight bias, and propose a theory of

---

Erin M. Harley, Department of Psychology, University of California, Los Angeles; Keri A. Carlsen and Geoffrey R. Loftus, Department of Psychology, University of Washington.

The writing of this article and the laboratory described in it were supported by National Institutes of Mental Health Grant MH41637 to Geoffrey R. Loftus. We thank Daniel Bernstein for valuable comments on drafts of this article. We thank Chris Dessert, Eunice Poon, Sarah Hayes, and Erik Gust for assistance in data collection.

Correspondence concerning this article should be addressed to Erin M. Harley, Department of Psychology, University of California, Los Angeles, Box 951563, 1285 Franz Hall, Los Angeles, CA 90095-1563. E-mail: harley@psych.ucla.edu

fluency misattribution to account for the bias. Finally, we present a third experiment, in which a prediction of the fluency-misattribution theory is tested and confirmed.

### Visual Hindsight Bias

To illustrate how hindsight bias might apply to visual perception, we return to the medical scenario described in the introduction to this article. In the field of radiology, it is common in medical malpractice litigation for a physician with outcome knowledge to judge whether another physician who did not have the benefit of such knowledge should have detected an abnormality on a medical image (Berlin, 2000). When making this judgment, it may, upon initial inspection, seem reasonable for the judging physician to simply assess his or her own ability to see the missed abnormality. However, the physician cannot use this method if the missed abnormality is more easily detected in hindsight.

There is evidence suggesting that medical abnormalities are, in fact, more visible when the physician has outcome knowledge. A striking example of tumors that are evident only in hindsight comes from Muhm, Miller, Fontana, Sanderson, and Uhlenhopp (1983), who conducted a screening program at the Mayo Clinic for men at high risk of lung cancer. The 4,618 members of the study group obtained chest radiographs every 4 months, and each radiograph was read by two to three radiologists or chest physicians. Over the course of the 6-year study, 92 tumors were detected in the study group. Of these, 75 (82%) were, as the authors termed it, “visible in retrospect” (p. 611). This means that when the physicians looked at the previous sets of radiographs, they found that they could detect the tumor on at least one, and often on multiple, radiographs (ranging from 4 to 53 months prior to diagnosis) in 82% of the cases that had initially been interpreted as normal.

When speculating on why the 75 tumors visible in hindsight were not detected on initial readings, Muhm et al. (1983) made the following comment: “The fact that some of the cancers had been overlooked was usually due to *perceptual errors* [italics added] by the observers” (p. 612). Does it make sense that two to three experienced radiologists made errors on 82% of the tumor-containing radiographs? We argue that the error may well lie not in the physicians’ initial interpretations of the radiographs but, rather, in their retrospective interpretations of the radiographs. The “overlooked” tumors became visible in hindsight not because they were more detectable to begin with, but because the physicians had the benefit of outcome knowledge.

Multiple factors contribute to allow a radiologist with outcome knowledge, or hindsight, to detect abnormalities that previously could not be seen. Since Bartlett’s (1932) pioneering work and Neisser’s (1967) landmark *Cognitive Psychology*, visual perception has been construed as an active and creative process involving far more than a direct translation of images projected onto the retina. What we perceive is influenced by top-down processes that include prior knowledge, expectations, context, and a great number of assumptions about how objects in the world behave (see, e.g., Palmer, 1999). The radiologists who initially read the films in the Mayo Clinic’s screening program had relatively low expectations for finding a tumor<sup>1</sup>, and if a tumor was present, they had no knowledge of what type of tumor to search for or where the tumor might be located. In contrast, the postdetection viewing conditions were quite different. During retrospective analysis of the radiographs, the physicians had high expectations for finding a tumor

and full knowledge of both tumor type and tumor location, and, because the outcome was known, no longer had to be cautious about avoiding false positives. These factors combined to make visible many of the abnormalities that had originally gone undetected.

There is some research indicating that participants who know what they are looking for will have higher detection rates in a difficult visual search task. Bruner and Potter (1964) reported that participants claim to be able to perceive a visual target at a more degraded state when it is viewed in a clear-to-blurry progression than when it is viewed in a blurry-to-clear progression. Given that target information can improve a participant’s ability to detect or identify a visual target, a participant asked to estimate the performance of a naive peer is faced with a difficult challenge. The participant cannot simply assess his or her current detection ability—it has been enhanced by outcome information—but instead must discount target information and imagine what a naive observer would perceive.

### Metaperception

When asked to estimate the visual performance of a naïve self or peer, a participant must attempt to gain insight into the strengths and limitations of his or her perceptual abilities as well as how these abilities are affected by experience and knowledge. We will term this insight *metaperception*. Although a large amount of research has been dedicated to the investigation of metacognition, and most specifically, metamemory (see Chambres, Izaute & Marescaux, 2002; Mazzoni & Nelson, 1998; Metcalfe & Shimamura, 1994, for reviews), very few researchers have explored metaperception.

In recent work, Levin (2002) and Levin, Momen, Drivdahl, & Simons (2000) have investigated metaperception in relation to *change blindness*, the counterintuitive finding that participants commonly fail to detect large visual changes in their environment (see, e.g., Blackmore, Brelstaff, Nelson, & Troscianko, 1995; Henderson, 1997; Pashler, 1988; Phillips, 1974; Rensink, O’Regan, & Clark, 1997; for a review, see Simons, 2000). The existence of change blindness suggests that participants do not retain many visual details in memory from one view to the next and that focusing attention on the changing item is critical for successful change detection. Most relevant to this study is that change blindness is a startling finding in that it contradicts people’s intuitions about their perceptual abilities. Prior to experience with change-detection tasks, most participants wrongly assume that they will have no trouble detecting changes in visual scenes (Levin, Drivdahl, Momen, & Beck, 2002). For example, Levin et al. (2002) found that 90% of participants believed they would notice a scarf disappear on an actor across a movie edit that 0% of participants in the original experiment had noticed. This metacognitive error has been termed *change blindness blindness* and is evidence that under some conditions, naïve observers have grossly inaccurate insights into their own and others’ perceptual abilities. This conflict between actual visual performance and estimated visual performance warrants further investigation of metapercep-

<sup>1</sup> Even though the patients were at high risk for lung cancer, over the 6-year screening only 92 tumors were detected in a group of 4,618 patients (roughly 2% of the sample).

tion. Specifically, under what conditions will participants make inaccurate estimates of their perceptual abilities? Judging in hindsight may be one such case.

Only two studies of which we are aware have examined hindsight bias in the visual domain. Winman, Juslin, & Bjorkman (1998) found a reverse visual hindsight bias, in other words, an underestimate of performance, when participants with outcome information estimated prior, naive performance in a two-alternative-forced-choice line-length discrimination task. In contrast, Harley, Carlsen, & Loftus (2001) found that participants show positive hindsight bias when predicting the performance of their peers in a digit-identification task, but that they only did so when the task was most difficult. The conflicting results of these two studies suggest that one can neither assume a priori that hindsight bias will exist for all perceptual tasks, nor assume that if it does, it will be a positive bias like the traditional verbal hindsight bias, rather than a negative bias like the one found by Winman et al (1998).

Another reason one cannot assume a priori that visual hindsight bias will mirror verbal hindsight bias is that confidence in foresight knowledge judgments appears to trend differently in the intellectual and sensory domains. People tend to be overconfident when making intellectual judgments, for example, "Who was the third president of the United States", but underconfident when making sensory judgments, for example, "What color are your colleague's eyes?" (e.g., Adams, 1957; Bjorkman, Juslin, & Winman, 1993; Dawes, 1980; Keren, 1988; Olsson & Winman, 1996; Winman and Juslin, 1993). If hindsight confidence ratings are based on the same information as foresight confidence ratings—Winman et al. (1998) suggest that the two are both based on an assessment of task difficulty—then a reverse hindsight bias may prove more prevalent in the visual domain. If, on the other hand, hindsight bias is a general phenomenon that affects decisions made in multiple domains, including the visual domain, then we expect to find evidence for a positive visual hindsight bias, in other words, an overestimation of naive performance after target identity is known.

In a series of experiments reported here, we tested the hypothesis that hindsight bias is a general phenomenon—an overestimation of foresight knowledge following the receipt of outcome knowledge—that is not restricted to the intellectual domain but occurs in the visual domain as well. To test this hypothesis, we examined whether a participant with knowledge of target identity, akin to outcome knowledge, could accurately predict the level of visual degradation at which a naive self or peer would be able to identify a celebrity face. In all experiments, participants exhibited visual hindsight bias. The bias was exhibited for judgments made about both self and others, and despite education and warnings to avoid the bias. We propose a fluency-misattribution theory to account for visual hindsight bias and provide confirmatory evidence of a prediction of the theory in Experiment 3.

### Experiment 1

In Experiment 1, participants identified degraded pictures of celebrity faces as they gradually became clearer. Later, in a surprise memory test, participants recalled the degree of degradation present at the time of original identification.

### Method

**Participants.** Forty-two University of Washington undergraduates, all with normal or corrected-to-normal vision, participated in exchange for course credit.


**Apparatus.** Data collection took place in a room equipped with four Macintosh eMac computers—each of which was equipped with a G4 processor and a 17-in. monitor—allowing for up to 4 participants to participate in each data-collection session. Curtains were hung between computers to prevent participants from viewing other monitors during the session. Each data collection session lasted approximately 30 min. The experiment was written and executed in MATLAB using the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997).

**Stimuli.** Stimuli were 36 grayscale pictures of celebrity faces. The set included well-known actors, musicians, politicians, and sports figures, for example, Jerry Seinfeld, Harrison Ford, Madonna, Hillary Clinton, and Michael Jordon (see Figure 7 in Harley, Dillon, & Loftus, 2004, for sample images of celebrity photos). Each face measured 500 pixels from the bottom of the chin to the top of the head and subtended a visual angle of about 21.8° vertically. The display monitor's background luminance measured 2.94 cd/m<sup>2</sup>.

Thirty successively more blurred versions of each face were created. Each blurred version was accomplished by Fourier-transforming the image from pixel space into spatial-frequency space, multiplying the resulting frequency amplitude spectrum by a low-pass filter and inverse-Fourier transforming the result back into pixel space. The low-pass filter was designed such that for each blur level, it passed frequencies perfectly; in other words, it had a value of 1.0 up to some value of  $f_0$  cycles per face height and then fell parabolically, reaching zero at the cutoff frequency value of  $f_1 = f_0 \times 3$  cycles per face height. As  $f_0$  and  $f_1$  are made smaller, the filter cuts off more spatial frequencies, and the resulting face becomes blurrier (see Loftus, 2001, for an additional explanation). See Figure 1 for samples of a filtered face. In each of the three experiments reported here,  $f_1$ , the filter cutoff frequency, was used as the dependent variable—a measure of the degree of blur present in the image when the participant was able to (or believed he or she would be able to) identify the celebrity. A blurrier picture is implied by a smaller  $f_1$  value.

**Design and procedure.** Phase 1 of the experiment, *baseline identification* (baseline ID), was a simple identification test. For each face, the 30 blurred images were displayed, in order, from most to least blurred, at a rate of 500 ms per image. To the participant, it appeared as if the celebrity's face were becoming clearer slowly over time. The following instructions were read aloud to participants:

Each celebrity will start very blurry and slowly become clear. Press the space bar as soon as you recognize who the celebrity is. After you press the space bar, type in this guess and hit RETURN. After you enter your guess, the picture will continue to become clear. If your guess changes, hit the space bar again and type in a new guess. You can guess as many times as you wish until you are certain you have



**Figure 1.** Sample of a celebrity face like those used in Experiments 1-3. Shown is a subset of the 30 low-pass filtered images created for Harrison Ford with corresponding  $f_1$  (cycles per face height) filter cutoff values. This image is in the public domain and was retrieved from [www.celebrity-walpaper.com/harrisonford1.html](http://www.celebrity-walpaper.com/harrisonford1.html)

identified the celebrity correctly. If you get all the way to the clearest image and you don't know who the celebrity is just type in 'don't know' or a question mark.

Participants were allowed to identify a celebrity in a number of ways, including any portion of the celebrity's name, the name of a character he or she plays on television, the name of a movie in which he or she has starred, and so on. Anything that indicated to the investigators that the celebrity was recognized was scored as an accurate response. At the completion of each trial, the participant was asked to verify his or her final identity guess. For each celebrity face, all 30 blurred images were displayed regardless of whether or when the face was identified. This was done to equate as best as possible the total time a participant viewed each face. The order in which the 36 celebrities appeared was randomized for each participant.

In Phase 2 of the experiment, participants were given a surprise memory test. The same celebrities were shown again in a different, random order. At the start of each trial, the face was shown in its most degraded form, and the participant's response from Phase 1 was printed on the screen below the face so that the celebrity's identity would be known despite having been introduced in a degraded state. The following instructions were read aloud to participants: "Now you are going to perform a memory test . . . Use the arrow keys to adjust the blurriness of the celebrity until it matches what the celebrity looked like when you correctly identified him or her in the first half of the experiment." Participants were allowed to range back and forth among the 30 filters until they were satisfied with their decisions; no time limits were imposed.

Participants completed two practice trials prior to each of the two tasks: baseline ID and memory test. Celebrities shown in practice trials did not appear in the experiment proper.

### Results and Discussion

We found evidence for visual hindsight bias in Experiment 1. When asked to recall how blurry each face looked when identified in Phase 1 of the experiment, participants systematically overestimated the degree of blur. Only trials for which an observer correctly identified the celebrity during baseline ID were included in the analysis. The mean baseline-ID point was  $f_1 = 25.68$ , whereas the mean *estimated* identification point from the memory test was  $f_1 = 20.73$ , implying that the participants remembered identifying the celebrity in a blurrier state ( $f_1 = 20.73$ ) than was actually the case ( $f_1 = 25.68$ ). The hindsight ratio (HR) can be quantified as the ratio of these two numbers: baseline-ID  $f_1$  divided by memory-test  $f_1$ . To the extent that participants show hindsight bias, HR will be greater than 1, as is the case here:  $HR = 1.28 \pm 0.07$  (see Figure 2A).<sup>2</sup> The bias was not driven by a few outlying participants; 37 of the 42 participants (88%) showed visual hindsight bias.

As we mentioned in the introduction to this article, there is evidence that verbal hindsight bias is greater for difficult than it is for easy items (Fischhoff, 1977; Wood, 1978). To assess the influence of task difficulty on the magnitude of visual hindsight bias, celebrities were divided into four difficulty quartiles based on baseline-ID point. Faces with lower  $f_1$  identification values, in other words, those that were identified at a more degraded state, were considered easier than those with higher  $f_1$  identification values, in other words, those that were identified at a less degraded state. This was done separately for each participant, and the means for each quartile were then averaged across participants.

Hindsight ratios for the four difficulty quartiles are shown in Figure 2A. No bias was found for the easiest faces,  $HR = 0.99 \pm 0.05$ . For the remaining three difficulty quartiles, hindsight ratios increase as difficulty increases, ranging from a small bias for faces in difficulty quartile 2,  $HR = 1.18 \pm 0.06$ , to a large bias for the most difficult faces,  $HR = 1.72 \pm 0.19$ .

### Experiment 2

In the verbal domain, many researchers have tried to eliminate or, at least, reduce the hindsight bias effect, and most attempts to do this have been unsuccessful. Fischhoff (1977) found that neither educating participants about hindsight bias nor warning them to do everything they could to avoid the bias were successful in reducing the effect. In Experiment 2, we used a similar approach to test whether visual hindsight bias is cognitively impenetrable. A mental function is said to be cognitively impenetrable if it cannot be influenced by purely cognitive factors such as goals, beliefs, and inferences (Pylyshyn, 1980). Experiment 2 was a replication of Experiment 1 with one addition: Immediately prior to the memory test, observers were educated about visual hindsight bias and were warned to avoid it and perform as accurately as possible.

### Method

*Participants.* Fifty-four University of Washington undergraduates, all with normal or corrected-to-normal vision, participated in exchange for course credit.

*Apparatus and Stimuli.* All equipment and stimuli were identical to those used in Experiment 1.

*Design and Procedure.* The design of Experiment 2 was identical to that of Experiment 1 with one difference. Following the baseline-ID task and instructions on how to perform the memory test, the following instructions were read aloud to the participants:

When remembering how blurry each celebrity was when you first recognized him or her, I would like you to be aware of hindsight bias. Hindsight bias is when someone who knows the outcome of an event thinks they would have predicted that outcome before it happened. In previous versions of this experiment, your peers tended to believe that they recognized celebrities earlier than they did in Phase 1 of the experiment. We believe that already knowing who the celebrity is before viewing the clarification is what causes this effect. The result is that your peers think they recognized celebrities earlier, at a blurrier point, than they actually did. Please try to avoid this bias and be as accurate as possible when performing the memory test.

### Results and Discussion

The instructions and warning given to participants in Experiment 2 did not reduce the size of the hindsight bias. As was the case with Experiment-1 participants, Experiment-2 participants systematically overestimated the degree of blur when asked to recall baseline-ID performance.

Data, averaged across 54 participants, are shown in Figure 2B. Only trials for which an observer correctly identified the celebrity during the baseline-ID task were included in the analysis. The mean baseline-ID point was  $f_1 = 29.13$ , whereas the mean esti-

<sup>2</sup> The notation " $x \pm y$ " refers to a mean plus or minus a 95% confidence interval.

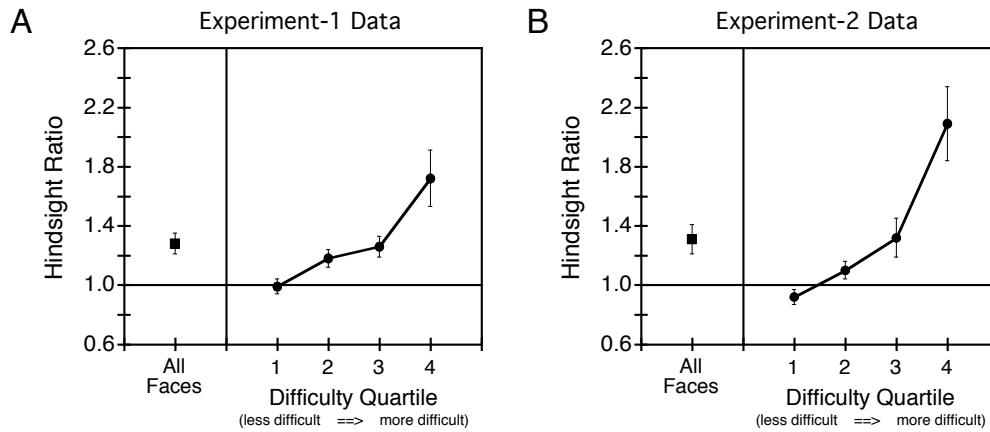


Figure 2. A: Experiment 1 data ( $N = 42$ ). Average hindsight ratios (baseline-ID  $f_i$  divided by memory-test  $f_i$ ) are plotted for all faces, and for faces divided into four difficulty quartiles. Error bars represent 95% confidence intervals. B: Experiment 2 data ( $N = 54$ ).

mated identification point from the memory test was  $f_i = 23.36$ . Recall that the HR can be quantified as the ratio of these two numbers: baseline-ID  $f_i$  divided by memory-test  $f_i$ ,  $HR = 1.31 \pm 0.10$ .

To assess the influence of task difficulty on the magnitude of the bias, we divided pictures of celebrities into four difficulty quartiles using the procedure used for the Experiment 1 data. The difficulty effect found in Experiment 1 was replicated in Experiment 2 (see Figure 2B). The pattern was identical to the pattern observed for the Experiment 1 data; no bias was found for the easiest faces,  $HR = 0.92 \pm 0.05$ , and for the remaining three difficulty quartiles, hindsight ratios increased as difficulty increased, ranging from a small bias for faces in Difficulty Quartile 2,  $HR = 1.10 \pm 0.06$ , to a large bias for the most difficult faces,  $HR = 2.09 \pm 0.25$ .

Note that there was no reduction in the hindsight bias effect observed in Experiment 2 when compared to the original Experiment 1 data. Educating observers about hindsight bias and warning them to avoid the bias and perform as accurately as possible were not effective in reducing bias. Experiment 2 data suggest that visual hindsight bias is similar to verbal hindsight bias in that they are both automatic, unconscious processes that participants are not easily able to control.

#### A Fluency-Misattribution Theory of Visual Hindsight Bias

When an observer views a visual stimulus, it is processed with some degree of *perceptual fluency* that reflects processing speed, effort, and accuracy. Stimulus variables such as clarity, long exposure duration, familiarity, and semantic relatedness can serve to enhance perceptual fluency. Jacoby and Whitehouse (1989) demonstrated that if participants are unaware of why fluency has been enhanced, they might misattribute the fluency to something else, for example, in their study, recent prior exposure to the stimulus. Since this work, research has shown that enhanced fluency can be misattributed to numerous factors in addition to prior exposure, such as stimulus duration, clarity, truth, and liking (for a review, see Winkielman, Schwarz, Reber, & Fazendeiro, 2003).

We propose that visual hindsight bias results from misattribution of enhanced perceptual fluency following the receipt of target-identity information. In the present experiments, processing of a degraded face after target identity was known (during the memory test) was enhanced compared with processing when it was not known (during the baseline-ID task). Participants had to ignore the enhanced fluency if they were to accurately estimate when a naive self or peer would make a correct identification. If participants failed to discount the enhanced fluency or did not discount enough, they may have mistakenly misattributed some or all of it to the predictability of the given outcome resulting in an overestimation of naive performance, in other words, hindsight bias (see Bernstein, Whittlesea & Loftus, 2002; Whittlesea & Williams, 2000; for similar accounts).

#### Experiment 3

The fluency-misattribution theory makes the following prediction: The more fluently a target is processed following the receipt of identity information, the larger the size of the hindsight effect should be. We designed Experiment 3 to test this prediction. Phase 1 was identical to the baseline-ID task used in Experiments 1 and 2: Participants viewed celebrity faces as they became clearer over time, and stopped the process when identification of the face was possible. Phase 2 differed from Experiments 1 and 2. Instead of a memory test, in Experiment 3 we used what is referred to as a “hypothetical hindsight design,” in which participants with outcome information predict the performance of a naive peer. Briefly, in the hindsight task, participants viewed an *outcome stimulus*—an unfiltered version of the face—followed by the same clarification process used during the baseline-ID task. Participants stopped the clarification when they thought that a naive peer, someone who did not see the outcome stimulus, would be able to identify the celebrity.

Critically, two sets of faces were shown in phase 2: the faces shown during Phase 1 (*old faces*), and *new*, previously unseen faces. We expected participants to show hindsight bias for both types of faces because processing fluency during the hindsight task would be enhanced as a result of viewing the outcome stim-

uli. However, we expected a larger bias for old faces compared with new faces, because processing fluency for old faces would receive additional enhancement from participants' viewing those faces during the baseline-ID task.

Experiment 3 also allowed for a test of the cognitive reconstruction theories proposed to account for verbal hindsight bias. Note that for the new faces shown in the hindsight task, participants viewed the outcome stimulus prior to any exposure to the degraded face. Cognitive reconstruction theories cannot account for bias in an outcome-first presentation because at the time at which outcome information is received, there is no prior evidence to be reworked or rejudged. If hindsight bias is observed for the new faces, rejudgment processes cannot be necessary for the production of visual hindsight bias.

### Method

**Participants.** Fifty-three University of Washington undergraduates, all with normal or corrected-to-normal vision, participated in exchange for course credit.

**Apparatus and Stimuli.** All equipment and stimuli were identical to those used in Experiment 1.

**Design and Procedure.** Each participant participated in a baseline-ID task followed by a hindsight task. The baseline-ID task was identical to that used in Experiments 1 and 2. Briefly, each face progressed from highly degraded to full clarity over 15 s, and participants stopped the resolution process as soon as they recognized the face.

Prior to completing the hindsight task, the following instructions were read aloud to participants:

In the second phase of the experiment, you are going to see the same celebrities again plus some new ones. Instead of indicating at what point you recognize the celebrity, your task this time will be to estimate at what point one of your peers would recognize the celebrity. So that you know the correct answer, we will show you a clear version of the celebrity at the beginning of each trial. If you know who they are, type that in. If you do not know who the celebrity is, type in a question mark, and it will skip to the next trial. After you identify the clear picture of the celebrity, the face will go blurry and slowly resolve—just like in Phase 1. Now imagine that a same-age peer is seeing the face for the first time, meaning they did not see the clear picture first. Press the space bar when you think your peer would recognize who the person is.

When the participant stopped the clarification process, the following question appeared on the screen, *Is this the point at which you think your peer would recognize this person?* If the participant entered “y” for yes, the trial ended. If the participant answered “n” for no, the clarification process continued. Participants were allowed to stop the process as many times as necessary until they were satisfied with the degree of blur present in the image.

Half (18 of 36) of the celebrities were shown in the baseline-ID task. The 18 baseline faces plus 18 new faces were mixed and shown in the hindsight task. The choice of which celebrities were shown twice (i.e., shown in both the baseline-ID and hindsight tasks) and the order of the 64 celebrities across trials were counterbalanced across participants. Participants completed two practice trials prior to each of the two tasks. Celebrities shown in practice trials did not appear in the experiment proper.

### Results and Discussion

Experiment 3 data are shown in Figure 3. Only trials for which a participant could correctly identify the celebrity were included in the analysis. Recall that the average  $f_1$  value for the baseline-ID condition represents the average degree of blur present at time of

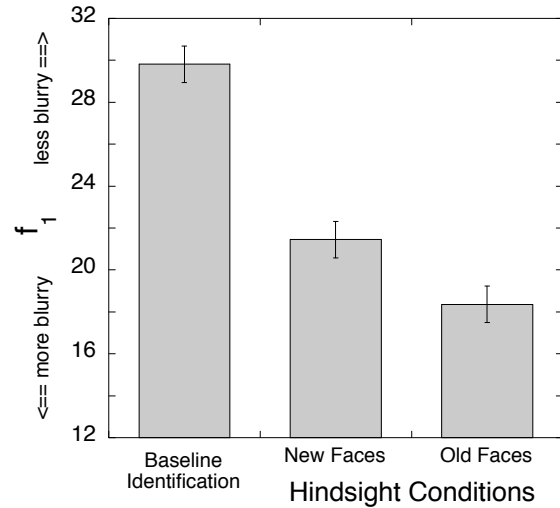


Figure 3. Experiment 3 data ( $N = 53$ ). Performance in the three conditions is plotted as a function of  $f_1$ , filter cutoff frequency expressed as cycles per face height. Error bars represent standard errors.

correct identification, whereas average  $f_1$  values for the two hindsight conditions represent the degree of blur present when participants predicted a naive peer would make an accurate identification. As predicted by the fluency-misattribution theory, average  $f_1$  identification point was greatest in the baseline-ID condition ( $M = 29.8$ ), followed by the new faces in hindsight ( $M = 21.4$ ), followed by old faces in hindsight ( $M = 18.4$ ).

Two orthogonal planned comparisons (C1 and C2) were tested to compare the three conditions: Baseline-ID performance was compared to the two hindsight conditions (C1), and the two hindsight conditions were compared with each other (C2). For each participant, a ratio of the weighted conditions was computed to test each comparison (the sum of positive-weighted conditions divided by the sum of negative-weighted conditions), and the mean of the ratios was computed. Compared with baseline-ID, participants demonstrated hindsight bias for both the old and new faces,  $M_{C1} = 1.45 \pm 0.11$ . Additionally, as predicted by the fluency-misattribution theory, the hindsight effect was larger for old than it was for new faces,  $M_{C2} = 1.08 \pm 0.02$ .

### General Discussion and Implications

Data from Experiments 1-3 provide evidence for visual hindsight bias; participants operating with the benefit of target-identity information may be biased when asked to estimate the performance of a naive peer or self. It was found for judgments made about both self (Experiments 1 and 2) and others (Experiment 3), and despite education and explicit instructions to participants to avoid the bias (Experiment 2). The bias was larger when postoutcome processing was more fluent and when preoutcome identification was more difficult. We now discuss these last two findings in more detail.

### *Fluency-Misattribution Theory*

We have proposed a theory of fluency misattribution to account for visual hindsight bias. Exposure to outcome information increases the perceptual fluency with which a degraded image is processed. This fluency must be discounted if a participant is to accurately predict identification performance of a naive self or peer. If the participant fails to fully discount the enhanced fluency, it is misattributed to the predictability of the outcome. The theory predicts larger hindsight bias for targets processed more fluently following the receipt of identity information. This prediction was confirmed in Experiment 3, in which a larger bias was found for old faces—those shown both in the baseline-ID task and the hindsight task—compared with new faces shown only in hindsight.

Jacoby and Whitehouse (1989) found that participants who were aware of why fluency had been enhanced were able to discount the fluency. Only when participants were made unaware of the source of the enhanced fluency did they misattribute it to something else. In the experiments reported here, participants may have had some awareness that processing fluency was enhanced via their learning the identities of the celebrities during the baseline-ID task. Although the data clearly indicate that participants did not successfully discount all of the enhanced fluency, it is possible that participants discounted some of the enhanced fluency. If the source of increased fluency were to be made more obscure, a larger bias might be found. In fact, decreased awareness of the source of enhanced fluency may have contributed to the larger bias observed for old faces compared with that observed for new faces in Experiment 3.

If hindsight bias is a general metacognitive error to which all modalities are vulnerable, as we believe it is, then the fluency-misattribution theory should account for verbal hindsight bias in addition to visual hindsight bias. There is some evidence that this may be the case. Using different terminology, Sanna et al. (2002) suggest that fluency plays a large role in verbal hindsight bias. The authors demonstrated that hindsight bias is reduced when evidence supporting alternative outcomes is easy to bring to mind, but it is increased when such evidence is difficult to bring to mind. They term the ease or difficulty with which these thoughts are brought to mind subjective accessibility, which, we argue, is really just another way of describing fluency. The outcome, and evidence for it, that is processed more fluently is seen as more likely and more predictable. In the case of verbal hindsight bias, fluency is conceptual rather than perceptual. This work suggests that a fluency-misattribution theory may account for verbal as well as visual hindsight bias. Further studies are warranted to examine this possibility.

### *Task Difficulty*

As mentioned in the introduction to this article, there is evidence that verbal hindsight bias is greater for difficult than it is for easy items and greater for events initially judged to be least plausible (Arkes et al., 1981; Fischhoff, 1977; Wood, 1978). In the visual domain, Harley et al. (2001) found that visual hindsight bias only occurred for the most difficult-to-detect targets. The results reported here are consistent with these findings. For Experiments 1 and 2, no hindsight effect was observed for the easiest faces, and the effect then increased monotonically as the degree of foresight identification difficulty increased.

One can account for the influence of task difficulty on the size of visual hindsight bias with the fluency-misattribution theory if one assumes that outcome information is more beneficial to the processing of difficult targets. To illustrate, prior to the receipt of target-identity information, some degraded targets will be more difficult to identify than others. Following the receipt of target-identity information, all of the images become identifiable at a more degraded state. If target-identity information is more beneficial for an item that was originally difficult to identify compared with one that was not difficult, then the discrepancy between baseline processing fluency and hindsight processing fluency will be larger for more difficult items. The greater this discrepancy, the more fluency a participant must discount to make an accurate prediction about a naive observer's ability. As evidenced by Experiment 1 and Experiment 2 data, participants do adjust appropriately for the easiest faces, those for which the outcome information presumably provides the least benefit, but the adjustment falls shorter and shorter as difficulty increases.

One could argue that the difficulty effect observed in Experiments 1 and 2 resulted from participants' failure to stop the clarification process at the point of recognition due to an inability to remember a celebrity's name. Recall that faces were divided into difficulty quartiles on the basis of the point at which they were identified during the baseline-ID task; faces identified later, in other words, in a clearer state, were categorized as more difficult than those identified earlier. It is possible that faces identified later (coded as more difficult) were faces for which participants recognized the celebrity but had trouble recalling a name or other identifying remark. In such cases, participants may have let the clarification process continue until a name could be recalled. Although this account is plausible, we have two reasons to believe that participants were not performing the task in this manner. First, participants were instructed to, "press the space bar as soon as you recognize who the celebrity is." They were also told that if their guess changed, they could stop the clarification process again and that there would be no penalty for guessing multiple times. Second, in watching participants perform the baseline-ID task, we noted that they appeared to stop the clarification process as soon as the face was recognizable and then to wait until a name or identifying remark could be generated before continuing. Given, however, that our second piece of evidence is anecdotal and that we cannot know whether all participants followed the directions, the alternative account of the difficulty effect cannot be ruled out.

### *Legal Implications*

Visual hindsight bias has a number of important legal implications. We return once again to the medical malpractice scenario described at the outset of this article. When asked to estimate whether Radiologist 1 should have detected a tumor at Time 1, Radiologist 2 should proceed with extreme caution. Not only will the tumor be more visible to an observer operating with the benefit of outcome knowledge (e.g., Muhm et al., 1983), but also, on the basis of the results reported here, it is highly likely Radiologist 2 will not fully discount this benefit, and as a result, will overestimate the detection ability of a naive observer.

Physicians judging the visibility of missed tumors are not the only legal players who may be susceptible to visual hindsight bias. Eyewitnesses to a crime are often asked to judge their perceptual abilities, and in doing so, may overestimate their ability to detect

or identify a visual target under poor viewing conditions. For example, if Warren Witness views someone fleeing from the scene of a crime 10 m away on a dark and rainy night, his memory of that person's face is likely to be poor. If Warren is later shown a clear picture of Sam Suspect in a police photo lineup (akin to outcome knowledge), Warren may integrate the clear picture of Sam with the original degraded image he has stored in memory. The result will be a "cleaned up" memory representation that more closely resembles Sam Suspect. This is clearly problematic, as Sam may or may not be the criminal Warren saw, but the trouble does not end here. When asked to testify against Sam Suspect, Warren Witness may overestimate his ability to have identified Sam under the original poor viewing conditions.

The notion that feedback can influence an eyewitness's estimate of original viewing conditions is not new. Verbal confirmatory postidentification feedback has been shown not only to inflate witnesses' confidence in the accuracy of their identification but also to distort their estimates of the original witnessing conditions (Bradfield, Wells, & Olson, 2002; Wells & Bradfield, 1998). Goodness of view, speed of identification, amount of attention paid to the suspect's face, and clarity of memory for the suspect are given inflated ratings by eyewitnesses who received confirmatory feedback following identification of a suspect. The data reported here add to the mounting evidence that outcome information, be it verbal or visual, can distort an eyewitness's beliefs about the original viewing conditions.

### Concluding Remarks

We have provided evidence that, under certain conditions, observers operating with the benefit of hindsight do not have accurate insights into the strengths and limitations of their perceptual abilities. This is evidenced by their failure to accurately predict the performance of a naive peer or self in visual identification tasks. Like verbal hindsight bias, visual hindsight bias appears to be moderated by task difficulty, with a greater bias occurring for more difficult (i.e., ambiguous) images, and appears to be cognitively impenetrable. We proposed a fluency-misattribution theory to account for the bias. The theory posits that exposure to target-identity information results in enhanced processing fluency of the degraded image. When asked to judge the performance of a naive observer, the increased fluency is not fully discounted and instead may be misattributed to the predictability of the given outcome. The theory predicts a larger hindsight effect for more fluently processed targets. This prediction was confirmed in Experiment 3, in which we observed a larger bias for targets made more familiar via repeated exposures.

### References

- Adams, J. K. (1957). A confidence scale defined in terms of expected percentages. *American Journal of Psychology*, *70*, 432-436.
- Arkes, H. R., Wortmann, R. L., Saville, P. D., & Harkness, A. R. (1981). Hindsight bias among physicians weighing the likelihood of diagnoses. *Journal of Applied Psychology*, *66*, 252-254.
- Bartlett, F. C. (1932). *Remembering: a study in experimental and social psychology*. Oxford, England: Macmillan.
- Berlin, L. (1996a). Malpractice issues in radiology: Perceptual errors. *American Journal of Roentgenology*, *167*, 587-590.
- Berlin, L. (1996b). Errors in judgment. *American Journal of Roentgenology*, *166*, 1259-1261.
- Berlin, L. (2000). Malpractice issues in Radiology: Hindsight bias. *American Journal of Roentgenology*, *175*, 597-601.
- Berlin, L., & Hendrix, R. W. (1998). Malpractice issues in radiology: Perceptual errors and negligence. *American Journal of Roentgenology*, *170*, 863-867.
- Bernstein, D. M., Whittlesea, B. W. A., & Loftus, E. F. (2002). Increasing confidence in remote autobiographical memory and general knowledge: Extensions of the revelation effect. *Memory and Cognition*, *30*, 432-438.
- Bjorkman, M., Juslin, P., & Winman, A. (1993). Realism of confidence in sensory discrimination: The underconfidence phenomenon. *Perception and Psychophysics*, *54*, 75-81.
- Blackmore, S. J., Brelstaff, G., Nelson, K., & Troscianko, T. (1995). Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception*, *24*, 1075-1081.
- Bradfield, A. L., Wells, G. L., & Olson, E. A. (2002). The damaging effect of confirming feedback on the relation between eyewitness certainty and identification accuracy. *Journal of Applied Psychology*, *87*, 112-120.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433-436.
- Bruner, J. S., & Potter, M. C. (1964). Interference in visual recognition. *Science*, *144*, 424-425.
- Chambres, P., Izaute, M., & Marescaux, P. (Eds.). (2002). *Metacognition: Process, function and use*. Boston, MA: Kluwer Academic Publishers.
- Christensen-Szalanski, J. J., & Willham, C. F. (1991). The hindsight bias: A meta-analysis. *Organizational Behavior and Human Decision Processes*, *48*, 147-168.
- Dawes, R. M. (1980). Confidence in intellectual vs. confidence in perceptual judgments. In E. D. Lanterman & H. Feger (Eds.), *Similarity and choice: Papers in honour of Clyde Coombs* (pp. 327-245). Bern, Switzerland: Huber.
- Fischhoff, B. (1975). Hindsight  $\neq$  foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, *1*, 288-299.
- Fischhoff, B. (1977). Perceived informativeness of facts. *Journal of Experimental Psychology: Human Perception and Performance*, *3*, 349-358.
- Fischhoff, B. (1982). For those condemned to study the past: Heuristics in hindsight. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge Univ. Press.
- Harley, E. M., Carlsen, K. A., & Loftus, G. R. (2001, November). *Is hindsight 20/20?: Evidence for hindsight bias in visual perception*. Poster session presented at the annual meeting of the Psychonomic Society, Orlando, FL.
- Harley, E. M., Dillon, A. M., & Loftus, G. R. (2004). Why is it difficult to see in the fog? How contrast affects visual perception and visual memory. *Psychonomic Bulletin & Review*, *11*(2), 197-231.
- Hawkins, S. A., & Hastie, R. (1990). Hindsight: Biased judgments of past events after the outcomes are known. *Psychological Bulletin*, *107*, 311-327.
- Henderson, J.M. (1997). Transsaccadic memory and integration during real-world object perception. *Psychological Science*, *8*, 51-55.
- Jacoby, L. L., & Whitehouse, K. (1989). An illusion of memory: False recognition influenced by unconscious perception. *Journal of Experimental Psychology: General*, *118*, 126-135.
- Keren, G. (1988). On the ability of monitoring non-veridical perceptions and uncertain knowledge: Some calibration studies. *Acta Psychologica*, *67*, 95-119.
- Levin, D. T. (2002). Change blindness blindness as visual metacognition. *Journal of Consciousness Studies*, *9*, 111-130.
- Levin, D. T., Drivdahl, S. B., Momen, N., & Beck, M. R. (2002). False predictions about the detectability of visual changes: The role of beliefs about attention, memory, and the continuity of attended objects in



- causing change blindness blindness. *Consciousness & Cognition: An International Journal*, *11*, 507-527.
- Levin, D. T., Momen, N., Drivdahl, S. B., & Simons, D. J. (2000). Change blindness blindness: The metacognitive error of overestimating change-detection ability. *Visual Cognition*, *7*, 397-412.
- Loftus, G. R. (2001, November). *It's Dennis Quaid; oops, it's Tom Cruise: Perceptual interference in face recognition*. Paper presented at the annual meeting of the Psychonomic Society, Orlando, FL.
- Mazzoni, G., & Nelson, T. O. (Eds.). (1998). *Metacognition and cognitive neuropsychology: Monitoring and control processes*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Metcalfe, J., & Shimamura, A. P. (Eds.). (1994). *Metacognition: Knowing about knowing*. Cambridge, MA: The MIT Press.
- Muhm, J. R., Miller, W. E., Fontana, R. S., Sanderson D. R., & Uhlenhopp, M. A. (1983). Lung cancer detected during a screening program using four-month chest radiographs. *Radiology*, *148*, 609-615.
- Neisser, U. (1967). *Cognitive psychology*. New York, NY: Prentice Hall.
- Olsson, H., & Winman, A. (1996). Underconfidence in sensory discrimination: The interaction between experimental setting and response strategies. *Perception and Psychophysics*, *58*(3), 374-382.
- Palmer, S. E. (1999). *Vision science: Photons to phenomenology*. Cambridge, MA: The MIT Press.
- Pashler, H. (1988). Familiarity and visual change detection. *Perception and Psychophysics*, *44*, 369-378.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437-442.
- Phillips, W.A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception and Psychophysics*, *16*, 283-290.
- Pylyshyn, Z. W. (1980). Computation and cognition: Issues in the foundations of cognitive science. *Behavioral and Brain Sciences*, *3*, 111-169.
- Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, *8*, 368-373.
- Sanna, L. J., Schwarz, N., & Small, E. M. (2002). Accessibility experiences and the hindsight bias: I knew it all along versus it could never have happened. *Memory and Cognition*, *30*, 1288-1296.
- Simons, D.J. (2000). Current approaches to Change Blindness. *Visual Cognition*, *7*, 1-15.
- Wells, G. L., & Bradfield, A. L. (1998). "Good, you identified the suspect." Feedback to eyewitnesses distorts their reports of the witnessing experience. *Journal of Applied Psychology*, *83*, 360-376.
- Whittlesea B. W. A., & Williams, L. D. (2000). The source of feelings of familiarity: The discrepancy-attribution hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 547-565.
- Winkielman, P., Schwarz, N., Reber, R., & Fazendeiro, T. A. (2003). Cognitive and affective consequences of visual fluency: When seeing is easy on the mind. In L. M. Scott, & R. Batra (Eds.), *Persuasive Imagery: A consumer response perspective* (pp. 75-89). Mahwah, NJ: Lawrence Erlbaum Publishers.
- Winman, A., & Juslin, P. (1993). Calibration of sensory and cognitive judgments: Two different accounts. *Scandinavian Journal of Psychology*, *34*, 135-148.
- Winman, A., Juslin, P., & Bjorkman, M. (1998). The confidence-hindsight mirror effect in judgment: An accuracy-assessment model for the knew-it-all-along phenomenon. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 415-431.
- Wood, G. (1978). The knew-it-all-along effect. *Journal of Experimental Psychology: Human Perception and Performance*, *4*, 345-353.

Received February 10, 2003

Revision received February 24, 2004

Accepted March 15, 2004