

Can Eyewitnesses Correct for External Influences on Their Lineup Identifications? The Actual/Counterfactual Assessment Paradigm

Steve D. Charman and Gary L. Wells
Iowa State University

Real-world eyewitnesses are often asked whether their lineup responses were affected by various external influences, but it is unknown whether they can accurately answer these types of questions. The witness-report-of-influence mental-correction model is proposed to explain witnesses' reports of influence. Two experiments used a new paradigm (the actual/counterfactual paradigm) to examine eyewitnesses' abilities to report accurately on the influence of lineup manipulations. Eyewitnesses were administered either confirming feedback or no feedback (Experiment 1, $n = 103$), or a cautionary instruction or no cautionary instruction (Experiment 2, $n = 114$). Eyewitnesses then gave actual responses (retrospective confidence, view, and attention measures in Experiment 1; identification decision in Experiment 2) as well as counterfactual responses stating how they would have responded in the alternative condition. Results across both studies showed an asymmetric estimation of influence pattern: Eyewitnesses who received an influencing manipulation estimated significantly less of a change in their responses than eyewitnesses who did not receive an influencing manipulation. A 48-hr delay between actual and counterfactual responses did not moderate any effects. Results are explained by witnesses' implicit theories of influence.

Keywords: eyewitness identification, eyewitness testimony, lineups, mental correction

Supplemental materials: <http://dx.doi.org/10.1037/1076-898X.14.1.5.supp>

Imagine that you are an eyewitness to a crime and are shown a lineup. You identify lineup member No. 3, after which the lineup administrator tells you "Good! That's our suspect!" You then tell the administrator that you are 90% confident that No. 3 is, in fact, the criminal. Based on your identification, formal charges are brought against No. 3. Months later, at trial, the defense attorney for the suspect argues to the judge that, because the lineup administrator told you that you had identified the suspect before you gave a confidence statement, you were unduly influenced and thus your confidence in your identification should be suppressed. The judge, not wanting to dismiss your confidence statement so readily, tells the defense attorney that he should simply ask you at trial whether the administrator's remark influenced your confidence statement. The defense attorney then asks you this question at trial and you begin to think. What confidence statement would you have given had you not been told that you identified the suspect? Would you have still been 90% confident? If this statement boosted your confidence, by how much did it boost it? Can

you give an accurate report on whether, and how, the lineup administrator's statement influenced you?

Self-reports of this type tend to be met with skepticism from people knowledgeable about the workings of human cognition and metacognition. A substantial body of psychological literature casts doubt on people's abilities to determine whether their judgments, attitudes, or behaviors were influenced by extraneous variables, as well as on their abilities to correct for any such influence so readily (Nisbett & Wilson, 1977; Stapel, Martin, & Schwarz, 1998; Wegener & Petty, 1995; Wilson & Brekke, 1994). According to many of these researchers, people have little if any introspective awareness into their cognitive processes. Because of this lack of introspective awareness, when people think that they have been influenced and must subsequently attempt to correct for that influence, they are forced to rely on their implicit theories—their beliefs about how objects and events are causally related to one another—to determine how they think the influencing variable affected them (Petty & Wegener, 1993; Wegener & Petty, 1995; Wilson, Centerbar, & Brekke, 2002; but see Martin & Stapel, 1998). These implicit theories about the causal relationships between variables may exist prior to people ever being influenced (e.g., people estimating how the weather affected their mood might have a preexisting implicit theory that rain causes bad moods), or they may be formed ad hoc at the time people attempt to assess how they were influenced (e.g., when asked to assess how a noisy power tool affected their liking of a film, people might spontaneously form the implicit theory that unpleasant noise decreases the evaluative rating of an unrelated stimulus). Because these implicit theories are often based on real-world observations and culture-dependent schemes, both of which are subject to many sources of

Steve D. Charman and Gary L. Wells, Department of Psychology, Iowa State University.

Steve D. Charman is now at the Department of Psychology, Florida International University, Miami, Florida.

Portions of this research were presented at Off the Witness Stand: Using Psychology in the Practice of Justice in New York, New York, March 2007. This research was supported by a National Science Foundation grant to the second author.

Correspondence concerning this article should be addressed to Steve Charman, Department of Psychology, Florida International University, DM 296A, Miami, FL 33199. E-mail: charmans@fiu.edu

bias, people are often objectively inaccurate. Consequently, relying on implicit theories can lead people to undercorrect, overcorrect, or unnecessarily correct (i.e., “correct” for influence that never existed in the first place) for the perceived influence (Stapel et al., 1998; Wegener & Petty, 1995; Wilson & Brekke, 1994).

There is, however, very little research within the eyewitness literature on whether witnesses can correctly estimate the influence of various lineup manipulations. Examining this question is especially important within a forensic context because the legal system appears to assume that witnesses can accurately report on how certain variables influenced them. In fact, much of the legal system’s resistance toward adopting various eyewitness research findings may ultimately stem from this assumption; after all, if it is true that witnesses can correct for influence, then the court can always ask witnesses whether (and how) they were influenced by various variables and thus “undo” the effects of influence. Consequently, avoiding influencing and suggestive procedures at the time of identification may seem unnecessary. However, if this commonly held assumption is false, then the legal system’s trust in witnesses’ self-reports of influence is misplaced and, to the extent that it reduces the perceived need for proper lineup reforms, counterproductive.

A model of how witnesses generate reports of influence is illustrated in Figure 1. Our witness-report-of-influence model is partly based on a more general model of mental contamination (Wilson & Brekke, 1994), but is slightly altered to reflect the witnessing experience in particular. According to the witness-report-of-influence model, witnesses’ abilities to generate accurate reports of influence are dependent on (a) their awareness of the influencing variable, and (b) their abilities to correct for its influence appropriately. Failures of awareness may occur for two reasons. First, witnesses may not consciously notice the influencing variable in the first place, as it may be too subtle to be consciously perceived even though it affects their behavior (e.g.,

Garrioch & Brimacombe, 2001). Second, even if witnesses initially notice the influencing variable, they may subsequently fail to remember its presence at a later time. In addition, witnesses may have to remember not only the presence of the influencing variable itself, but also how they were affected by the influencing variable. Memory for the cognitive processes that led to a particular influenced response may not survive a long period of time. This is especially problematic given that eyewitnesses are often asked about the effects of an influencing variable at trial, usually many months or even years after the identification procedure itself (e.g., in a recent DNA exoneration case, a witness was asked about the identification procedure that took place 15 years prior; Newsome v. McCabe, 2002). Such failures of awareness will lead witnesses to misestimate the effects of that variable.

Assuming that witnesses are aware of having been exposed to an influencing variable, they must then properly correct for its influence. Failures of proper correction may occur for a number of reasons, and are largely predicated on the assumption that witnesses’ corrections are generally based on their own (sometimes incorrect) implicit theories of influence (Nisbett & Wilson, 1977; Wegener & Petty, 1995). First, witnesses might acknowledge the presence of an influencing variable but then erroneously claim, because of an incorrect implicit theory, that it did not influence their lineup responses. Conversely, witnesses might erroneously claim, again due to an incorrect implicit theory, that a noninfluencing variable in fact did influence their lineup responses. Second, even if witnesses do correctly claim to have been influenced by an influencing variable, they might fail to estimate the direction and magnitude of that variable’s influence appropriately, an especially likely occurrence if witnesses’ implicit theories of influence are incorrect. Overall, because corrective processes tend to be cognitively effortful and deliberative, witnesses’ motivation is also likely to affect their ability to correct for influence appropriately (Gilbert, Pelham, & Krull, 1988; Wilson & Brekke, 1994). Witnesses who are highly motivated to determine the extent of their bias accurately are more likely to use their implicit theories and engage in more deliberative corrective processes than nonmotivated witnesses. Various failures of correction will lead witnesses to misestimate the effects of the influencing variable, leading witnesses to overcorrect, undercorrect, or even correct for the influence in a direction opposite its actual influence.

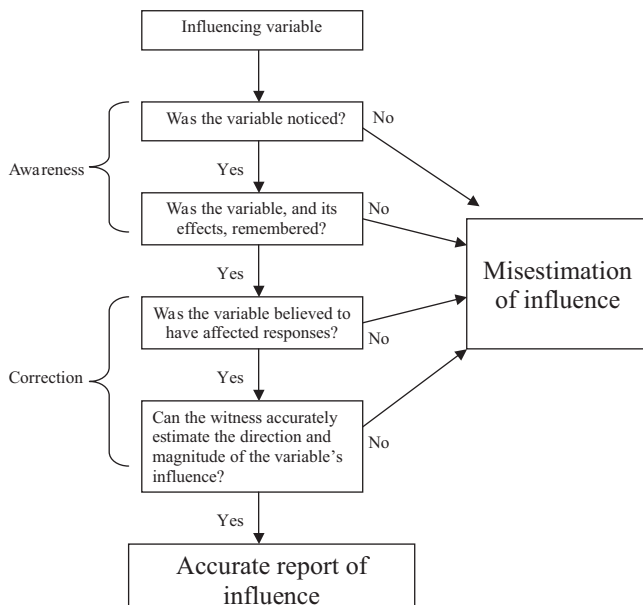


Figure 1. The witness-report-of-influence model.

Previous Paradigms

How does one ascertain whether witnesses’ *believed* effects of influence accurately match the *actual* effects of said influence? Within the more general mental-correction literature, this is generally accomplished by delivering a warning to experimental participants not to let a previously administered influencing stimulus affect their subsequent responses. These responses are then compared to the responses of control participants who never received the influencing stimulus. If people can accurately correct for the influence, the experimental group’s responses should not differ from the control group’s responses.

The warning paradigm, however, is not representative of the problem at hand. Real-world eyewitnesses have already given a response that was previously affected by the influencing variable by the time they are asked to report on the influence. Even if witnesses were able to avoid being influenced if given a warning

prior to their responses, they might be quite poor at retrospectively assessing how their previously generated responses were influenced. Once a response is given, witnesses may feel a stronger need to believe that their freely chosen action was not influenced (Taylor & Brown, 1988). In addition, an overt response may anchor and limit corrective processes (Epley & Gilovich, 2001).

Furthermore, the warning paradigm only assesses people's abilities to subtract or *remove* the effects of an influencing variable. In a forensic context, an eyewitness might be asked to *add* the effects of a variable. For example, a lineup administrator might neglect to give a key instruction to a witness prior to showing that witness a lineup. The witness may subsequently be asked at trial how she would have responded had she in fact received that instruction. For this reason, the term *correction* is used slightly differently from previous mental-correction models, as we use it to refer not only to witnesses' abilities to *subtract* the effects of influence, but also to witnesses' abilities to *add* the effects of influence.

Wells and colleagues developed one paradigm (dubbed here the *yes/no paradigm*) to examine witnesses' corrective abilities (Wells & Bradfield, 1998; Wells, Olson, & Charman, 2003). Participants viewed a mock crime and then attempted an identification of the criminal from a target-absent lineup. Following their identification, the lineup administrator randomly gave witnesses either confirming feedback ("Good, you identified the suspect") or disconfirming feedback ("Actually the criminal was someone else"). Confirming feedback dramatically inflated witnesses' retrospective confidence (i.e., how confident they reported being at the time of the identification) relative to disconfirming feedback. Participants were then asked whether they thought their confidence responses had been influenced by the feedback they had received. Despite witnesses' reports, those who believed that they were influenced were no more influenced than those who believed that they were not influenced.

Unfortunately, the *yes/no paradigm* fails to test an important question in the current work, namely, whether witnesses can properly estimate the *magnitude* of the influence. Dichotomizing a continuous variable, such as estimates of influence, results in a loss of information and may be misleading (Charman & Wells, 2007; Weber, Brewer, Wells, Semmler, & Keast, 2004). For example, imagine that future experiments using the *yes/no paradigm* do indeed find that witnesses who respond that they were influenced were in fact influenced, and those who respond that they were not influenced were in fact not influenced. Even with such findings, we would not know whether witnesses could accurately estimate the *magnitude* of influence. Perhaps witnesses who accurately assign themselves into the "Yes I was influenced" category nonetheless systematically underestimate (or overestimate) the effects of influence. Consequently, we would be no closer to answering the question of whether witnesses can estimate (and appropriately correct for) the effects of influence than we were before conducting the experiments. A proper answer to this question requires a paradigm that can assess not only witnesses' beliefs about whether or not they were influenced, but also the accuracy of witnesses' *estimated magnitude* of influence.

The Actual/Counterfactual Paradigm

To assess the accuracy of corrections, one must be able to compare witnesses' estimates of the magnitude of influence

against the actual magnitude of influence. Although the actual magnitude of influence may be empirically ascertained at the group level, witnesses' estimates of the magnitude of influence requires each witness to, in effect, calculate the imagined difference between their response having received the influencing variable and their response having not received the influencing variable. Thus, the technique employed to ascertain witnesses' corrective abilities was to have all witnesses give not only an *actual* response to the influence, but also a *counterfactual* response. The term *counterfactual* refers to an alternate version of the past or present that could have occurred, but did not (Roese & Olson, 1995). Counterfactual assessment requires a mutable antecedent that, had it been different, could have conceivably led to an alternate outcome. Mentally mutating an antecedent can consequently lead people to think about what might have been.

For the present purposes, a counterfactual response refers to witnesses' reports of how they would have responded had they been in an alternate condition. For example, witnesses who receive confirming feedback following an identification and then report their retrospective confidence (an actual response) would then be asked to indicate how they think they would have responded to the confidence question had they not received the confirming feedback (a counterfactual response). Similarly, witnesses who receive no feedback following an identification and then report their retrospective confidence would subsequently give a counterfactual response indicating how they think they would have responded to the confidence question had they received confirming feedback. Thus, all witnesses are required to report an *actual* response to influence (or noninfluence) and a *counterfactual* response to noninfluence (or influence).

Note that although there is a certain amount of arbitrariness as to what is called an "influencing manipulation" and what is called an "absence of an influencing manipulation" (just as there is some arbitrariness as to what is called an "event" and what is called a "nonevent"), witnesses will nonetheless be referred to throughout the manuscript as having received or not received an influencing manipulation, for economy-of-language purposes. More formally, for purposes of the current manuscript, an influencing manipulation will be defined as one that provides witnesses with some sort of additional information (i.e., feedback or an instruction).

The actual/counterfactual paradigm introduced in this work provides two main advantages over the *yes/no paradigm*. First, by collecting both actual and counterfactual responses from witnesses who receive an influencing manipulation and from witnesses who do not receive an influencing manipulation, the paradigm provides an objective way to compare the accuracy of witnesses' counterfactual responses. Second, instead of using a simple "Yes, I was influenced"/"No, I was not influenced" judgment, the actual/counterfactual paradigm uses continuous, quantitative outcome measures, and can therefore measure the *magnitude* of estimated influence.

Possible Patterns of Data

There were a number of distinct patterns that the actual/counterfactual data could take; however, four distinct possibilities seemed most likely given their theoretical grounding. These possibilities are illustrated in the four panels of Figure 2.

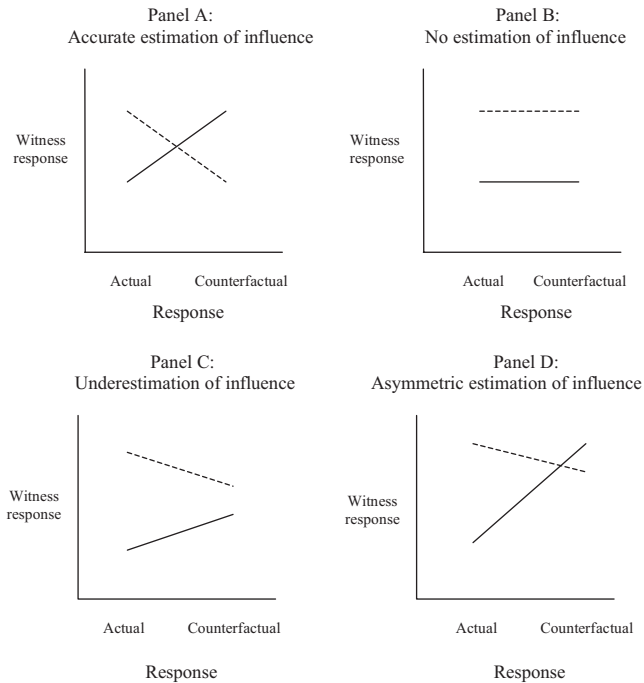


Figure 2. Possible patterns between witnesses' responses (actual vs. counterfactual) and experimental manipulation. Dotted and solid lines represent witnesses who received and did not receive the influencing variable, respectively. "Witness response" is retrospective confidence for Experiment 1 and probability of a "not there" response for Experiment 2.

Accurate estimation of influence. First, at one extreme, if witnesses can accurately estimate the effects of influence, we would expect to see a complete crossover interaction with no main effects (A: Accurate estimation of influence). Witnesses who receive the influencing variable should be able to subtract the influence from their responses completely, and witnesses who do not receive the influencing variable should be able to add the influence to their responses. Thus the counterfactual responses of witnesses who receive the influencing manipulation should match the actual responses of witnesses who do not receive the influencing manipulation, and vice versa.

No estimation of influence. Second, at the opposite extreme, witnesses may deny that a variable influenced (or would have influenced) their responses (B: No estimation of influence). This would be evidenced by no interaction, no main effect for actual versus counterfactual responses, and a main effect for the influencing variable. In other words, witnesses' counterfactual responses would not differ from their actual responses. This might happen, for instance, if witnesses perceive attempts to influence them as threats against their freedoms, which should lead them to engage in strategies that attempt to reexert those freedoms (Brehm, 1966). Denying that a variable influenced one's responses may be such a strategy. Additionally, people generally like to believe that their choices are under their own control, and may consequently tend to dismiss the idea that external influences affected their behavior (Langer, 1975; Taylor & Brown, 1988). More generally, if people's claims of influence are strictly the result of their implicit causal theories, the lack of a causal theory about how a

specific manipulation affects a specific response, or a causal theory that specifically denies the influence of such a manipulation, will result in a denial of influence (Nisbett & Wilson, 1977). There are data within the eyewitness psychology field suggesting that witnesses may fail to appreciate that a demonstrably influential variable affected their responses (specifically, postidentification feedback; Semmler, Brewer, & Wells, 2004, Study 1).

Underestimation of influence. In many domains, cognitive correction processes tend to be plagued by systematic undercorrection (e.g., Gilbert & Malone, 1995). If this general finding is true of witnesses, we would expect witnesses to underestimate the effects of the influencing manipulation (C: Underestimation of influence). This would be evidenced by the presence of an incomplete crossover interaction, a main effect for the influencing variable, and no main effect for actual versus counterfactual response. This possibility is based in part on the anchoring and adjustment heuristic, whereby people, when estimating an unknown value, tend to "anchor" on a known value and adjust insufficiently from that anchor (Epley & Gilovich, 2001; Kahneman & Tversky, 1973). Thus, witnesses might anchor on their actual response and adjust insufficiently, leading to an underestimation of influence. Consistent with this possibility, some evidence suggests that witnesses may in fact underestimate the magnitude of influence (Semmler et al., 2004, Study 2).

Asymmetric estimation of influence. The most complex interaction pattern is the fourth possibility, which is that the ability of witnesses to estimate the effects of influence accurately may depend on the presence or absence of the manipulated variable (D: Asymmetric estimation of influence). Specifically, witnesses may more readily add the effects of an influencing variable than subtract those effects. This would lead to an asymmetric interaction pattern. This possibility has support from two distinct literatures.

An asymmetric pattern is consistent with research on the hindsight bias, which states that once an outcome is known, people tend to overestimate the likelihood that they knew that outcome would occur all along (Fischhoff, 1975; Guilbault, Bryant, Brockway, & Posavac, 2004). Because an influencing variable transmits information to the witness, witnesses who receive an influencing manipulation may overestimate the likelihood that they would have known that information anyway when generating counterfactual responses. For example, confirmatory feedback following a lineup identification provides information to the witness that he or she identified the correct lineup member. When generating counterfactual responses, these witnesses may tend to think that even if they had not received confirming feedback, they would have nonetheless known that they identified the criminal. This hindsight bias should mitigate witnesses' ability to subtract the effects of the influencing variable accurately, and they should consequently tend to underestimate its influence. Witnesses who do not receive the influencing manipulation, on the other hand, should not exhibit any hindsight bias because they did not receive any extra information. Witnesses who do not receive the influencing manipulation should therefore estimate a greater influence than witnesses who do receive the influencing manipulation.

An asymmetric pattern is also consistent with research on counterfactual thinking. Specifically, it has been found that framing counterfactuals in terms of mental addition, in which people must assess whether an alternative antecedent would produce a greater outcome, has more impact than framing counterfactuals in terms of

mental subtraction, in which people assess whether an alternative antecedent would produce a lesser outcome (Dunning & Pappal, 1989). When generating counterfactuals, witnesses who did not receive the influencing variable should self-generate an additive frame (e.g., “Had I gotten confirming feedback, how much *more* confident would I be?”). Conversely, witnesses who did not receive the influencing variable should self-generate a subtractive frame (e.g., “Had I not gotten confirming feedback, how much *less* confident would I be?”). Accordingly, witnesses who do not receive an influencing manipulation should estimate a greater influence than witnesses who do receive an influencing manipulation.

Time Delay

If the accuracy of witnesses’ counterfactual responses relies on their abilities to “retrace” their cognitive processes, it should be harmed by time delay because access to those cognitive processes would decay with time. If witnesses use implicit theories for correction, however, the accuracy of the counterfactual responses should not be affected by time delay. Thus, a time-delay manipulation between witnesses’ actual responses and their counterfactual responses is also included in the design of the current experiments. Some witnesses gave their counterfactual responses immediately following their actual responses, whereas other witnesses gave their counterfactual responses 48 hr after their actual responses.

Experiment 1: Postidentification Feedback

Experiment 1 examined witnesses’ abilities to report accurately on the influence of confirming feedback on their retrospective confidence. Because other self-reports, such as how good a view witnesses had of the criminal and how much attention they paid to the criminal’s face, have also been shown to be affected by feedback (e.g., Wells & Bradfield, 1998, 1999; Wells et al., 2003), these measures were also assessed.

Method

Participants. Participants ($N = 117$) were drawn from the research pool of undergraduate psychology students at a large Midwestern university and received course credit for their participation. Fourteen participants’ data could not be used for various reasons (eight delay participants did not return 48 hr later to complete the study; data from two participants had to be thrown out due to experimenter error, four participants expressed previous knowledge about the study), leaving 103 participants in the experiment.

Design and Overview

Experiment 1 was a 2 (feedback: confirming vs. none) \times 2 (response: actual vs. counterfactual) \times 2 (timing of counterfactual response: immediate vs. delayed) mixed-factorial design, with feedback and timing of counterfactual response being between-subjects variables, and response being a within-subjects variable. All participants engaged in the following order of tasks: They witnessed a mock crime, attempted an identification of the criminal from a lineup, received either confirming feedback or no feedback about their decision, responded to testimony-relevant

questions (*actual* responses concerning their retrospective confidence, their view of the criminal, and the amount of attention paid to the criminal), answered questions about their memory for the feedback they were given (if any) as well as for their actual responses, and reported how they would have answered the testimony-relevant questions had they been in the alternative feedback condition (*counterfactual* responses). Responses were made on a 0–100% scale in 10% increments for retrospective confidence, and a 1–10 scale for view and attention. Participants in the delay condition experienced a 48-hr delay before answering questions about their memory for influence and giving their counterfactual responses.

Materials

Mock crime. Participants viewed a mock crime approximately 45 s long. The video showed a male fiddling with a bomb on the roof of a campus building. The video focuses on the criminal for a few seconds, after which the criminal realizes he is being watched and escapes down a hallway. The criminal’s face is visible on at least three occasions for a total of approximately 15 s.

Lineups. All participants viewed the same lineup that did not contain the criminal. The lineup was composed of six pictures of males who were similar in appearance to the criminal. Lineup members were shown straight on from the shoulders up in a simultaneous 2 (down) \times 3 (across) display. Each picture measured approximately 5.1 \times 7.6 cm.

Procedure

Participants signed a consent form and were told that they would be watching a short video of some people and that we were interested in their impressions of these people. They were led individually to private cubicles and were seated in front of a computer. The experimenter began the mock crime video and left the room. Immediately following the video, the experimenter returned and informed participants (henceforth called “witnesses”) that they were now witnesses to a crime and asked them to attempt to make an identification from a photo lineup. The experimenter showed witnesses a six-person lineup that did not contain the actual criminal. Criminal-absent lineups were used for two reasons: First, to ensure that all witnesses were inaccurate and thus to eliminate any noise associated with having a mixture of accurate and inaccurate witnesses, and second, to maximize the effect of postidentification feedback, which has been shown to have stronger effects for inaccurate witnesses than accurate witnesses (Bradfield, Wells, & Olson, 2002). Because it was important to have all witnesses make an identification, the experimenter instructed witnesses to “select the person who you think planted the bomb on the roof.” This instruction has been shown to cause almost all witnesses to identify someone from a lineup (Wells & Bradfield, 1998). If a witness failed to make an identification, the experimenter told the witness “I need you to try to identify the bomber.”

Following the lineup identification, the experimenter gave either confirming feedback (“Good, you identified the suspect,” $n = 52$) or no feedback ($n = 51$) to witnesses, depending on the specific witness’s condition, which was determined randomly. Witnesses then responded to questions about their retrospective confidence in their identification, how good a view they had of the criminal, and

how much attention they paid to the criminal's face. Witnesses in the no-delay condition ($n = 54$) immediately proceeded with the rest of the experiment; witnesses in the delay condition ($n = 49$) were excused at this point and returned exactly 48 hr later. When delay witnesses returned, they followed the same procedure as no-delay witnesses.

Witnesses were then given questions concerning their memory for the influence and for their actual responses. Specifically, witnesses were asked (a) to give an open-ended response as to what the experimenter had told them immediately following their identification (the cued recall question); (b) to choose, from a list of five options, which statement the experimenter had given them immediately following the identification (the recognition question); (c) to select whether the statement that the experimenter gave them (if they were given a statement) occurred before or after they responded to the confidence question; and (d) to re-respond from memory how they had originally responded to the confidence, view, and attention questions. Once witnesses responded to each question, they were told what the correct response was (except after the first open-ended question), and then allowed to see the next question.

In order to ensure that any effect of time delay on counterfactual responses was not confounded with witnesses' memories for their actual responses, participants were then shown both the lineup and their original responses to the confidence, view, and attention questions. Participants were instructed to imagine that they were in the alternative feedback condition, and were asked to re-respond to the confidence, view, and attention questions as they would have had they been in the alternative condition. These were witnesses' counterfactual responses. Following completion of these questions, all participants were debriefed and excused.

Results

Overview. Only two participants (1.9%) did not immediately make an identification from the lineup. However, both participants made an identification after the additional prompt by the experimenter (to "try and identify the bomber"), and thus their data are included in all analyses. The results section is divided into two sections that follow the logic of the witness-report-of-influence model in Figure 1. The first section examines witnesses' awareness of the confirming feedback; specifically, whether they noticed and remembered the confirming feedback, and whether they remembered their actual responses. This is ascertained by examining (a) whether witnesses can recall what sort of feedback, if any, they received (as determined through open-ended responses); (b) whether witnesses can recognize what the experimenter told them immediately following their identification (as determined through multiple-choice responses); and (c) whether witnesses can remember how they originally answered the actual confidence, view, and attention questions. Time delay was also examined as a possible moderator of any effects.

The second section examines witnesses' abilities to correct for influence by examining whether they could accurately estimate the magnitude of any such influence. This was done by comparing the counterfactual responses given by witnesses who received confirming feedback to the actual responses made by witnesses who did not receive feedback (and vice versa). These comparisons were used to determine (a) whether witnesses show any ability to

estimate the effects of influence accurately, (b) whether the estimation of influence is equal for witnesses who received confirming feedback and witnesses who did not receive confirming feedback, and (c) whether witnesses tended to overestimate, underestimate, or estimate accurately the effects of confirming feedback. Time delay was also examined as a possible moderator of estimates of influence. In order to use a common metric, effect sizes for t tests and for single-degree-of-freedom F tests were converted to Cohen's d , which was defined as a small, medium, and large effect size for $d = 0.20, 0.50,$ and $0.80,$ respectively (Cohen, 1988). In the case of single-degree-of-freedom chi-square tests, the effect size w was used, for which small, medium, and large effect sizes are $0.10, 0.30,$ and $0.50,$ respectively (Rosenthal & Rosnow, 1991). Confidence, view, and attention scores were all significantly correlated with one another on both the actual measures and the counterfactual measures (see Table 1). Unlike view and attention, however, confidence is a central construct in the eyewitness-identification literature. Eyewitness confidence has been directly tied to whether jurors will believe the eyewitness, and confidence is the most researched self-report measure in the eyewitness-identification literature. Therefore, confidence was examined independently with univariate ANOVAs, whereas view and attention were analyzed with the use of MANOVAs. Means of all scores are displayed in Table 2. Patterns of data are displayed in Figure 3.

Did Witnesses Notice and Remember the Confirming Feedback and Their Actual Responses?

Open-ended responses. Witnesses' reports of what they were told immediately following their identification were content-analyzed for whether they reported having received confirmation about their identification decision. Two blind raters independently analyzed each witness's response; the two raters agreed on 100% of the responses. Of witnesses who received confirming feedback, 80.8% correctly reported in the open-ended responses that they received some sort of experimenter confirmation of their identification. Witnesses in the delay condition (who answered this question 48-hr after having received feedback) did not significantly differ in their likelihood of reporting having received confirming feedback from witnesses in the no-delay condition (80% vs. 81%, respectively), $\chi^2(1, N = 52) = 0.02, p = .89, w = 0.02.$ Two

Table 1
Pearson Correlation Coefficients Between Confidence, View, and Attention on Actual and Counterfactual Responses for Experiment 1

Measure	Confidence	View	Attention
	Actual responses		
Confidence	—	.38**	.28**
View		—	.25*
Attention			—
	Counterfactual responses		
Confidence	—	.54**	.36**
View		—	.42**
Attention			—

* $p < .05.$ ** $p < .001$

Table 2

Means and Standard Deviations (in Parentheses) of Confidence, View, and Attention Scores as Functions of Feedback, Delay, and Response for Experiment 1

Measure	Delay			No delay			Overall		
	Actual	Remembered	Counterfactual	Actual	Remembered	Counterfactual	Actual	Remembered	Counterfactual
Confirming feedback									
Confidence	6.2 (2.8)	6.1 (2.7)	5.1 (2.6)	5.7 (2.5)	5.6 (2.5)	4.3 (2.3)	5.9 (2.6)	5.8 (2.6)	4.7 (2.5)
View	7.3 (2.1)	7.4 (2.3)	6.6 (2.1)	6.7 (2.4)	6.6 (2.4)	5.6 (2.7)	7.0 (2.2)	6.9 (2.4)	6.2 (2.4)
Attention	7.1 (2.0)	6.8 (2.1)	6.2 (2.4)	6.8 (2.1)	6.8 (2.2)	6.0 (2.5)	6.9 (2.1)	6.8 (2.1)	6.1 (2.5)
No feedback									
Confidence	5.2 (2.6)	5.2 (2.5)	7.5 (2.1)	4.9 (2.5)	4.9 (2.4)	7.7 (1.9)	5.0 (2.5)	5.0 (2.4)	7.7 (2.0)
View	6.3 (1.8)	6.3 (2.3)	7.7 (1.6)	6.1 (2.0)	6.0 (2.0)	7.5 (1.9)	6.2 (1.9)	6.2 (2.1)	7.6 (1.8)
Attention	6.3 (2.1)	6.3 (1.8)	7.3 (1.8)	5.9 (2.5)	5.9 (2.4)	7.4 (2.2)	6.1 (2.3)	6.1 (2.1)	7.4 (2.0)

Note. Confidence scores were converted to a 0–10 scale to match the other measures. Higher scores indicate greater confidence, a better view, and more attention paid to the criminal's face. Overall actual and counterfactual responses are also plotted in Figure 3.

witnesses (3.9%) in the no-feedback condition reported having received some confirmation of their identification from the experimenter.

Multiple-choice recognition responses. Overall, 90.3% of witnesses correctly selected the statement that had been given to them by the experimenter immediately following their identification from a list of five options. This did not vary significantly as a function of whether witnesses received feedback, $\chi^2(1, N = 103) = 1.69, p = .19, w = 0.13$, or as a function of delay, $\chi^2(1, N = 103) = 3.38, p = .07, w = 0.18$.

Of those witnesses who received confirming feedback, 65.4% correctly recalled having received the feedback *before* answering the question about their confidence. The remaining 34.6% of witnesses incorrectly reported having received the feedback *after* answering the question about their confidence. This did not vary significantly as a function of delay, $\chi^2(1, N = 52) = .15, p = .70, w = .05$.

Remembered responses. Paired-samples *t* tests indicated that witnesses' remembered reports did not differ significantly from their actual reports on any of the three variables: $t(102) = 1.00$,

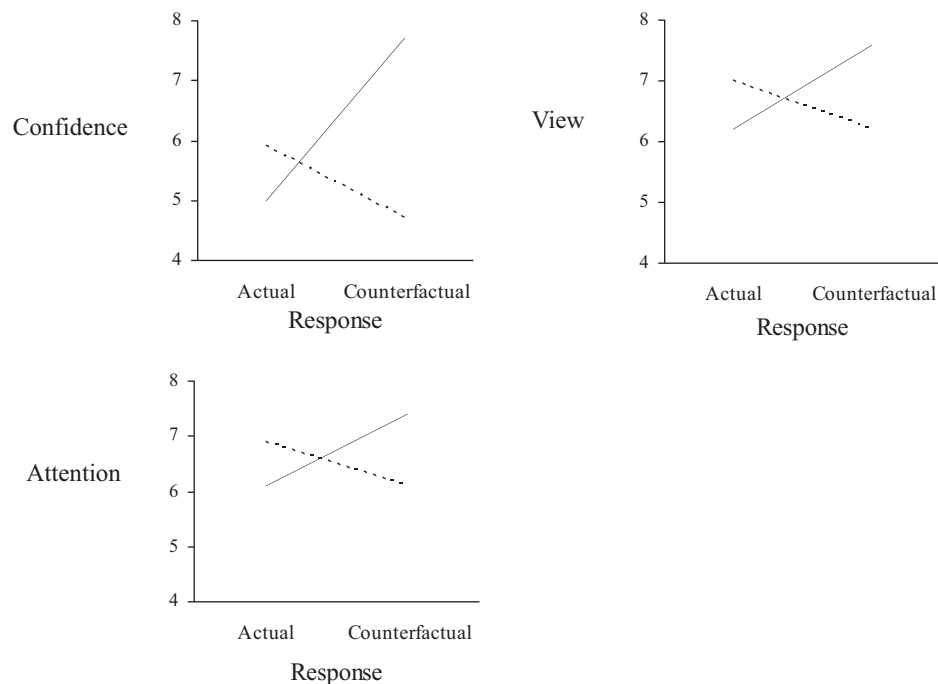


Figure 3. Actual and counterfactual responses to confirming feedback for Experiment 1. Dotted lines represent witnesses who received confirming feedback; solid lines represent witnesses who did not receive confirming feedback.

$p = .32$, $d = 0.20$ for confidence; $t(102) = 0.21$, $p = .83$, $d = 0.04$ for view; $t(102) = 1.09$, $p = .28$, $d = 0.22$ for attention. Difference scores were calculated by subtracting witnesses' actual responses to the confidence, view, and attention measures from their respective remembered responses to these measures. These difference scores equaled the degree to which witnesses misremembered their initial responses. A 2 (delay vs. no delay) \times 2 (confirming feedback vs. no confirming feedback) ANOVA was conducted on the confidence difference score and a 2 \times 2 MANOVA was conducted on the remaining testimony-relevant judgments (i.e., view and attention). Delay did not significantly affect the degree to which witnesses misremembered their confidence, $F(1, 99) = 0.03$, $p = .87$, $d = 0.03$, nor their other testimony-relevant judgments, $F(2, 98) = 0.44$, $p = .65$. Feedback did not significantly affect the degree to which witnesses misremembered their confidence, $F(1, 99) = 1.95$, $p = .17$, $d = 0.28$, nor their other testimony-relevant judgments, $F(2, 98) = 0.67$, $p = .52$. Likewise, no interaction between delay and feedback existed for confidence, $F(1, 99) = 0.07$, $p = .80$, $d = 0.06$, or for the other testimony-relevant judgments, $F(2, 98) = 1.55$, $p = .22$. Thus, witnesses tended to be quite accurate when remembering their initial responses, a finding that was not moderated by either time delay or feedback. Because preliminary analyses showed no effect of delay, data were initially collapsed across this variable.

Did Witnesses Accurately Correct for the Effects of Confirming Feedback?

Feedback effect on actual responses. Based on the wealth of previous research concerning the feedback effect, including a recent meta-analysis (Douglass & Steblay, 2006), as well as its demonstrated consistency and magnitude, one-tailed tests were used to assess the feedback effect. Witnesses who received confirming feedback reported being significantly more confident, $t(101) = 1.86$, $p = .03$, $d = 0.37$, and reported significantly higher scores on testimony-relevant judgments, $F(2, 100) = 3.20$, $p = .05$, compared to witnesses who did not receive feedback. The typical feedback effect was thus replicated.

Test of estimation of influence. If witnesses show any ability to estimate the effects of feedback accurately, then those witnesses who received feedback should have *decreased* their reports of confidence, view, and attention when generating counterfactual responses. Conversely, those witnesses who did not receive feedback should have *increased* their reports of confidence, view, and attention when generating counterfactual responses. Thus, any ability of witnesses to estimate the effects of influence accurately should appear as a significant response (actual vs. counterfactual) \times feedback (confirming vs. none) interaction. A 2 \times 2 mixed ANOVA on confidence revealed a significant interaction, $F(1, 101) = 137.11$, $p < .001$, $d = 2.34$. A MANOVA also revealed a significant response \times feedback interaction for the other testimony-relevant judgments, $F(2, 100) = 59.85$, $p < .001$. The presence of significant interactions suggests that witnesses estimated that confirming feedback had (or would have had) at least some influence on their responses.

Test of equality of estimation of influence. The previous analysis does not, however, indicate whether the estimated effect of feedback was the same for witnesses who received feedback as it was for witnesses who did not receive feedback. The asymmetric

estimation of influence hypothesis predicted that witnesses who received the influencing variable (confirming feedback) would estimate less of an influence than witnesses who did not receive the influencing variable. To test this prediction, difference scores were calculated for each of the individual measures. For witnesses who received feedback, this difference score was calculated as their actual measure score minus their counterfactual measure score. For witnesses who did not receive feedback, this difference score was calculated as their counterfactual measure score minus their actual measure score. Thus, positive difference scores indicate that the participant changed his or her response in a direction consistent with predictions, whereas negative difference scores indicate that the participant changed his or her response in a direction opposite to the predictions. These scores are displayed in Panel A of Table 3.

Consistent with the predicted asymmetric estimation of influence hypothesis, witnesses who received confirming feedback estimated significantly less of an effect of feedback than did witnesses who did not receive confirming feedback on the confidence measure, $t(101) = 4.26$, $p < .001$, $d = .85$, and on the other testimony-relevant judgments, $F(2, 100) = 4.46$, $p = .01$.

According to the witness-report-of-influence model, there are two possible ways in which witnesses who received confirming feedback could have estimated a smaller effect of feedback than witnesses who did not receive confirming feedback. First, witnesses who received confirming feedback might simply have been less likely to acknowledge any effect at all of feedback (i.e., their counterfactual responses would be more likely to be the same as their actual responses). Alternatively, among those witnesses who did acknowledge an effect of feedback, the estimate of the magnitude of the feedback effect may have been less among witnesses who received feedback than among witnesses who did not receive

Table 3
Means and Standard Deviations (in Parentheses) of Witnesses' Estimates of the Magnitude of the Feedback Effect (Panels A and C) and Percentage of Witnesses Who Acknowledged an Effect of Feedback (Panel B) as a Function of Feedback Condition and Measure for Experiment 1

Measure	Confirming feedback	No feedback
Panel A		
Estimates of the magnitude of the feedback effect among all witnesses		
Confidence	1.23 (1.10)	2.64 (2.11)
View	0.77 (1.04)	1.44 (1.38)
Attention	0.81 (1.07)	1.27 (1.34)
Panel B		
Percentage of witnesses who acknowledged an effect of feedback		
Confidence	73	88
View	44	75
Attention	50	73
Panel C		
Estimates of the magnitude of the feedback effect among witnesses who acknowledged an effect of feedback		
Confidence	1.68 (0.93)	2.99 (2.00)
View	1.74 (0.86)	1.93 (1.26)
Attention	1.62 (0.98)	1.76 (1.28)

Note. Confidence scores were converted to a 0–10 scale to match the other measures.

feedback. Data for these two possibilities are displayed in Panels B and C, respectively, of Table 3. Consistent with the former interpretation, witnesses who received confirming feedback were significantly less likely to acknowledge any effect of feedback than witnesses who did not receive confirming feedback on the confidence measure, $\chi^2(1, N = 103) = 3.78, p = .05, w = 0.19$; the view measure, $\chi^2(1, N = 103) = 9.78, p = .002, w = 0.31$; and the attention measure, $\chi^2(1, N = 103) = 5.51, p = .02, w = 0.23$. Partial support was found for the latter interpretation: Looking only at witnesses who acknowledged an effect of feedback revealed that witnesses who received confirming feedback estimated the magnitude of the effect to be significantly less than witnesses who did not receive confirming feedback on the confidence measure, $t(81) = 3.70, p < .001, d = 0.82$, but not on either the view measure, $t(59) = 0.65, p = .52, d = 0.17$, or the attention measure, $t(61) = 0.47, p = .64, d = 0.12$.

Recall that a substantial proportion of witnesses who received confirming feedback (34.6%) incorrectly reported having received confirming feedback *after* having answered the confidence question. Oddly, however, witnesses' estimates of the influence of confirming feedback on their confidence reports did not significantly differ as a function of whether they reported having received the feedback before ($M = 13.2$) versus after ($M = 10.6$) having answered the confidence question, $t(50) = 0.84, p = .41, d = 0.24$.

Test of accuracy of estimation. Although witnesses who received confirming feedback estimated the influence of feedback to be significantly less than did witnesses who did not receive confirming feedback, a crucial question is whether either of those estimates was accurate. To assess the accuracy of witnesses' estimates of influence, the counterfactual scores of witnesses who received feedback were compared to the actual scores of witnesses who did not receive feedback. Similarly, the counterfactual scores of witnesses who did not receive feedback were compared to the actual scores of witnesses who did receive feedback. If witnesses were accurately estimating the effects of influence, these differences should be nonsignificant.

The counterfactual responses of witnesses who received feedback did not differ significantly from the actual responses of witnesses who did not receive feedback on the confidence measure, $t(101) = 0.61, p = .54, d = 0.12$, or on the other testimony-relevant judgments, $F(2, 100) = 0.01, p = .99$. Thus, witnesses who received feedback were relatively accurate in their estimates of the influence of feedback. However, the counterfactual responses of witnesses who did not receive feedback were significantly higher than the actual responses of witnesses who did receive feedback on the confidence measure, $t(101) = 3.72, p < .001, d = 0.74$, but not on the other testimony-relevant measures, $F(2, 100) = 1.27, p = .28$. Thus, witnesses who did not receive feedback overestimated the influence feedback would have had on their confidence.

Test of a moderating influence of time delay. Because all witnesses, regardless of delay condition, responded to the confidence, view, and attention questions immediately following the administration of feedback, delay could only have affected witnesses' counterfactual responses. Thus, to test the moderating effect of time delay, 2 (timing of counterfactual responses: delayed vs. immediate) \times 2 (feedback: confirming vs. none) between-subjects ANOVAs were conducted on the counterfactual re-

sponses. Delay had no significant main effects on witnesses' counterfactual confidence, $F(1, 99) = 0.44, p = .51, d = 0.12$, or on the counterfactual responses of the other testimony-relevant judgments, $F(2, 97) = 0.61, p = .55$. Delay also did not significantly interact with feedback on witnesses' counterfactual confidence, $F(1, 99) = 1.23, p = .27, d = 0.22$, or on the counterfactual responses of the other testimony-relevant judgments, $F(2, 97) = 0.25, p = .78$.

However, the previous analyses do not control for witnesses' initial scores on these measures. Although they demonstrate that delay did not moderate witnesses' counterfactual responses, they do not speak as to whether time delay moderated the *magnitude* of witnesses' estimates of influence. Hence, the previous analyses were repeated with witnesses' difference scores (which represent the magnitude of estimation of influence) used as the dependent variable. Again, delay had no significant main effects on the estimated magnitude of influence of feedback on confidence, $F(1, 99) = 1.55, p = .22, d = 0.24$, or on the other testimony-relevant judgments, $F(2, 98) = 0.11, p = .90$. Delay also did not significantly interact with feedback on the estimated magnitude of influence of feedback on confidence, $F(1, 99) = 0.14, p = .71, d = 0.06$, or on the other testimony-relevant judgments, $F(2, 98) = 0.69, p = .50$. Therefore, delay did not significantly moderate any of the witnesses' responses.

Discussion

The purpose of Experiment 1 was to determine whether witnesses could accurately report on the influence of confirming feedback. According to the witness-report-of-influence model, misestimation of influence may occur because of lack of awareness for the influencing variable or inappropriate correction of the influencing variable. Concerning witnesses' awareness of the confirming feedback, the results were clear: The majority of witnesses who received confirming feedback accurately reported having received confirming feedback, effects that were not moderated by a 48-hr time delay.

Concerning witnesses' corrective abilities, the results are clearly consistent with the asymmetric estimation of influence hypothesis. Witnesses who did not receive confirming feedback estimated the influence of confirming feedback to be significantly greater in magnitude than witnesses who did receive confirming feedback. This asymmetry on the confidence measure resulted from the fact that, compared to witnesses who did not receive confirming feedback, witnesses who did receive confirming feedback (a) were less likely to report that confirming feedback had (or would have had) any influence whatsoever, and (b) estimated less of an impact of confirming feedback among those who did make a correction. The asymmetry on the view and attention measures resulted from the fact that witnesses who received confirming feedback were less likely than witnesses who did not receive confirming feedback to report that the feedback had (or would have had) any influence whatsoever on their responses. A 48-hr time delay had no significant effect on this pattern. Thus, the ability of witnesses to estimate the influence of confirming feedback accurately was dependent on whether those witnesses actually received feedback or not.

Would this asymmetric estimation of influence pattern generalize across another influencing variable or is it idiosyncratic to

postidentification feedback? To test whether this observed asymmetric estimation of influence would generalize across a different influencing manipulation, Experiment 2 used the same paradigm as Experiment 1, this time using prelineup instructions as the influencing variable.

Experiment 2: Prelineup Instructions

It has been shown repeatedly that instructions that imply that the criminal is in the lineup significantly increase the probability of an identification attempt, and instructions that suggest that the criminal may not be in the lineup significantly decrease the probability of an identification attempt (Malpass & Devine, 1981; Steblay, 1997). This prelineup instruction manipulation was chosen to be the influencing variable for Experiment 2 because it differed from postidentification feedback in several important respects, and thus any common patterns of results between the two experiments are less likely to be idiosyncratic artifacts of the specific influencing manipulations.

First, the two manipulations exert their main effects on different measures: postidentification feedback affects retrospective confidence (a subjective response), whereas prelineup instructions affect the probability of an identification attempt (an objective behavior). Second, the timing of the two manipulations differs; postidentification feedback occurs *after* a lineup identification task, whereas administration of the cautionary instruction occurs *before* a lineup identification task in Experiment 2. Third, the influencing manipulations differ in terms of their effects on the accuracy of witnesses' reports. Whereas the presence of confirming feedback makes witnesses *less* accurate in their confidence assessments, the presence of the cautionary instruction makes witnesses *more* accurate in their identification decisions.

Method

Participants. Participants ($N = 141$) were drawn from the research pool of undergraduate psychology students at a large Midwestern university and received course credit for their participation. Data from 18 participants could not be used for various reasons (16 delay participants did not return 48 hr later to complete the study; two participants expressed previous knowledge about the study), leaving 123 participants in the experiment.

Design and overview. Experiment 2 was a 2 (cautionary instruction: given vs. not given) \times 2 (response: actual vs. counterfactual) \times 2 (timing of counterfactual response: immediate vs. delayed) mixed factorial design, with feedback and timing of counterfactual response being between-subjects variables, and response being a within-subjects variable. All participants engaged in the following order of tasks: They witnessed a mock crime, were given either the cautionary instruction by the experimenter or not, attempted their identification (their *actual* identification response), answered questions about their memory for the influence and for their actual responses, and reported how they would have responded to the lineup had they been in the alternative instruction condition (their *counterfactual* identification response). Participants in the delay condition experienced a 48-hr delay before answering questions about their memory for influence and giving their counterfactual responses.

Materials and procedure. Both the video of the mock crime and the lineup were identical to those used in Experiment 1.

Participants viewed the mock crime and lineups individually. Participants were randomly assigned to both instruction and delay conditions. Participants followed the same general procedure as participants in Experiment 1. Following the mock crime, however, participants (now called "witnesses") were given one of two instructions. Before viewing the lineup, witnesses assigned to the cautionary-instruction condition ($n = 63$) were told to "keep in mind that the actual criminal may or may not be in the lineup. Look at the people in the photo lineup and tell me if you see the criminal." Witnesses in the no-cautionary-instruction condition ($n = 60$) were told to "look at the people in the photo lineup and try to identify the criminal you saw in the video." Witnesses could then identify a lineup member or choose to identify no one. They were then asked about their confidence in their decision. Their responses were recorded by the experimenter. No feedback was given to the witnesses following the identification decision. Witnesses in the no-delay condition ($n = 71$) immediately proceeded with the rest of the experiment; witnesses in the delay condition ($n = 52$) were excused at this point and returned exactly 48 hr later. When delay witnesses returned, they followed the exact same procedure as no-delay witnesses.

Witnesses were given questions concerning their memory for the influence and for their actual responses. Specifically, witnesses were asked (a) to give an open-ended response as to what the experimenter had instructed them immediately before showing them the lineup (the recall question); (b) to choose, from a list of five options, which instruction, if any, the experimenter had given them immediately before showing them the lineup (the recognition question); and (c) to re-respond from memory how they had originally responded to the lineup. After witnesses responded to each question, they were told what the correct response was (except for the first open-ended question), and then allowed to see the next question.

Witnesses were then shown both the lineup and the sheet on which the experimenter had recorded their lineup decision, were instructed to imagine that they were in the alternative instruction condition, and were asked to re-respond to the lineup as if they had been in the alternative instruction condition. These were witnesses' *counterfactual* responses. Following completion of these questions, all participants were debriefed and excused.

Results

Overview. Following the witness-report-of-influence model, results are divided into two sections and are organized in the same format as they were for Experiment 1. The first section examines witnesses' awareness of the prelineup instruction. This is ascertained by examining (a) whether witnesses can recall what prelineup instructions, if any, they received (as determined through open-ended responses); (b) whether witnesses can recognize what the experimenter told them immediately prior to the identification attempt (as determined through multiple-choice responses); and (c) whether witnesses can remember how they originally responded to the lineup. Time delay was also examined as a possible moderator.

The second section examines witnesses' abilities to correct for influence by examining whether they could estimate the extent to which they were influenced by the cautionary instruction. This was done by comparing the actual responses of witnesses who received the cautionary instruction to the counterfactual responses of wit-

nesses who did not receive the cautionary instruction (and vice versa). The comparisons were used to determine (a) whether witnesses show any ability to estimate the effects of influence accurately, (b) whether estimates of the effect of influence are equal for witnesses who received the cautionary instruction and witnesses who did not receive the cautionary instruction, and (c) whether witnesses tended to overestimate, underestimate, or estimate accurately the influence of prelineup instructions. Again, time delay was examined as a possible moderator.

Did Witnesses Notice and Remember the Cautionary Instruction and Their Actual Responses?

Open-ended responses. Witnesses' responses to open-ended questions concerning what instructions they were given immediately before being shown the lineup were content analyzed for whether they reported receiving a cautionary instruction that the criminal may not be in the lineup. Two blind raters independently analyzed each witness's response; the two raters agreed on 100% of the responses. Of witnesses who received the cautionary instruction, 82.5% correctly reported having been instructed that the criminal may not be in the lineup. Witnesses in the delay condition (who answered this question 48 hr after having received the instruction) were as likely to report having received a cautionary instruction (74%) as witnesses in the no-delay condition (89%), $\chi^2(1, N = 63) = 2.35, p = .13, w = 0.19$. One person (1.7%) in the no-instruction condition reported having received an instruction that the criminal may not be in the lineup.

Multiple-choice recognition responses. Overall, 84.6% of witnesses correctly selected the instruction that was given to them from a list of five options. This did not vary significantly as a function of either instruction, $\chi^2(1, N = 123) = 0.02, p = .89, w = 0.01$, or delay, $\chi^2(1, N = 123) = 1.21, p = .27, w = 0.10$.

Remembered responses. Overall, 98.4% of witnesses correctly remembered their identification decision. This did not vary significantly as a function of delay, $\chi^2(1, N = 123) = 0.04, p = .84, w = 0.02$. Difference scores, which equaled the degree to which witnesses misremembered their confidence, were calculated by subtracting witnesses' reported confidence from their remembered confidence. These difference scores did not differ significantly from zero, indicating that witnesses' reported confidence in their identification did not differ significantly from their subsequent remembered confidence, $t(122) = 1.47, p = .15, d = 0.27$. In fact, actual confidence and remembered confidence were highly correlated, $r = .98, p < .001$. The difference scores did not differ significantly as a function of delay, $t(121) = 0.85, p = .40, d = 0.15$ or instructions, $t(121) = 0.66, p = .51, d = 0.12$. Thus, witnesses tended to be quite accurate when remembering their

initial responses, a finding that was not moderated by time delay or instructions.

Did Witnesses Accurately Correct for the Influence of the Cautionary Instruction?

Proportions of correct lineup rejections (i.e., "not there" responses) as a function of condition are displayed in Table 4. Patterns of data are displayed in Figure 4.

Instruction effect on actual responses. Consistent with past research, the cautionary instruction to "keep in mind that the criminal may or may not be present in the lineup" led to more correct "not there" responses (47.6%) than an instruction that did not include this phrase (6.7%), $\chi^2(1, N = 123) = 25.77, p < .001, w = 0.46$.

Test of estimation of influence. Given that the previous analysis demonstrated that instructions indeed influenced identification decisions, if witnesses show any ability to estimate the effects of the cautionary instruction accurately, then some proportion of witnesses who received the cautionary instruction and subsequently rejected the lineup (i.e., said "not there") should report that had they instead not received the cautionary instruction they would have made an identification. Conversely, some proportion of witnesses who did not receive the cautionary instruction and made an identification should report that had they instead received the cautionary instruction they would have rejected the lineup. Therefore, any tendency to estimate the effects of the cautionary instruction accurately should appear as a significant response (actual vs. counterfactual) \times instruction (cautionary vs. no cautionary) interaction. Because all identifications were incorrect, no distinction was made between witnesses who identified different lineup members; all witnesses were classified as either having made an identification or having rejected the lineup. A difference score was calculated for each participant that represented whether the counterfactual response was the same as the actual response score (i.e., both identifications or both nonidentifications), was greater than the actual response score (i.e., an identification following a non-identification), or was less than the actual response score (i.e., a nonidentification following an identification). These difference scores were compared across the instruction condition. Consistent with the idea that witnesses were at least partially estimating for the effects of the cautionary instruction, witnesses who received the cautionary instruction exhibited a significantly different pattern of change in responses than witnesses who did not receive the cautionary instruction, $\chi^2(2, N = 123) = 42.84, p < .001$.

Test of equality of estimation of influence. The asymmetric estimation of influence hypothesis predicts that witnesses who receive a manipulation (in this experiment, a cautionary instruc-

Table 4
Proportion of Correct Lineup Rejections as a Function of Instruction Condition, Delay, and Response for Experiment 2

Instruction	Delay		No delay		Overall	
	Actual	Counterfactual	Actual	Counterfactual	Actual	Counterfactual
Cautionary instruction	.52	.30	.44	.14	.48	.21
No cautionary instruction	.04	.50	.09	.56	.07	.53

Note. Overall actual and counterfactual proportions are also plotted in Figure 4.

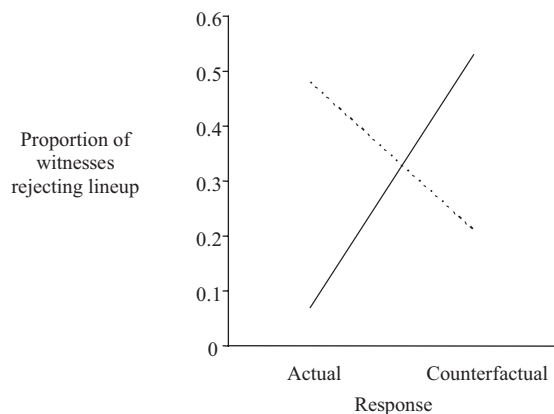


Figure 4. Proportion of actual and counterfactual lineup rejections for Experiment 2. The dotted line represents witnesses who received the cautionary instruction; the solid line represents witnesses who did not receive the cautionary instruction

tion) should estimate less of an influence than should witnesses who do not receive a manipulation. To test this prediction, a chi-square test was performed on the percentage of witnesses who gave a counterfactual response that was different from their actual response and in the direction of actual influence. Consistent with the predicted asymmetric estimation of influence hypothesis, witnesses who received the cautionary instruction were less likely to change their response in the predicted direction (30.2%) than witnesses who did not receive the cautionary instruction (46.7%), although this effect fell slightly short of attaining statistical significance, $\chi^2(1, N = 123) = 3.55, p = .06, w = 0.17$.

Test of accuracy of estimation. If witnesses can accurately estimate the effects of the cautionary instruction, then the proportion of witnesses who received the cautionary instruction and rejected the lineup should not differ significantly from the proportion of witnesses who did not receive the cautionary instruction but reported that they would have rejected the lineup had they received the cautionary instruction. Similarly, the proportion of witnesses who did not receive the cautionary instruction and rejected the lineup should not differ significantly from the proportion of witnesses who received the cautionary instruction but reported that they would have rejected the lineup had they not received the cautionary instruction. In fact, witnesses who received the cautionary instruction underestimated the influence of the instruction: they were more likely than witnesses who did not receive the cautionary instruction to report that they would have rejected the lineup (20.6% vs. 6.7%, respectively), $\chi^2(1, N = 123) = 5.03, p = .03, w = 0.20$. In contrast, those witnesses who did not receive the cautionary instruction did not significantly over- or underestimate the influence of the cautionary instruction: They were as likely as witnesses who actually received the cautionary instruction to report that they would have rejected the lineup (53.3% vs. 47.6%, respectively), $\chi^2(1, N = 123) = 0.40, p = .53, w = 0.06$.

Test of a moderating influence of time delay. Because all witnesses, regardless of delay condition, gave their actual responses to the lineup on Day 1, delay could only have affected witnesses' counterfactual responses. Delay did not significantly affect the proportion of witnesses who reported that they would

have rejected the lineup, $\chi^2(1, N = 123) = 0.37, p = .54, w = 0.06$. A log-linear analysis indicated that delay did not interact with instructions on the likelihood of participants' rejecting a lineup, $\chi^2(1, N = 123) = 2.61, p = .11, w = 0.15$. Delay also had no significant effect on the likelihood that witnesses would correctly change their responses, $\chi^2(1, N = 123) = .01, p = .93, w = 0.01$. Therefore, time delay did not significantly moderate any of the witnesses' responses.

Discussion

The purpose of Experiment 2 was to determine whether witnesses could accurately report on the influence of the cautionary instruction, and more specifically, whether the results of Experiment 1—in particular the asymmetric estimation of influence pattern—would generalize to a prelineup instruction manipulation. According to the witness-report-of-influence model, misestimation of influence may occur as a result of lack of awareness for the influencing variable or because of inappropriate correction for the influencing variable. Concerning witnesses' awareness for the influencing variable, the majority of witnesses who received the cautionary instruction accurately reported having received an instruction that the criminal may or may not be in the lineup, an effect that was not moderated by a 48-hr time delay. Concerning witnesses' corrective abilities, results are clearly consistent with the asymmetric estimation of influence hypothesis. Witnesses who received the cautionary instruction estimated significantly less of an effect of the instruction than did witnesses who did not receive the cautionary instruction. A 48-hr time delay had no significant effect on this pattern. Thus, the ability of witnesses to estimate the influence of the cautionary instruction accurately was dependent on whether they received that instruction or not.

General Discussion

Results of the current studies indicate that witnesses' abilities to report on the influence of various lineup manipulations accurately may not be as poor as some have thought. According to the witness-report-of-influence model, accurate reports of influence are based on (a) an awareness of the influencing variable, and (b) an ability to correct for the influencing variable appropriately. Concerning the first point, witnesses across two experiments clearly exhibited an awareness of the influencing variable. They correctly tended to notice and to remember subsequently having been presented with the influencing variable in question, even after a 2-day delay. Concerning the second point, when witnesses were asked to correct for the influence of that variable, they tended to adjust their responses in a direction consistent with the actual influence of the variable.

Nonetheless, these corrections exhibited a systematic bias. When asked to provide an estimate of the magnitude of influence produced by the influencing variable, witnesses who received the influencing manipulation (postidentification feedback in Experiment 1; cautionary instruction in Experiment 2) consistently estimated less of an effect of that variable than did witnesses who did not receive the influencing manipulation. This asymmetric estimation of influence resulted in witnesses who received feedback accurately estimating its influence (compared to witnesses who did not receive feedback who overestimated its influence) in Experi-

ment 1, and witnesses who received a cautionary instruction underestimating its influence (compared to witnesses who did not receive the cautionary instruction who accurately estimated its influence) in Experiment 2.

How Do Witnesses Estimate Influence?

A large corpus of psychological research indicates that people have little insight into their cognitive processes (Nisbett & Wilson, 1977; Wegener & Petty, 1997; Wilson et al., 2002). When estimating whether, and how, they were influenced by some variable, people are forced to rely on their implicit theories. Witnesses in the current experiments almost certainly did not have a prior implicit theory upon entering the lab about how postidentification feedback, or a cautionary instruction, would affect their lineup responses. When asked how they were influenced, these witnesses were likely forced to generate an implicit theory about the relationship between those variables ad hoc, the result of which was used to estimate the magnitude of the influencing variable.

That these implicit theories led to an asymmetric estimation of influence pattern, depending on whether witnesses received or did not receive the influencing variable, was predicted a priori on the basis of both the hindsight-bias literature (Fischhoff, 1975) and the counterfactual-thinking literature (Dunning & Parpal, 1989). Consistent with results, both of these literatures predict that the implicit theories generated by people who receive some variable will be of a lesser magnitude than the implicit theories generated by people who do not receive the variable.

Notably, the observed asymmetric estimation of influence pattern simultaneously supports this implicit-theories interpretation, and also eliminates potential alternative explanations for the data. For example, perhaps the estimated influence of a variable (as assessed via counterfactual responses) that we are claiming is the result of witnesses' implicit theories was instead the result of a shifting decision criterion. A witness who received confirming feedback in Experiment 1 may have adopted a relatively lax criterion, resulting in high retrospective confidence, view, and attention scores. When asked to estimate how they would have responded had they not received feedback, the witnesses raised their decision criterion, resulting in lower scores. Conversely, a witness who received a cautionary instruction in Experiment 2 may have adopted a strict decision criterion, resulting in few false identifications. When asked to estimate how they would have responded had they not received the instruction, the witnesses lowered their decision criterion, resulting in more false identifications.

There is reason to be skeptical of this alternative explanation. A shifting decision criterion may affect an individual's likelihood of responding, but it is unclear how it could also affect retrospective reports of confidence, view, and attention, as observed in Experiment 1. More critically, however, this decision-criterion explanation does not predict that the magnitude of the shift would be greater for witnesses who received the variable as compared to witnesses who did not receive the variable, and is, consequently, inadequate as an alternative explanation for the observed results.

A second alternative explanation to the implicit-theories perspective is the following: Perhaps witnesses who received feedback were certain in their identification, which was reflected in their responses. When later generating their counterfactual re-

sponses, they reasoned that had they not received feedback, they would have been less certain, and they appropriately lowered their responses. Conversely, perhaps witnesses who received a cautionary instruction were full of doubt about a possible identification, resulting in few false identifications. When generating their counterfactual responses, they reasoned that had they not received the instruction, they would have been more certain, and appropriately altered their lineup responses to reflect that.

This doubt-certainty hypothesis, however, suffers from a similar problem as the decision-criterion explanation: It does not on its own predict an asymmetric estimation of influence pattern. In fact, the only way that it can explain this asymmetry is to resort to the idea that witnesses use implicit theories. To reason that one would have been more certain (or more full of doubt) had an influencing variable been received (or not received) is to say that one has an implicit theory about the effects of the influencing variable on one's feelings of certainty/doubt. How else, other than with an ad hoc implicit theory, could witnesses *know* that they would have been more or less certain? The doubt-certainty hypothesis is not an alternative explanation at all, but simply a reformulation of the implicit theories explanation.

A final alternative explanation is that witnesses relied on introspective awareness when estimating the effects of influence. If witnesses have access to their higher-order cognitive processes, they could simply retrace their cognitive steps and add or subtract the influence of a given variable. Consequently, there would be no need for implicit theories. Again, however, this alternative explanation falls prey to the observed asymmetric estimation of influence pattern. If witnesses could retrace their cognitive steps, why would some witnesses have overestimated, some underestimated, and some accurately estimated the magnitude of influence, and why would this tendency depend on whether or not they received the influencing variable? Additionally, access to their preinfluenced cognitive state would be necessary if witnesses could introspect; however, the absence of any effects of time delay on witnesses' memory for influence suggests that witnesses did not rely on their memory for their preinfluenced state. An implicit-theories perspective, on the other hand, can easily explain the lack of effects of time delay, because the generation of an implicit theory in no way relies on memory for a preinfluenced state, and thus should be insensitive to time delay. This skepticism toward introspective access among witnesses is consistent with recent work on the postidentification feedback effect, from which a consensus is growing that the feedback effect occurs because witnesses do not have cognitive access to their prefeedback states (e.g., Neuschatz et al., 2005; Wells & Bradfield, 1999; Wells et al., 2003).

How Accurate Are Witnesses' Implicit Theories?

The implicit-theories approach is currently the only explanation able to handle the asymmetric estimation of influence pattern. Furthermore, it can also explain an interesting difference across the two experiments. Specifically, witnesses who accurately estimated the magnitude of influence in Experiment 1 were those who received postidentification feedback and had to subtract its influence. Witnesses who did not receive feedback overestimated its influence. However, witnesses who accurately estimated the magnitude of influence in Experiment 2 were those who did *not* receive

the cautionary instruction had to add its influence. Witnesses who received the cautionary instruction underestimated its influence. Although both studies demonstrated an asymmetric estimation of influence pattern, they differed in the specifics of how that pattern was exhibited.

This interesting difference between the two studies is easily explained via the specific implicit theories that witnesses hold, and is partly based on the claims of both the hindsight-bias and counterfactual-thinking literatures that the cause of the asymmetry is that people who receive a variable do not sufficiently partial out the effects of the variable. Thus, if witnesses hold relatively accurate implicit theories about the influence of a variable on their responses, they should accurately estimate the variable's influence when adding it to their responses, but underestimate the variable's influence when subtracting it from their responses, a pattern observed in Experiment 2. If, on the other hand, witnesses hold implicit theories that overestimate the influence of a variable on their responses, they should overestimate the variable's influence when adding it to their responses, but more accurately estimate the variable's influence when subtracting it from their responses, a pattern observed in Experiment 1. We would expect a third pattern if witnesses hold implicit theories that underestimate the influence of a variable on their responses: They should underestimate the variable's influence regardless of whether they received the influencing variable or not, but those who did receive it should underestimate to a greater extent.

Although some studies have demonstrated that peoples' theories about the effects of influencing lineup manipulations are often inaccurate (e.g., Hasel, 2006; Kassir & Barndollar, 1992), little is known about what specific implicit theories witnesses actually hold or how witnesses build those theories on an ad hoc basis. However, the likelihood that a witness will generate a correct implicit theory about the causal effects of an influencing variable will likely be a function of a number of factors, including the salience of the variable, the ease with which the variable can be imagined to influence the response in question, the perceived malleability of the response, and so forth. Witnesses may, for example, readily assume that postidentification feedback would affect their retrospective confidence reports because they perceive retrospective confidence to be malleable. They may not think that feedback could affect their reports of view and attention, on the other hand, because they perceive these variables to be more objective and hence less malleable. If this is true, we would expect witnesses to estimate a greater effect of confirming feedback on their confidence responses than on their view and attention responses, a finding that was in fact observed in Experiment 1.

Memory error. Although most of the observed errors in witnesses' estimation of influence can be accounted for by witnesses' reliance on erroneous implicit theories and not by memory errors, Experiment 1 provides some interesting data that undercuts any premature suggestions that witnesses' memory for the influence variable itself is pristine. Witnesses in Experiment 1 were asked to indicate the point in the experimental procedure during which they received feedback. Although most witnesses correctly reported having received postidentification feedback *before* answering questions about their confidence, view, and attention, over one third incorrectly reported having received the feedback *after* answering those questions, suggesting that witnesses misremembered the sequence of events. This finding is especially intriguing

in light of the fact that such a temporal order of events would have made it impossible for feedback to have influenced witnesses' responses, and yet these witnesses nonetheless estimated having been just as influenced by the feedback as witnesses who correctly reported having received feedback before giving their responses.

It should be noted, however, that all witnesses were told the correct temporal order of events prior to generating their counterfactual responses, and this reminder may have led witnesses to assume that feedback could have influenced them. The current data do not allow an assessment of whether a similar finding would occur among witnesses who are not told the correct temporal order before generating counterfactual responses. If this finding is replicable, it would be more than simply an intellectual curiosity, as misremembering when one was presented with an influencing variable has important real-world implications. A witness who incorrectly claims that she did not receive feedback until after giving a confidence statement might lead a defense attorney to forgo the possibility that the feedback influenced her confidence statement.

Contributions of This Research

Introduction of a New Paradigm

We developed the actual/counterfactual paradigm with the hope that it would become a new tool for eyewitness researchers to test propositions about whether eyewitnesses can estimate external influences on their identification decisions and identification confidence. Prior to collecting the data, we had several concerns about whether this paradigm (or any other paradigm) would be able to accomplish this goal. For instance, we were concerned that the powerful human need for consistency or the natural tendency of people to deny external influences on their behavior would prevent them from reporting counterfactual answers that were different from their previous actual answers. We found that participant-witnesses not only gave counterfactual answers that were different from their actual answers, but in fact even overcorrected in one condition. Hence, we believe that the paradigm is not artificially anchored by consistency drives or by denial of external influence.

We also had the opposite concern, namely, that the counterfactual procedure might have an inherent demand that people should give a different answer for their counterfactual responses than for their actual responses. We found, however, that the percentage of participant-witnesses who did *not* give a different answer for the counterfactual than they gave for the actual was as high as 56% and as low as 12% for Experiment 1 and as high as 70% and as low as 53% for Experiment 2, depending on condition. Hence, people felt free not to change their responses. Furthermore, these rates of change across conditions followed meaningful patterns related to whether counterfactual generation was an addition or subtraction process, thereby providing support for the idea that the participant-witnesses were not simply responding to demands of the counterfactual paradigm. We are further encouraged that the results showed evidence of all three possible results, namely, overestimation, underestimation, and accurate estimation, depending on condition. This suggests that the paradigm itself is not fundamentally biased toward underestimation or overestimation and should be sensitive to detect both outcomes.

Theoretical Contributions.

Prior to the current work, no theory-based model existed for predicting or understanding eyewitnesses' estimates of how variables influenced them. The witness-report-of-influence model, in conjunction with theoretical processes involved in counterfactual thinking and hindsight processes, permits the prediction that such estimates will follow an asymmetric pattern in which eyewitnesses will estimate a greater impact for adding an influence than they will for subtracting that same influence. Virtually any influencing variable is either present or absent in a given situation. Hence, it is important to know how people who experienced a variable's presence estimate their response had the variable been absent versus how people who experienced the variable's absence estimate their response had the variable been present. The asymmetric-influence model clearly makes predictions about other variables not yet tested, such as cowitness influence, postevent information, prior exposure to mug shots, within-race versus cross-race identifications, and scores of other variables, thereby attesting to the heuristic value of the theoretical model. Although the asymmetric-influence model does not yet allow for precise predictions of the exact amount of influence that witnesses will estimate for a given variable, it specifies a conceptual framework based on ad hoc implicit theories and does not rely on the largely discounted notion of introspection.

Practical Implications

It is not uncommon for witnesses to be asked in court whether they were influenced by variable X, or whether they would have responded to the lineup differently had variable Y been present or not present. Results from the current experiments afford a tentative conclusion concerning witnesses' abilities to answer such questions: Witnesses' estimations of influence are largely based on their implicit theories of influence, which lead them to estimate the effects of influencing variables asymmetrically, depending on whether they received the variable or not. This asymmetric pattern should give the legal system pause in assuming that witnesses' reports of influence reflect the extent to which they were actually influenced. Witnesses' responses to these types of questions probably tell us more about how witnesses *think* they were influenced than how they actually *were* influenced. This subtle yet important distinction should increase the skepticism with which the legal system interprets witnesses' self-reports: To the extent that witnesses' implicit theories of influence do not match the actual influence, they will be inaccurate. But if witnesses themselves cannot accurately answer questions about influence, then who can? How can we ever truly discover whether a given real-world witness really was influenced (or would have been influenced) by some lineup practice? The simple answer is we may not be able to at all. Once influenced, one cannot be uninfluenced, despite the assumptions or desires of the legal system. If this sounds pessimistic, it likely is. And it is also realistic. But rather than focusing on what the legal system cannot do, it is important to emphasize what they can do, and to that end we can make one unequivocal recommendation in light of the current experiments: The best way to avoid influenced responses among eyewitnesses is to avoid influencing them in the first place.

References

- 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.
- Brehm, J. (1966). *A theory of psychological reactance*. Oxford, England: Academic Press.
- Charman, S., & Wells, G. L. (2007). Applied lineup theory. In R. C. L. Lindsay, D. Ross, D. Read, & M. Togli (Eds.), *Handbook of eyewitness psychology* (pp. 219–254). Mahwah, NJ: Erlbaum.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.
- Douglass, A. L., & Steblay, N. (2006). Memory distortion in eyewitnesses: A meta-analysis of the post-identification feedback effect. *Applied Cognitive Psychology, 20*, 859–869.
- Dunning, D., & Pargal, M. (1989). Mental addition versus subtraction in counterfactual reasoning: On assessing the impact of personal actions and life events. *Journal of Personality and Social Psychology, 57*, 5–15.
- Epley, N., & Gilovich, T. (2001). Putting adjustment back in the anchoring and adjustment heuristic: Differential processing of self-generated and experimenter-provided anchors. *Psychological Science, 12*, 391–396.
- Fischhoff, B. (1975). Hindsight ≠ foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance, 1*, 288–299.
- Garrioch, L., & Brimacombe, C. A. E. (2001). Lineup administrators' expectations: Their impact on eyewitness confidence. *Law and Human Behavior, 25*, 299–315.
- Gilbert, D., & Malone, P. (1995). The correspondence bias. *Psychological Bulletin, 117*, 21–38.
- Gilbert, D., Pelham, B., & Krull, D. (1988). On cognitive busyness: When person perceivers meet persons perceived. *Journal of Personality and Social Psychology, 54*, 733–740.
- Guilbault, R. L., Bryant, F. B., Brockway, J. H., & Posavac, E. J. (2004). A meta-analysis of research on hindsight bias. *Basic and Applied Social Psychology, 26*, 103–117.
- Hasel, L. E. (2006). *Do people discount testimony by eyewitnesses who received testimony-bolstering feedback?* Unpublished master's thesis, Iowa State University, Ames.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80*, 237–251.
- Kassin, S. M., & Barndollar, K. A. (1992). On the psychology of eyewitness testimony: A comparison of experts and prospective jurors. *Journal of Applied Social Psychology, 22*, 1241–1249.
- Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology, 32*, 322–328.
- Malpass, R. S., & Devine, P. G. (1981). Eyewitness identification: Lineup instructions and the absence of the offender. *Journal of Applied Psychology, 66*, 482–489.
- Martin, L. L., & Stapel, D. A. (1998). Correction and metacognition: Are people naïve dogmatists or naïve empiricists during social judgments? In V. Y. Yzerbyt, G. Lories, & B. Dardenne (Eds.), *Metacognition: Cognitive and social dimensions* (pp. 228–247). Thousand Oaks, CA: Sage.
- Neuschatz, J. S., Preston, E. L., Burkett, A. D., Togli, M. P., Lampinen, J. M., Neuschatz, J. S., Fairless, A. H., Lawson, D. S., Powers, R. A., & Goodsell, C. A. (2005). The effects of post-identification feedback and age on retrospective eyewitness memory. *Applied Cognitive Psychology, 19*, 435–453.
- Newsome v. McCabe, 256 F.3d 747.
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84*, 231–259.
- Petty, R. E., & Wegener, D. T. (1993). Flexible correction processes in social judgment: Correcting for context-induced contrast. *Journal of Experimental Social Psychology, 29*, 137–165.
- Roese, N. J., & Olson, J. M. (1995). *What might have been: The social psychology of counterfactual thinking*. Hillsdale, NJ: Erlbaum.

- Rosenthal, R., & Rosnow, R. L. (1991). *Essentials of behavioral research: Methods and data analysis*. New York: McGraw-Hill.
- Semmler, C., Brewer, N., & Wells, G. L. (2004). Effects of postidentification feedback on eyewitness identification and nonidentification confidence. *Journal of Applied Psychology, 89*, 334–346.
- Stapel, D. A., Martin, L. L., & Schwarz, N. (1998). The smell of bias: What instigates correction processes in social judgments? *Personality and Social Psychology Bulletin, 24*, 797–806.
- Stebly, N. M. (1997). Social influence in eyewitness recall: A meta-analytic review of lineup instruction effects. *Law and Human Behavior, 21*, 283–297.
- Taylor, S. E., & Brown, J. D. (1988). Illusion and well-being: A social psychological perspective on mental health. *Psychological Bulletin, 103*, 193–210.
- Weber, N., Brewer, N., Wells, G., Semmler, C., & Keast, A. (2004). Eyewitness identification accuracy and response latency: The unruly 10–12-second rule. *Journal of Experimental Psychology: Applied, 10*, 139–147.
- Wegener, D. T., & Petty, R. E. (1995). Flexible correction processes in social judgment: The role of naïve theories in corrections for perceived bias. *Journal of Personality and Social Psychology, 69*, 5–15.
- Wegener, D. T., & Petty, R. E. (1997). The flexible correction model: The role of naïve theories of bias in bias correction. In M. Zanna (Ed.), *Advances in experimental social psychology* (pp. 141–208). San Diego, CA: Academic Press.
- 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.
- Wells, G. L., & Bradfield, A. L. (1999). Distortions in eyewitnesses’ recollections: Can the postidentification feedback effect be moderated? *Psychological Science, 10*, 138–144.
- Wells, G. L., Olson, E., & Charman, S. (2003). Distorted retrospective eyewitness reports as functions of feedback and delay. *Journal of Experimental Psychology: Applied, 9*, 42–52.
- Wilson, T. D., & Brekke, N. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin, 116*, 117–142.
- Wilson, T. D., Centerbar, D. B., & Brekke, N. (2002). Mental contamination and the debiasing problem. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 185–200). New York: Cambridge University Press.

Received May 2, 2007

Revision received December 11, 2007

Accepted December 17, 2007 ■