But where's the evidence? The effect of explanatory corrections on inferences about false information

Saoirse Connor Desai (saoirse.connor-desai@city.ac.uk) & Stian Reimers (stian.reimers.1@city.ac.uk)
City, University of London, Department of Psychology, Whiskin Street
London, EC1R 0JD

Abstract

Research on the *continued influence effect* has consistently shown that people continue to rely on false causal information despite being corrected by more recent information. Corrections are most effective when paired with an alternative explanation that 'fills the causal gap' left by the correction. However, it may not always be possible provide an alternative explanation. Previous research suggests people more readily discount unreliable information. Two experiments examined whether corrections to false causal information in a news report are more effective when the correction explains why the source of the false information was unreliable. The results showed that a correction did not fully eliminate reliance on false information and that an explanatory correction was no more effective than a non-explanatory correction. People also continued to rely on false information when there was limited information to support its validity. Possible explanations for the ineffectiveness of explanatory corrections are discussed.

Keywords: False Information; Continued Influence; Corrections; Inference; Explanations; Reasoning; Memory

Introduction

Many news organizations now report breaking news via social media platforms. Although social media helps to keep people up to date, breaking news can be based on mistaken, inaccurate, or incomplete information. When false information[1] is reported news organizations would ordinarily issue a correction, revising their original account. Provided that the false information has not proliferated, being aware of a correction should normatively neutralize belief in the false information. In contrast, numerous experiments on the *continued influence effect* (CIE) have shown that causal false information continues to be influential beyond a correction (Ecker, Lewandowsky, & Apai, 2011; Ecker, Lewandowsky, Swire, & Chang, 2011; Ecker, Lewandowsky, & Tang, 2010; Johnson & Seifert, 1994; Wilkes & Leatherbarrow, 1988).

The standard experimental paradigm for studying the CIE involves reading a series of messages describing a fictional news story over time. Explanatory target information is presented and subsequently corrected for one group of participants, but remains uncorrected for a control group. Inferences and memory for the news report are then assessed through a series of open-ended questions. In Johnson and Seifert (1994), participants read a story about a warehouse fire wherein target information implies that carelessly stored flammable materials (oil paint and gas cylinders), were a likely cause of the fire. Later in the story, some participants learn that no such materials had actually been found. A comprehension test follows, which includes indirect inference questions (e.g., "what could have caused the explosions?"), and questions assessing recall of basic facts (e.g., "what was the cost of the damage done?"). Inference responses are then coded in order to measure the extent to which the target information (oil paint and gas cylinder) has been discounted. Responses are coded according to whether they are consistent with the explanatory theme implied by the target information (e.g., "exploding gas cylinders") or not (e.g., "electrical short circuit").

The key finding from CIE studies (for reviews see (Lewandowsky, Ecker, Seifert, Schwarz, & Cook, 2012; Seifert, 2014) is that corrections do not fully eliminate reliance on false information. People continue to rely on false information despite recalling the correction, when given prior warnings about false information in news reports; whether corrections are repeated, or appear immediately after the false information (Ecker, Lewandowsky, & Apai, 2011; Ecker et al., 2011; Ecker et al., 2010; Johnson & Seifert, 1994). The CIE has been replicated using various types of news stories, types of false information (e.g., Ecker et al., 2011), and using direct or indirect measures of reliance on false information (Connor Desai & Reimers, 2016; Rich & Zaragoza, 2016). Identifying the cognitive mechanisms underlying the successful correction of false information has timely, real-world implications in a wide variety of domains (e.g., news stories, public health information, in the courtroom).

Filling the causal gap One explanation for the CIE is that corrections are ineffective because a correction alone leaves a causal gap in a person's mental model of the reported event (e.g., Johnson & Seifert, 1994; Wilkes & Leatherbarrow, 1988; Lewandowsky et al., 2012). In this view, people maintain the false information because they prefer an inconsistent to an incoherent mental event model. In the warehouse fire example, an individual might infer that a fire started by an electrical short circuit was a result of negligence, based on information suggesting that flammable liquids were carelessly stored. Correcting a key piece of causal information (i.e., no flammable liquids) results in an incoherent mental model. People might continue to draw causal inferences from the false information because it is the only explanation available to them. In line with the mental models account, combining a correction with an alternative

[1] In this context the term 'false information' refers to incorrect or inaccurate information that is initially presented as true.

explanation to 'fill the causal gap' considerably reduces the degree to which people rely on false information (e.g., there was evidence the fire was caused by arson; Ecker et al., 2011; Johnson & Seifert, 1994a; Rich & Zaragoza, 2016; Tenney, Cleary, & Spellman, 2009).

In the real world it is not always possible to provide a single, coherent, alternative explanation to replace corrected false information (e.g., the true cause the Flight MH370 disappearance still remains unknown). Due to the fact that alternative explanations are not always available, it is important to identify other means of increasing the effectiveness of corrections.

Explanatory corrections One way of increasing the impact of corrections is to explain why the original information is no longer relevant or useful. For example, Bush, Johnson, and Seifert (1994) found that explaining that target information had been of poor quality (*the storeroom actually contained cans of coffee and soda canisters*), or was no longer relevant (*a delivery of paint and gas cylinders was expected but never arrived*), enhanced the effectiveness of the correction statement compared to a correction alone, but an explanatory correction was still not as effective an offering an alternative explanation. Bush et al., also found that ruling out the involvement of the corrected information (*there was clear evidence that no paint or gas were ever on the premises*) without providing an explanation actually decreased the effectiveness of the correction. These findings can be understood by the pragmatic inferences people draw about the conversational implications of the original statement (cf. Seifert, 2014). The validity of corrected information might be reinforced because people assume that a speakers only offer true (*maxim of quality*) and relevant (*maxim of relevance*) information (Grice, 1975). Bush et al's findings suggest that the person issuing the correction must explain why the original information should no longer be believed in order ensure the correction is understood. Legal decision-making studies support the idea that explaining why initial information is unreliable can enhance the effectiveness of a correction. For example, Kassin and Sommers (1997) found that mock-jurors who learned a key piece of incriminating evidence was inadmissible because it was unreliable (*a taped confession secured without a warrant*) were more likely to convict a defendant than mock-jurors who were told that the evidence was unreliable (*the tape was inaudible*). Similarly, Fein, McCloskey and Tomlinson (1997) found mock-jurors discounted inadmissible incriminating testimony when its reliability was called into question. Finally, Lagnado and Harvey (2008) showed that people providing evidence that an eyewitness has a 'longstanding grudge' against the suspect resulted in participants discounting that testimony. These studies suggest that explanatory corrections could be as effective as combining a correction with an alternative explanation when the correction explains why the initial source of the false information is unreliable.

Pilot study

The pilot study tested whether explanatory corrections are more effective than a correction alone when the correction explains why the original source of the false information is unreliable (i.e., mistaken or intentionally deceptive). There were two main predictions: 1) Explanatory correction groups would produce fewer target information consistent inferences than the correction alone group, and 2) Correction only group would produce fewer target information consistent inferences than a group who was never exposed to a correction.

Methods

Participants Forty-five U.S. based participants were recruited from Amazon Mechanical Turk (17 female, age 36.3±9.70). Participants were paid $1 and took an average of 14 minutes to complete the experiment.

Design Participants were randomly assigned to either the no correction (11), correction only (10), explanatory correction error (10), or explanatory correction lie (14) correction groups. There were four main dependent measures: 1) references to target information on inference questions, 2) recall on filler items, and 3) awareness of the correction.

Materials and Procedure Participants read a news story describing a warehouse fire, displayed as a series of sequentially presented short messages. Materials were reconstructed from an experiment by (Johnson & Seifert, 1994; Exp 3a). There were 12 discrete messages (1, target message, 1 critical message, 1 causal detail message, 9 additional messages), in the style of 'Tweets' from the social media platform Twitter, an approach inspired by Hardwicke, Manning and Shanks (2016). The 'Tweets' originated from the same fictional news outlet, called "news now" and each message was no longer than 140 characters. Messages appeared one a time for a minimum of 5 seconds each; there was no maximum time. Participants clicked a button to proceed to the next message; they were unable to return and view previous messages.

Participants completed an instructional attentional check (e.g., Oppenheimer, Meyvis, & Davidenko, 2009) before starting the experiment. The explanatory theme implied by the target message was that flammable materials had been carelessly stored in a storeroom. The *target message*, containing information about a possible cause of the fire (*there were cans of oil paint and gas cylinders present in a storeroom*), was presented at Message 5. The *causal detail* containing information consistent with the explanatory theme implied in the target message (*thick, oily smoke + sheets of flames hinder firefighters efforts, intense heat has made the fire difficult to bring under control*) appeared at Message 8. The *critical message* varied depending on condition and appeared at Message 11. The remaining (filler) messages provided event information, which was neutral with respect to the explanatory theme implied by the target message (e.g., Three warehouse workers working overtime, have been taken to St Columbus Hospital, due to smoke inhalation).

Table 1 Example questions and responses from pilot study

Question	Example response to receive score of 1 on false information measure
What aspect of the fire should the police focus on in their investigation?	They should focus on the chemical aspects because it seemed to have started from paint or gas.
What was the most likely overall cause of the fire?	The oil paint cans and pressurized gas cylinders.
Is there any evidence of careless management in relation to this fire?	Yes the pressurized cylinders should not have been kept indoors next to paint cans.

In the three correction conditions Message 11 corrected earlier information about the contents of the storeroom; participants in the no correction condition learned instead that warehouse workers taken to hospital had been released. The explanatory correction groups either learned that the target information had been corrected because an employee confused the soda canisters and coffee cans for paint and gas (*error*) or that an employee lied that there were flammable materials in the storeroom (*lie*). There were four narrative versions in total.

After reading all of the 'Tweets' participants completed a questionnaire consisting of seven inference questions, seven filler questions and two questions assessing awareness and understanding of the correction. Inference and filler questions were presented in a random order. Inference questions asked participants about information not explicitly mentioned in the news report (e.g., "Is there any evidence of careless management in relation to this fire?"), and included a question querying participants about what they thought the most likely cause of the fire was. Filler questions enquired about the explicit details included in additional (filler) messages included in the news story (e.g., "Which hospital were the workers taken to?"). Two further questions assessed awareness and understanding of the correction message. Participants typed a response to each of 16 questions in a text box, were required to use a minimum of 25 characters, and encouraged to answer using full sentences.

Pilot study: Results

Coding of Responses

The main dependent variable extracted from inference question responses was 'references to target information'. References that explicitly stated, or strongly implied, that the fire was caused by gas and oil paints were scored a 1 on the target information measure, and were otherwise scored as 0. The maximum individual score for inference questions was 7. Filler question responses were scored for accuracy. Correct or partially correct responses were scored 1 and a score 0 was given for an incorrect response. The maximum individual score for filler questions was also 7. Awareness

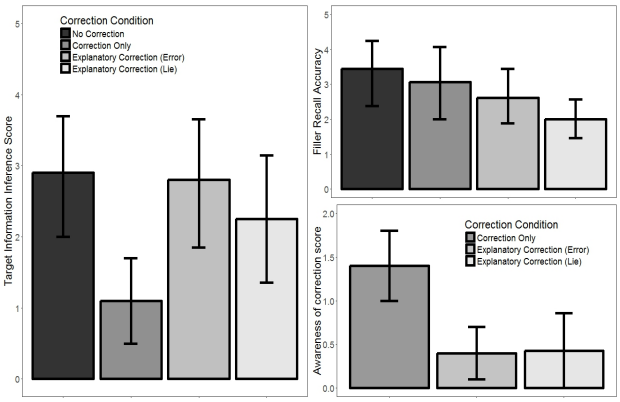

Figure 1: Mean target information inference scores (*left panel*), filler accuracy scores (*top right panel*), and awareness of correction scores (*bottom right panel*) as a function of correction. Error bars represent 95% confidence interval of the mean.

of correction scores were computed using the same criteria; the maximum individual awareness of correction score was 2. One-way ANOVA analyzed differences between the correction conditions for all three measures. [2] Fig 1 shows mean inference, filler and awareness of correction as a function of correction condition.

Inference scores There was significant effect of correction on the number references to target information, F (3, 41) = 3.32, p <. 05, η^2 = .20. Planned contrasts showed a correction reduced references to target information compared to no correction, t (19) = -2.98, p < .01, d = 1.46. However, neither an error explanatory correction, t (19) = -0.55, d = 0.23, nor a lie explanatory correction, t (23) = -1.45, d = 0.53, reduced references to target information compared to no correction.

Filler recall accuracy There was a significant effect of correction on filler recall accuracy, F (3, 41) = 4.39, p < .01, η^2 = .24. Tukey's tests showed the no correction group recalled significantly more filler details than the lie explanatory correction group, t (23) = -3.34, p =. 009, d =1.14. None of the other differences were significant (p's > .05).

Awareness of correction There was a significant effect of correction on awareness of correction scores, F (2, 31) = 6.52, p < .01, η^2 = .30. Tukey's tests revealed that the correction only group showed more awareness of the correction than the error, t (18) = -3.09, p = .01 d = 1.62, or lie, t (22) = -3.24, p < .01, d = 1.22 explanatory correction groups. The two explanatory correction groups did not significantly differ, p = .10.

[2] Planned contrasts are reported for predicted differences. Tukey's post-hoc tests are reported when no difference between conditions was predicted.

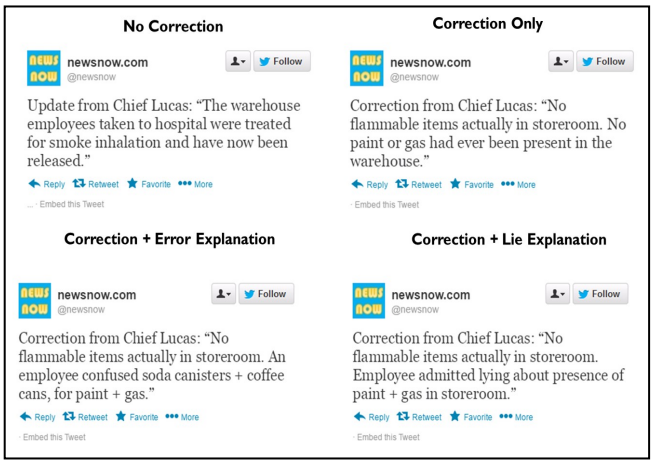

Figure 2: Content of critical messages in main experiment. In contrast to previous studies, the critical message in each of the correction conditions explicitly stated that the message was a correction.

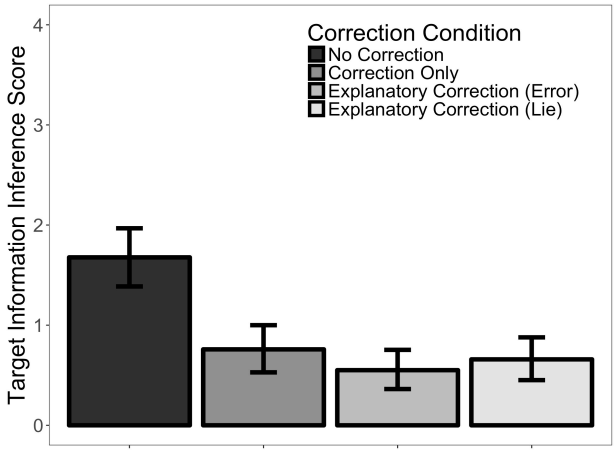

Figure 3: Mean target information inference score as a function of condition in main experiment. Error bars represent the 95% confidence interval of the mean.

Interim Discussion

A correction alone reduced, but did not fully eliminate, target information references when compared to no correction. Pilot results also showed that an explanatory correction did not reduce references to target information compared to no correction. On average both explanatory correction groups made more target information references than the correction only group, although these differences were not significant. These results are inconsistent with previous findings showing that an explanatory correction was more effective at reducing reliance on target information than a correction alone (cf. Bush, Johnson & Seifert, 1994).

The most likely reason that an explanatory correction was less effective than a correction alone is that participants in the explanatory correction groups showed poorer awareness and understanding of the correction, than the correction only group. Only 40% of the error explanatory correction group, and 21% of the lie explanatory correction group understood and were aware of the correction, compared to 90% of the correction only group. Both explanatory correction groups also recalled fewer story details on average than the correction only group. Some participants' responses indicated doubts about the credibility of the correction message (e.g., questioning whether the employee really lied about the contents of the storeroom), and other responses suggested misunderstanding of the correction message (e.g., the employee thought there was soda and coffee but there was actually paint and gas). A lack of clarity of the explanatory correction messages could explain poorer awareness and understanding in explanatory correction conditions could also explain why the current results do not replicate previous findings (cf. Bush et al., 1994). The main experiment sought to rectify these issues by enhancing the clarity of the correction messages.

Main Experiment

The same general setup was employed in the main experiment except that a number of changes were made to rule out explanations identified in the interim discussion. The hypotheses and predictions were also the same as the pilot study.

Participants Three-hundred and twelve U.S. based participants were recruited from Amazon Mechanical Turk (146 female, age 39.67±12.31). Participants were paid $1 and took an average of 20 minutes to complete the experiment.

Design, materials and procedure Participants were randomly assigned to either the no correction (71), correction only (87), explanatory correction error (71), or explanatory correction lie (83) groups. Dependent measures were the same as in the pilot study.

Content of the critical messages was modified from the pilot study in order to make it unequivocally clear that the target information was being corrected (see Fig 2). Unlike previous studies, the critical message for the correction conditions explicitly stated that the target information was being corrected.

Results

Additional coding and analysis was performed on one of the filler questions to the total number of references indicating flammable substances had been in the storeroom before the fire. The additional 'discounting' measure further assessed the extent to which the false information had been disregarded. Responses were scored 1 if the response indicated there were flammable substances in the storeroom before the fire and 0 otherwise. One-way ANOVA analyzed differences between conditions for all four dependent measures.

Inference scores There was a significant effect of correction on references to the target information, F (3, 308) = 23.23, p < .001, η^2 = .18. Planned contrasts revealed a correction significantly reduced the number of references to target information, t (156) = -6.84, p < .001, d = 0.98.

Likewise, the error, t (140) = -6.90, p < .001, d = 1.02, and lie, t (152) = -6.99, p <. 001, d = 1.01., explanatory correction groups, made significantly fewer references to the target information than the no correction group. Mean target information inference scores are shown in Fig 3.

Filler recall accuracy There was no effect of correction on filler recall accuracy, F (3, 308) = 0.64, p = .90, η^2 = .01, so it was not necessary to perform contrast analysis. Mean filler recall scores ranged from 4.58 to 4.96 (out of 7).

Awareness of correction No correction group responses were excluded from analysis because their responses to awareness of correction questions were meaningless. There was a significant effect of correction condition on awareness of correction scores**,** F (2, 238) = 3.76, p < .05, η^2 = .03. Tukey's tests showed a significant difference between the correction only and explanatory correction error group, t (168) = -2.39, p = .05, d = 0.37. There were non-significant differences between the explanatory correction error and correction only groups (p =.06), and both the explanatory corrections groups (p = 1). Given the small effect size the difference is considered negligible.

Discounting false information The inference scores suggest that explanatory corrections are treated the same as a correction alone. If this is the case, then the number of references indicating the storeroom contained flammable substances before the fire should be equivalent to inference scores. There was a significant effect of correction on the number of responses indicating the storeroom contained flammable substances before the fire, F (3, 308) = 57.25, p < .001, $\eta 2$ = .36. Planned contrasts confirmed the same pattern of results as inference scores; there were significantly higher number of references stating that flammable substances had been in the storeroom before the fire in the no correction than the correction only group, t (156) = 9.22, p < .001, d = 1.49, the error correction group, t (140) = -12.27, p < .001, d = 2.81, or the correction lie group, t (152) = -9.99, p < .001, d = 1.68. A closer inspection of the responses suggested that explanatory corrections were not treated the same as a correction alone. Fig 4 shows the mean number of responses indicating that flammable substances were in storeroom before the fire. Tukey's tests showed that the explanatory correction error group significantly differed to the correction only, t (156) = -3.66, p < .01, d = 0.55, and explanatory correction lie group, t (152) = -2.75, d = -.45. The difference between the correction only and explanatory correction lie groups was not significant, p = .80.

Discussion

The results show a clear continued influence effect; a correction significantly reduced, but did not eliminate, references to target information. A correction appeared to have a similar impact on inferences whether accompanied by an explanation as to why the original source of the false information should not be trusted, or not. A closer inspection of responses suggested fewer people continued to think that flammable substances had been in the storeroom before the fire when the correction replaced the contents of the storeroom (i.e., there were soda cans and gas canisters in the storeroom) than when the correction left the storeroom empty before the fire (i.e., the employee had lied about flammable materials in the storeroom). In addition, the continued influence effect was still observed despite the fact that the correction to target information was explicitly stated in the correction message.

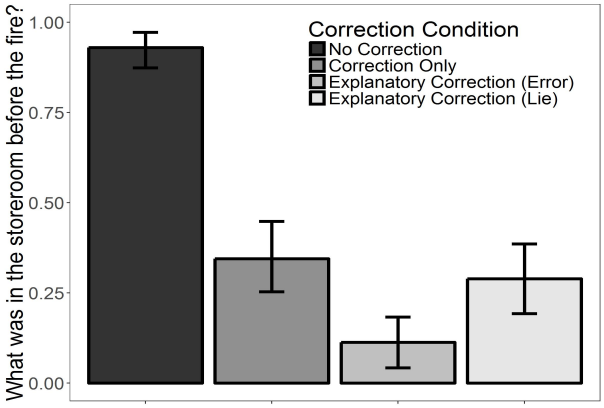

Figure 4: Mean number of references to presence of flammable substances in the storeroom before the fire as a function of correction condition. Error bars represent 95% confidence interval of the mean.

General Discussion

The experiments reported in this paper examined the impact of explanatory corrections on inferences about false information in the context of breaking news reports on social media. The findings reported here are consistent with previous studies showing that corrections do not fully eliminate reliance on false information (Ecker, Lewandowsky, Swire, et al., 2011; Ecker et al., 2010; Johnson & Seifert, 1994; Wilkes & Leatherbarrow, 1988). These results also provide a novel contribution to the literature on the continued influence effect. Specifically, a correction that explained why the original source of the false information was unreliable was no more effective in reducing reliance on false information, than a correction alone. Participants made an equivalent number of references to target information whether the correction provided an explanation for why the target information should no longer be believed (i.e., the current inaccuracy of the target information was directly attributed to a mistaken or a deceptive individual), or not.

Despite this finding there was evidence to suggest that corrections were not treated equally. People were less likely to say that flammable substances (*oil paint and gas canisters*) were in the storeroom before the fire when the contents of the storeroom were replaced with other objects (*soda canisters and coffee cans*) than when the contents of the storeroom were not replaced (i.e., the employee lied that there were flammable items in the storeroom). One explanation for this inconsistency between inferences and

memory of the storeroom contents is that people who received the error correction updated their representation of the contents of the room whilst maintaining an inconsistent mental model of the event. In contrast, people did not update their representation of the storeroom when the contents were not replaced with alternative materials. These results do not support previous findings showing that explanatory corrections are more effective than a correction alone (Bush, Johnson, & Seifert, 1994).

The main methodological difference between the current study and previous study is that the explanatory correction conditions in this study involved an additional source of information. In addition to making a judgment about whether the correction sufficiently negated the false information participants had to establish why or how the original source of the information (i.e., the employee) provided the information in the first place. Without knowing why the employee lied about the flammable materials or how the employee was able to confuse flammable for non-flammable substances, people might still assume the false information is relevant. Another possible reason for the inconsistent results could be that in at least one of Bush et al's conditions the correction made it logically impossible to continue to rely on the false information whereas this was not the case in the current study. These findings further demonstrate that pragmatic inferences play an important role in successfully correcting false information.

The current studies also showed evidence of the continued to rely on false information even though the report only contained one piece of information that reinforced the false information explanatory theme. This suggests that people construct a mental model of the incident on the basis of limited causal information. If there is no information to indicate an alternative explanation then people fall back on the only explanatory information available to them. It is also possible that people interpret (or re-interpret) information as supporting their leading hypothesis (e.g., Carlson & Russo, 2001). Future studies are necessary to address whether people re-interpret neutral information to fit false causal information or whether people construct their mental event model based on limited information.

While the current study provides initial steps, there is a lot more left to explore. It will be necessary to further explore why explaining why the information was unreliable was no more effective than withdrawing the false information and why the current findings are discrepant with previous continued influence (Bush et al., 1994) and legal decision making studies (e.g., Lagnado & Harvey). Future studies will further investigate the role of source reliability in correcting false information, and use a wider range of scenarios as well as types of false information.

Acknowledgments

We thank Dr. Tom Hardwicke for his insights and expertise regarding the continued influence effect paradigm.

References

Bush, J. G., Johnson, H. M., & Seifert, C. M. (1994). The implications of corrections: Then why did you mention it? In *Proceedings of the Sixteenth Annual Cognitive Science Society Conference, Atlanta, Georgia* (Vol. 27, pp. 112–117). Atlanta.

Carlson, K. a, & Russo, J. E. (2001). Biased interpretation of evidence by mock jurors. *Journal of Experimental Psychology. Applied*, *7*(2), 91–103. http://doi.org/10.1037/1076-898X.7.2.91

Connor Desai, S., & Reimers, S. (2016). *Direct and indirect measures of the Continued Influence Effect*. Unpublished manuscript.

Ecker, U. K. H., Lewandowsky, S., & Apai, J. (2011). Terrorists brought down the plane!--No, actually it was a technical fault: processing corrections of emotive information. *Quarterly Journal of Experimental Psychology (2006)*, *64*(2), 283–310. http://doi.org/10.1080/17470218.2010.497927

Ecker, U. K. H., Lewandowsky, S., Swire, B., & Chang, D. (2011). Correcting false information in memory: manipulating the strength of misinformation encoding and its retraction. *Psychonomic Bulletin & Review*, *18*(3), 570–578. http://doi.org/10.3758/s13423-011-0065-1

Ecker, U. K. H., Lewandowsky, S., & Tang, D. T. W. (2010). Explicit warnings reduce but do not eliminate the continued influence of misinformation. *Memory & Cognition*, *38*(8), 1087–1100. http://doi.org/10.3758/mc.38.8.1087

Fein, S., McCloskey, a. L., & Tomlinson, T. M. (1997). Can the Jury Disregard that Information? The Use of Suspicion to Reduce the Prejudicial Effects of Pretrial Publicity and Inadmissible Testimony. *Personality and Social Psychology Bulletin*, *23*(11), 1215–1226. http://doi.org/10.1177/01461672972311008

Grice, H. P. (1975). Logic and conversation. In J. L. Cole, P., & Morgan (Ed.), *Syntax and semantics, Vol 3: Speech acts* (pp. 41–58). New York: NY: Academic Press.

Johnson, H. M., & Seifert, C. M. (1994). Sources of the Continued Influence Effect: When Misinformation in Memory Affects Later Inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(6), 1420–1436. http://doi.org/10.1037/0278-7393.20.6.1420

Kassin, S. M., & Sommers, S. R. (1997). Inadmissible Testimony, Instructions to Disregard, and the Jury: Substantive Versus Procedural Considerations. *Personality and Social Psychology Bulletin*. http://doi.org/10.1177/01461672972310005

Lagnado, D. a, & Harvey, N. (2008). The impact of discredited evidence. *Psychonomic Bulletin & Review*, *15*(6), 1166–1173. http://doi.org/10.3758/PBR.15.6.1166

Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and Its Correction: Continued Influence and Successful Debiasing. *Psychological Science in the Public Interest*, *13*(3), 106–131. http://doi.org/10.1177/1529100612451018

Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, *45*(4), 867–872. http://doi.org/10.1016/j.jesp.2009.03.009

Rich, P. R., & Zaragoza, M. S. (2016). The continued influence of implied and explicitly stated misinformation in news reports. *Journal of Experimental Psychology. Learning, Memory & Cognition*, *42*(1), 62–74. http://doi.org/http://dx.doi.org/10.1037/xlm0000155

Seifert, C. M. (2014). *The Continued Influence Effect: The Persistence of Misinformation in Memory and Reasoning Following Correction*. (Rapp, David N., Braasch, Jason L.G., Ed.). Cambridge, Massachusetts: MIT Press.

Tenney, E. R., Cleary, H. M. D., & Spellman, B. A. (2009). Unpacking the doubt in "beyond a reasonable doubt": Plausible alternative stories increase not guilty verdicts. *Basic and Applied Social Psychology*, *31*(1), 1–8. http://doi.org/10.1080/01973530802659687

Wilkes, A. L., & Leatherbarrow, M. (1988). Editing episodic memory following the identification of error. *The Quarterly Journal of Experimental Psychology Section A*, *40*(2), 361–387. http://doi.org/10.1080/02724988843000168