



Published in final edited form as:

Psychon Bull Rev. 2009 February ; 16(1): 88–92. doi:10.3758/PBR.16.1.88.

Surprising feedback improves later memory

Lisa K. Fazio and Elizabeth J. Marsh

Duke University

Abstract

The hypercorrection effect is the finding that high-confidence errors are more likely to be corrected after feedback than are low-confidence errors (Butterfield & Metcalfe, 2001). In two experiments we explored the idea that the hypercorrection effect results from increased attention to surprising feedback. In Experiment 1, participants were more likely to remember the appearance of the presented feedback when the feedback did not match expectations. In Experiment 2, we replicated this effect using more distinctive sources, and also demonstrated the hypercorrection effect in this modified paradigm. Overall, participants better remembered both the surface features and the content of surprising feedback.

People do not have perfect knowledge about the world around them. As we go about our lives and interact with the world, we discover errors in our knowledge that we have to correct. How do we correct these errors? An examination of which false beliefs are easier versus more difficult to update should inform us about the mechanisms of correction. Many theories of memory would hold that errors that you believe very strongly in, those made with high confidence, would be the most difficult to later correct (i.e. McGeogh, 1942; Raaijmakers & Shiffrin, 1981). The argument is that errors made with high confidence are firmly established in our memories, and thus difficult to eradicate from our knowledge base.

Intriguingly, several studies have shown that high-confidence errors are actually more likely to be corrected after feedback than are low-confidence errors. In an early demonstration, participants read short paragraphs about the eye, then answered multiple-choice questions, rated their confidence in each answer, and received feedback about the correct answers (Kulhavy, Yekovick, & Dyer, 1976). On a final multiple-choice test that repeated the same 30 questions, participants corrected more of their high-confidence errors than their low-confidence errors. More recently, Butterfield and Metcalfe (2001) found the same effect with different stimuli. In their experiment, participants answered general world knowledge questions such as “What poison did Socrates take at his execution?” Participants rated their confidence in each response and then were told the correct answer to each question. Similar to Kulhavy et al., Butterfield and Metcalfe (2001) found that high-confidence errors were more likely to be corrected on a retest than were low-confidence errors. The authors named this finding the *Hypercorrection Effect*.

Why is it that these high-confidence errors, which should be firmly established in memory and difficult to update, are instead more likely to be corrected than are low-confidence errors? One possibility is that participants attend more to unexpected feedback, with positive consequences for memory. In other words, when a participant makes an error with high confidence, the feedback is surprising, leading the learner to more deeply encode the

feedback. This hypothesis is similar Kulhavy's model of how feedback affects learning (Kulhavy, 1977; Kulhavy et al., 1976), and owes a debt to Rescorla and Wagner's (1972) model of animal learning (which stated that learning occurs fastest when events violate the organism's expectations). Kulhavy proposed that a large discrepancy between the participant's initial beliefs and the correct answer leads the participant to expend more effort to correct the misunderstanding.

One prediction of this model is that participants should choose to spend more time studying the feedback after a high-confidence error; this was confirmed by Kulhavy (1977). However, the hypercorrection effect occurs even when the duration of the feedback is held constant (as in Butterfield & Metcalfe, 2001), and so the challenge is to find evidence for surprise when differential study times are not possible. Some support comes from neuroimaging data; for example, Butterfield and Mangels (2003) used ERPs to show that high-confidence errors elicited activity in frontal areas that have been linked to novelty in other studies (see Butterfield (2003) for a similar result using fMRI).

When answering questions, the feedback may be surprising in two different situations. In addition to high-confidence errors, an individual should also be surprised when he or she believes a response to be a guess (and rates confidence as low), and yet finds out that the guess was correct. A test-feedback-retest paradigm only allows examination of errors, as low-confidence correct answers do not need to be corrected. But the surprise hypothesis predicts that both situations should have consequences for attention, and in turn for later memory. Butterfield and Metcalfe (2006) used this logic in a pair of experiments in which participants did a tone detection task while answering the initial general knowledge questions, and then were retested with full attention. During the initial test, participants simply had to press a key whenever they heard a tone; critical was participants' ability to detect tones played concurrently with feedback. Surprising feedback was presumed to divert attention from the tone detection task. Consistent with this, participants missed more tones when the feedback revealed an error made with high confidence. For correct answers, tone detection was better for high-confidence responses than for correct guesses. Overall, tone detection was negatively related to performance on the retest, suggesting that participants encoded the feedback at the expense of detecting the tone.

Our research also takes advantage of the fact that the surprise hypothesis predicts increased attention (and memory) for both high-confidence errors and low-confidence correct answers. However, instead of using distraction from another task to infer attention to the feedback, we chose a more direct measure of attention to the feedback, one that could be measured for both correct answers and errors: memory for the feedback's appearance. This dependent measure has been used in emotion research, with the result that memory is better for the surface features (e.g., font colors) of attention-grabbing emotional and taboo words than neutral words (Doerkenson & Shimamura, 2001; MacKay & Ahmetzanov, 2005). This is a measure of source memory or memory for the "conditions under which a memory is acquired" (Johnson, Hashtroudi, & Lindsay, 1993, p. 3). We are using a broad definition of source memory that includes everything that gets encoded about the feedback other than its content. Our argument is that source memory will be better when feedback is surprising, for both correct and incorrect answers.

In Experiment 1, participants answered general world knowledge questions, rated their confidence in each answer, and received feedback in the form of the correct answer to each question. Critically, the feedback appeared either in red or green font. After a short delay, participants identified whether each correct answer had been presented in green or red during the feedback phase. If a discrepancy between the participant's expectation and the feedback leads to a deeper encoding of the feedback, then source memory should be better for high-confidence errors *and* low-confidence correct responses, as compared to low-confidence errors or confident correct answers. In other words, when the feedback confirmed participants' beliefs, they should have paid less attention to it, resulting in lower memory for the feedback's appearance. This same relationship was expected in Experiment 2, where male and female voices delivered the feedback. One group of participants completed a source test (as in Experiment 1): for each correct answer, they identified whether it had been spoken in a male or female voice. Other participants were simply retested on the general knowledge questions. Thus Experiment 2 was designed to generalize the relationship between confidence and appearance memory to a different source judgment, as well as to demonstrate the standard hypercorrection effect in our modified paradigm. Because of the similarities between the two experiments, they will be discussed together in a single general discussion.

Experiment 1

Method

Participants—Forty-six Duke University undergraduates participated in the experiment for partial fulfillment of a course requirement. Seventeen additional participants were tested but performed at chance on the source discrimination task; thus, their data was excluded from the analyses. Chance was defined as answering less than 55% of the source questions correctly. None of the participants was color-blind.

Materials—One hundred forty general knowledge questions were selected from the Nelson and Narens (1980) norms. Questions ranged in difficulty; on average, 40% of participants in the norming study answered these items correctly (ranging across items from 0% to 92% correct). The feedback appeared in Times New Roman font. It was either red and italicized in 64 pt. font or green, underlined and bolded in 12 pt. font. All other text was presented in light blue 24 pt. font.

Procedure—The experiment began with a general knowledge test. Participants were told they were to answer a series of questions and rate their confidence in each answer. They were warned that some of the questions would be difficult and that they should make educated guesses, or else respond "I don't know." Furthermore, they were told they would receive feedback on their answers and that they would later take a second test. Critically, the nature of the second test was never mentioned.

Participants typed their response to each question, and rated their confidence using a 7-point scale. Following Butterfield & Metcalfe (2001), the scale ranged from 1 (sure wrong) to 4 (unsure) to 7 (sure correct). The correct answer appeared for 5 seconds after each confidence rating was recorded. This feedback took the form of a sentence, and was presented

regardless of whether the question was answered correctly or not. For example, if the question was “What’s the longest river in South America?” then the feedback was “Amazon is the longest river in South America.” For half of the items, the feedback was presented in the red font, whereas for the other half feedback appeared in the green font.

Immediately following the general knowledge test, participants completed a source test on their memory for the feedback’s appearance. The feedback sentences were tested one at a time, in random order, in the light blue font. For each item, participants identified whether the feedback had been presented previously in red or green font. After the source test, participants were debriefed and thanked for their participation.

Results

Unless otherwise noted, differences were significant at the .05 level.

Initial Test—On the initial test, participants answered an average of 43% of the questions correctly, and their average confidence was 4.11. Participants were well calibrated in their use of the confidence scale; the average within-participant gamma correlation between initial test accuracy and confidence was .78.

Source Test—Participants correctly identified the prior color of the feedback for 69% of the facts.

Of primary interest was the relationship between confidence on the initial test and performance on the source test. The surprise hypothesis predicts a different relationship between confidence and source memory for items answered correctly vs. incorrectly on the initial general knowledge test. For general knowledge questions answered *correctly*, the feedback would have been unexpected for guesses, thus predicting better source memory for low-confidence correct answers than for high-confidence ones. In contrast, for general knowledge questions answered *incorrectly*, the feedback would have been surprising for high-confidence errors, thus predicting better source memory for high-confidence errors than low-confidence ones. In short, the surprise hypothesis predicts a negative relationship between source memory and confidence for items answered correctly on the initial general knowledge test, but a positive relationship for errors.

Figure 1 shows the relationship between source memory and confidence as a function of correctness on the initial general knowledge test. As predicted, source memory was highest when participants’ confidence was mismatched with the accuracy of their original responses. For correct answers, lower confidence on the initial test was associated with better source memory. The mean within-subject gamma correlation between initial confidence and later source memory was significantly negative, $\gamma = -.19$, $t(45) = 2.61$, $SEM = .07$. For incorrect answers, higher confidence was associated with better source memory, $\gamma = .12$, $t(45) = 2.23$, $SEM = .05$.

A series of additional analyses were conducted to ensure that the key results were not due to differential memory for the red font ($M = .72$), which turned out to be more memorable than the green font ($M = .65$), $t(45) = 3.07$, $SEM = .02$. The reader will remember that half of the

feedback statements were presented in red and half in green; but because we could not predict a priori an individual's responses nor their confidence in these responses, we could not counterbalance the font across the 14 cells. Critically, for errors, red and green feedback were not unequally distributed across the seven levels of confidence, $F(6, 270) = 1.05$, $Mse = 2.65$, $p > .3$, and thus better memory for red feedback could not explain the positive relationship between confidence and source memory found for errors. For correct responses, disproportionately more red feedback occurred in the high-confidence cells, $F(6, 270) = 6.95$, $Mse = 2.42$, but this is not concerning as the result predicted (and obtained) was in the opposite direction.

In short, the results were consistent with the surprise hypothesis: the relationship between confidence and source memory was positive for errors but negative for correct answers.

Because Experiment 1 focused on source memory, there was no measure of error correction. That is, a second general knowledge test was not administered after the source test, as the source memory test effectively presented the feedback for a second time. In Experiment 2, one group of participants took the source test and another group was retested on the general knowledge questions, to ensure that the changed paradigm did not eliminate the basic hypercorrection effect. In addition, the sources were made more distinctive in Experiment 2, to minimize loss of participants due to chance performance on the source test.

Experiment 2

Method

Participants—Seventy-two undergraduates participated in the experiment for partial fulfillment of a course requirement. In the final test phase, fifty participants took the source test. Six additional participants were tested in this condition but were excluded because they performed at chance on the source test (chance was defined as in Experiment 1). Twenty-two participants were in the retest condition; one additional participant was tested but excluded because he corrected all of his initial errors on the second test, making it impossible to calculate the relationship between his confidence in the initial error and the probability of it being corrected on the retest.

Materials—We used 120 of the original 140 questions from Experiment 1. Feedback was presented in one of two ways. For half of the items, a female voice read the feedback aloud, while a woman's picture appeared on the left side of the computer screen and the feedback printed in pink lettering appeared on the right. For the other half of the items, the voice was male and the computer screen showed a man's picture on the right and the feedback in blue lettering on the left.

Procedure—As in Experiment 1, participants answered a series of general knowledge questions, rated their confidence in each response, and then received feedback. To improve source memory, the feedback appeared for 6 seconds instead of 5 seconds as in Experiment 1.

After the general knowledge test, participants in the source memory condition immediately began the source test. The feedback sentences appeared on the screen in a neutral font and the participants identified whether the male or the female source had presented the feedback.

Participants in the retest condition solved visuo-spatial puzzles for 4 minutes before taking their final test, as pilot testing showed that participants were at ceiling without a short filler task. These participants then retook the general knowledge test, which was identical to the first test except no feedback was provided.

Results

Unless otherwise noted, differences were significant at the .05 level.

Initial Test—Performance on the initial test did not differ across the two conditions. Participants correctly answered 42% of the initial questions in the source condition and 43% in the retest condition, $t < 1$. Confidence on the initial test was also similar across the conditions, averaging 4.01 in the source condition and 4.21 in the general knowledge retest condition, $t < 1$. These values are similar to what was observed in Experiment 1, as were participants' confidence-accuracy correlations. The average within-subject gamma correlation between proportion correct and confidence on the initial test was .81 in the source condition and .76 in the retest condition, $t(70) = 1.35$, $SEM = .04$, $p = .18$.

General Knowledge Retest—For the participants in the general knowledge retest condition, we compared performance on the initial test to performance on the final test. Feedback improved performance, with participants answering 80% of the questions correctly on the second test as compared to 43% on the initial test, $t(21) = 26.17$, $SEM = .01$.

Of primary interest was whether the hypercorrection effect occurred. For each of the seven confidence levels on the first test, we examined the proportion of errors that were successfully corrected on the second test. Figure 2 shows hypercorrection: participants corrected more of the errors that had been committed with high confidence than those made with low confidence. The mean within-subject gamma correlation between initial confidence and proportion of errors later corrected was significantly positive, $\gamma = .23$, $t(21) = 2.27$, $SEM = .10$.

Source Test—For the participants in the source condition, we examined memory for the source of the feedback. On average, participants correctly identified the source for 68% of the facts.

As in Experiment 1, of primary interest was the relationship between confidence on the initial test and later memory for the source of the feedback. Replicating Experiment 1, there was a negative relationship between confidence and source memory for correct answers and a positive relationship between confidence and source memory for errors, as shown in Figure 3. After answering a question correctly, participants were more likely to remember the source of the feedback if they had answered with low confidence than if they had answered with high confidence $\gamma = -.28$, $t(49) = 4.24$, $SEM = .07$. The pattern was opposite for errors; participants were more likely to remember the source of the feedback if they had

answered with high confidence than if they had answered with low confidence $\gamma = .12$, $t(49) = 2.18$, $SEM = .06$.

As in Experiment 1, we conducted additional analyses to ensure that our results were not due to one source being more memorable than the other. We found that the male source was more likely to be accurately identified ($M = .70$) than the female source ($M = .66$), $t(49) = 2.53$, $SEM = .02$. However, the male and female feedback was not unequally distributed across the confidence levels for correct answers, $F(6, 294) = 1.80$, $Mse = 2.29$, $p > .1$ or for errors, $F < 1$. Thus better memory for the male source cannot explain our results.

General Discussion

In two experiments surprising feedback improved memory for both the surface features and the content of presented feedback. In Experiment 1, participants were better able to remember the color of feedback when it was incongruent with their expectations. That is, source memory was better for feedback that had been presented in response to correct guesses or errors made with high confidence. In Experiment 2, participants showed improved memory for both the content and the source of the feedback. Participants were more likely to correct high-confidence errors than low-confidence errors and they were more likely to remember the source of the feedback when it was unexpected.

While the observed relationships between initial confidence and source memory were relatively small, they were as predicted in both experiments and occurred for both correct and erroneous answers. It is not surprising that the effects were smaller ones given that remembering the appearance of the feedback was not participants' main task. The participants in both experiments were lead to believe that they would be retested on the general knowledge questions - the source memory test was unexpected. Thus, most of the participants' additional attention should have been, and was, directed towards the content of the surprising feedback, rather than its surface features. This can be seen most clearly in Experiment 2 where memory for the content of the feedback (the correct answer) increased more than 10% across the confidence levels, while source memory increased less than 5%.

These experiments support the surprise hypothesis, which states that unexpected feedback leads to a greater expenditure of effort to encode that feedback, with positive consequences for memory. Data across laboratories are converging in support of the surprise hypothesis. Putting these results together, a consistent picture is emerging: feedback can be surprising (Butterfield & Mangels, 2003), leading to a focus on the feedback (the present studies) at the expense of other tasks (Butterfield & Metcalfe, 2006).

In addition to the surprise hypothesis, there is at least one other possible explanation of the hypercorrection effect. The knowledge hypothesis posits that confidence tends to be correlated with how much a participant knows generally about the target domain (Butterfield & Metcalfe, 2001). The argument is that if participants have little knowledge about a domain, then they have nothing with which to associate the incoming information. In other words, it will be more difficult to integrate the correct answer into their semantic memory if it is an unfamiliar domain. Although our experiments were not designed to test the knowledge hypothesis, it is not immediately clear what the knowledge hypothesis would

predict about memory for the source of the feedback. In particular we doubt that the knowledge hypothesis would predict a negative relationship between source memory and confidence in correct answers. Of course our data do not rule out the knowledge hypothesis, as the two hypotheses are not mutually exclusive. It is quite plausible that knowledge updating requires both deep encoding of the feedback and a knowledge structure that allows the new information to be easily assimilated – but our data suggest that differences in domain knowledge are unlikely to be solely responsible for the hypercorrection effect.

It is important that people are able to accurately update their general world knowledge. We believe that confidence judgments play an important role in dictating which errors are most essential to correct. Because confidence judgments are in general a valid indicator of overall accuracy (Brewer, Sampaio, & Barlow, 2005; Perfect, Watson, & Wagstaff, 1993), it is informative when there is a conflict between a person's confidence and the actual answer. The contradictory information tells the person that something is seriously wrong with his or her knowledge structure. It is the importance of this miscalibration that causes the feedback to be better processed and better remembered.

Acknowledgments

This work was supported by a collaborative activity award from the James S. McDonnell foundation. We thank Barbie Huelser for help with data collection.

References

- Brewer WF, Sampaio C, Barlow MR. Confidence and accuracy in the recall of deceptive and nondeceptive sentences. *Journal of Memory & Language*. 2005; 52:618–627.
- Butterfield B. The hypercorrection effect and its neural correlates. *Dissertation Abstracts International*. 2003; 66(05)
- Butterfield B, Mangels JA. Neural correlates of error detection and correction in a semantic retrieval task. *Cognitive Brain Research*. 2003; 17(3):793–817. [PubMed: 14561464]
- Butterfield B, Metcalfe J. Errors Committed With High Confidence Are Hypercorrected. *Journal of Experimental Psychology: Learning, Memory, & Cognition*. 2001; 27(6):1491–1494.
- Butterfield B, Metcalfe J. The Correction of Errors Committed with High Confidence. *Metacognition and Learning*. 2006; 1(1):69–84.
- Doerkenson S, Shimamura AP. Source memory enhancement for emotional words. *Emotion*. 2001; 1(1):5–11. [PubMed: 12894807]
- Johnson MK, Hashtroudi S, Lindsay DS. Source monitoring. *Psychological Bulletin*. 1993; 114:3–28. [PubMed: 8346328]
- Kulhavy RW. Feedback in Written Instruction. *Review of Educational Research*. 1977; 47(1):211–232.
- Kulhavy RW, Yekovick FR, Dyer JW. Feedback and Response Confidence. *Journal of Educational Psychology*. 1976; 68(5):522–528.
- MacKay DG, Ahmetzanov MV. Emotion, Memory, and Attention in the Taboo Stroop Paradigm An Experimental Analogue of Flashbulb Memories. *Psychological Science*. 2005; 16(1):25–32. [PubMed: 15660848]
- McGeogh, JA. *The psychology of human learning*. New York: Longmans, Green; 1942.
- Nelson TO, Narens L. Norms of 300 general-information questions: Accuracy of recall, latency of recall, and feeling-of-knowledge ratings. *Journal of Verbal Learning & Verbal Behavior*. 1980; 19:338–368.
- Perfect TJ, Watson EL, Wagstaff G. Accuracy of confidence ratings associated with general knowledge and eyewitness memory. *Journal of Applied Psychology*. 1993; 78:144–147.

- Raaijmakers JGW, Shiffrin RM. Search of associative memory. *Psychological Review*. 1981; 88:93–134.
- Rescorla, RA.; Wagner, AR. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Black, AH.; Prokasy, WF., editors. *Classical Conditioning II: Current Research and Theory*. New York: Appleton-Century-Crofts; 1972. p. 64-99.




Figure 1. Average proportion correct on the source test for each confidence level, as a function of whether the answer on the initial test was correct or incorrect. Bolded lines are the best fitting trend lines. (Experiment 1)

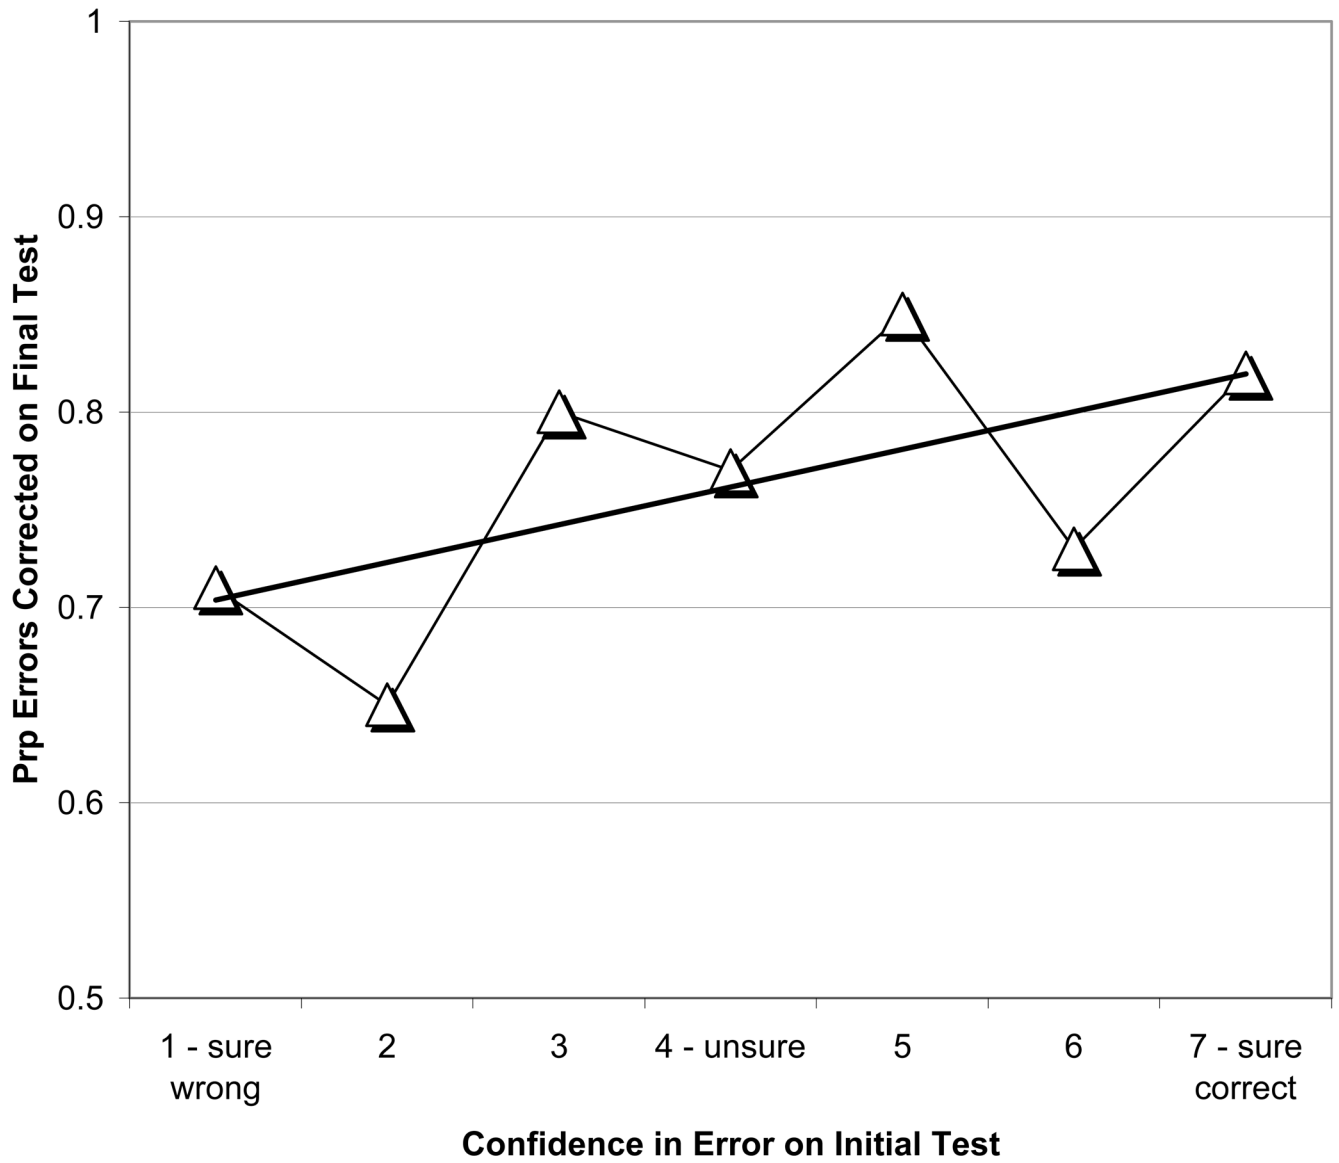


Figure 2. Average proportion of the errors on the first test that were corrected on the second test for each confidence level. Bolded line is the best fitting trend line. (Experiment 2)

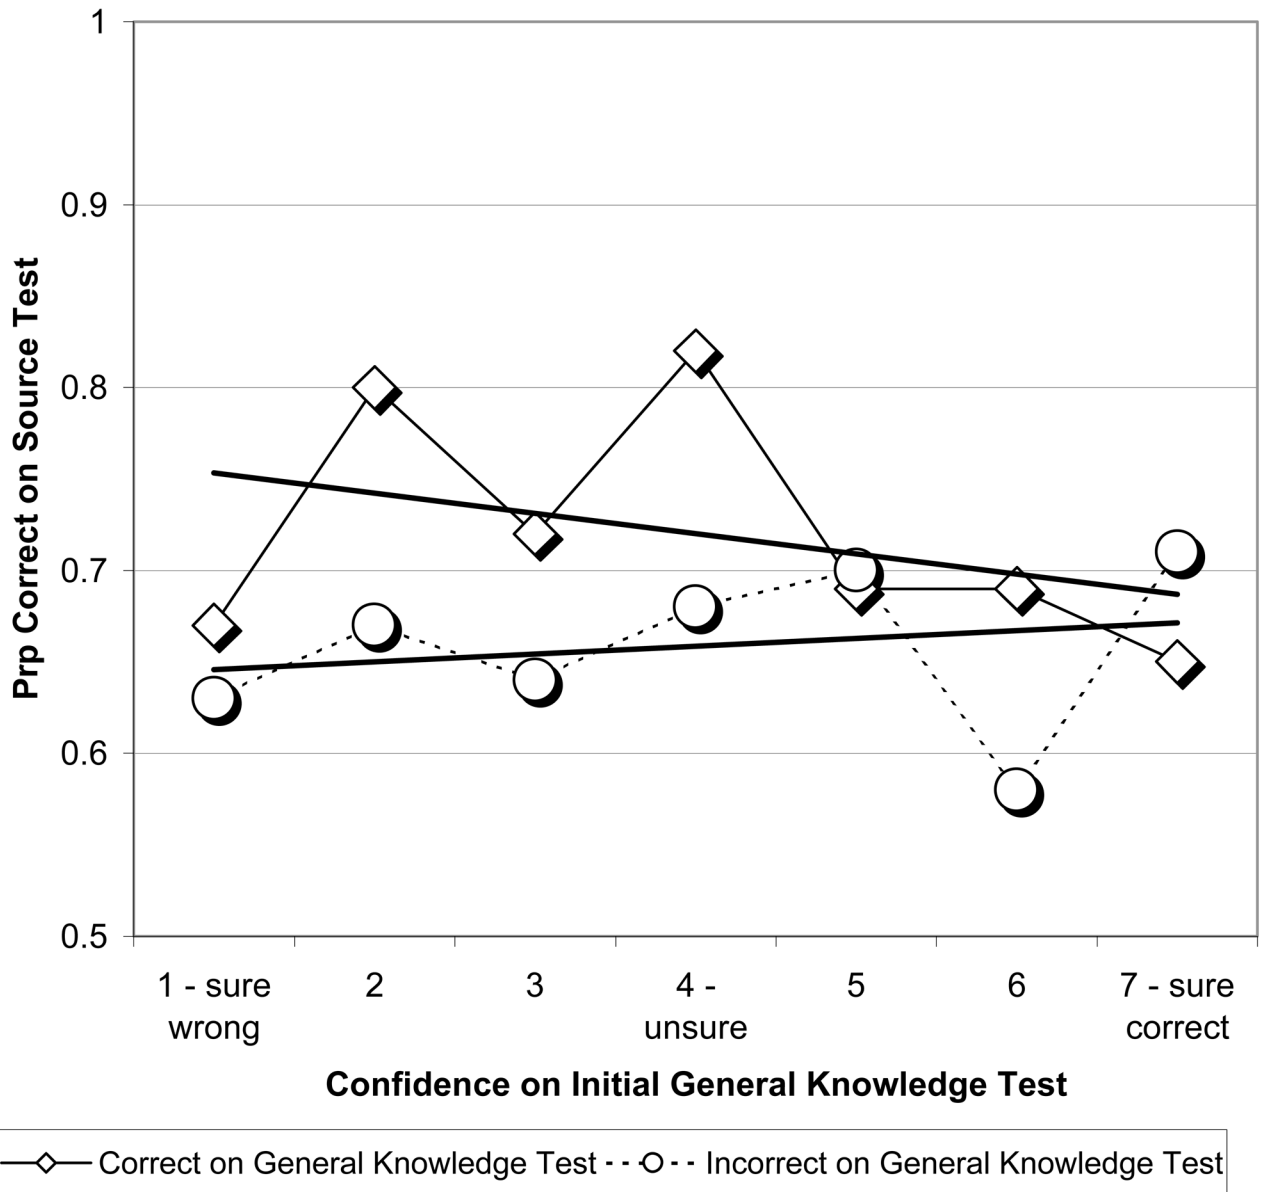


Figure 3. Average proportion correct on the source test for each confidence level split by whether the answer on the initial test was correct or incorrect. Bolded lines are the best fitting trend lines. (Experiment 2)