The Effects of Collaborative Discussion and Exposure to Misinformation on the Retrieval of General Knowledge

A Dissertation Presented

by

Sarah Pociask

to

The Graduate School

in Partial Fulfillment of the

Requirements

for the Degree of

Doctor of Philosophy

in

Cognitive Science

Stony Brook University

August 2015

Stony Brook University

The Graduate School

Sarah Pociask

We, the dissertation committee for the above candidate for the

Doctor of Philosophy degree, hereby recommend

acceptance of this dissertation.

Dr. Suparna Rajaram – Dissertation Advisor
Professor, Department of Psychology

Dr. Christian Luhmann – Chairperson of Defense
Associate Professor, Department of Psychology

Dr. Bonita London – Dissertation Committee Member (Outside Area)
Associate Professor, Department of Psychology

Dr. Elizabeth Marsh – Dissertation Committee Member (Outside Program)
Associate Professor, Department of Psychology and Neuroscience, Duke University

This dissertation is accepted by the Graduate School

Charles Taber
Dean of the Graduate School

Abstract of the Dissertation

The Effects of Collaborative Discussion and Exposure to Misinformation on the Retrieval of General Knowledge

by

Sarah Pociask

Doctor of Philosophy

in

Cognitive Science

Stony Brook University

2015

Recent research from our laboratory demonstrated that collaboration has a positive impact on knowledge retrieval—increasing accuracy, reducing omissions and errors—highlighting the error pruning mechanisms of group discussion. The present experiments put the error pruning mechanisms of collaboration to test by investigating how collaborative discussion influences the transmission of misinformation introduced via fictional story reading. Given the positive impact of collaborative discussion on knowledge retrieval, we expected that collaboration would reduce the negative effects of misinformation exposure. In two experiments, participants completed three general knowledge tests in one of the following sequences: individually, collaboratively, and then individually (I-C-I), or in three successive individual tests (I-I-I). The timing of collaboration relative to misinformation exposure varied: in Experiment 1, misinformation exposure occurred before the second test, while in Experiment 2, it occurred after. Consistent with previous findings, both experiments demonstrated that reading misinformation reduced correct responding and increased misinformation production, and this outcome was particularly salient for participants who completed repeated individual tests. Although participants who worked collaboratively also experienced some of the negative effects of misinformation, there were benefits associated with collaboration; I-C-I participants outperformed those that worked individually and were more confident in their responses. Moreover, on questions for which participants demonstrated accurate prior knowledge, a significant reduction in correct responses was observed only for participants who repeatedly worked individually. Together, these experiments show that although collaborative discussion did not eliminate the transmission of misinformation, it significantly reduced its impact.

Keywords: collaboration, misinformation, retrieval, general knowledge

Acknowledgements

I want to sincerely thank everyone who has guided and supported me—my advisor, fellow lab mates, dedicated research assistants, friends, and family members—all of you have helped me complete this dissertation in one way or another, and for that I am truly grateful.

Table of Contents

List of Tables

List of Figures

I. Introduction

The Effects of Collaborative Discussion and Exposure to Misinformation on the Retrieval of General Knowledge

Memory is a reconstructive process. With each attempt at retrieval, new details may be added to our memories, while others may be changed or omitted. A good example of this comes from thinking generally about the process of learning. Throughout our lives, we constantly update our knowledge bases by learning new things and relearning things we have encountered in the past. Being able to change and update our memories often works to our advantage by allowing us to reinforce and connect prior knowledge, correct previous misconceptions, and fill in gaps in our understanding. However, sometimes the reconstructive nature of memory leads to memory errors. For instance, in situations where sources of information are incomplete, inaccurate, or conflicting, updating our memories does not always produce a more accurate or more complete record of past events or knowledge. This turns out to be the case for many different kinds of memories, including memory for general knowledge (Marsh, Butler, & Umanath, 2012), autobiographical memories (Hyman & Loftus, 1998), and a whole host of other episodic memories (Lewandowsky, Ecker, Seifert, Schwarz, & Cook, 2012).

The present experiments focus on the malleability of general knowledge, and in particular, the shifts that occur in general knowledge as a result of collaborative discussion and exposure to misinformation. The present experiments were motivated by our recent work (Pociask, Rajaram, & Marsh, 2014) where we found that collaboration significantly improved general knowledge retrieval relative to repeated individual testing. In our previous research, we did not manipulate participant knowledge; participants came to the experiment with varying levels of prior knowledge, discussed this knowledge freely as they decided upon their final answers, and in doing so, answered questions more accurately. In the present experiments we put the error pruning mechanisms of collaborative discussion to test by exposing participants to misinformation using a fictional story reading manipulation that has been shown to reduce accuracy of knowledge retrieval (Marsh, Meade, & Roediger, 2003). Critically, by including this misinformation manipulation and by comparing the performance of participants who were given the opportunity to engage in collaborative discussion to those who worked on their own, the present experiments allowed us to test whether the benefits of collaboration are upheld when people are confronted with misinformation that was systematically presented to them, rather than just sporadically generated during group discussion.

Extensive research has been conducted in order to understand how, why, under what circumstances, and to what extent our memories change such that incorrect information becomes integrated in memory and false memories are recalled (for reviews and relevant comentaries, see DePrince, Allard, Oh, & Freyd, 2004; Frenda, Nichols, & Loftus, 2011; Gallo, 2010; Loftus, 2005; Pezdek & Lam, 2007; Roediger & McDermott, 2000; Schacter, 1999; Wade et al., 2007). In order to study the various real-life scenarios in which memory errors occur, researchers have developed a vast array of laboratory paradigms that model different variations of memory errors and the ways in which they emerge. This paradigmatic diversity has allowed for the differentiation between different types of memory errors and the mechanisms that underlie their formation. To provide a broad view of this literature, we first present an overview of two of the most commonly used and adapted paradigms for studying false memories in the cognitive literature, and some of the key findings pertaining to false memories and memory errors more broadly. We then move to the description of the paradigm we selected for the present experiments to target a test of the shifts that occur in a person's general knowledge representations (e.g., Marsh, et al., 2003), followed by a discussion of the collaborative memory literature, as the present research represents a merger of these two research areas.

Overview of False Memory Literature

Over the past several decades, one paradigm that has been repeatedly used and adapted to study how memories change is the three-phase misinformation paradigm (Loftus, 1975, 1979, 2005; Loftus, Miller, & Burns, 1978). This paradigm is particularly well suited to study the memory errors that might occur in situations where someone witnessed a crime and was subsequently exposed to misleading information from outside parties like co-witnesses or police officers. As such, this paradigm has contributed significantly to the legal domain, particularly in the evaluation of the accuracy of eyewitness testimony (Zaragoza, Belli, & Payment, 2006). As suggested in the name, this laboratory paradigm involves three key phases: the initial encoding of critical stimuli, exposure to post-event information, and a final memory test. During the encoding phase, participants watch a video, look at pictures, or witness an event unfolding before them in the lab, and then in the post-event information phase, they are asked questions about what they saw. Critically, some participants are asked questions that contain misinformation that contradicts what was initially presented. Finally, participants are tested on the details of the original event. Here, the measure of interest is whether participants correctly remember the events that they had initially seen, or instead, if they report the details suggested with the misinformation.

The overarching finding from this body of research is that people integrate information from multiple sources when retrieving information from memory, and when exposed to misinformation following an event, accuracy for the original details can be reduced, a phenomenon that has become known as the misinformation effect (Tousignant, Hall, & Loftus, 1986). Researchers have developed many interesting variations of the original three-phase misinformation paradigm, for example, using different stimuli at encoding (e.g. film clips, pictures, stories, experienced events); different ways of exposing people to misinformation (e.g. misleading questions, videos, fictitious co-witness reports, discussion with another person); and different ways of testing memory (e.g. recall, cued recall, recognition, modified recognition, source monitoring tests), for a review, see Loftus (2005). Identifying and comparing the circumstances under which the misinformation effect is reduced (e.g., Clifasefi, Garry, Harper, Sharman, & Sutherland, 2007; Eakin, Schreiber, & Sergent-Marshall, 2003; MacLeod & Saunders, 2005; Parker, Garry, Engle, Harper, & Clifasefi, 2008; Tousignant, et al., 1986; Zaragoza, McCloskey, & Jamis, 1987) versus enhanced (e.g., Assefi & Garry, 2003; Eakin, et al., 2003; Zhu et al., 2010), has allowed researchers to pose and address important theoretical questions, like why the effect occurs, how it can be avoided, and what happens to the original event memory. These are some of the key questions that underlie all false memory research, and despite the vast amount of research conducted on this topic, many of these questions have yet to be fully resolved (Michael, Garry, & Kirsch, 2012).

The Deese/Roediger-McDermott (DRM) paradigm (Deese, 1959; Roediger & McDermott, 1995) has also been widely used to study false memories in the laboratory. Within the classic version of this paradigm, participants study lists of associatively related words (e.g. *sandwich, toast, butter, loaf*), and then some time later they take a memory test that requires them to recognize or recall the words that they had previously seen. The measure of false memory in this paradigm is the rate at which participants report the critical, non-studied lures that are associated with the lists that they previously studied, in this case, remembering the word *bread*, even though it was not presented among the initial set of items. Not only do participants incorrectly remember critical items, but their recollections often include contextual details and high levels of confidence in these responses (Roediger, McDermott, Pisoni, & Gallo, 2004; Roediger & McDermott, 1995). Also, as observed with the misinformation effect, the associative illusions observed within the DRM paradigm also persist, but are reduced following warnings about the nature of the illusion (e.g., Gallo, Roediger, & McDermott, 2001; McCabe & Smith, 2002; Multhaup & Conner, 2002; Westerberg & Marsolek, 2006). One explanation of why false memories arise in DRM paradigm comes from the activation/monitoring framework (Roediger & McDermott, 1995, for a review, see Gallo, 2010). Very briefly, activation refers to the idea that there are connections and relationships among the words in the lists and the lure itself, and these connections prompt the viewer to think of the critical lure. Monitoring is the process by which we make decisions about what we saw by evaluating the qualities of the memories and deciding, for instance, if a given word was seen or something we merely thought of. Thus, considering both processes within this broad framework, the DRM illusion can be explained by the activation of the critical lure via associative processing and the failure to adequately monitor the source of the memory.

Malleability of General Knowledge

Whereas the three-phase misinformation paradigm and the DRM paradigm focus on false memories for words, pictures, and experienced events (laboratory or otherwise), the present experiments will focus on memory errors for general world knowledge. Throughout our lives we learn about the world, its history, and our place in it not only from reference books and schooling, but also through everyday experiences, conversations, written accounts (fiction and nonfiction), and other types of popular media sources. The accuracy of these sources varies considerably, and is of particular concern today where we have the ability to be constantly connected to media sources that provide access to rapidly generated and often unregulated content (Lewandowsky, et al., 2012). To address this question of how misinformation affects retrieval of general knowledge several different paradigms that have been used (e.g., Bottoms, Eslick, & Marsh, 2010; Gerrig & Prentice, 1991; Kelley & Lindsay, 1993; Lewis & Anderson, 1976; Marsh, et al., 2003). Although the specific types of misinformation and modes of exposure differ across these experimental paradigms, what is consistent is that participants are in some way exposed to misinformation pertaining to factual knowledge and then following this exposure, they are tested on their knowledge for this information.

Intuitively, it is reasonable to predict that many general knowledge facts are well consolidated in memory, and as a result, people might be less likely to be negatively influenced by exposure to false general knowledge information. This, however, does not appear to be the case, as it has been repeatedly demonstrated that exposure to incorrect information disrupts people's abilities to retrieve general knowledge. For example, reading "fantasy facts" (e.g. "Napoleon Bonaparte was from India.") led to more errors on a fact verification task and slower response times when verifying true facts (Lewis & Anderson, 1976). Similarly, after reading incorrect assertions in fictional narratives, people were slower to reject these false facts on a subsequent verification task (Gerrig & Prentice, 1991). Simply presenting people with lists of incorrect answers prior to testing reduced correct responding on a subsequent general knowledge test, even when people were warned about the errors (Kelley & Lindsay, 1993). Based on these findings, it is clear that exposure to misinformation, in a variety of forms, can have a negative impact on our ability to retrieve and report general knowledge.

In the context of these findings, and given the widespread use of fictional materials in classrooms as a way to supplement the teaching of factual information (Marsh, et al., 2012), one of the many important applications of this work involves understanding the implications of exposing students to incorrect information via fictional sources like films and novels. Including works of fiction in classes might improve student engagement and help students understand and relate to the material more effectively, but on the other hand, these materials are not entirely accurate depictions of history, and thus it is also necessary to consider the potential costs of using them as teaching tools. To investigate the extent to which people use and incorporate information from fictional sources into their general knowledge base, Marsh and colleagues have developed a paradigm that uses fictional stories (Marsh, et al., 2003) and films (Butler, Zaromb, Lyle, & Roediger, 2009) as a means of exposing people to correct and incorrect information about the world. A number of critical findings have emerged from this work. For instance, reading incorrect facts increased the likelihood that people would answer general knowledge questions with misinformation (i.e. the same incorrect information they had been exposed) and reduced their production of correct answers to those questions (Barber, Rajaram, & Marsh, 2008; Eslick, Fazio, & Marsh, 2011; Fazio, Barber, Rajaram, Ornstein, & Marsh, 2013; Fazio & Marsh, 2008; Marsh, et al., 2003). Surprisingly, the negative consequences of reading misinformation in stories occurred even when monitoring of facts was encouraged through slower story presentation (Fazio & Marsh, 2008), multiple story readings (Marsh, et al., 2003), warnings (Experiment 1, Marsh & Fazio, 2006), reducing cognitive load (Experiment 2, Marsh & Fazio, 2006), and when given explicit error detection instructions (Experiment 3, Eslick, et al., 2011; Experiment 3, Marsh & Fazio 2006). In fact, slowing down the reading time of the story actually increased the likelihood of misinformation production (Fazio & Marsh, 2008), which is the opposite of what might be expected based on the episodic false memory literature. Slower presentation speed would have been expected to increase monitoring and thus reduce illusions of knowledge, but that is not what was observed in this instance.

Overview of Collaborative Memory Literature

Each of the aforementioned laboratory paradigms represents different ways that errors can be encoded in memory in everyday life. We may incorrectly remember seeing something because it was closely related to what we actually saw; we may remember details that someone else recounted as our own personal experience; and we might mistake the creative license exercised by an author of fiction as an actual historical fact that we have always known to be true. Likewise, people view events from different perspectives and know different things about the world, and by engaging with others we are exposed to alternative ideas and perspectives. Any given person is going to have different ways of thinking about and organizing things based on their own preferences and experiences. With this in mind, what happens when people share and discuss information in groups? Or put another way, when we interact with others and share information, how does this affect what we remember later on? The impact of such discussions on one's knowledge retrieval, particularly when confronted with misinformation, is at the crux of the research questions tested in this dissertation research.

Over the past few decades, memory researchers have become increasingly interested in studying memory at the group level. Whereas prior research had focused almost exclusively on the memory of individuals in isolation, collaborative memory research focuses on understanding what happens when people remember information as a group, and what the consequences of group discussion are for individual memory later on (for a review, see Rajaram & Pereira-Pasarin, 2010). The typical collaborative memory paradigm involves at least three phases. The first phase is an encoding phase where participants individually study a given set of experimental stimuli (e.g. words, pictures, narratives, etc.). Following the encoding phase and a distractor period, participants are tested on the material, either alone or collaboratively in a small group. Finally, all participants typically complete one final individual test. This is just one variation of the classic collaborative memory paradigm, as it is often adapted to include different numbers and sequences of study or test phases (collaborative or individual) as motivated by the research questions at hand (e.g. for variations, see Blumen & Rajaram, 2008; Choi, Blumen, Congleton, & Rajaram, 2014; Congleton & Rajaram, 2011).

Based on this design, several comparisons can be made to evaluate the consequences of remembering information with others. To highlight a few key findings that have emerged from this literature, we first discuss the robust, counterintuitive phenomenon known as collaborative inhibition. As per the saying, 'two heads are better than one,' it is commonly believed that performance on a memory task will be improved when working as a group. Although groups do tend to outperform individuals, comparing group performance to the pooled, non-redundant responses of individuals in nominal groups[1] reveals that groups often do not perform up to their full potential (B. H. Basden, D.R. Basden, Bryner, & Thomas, 1997; Weldon & Bellinger, 1997). One theory as to why this counterintuitive outcome occurs is the retrieval disruption hypothesis (for a discussion of retrieval disruption and other mechanisms underlying collaborative inhibition, see Barber, Harris, & Rajaram, 2014; Hyman, Cardwell, & Roy, 2013). Retrieval disruption is similar to the process that underlies the part-set cueing effect, where seeing a subset of the previously studied list items impairs recall of the items that remain (D.R. Basden, B.H.

[1] Nominal groups consist of the pooled, non-redundant responses of participants who worked individually; they did not work as a group during the experiment, but their data are pooled as if they had, thus they are a group in name only. Nominal groups provide a way of understanding what the group might have been able to produce had the group members worked independently and simply pooled their responses.

Basden, Galloway, 1977; Finlay, Hitch, & Meudell, 2000; Nickerson, 1984; Roediger & Neely, 1982). Remembering information as a group is disruptive because the items recalled by fellow group members inhibit one's own individual recall, and moreover, because each person has their own way of organizing the material, retrieving the material as a group disrupts the optimal retrieval strategy of each person, and thus the total output for the group is reduced (B. H. Basden, et al., 1997).

Following collaboration there are both positive and negative consequences for individual memory (B. H. Basden, D. R. Basden, Henry, 2000; e.g., Hirst & Echterhoff, 2012). For instance, when working alone after remembering with a group, the costs of group recall can be mitigated, that is, individuals are able to recall some of the items that may have been blocked during group recall (Finlay, et al., 2000). In addition to the release from collaborative inhibition, post-collaborative benefits are also observed as a result of re-exposure to items during group recall (Blumen & Rajaram, 2008), however, post-collaborative forgetting has also been observed for items that are not discussed as a group (Barber, Harris, & Rajaram, 2014; B. H. Basden, et al., 2000; Coman, Manier, & Hirst, 2009; Hirst & Echterhoff, 2012).

It is also worth noting that accuracy is not the only metric used to evaluate group memory. For example, researchers have also measured the consistency of memory reports and the extent to which representations are shared across individuals within a given group, as in collective memory (Hirst & Manier, 2008). In fact, repeated collaboration has been found to increase the overlap that is observed across post-collaborative individual memory (Congleton & Rajaram, 2014). Finally, although less commonly observed, it is also important to note that under certain conditions, collaborative group performance has been found to exceed what would be expected based on nominal group performance, a finding that is known as collaborative facilitation (Clark, Hori, Putnam, & Martin, 2000; Harris, Barnier, & Sutton, 2012; Harris, Keil, Sutton, Barnier, & McIlwain, 2011; Meade, Nokes, & Morrow, 2009).

Researchers have investigated the effects of collaboration on group memory using a variety of study materials, including associatively related DRM word lists and other stimuli that allow for the study of the transmission of false memories among group members. Considering the underlying processes of activation and monitoring that are thought to be involved in the formation of associative false memories, collaboration might be predicted to increase the occurrence of false memories to the extent that group members collaboratively endorse the critical lures. Alternatively, collaboration might reduce false memories through increased monitoring. Demonstrated across two studies using DRM lists, as group size increased from dyads to quartets, recall and recognition of critical lures and studied words increased (Thorley & Dewhurst, 2007; Thorley & Dewhurst, 2009). When participants were under higher pressure collaborative conditions (i.e. turn-taking free recall) compared to relatively lower pressure (i.e. free-for-all recall), more critical lures were recalled (Takahashi, 2007; Thorley & Dewhurst, 2007). These same costs and benefits were observed post-collaboratively, as participants that had previously worked in groups again recognized more lures and also more study words. In contrast, evidence for group filtering has been found during collaborative recall of DRM lists (Weigold, Russell, & Natera, 2014), and using visual stimuli instead of the classic DRM word lists (i.e. household scenes, see Meade & Roediger, 2002), collaboration has also been found to reduce the production of errors (Ross, Spencer, Blatz, & Restorick, 2008).

To briefly summarize, when working to remember as a group there are a number of different processes that interact to determine the net outcome of collaboration, some processes that impact memory in a positive way (i.e. re-exposure, re-learning through retrieval, crosscuing, and error pruning) and others that impact memory in a negative way (i.e. social contagion of errors, blocking, and retrieval disruption (Rajaram & Pereira-Pasarin, 2010). Depending on these processes, and their interactions with a number of other factors related to the collaborative task itself—like for instance, the materials being recalled, type of memory task, type of collaborative instructions, and characteristics of the individuals in a group—collaboration can either enhance or impair memory (Andersson, Helstrup, & Rönnberg, 2007; Rajaram, 2011). The effects of collaboration have largely been explored with the recall or recognition of verbal and visual study materials and autobiographical materials as opposed to general knowledge. Interestingly, the few studies that have directly compared the effects of collaboration when participants recalled details from recently read stories (i.e. episodic information) to the effects of collaboration for recalling general knowledge found that collaborative inhibition was not observed for general knowledge, but it was observed for recalling details for the stories (Andersson & Rönnberg, 1996; Johansson, Andersson, & Rönnberg, 2005). More recently, researchers have continued investigating how the effects of collaboration might be different in the context of retrieval of general knowledge by specifically focusing on the costs and benefits.

One set of studies investigating how collaboration affects the retrieval of general knowledge looked at the effects of collaboration when people work together to solve crossword puzzle clues (Szary & Dale, 2013, 2014). In this task, participants worked to come up with answers to a variety of different types of crossword puzzle clues (e.g. sayings, logic, general knowledge, fill in the blank, etc.) within a timed round, where 30 seconds were provided for each question. Participants were allowed unlimited chances to correctly answer to each clue within the 30 second time limit, thus correct answer and the number of errors made before arriving at this answer was documented within this design. When working collaboratively, participants could discuss the clue and problem solve with their partner. Across two studies, both costs and benefits of collaboration were identified. When comparing the average performance of the two participants working in dyads versus the average performance of two participants working individually, working collaboratively on this type of task was beneficial in that the average performance of participants in dyads solved more puzzle questions and did so more accurately and quickly than peers working individually (Szary & Dale, 2013). In a second study, both costs and benefits were observed as collaborative groups answered fewer trivia questions and did so more slowly compared to nominal groups, but on the other hand, groups made fewer errors before ultimately arriving at the correct solution (Szary & Dale, 2014).

In a different collaborative task that is more directly related to the present studies (Pociask, et al., 2014), participants worked in triads to answer general knowledge questions (e.g., see Marsh, 2004; Nelson & Narens, 1980; Tauber, Dunlosky, Rawson, Rhodes, & Sitzman, 2013). The study was not concerned with collaborative inhibition; rather it tested whether collaboration under these circumstances promotes error contagion or error pruning when retrieving general knowledge. Across two experiments, findings showed that collaboration improved performance by pruning errors and reducing omissions relative to participants who worked individually. Collaboration also led to increased confidence in one's answers. Taken together, these findings suggest that there are benefits associated with retrieving general knowledge in groups.

Theoretical explanations from within the collaborative memory literature offer some possible explanations for why the positive effects of collaboration (i.e. error pruning, re-exposure, relearning via retrieval) emerged under these conditions as opposed to social contagion of errors. As previously discussed, it is important to consider that the effects of collaboration depend on a number of factors including the nature of the information being recalled as well as the goals of the recall context. When working in a group to answer general knowledge questions, the inherent goal is to arrive at the single correct answer. This is in contrast to other collaborative memory situations (e.g. recalling all of the words on a recently studied word list or discussing various perspectives of an experienced event), where there is not necessarily one answer that is correct, and instead, the goal is to maximize all group members' memories and coordinate the narratives across collaborators (e.g., Hirst & Echterhoff, 2012; Hirst & Manier, 2008). Furthermore, research suggests that error pruning tends to prevail in free-flowing collaboration as opposed to turn-taking (e.g., Barber, Rajaram, & Aron, 2010), and consensus collaboration (i.e. where all group members must agree on a single correct answer) is also associated with improved accuracy (Harris, et al., 2012). It is also important to consider the characteristics of general knowledge, and how retrieving this type of information is different from retrieving more recently learned information. As posited by Andersson and Ronnberg (1996) retrieval cues generated by group members in effort to recall answers to general knowledge questions might be more effective because this material is more organized and because it is "overlearned," meaning that participants have been exposed to the material and learned and relearned it many times (see also Meade, et al., 2009). This explanation is consistent with other findings within the collaborative memory literature that show that prior repeated retrieval of word lists strengthened retrieval organization, and as a result, collaborative inhibition was not observed (Congleton & Rajaram, 2011).

The Present Experiments

Merging the collaborative memory paradigm (Basden, et al., 1997; Weldon & Bellinger, 1997) with the illusions of general knowledge paradigm (Marsh, et al., 2003), the present experiments allowed us to investigate whether illusions of knowledge were spread or dispelled by group discussion. Experiments 1 and 2 leveraged the findings from our recent work (Pociask et al., 2014) and explored the effects of collaboration on the retrieval of general knowledge by focusing on how collaborative discussion influences the transmission of misinformation that is introduced via the reading of fictional stories. In Experiment 1, the focus was on exploring the error pruning versus error propagation processes involved in collaboration. After completing a baseline general knowledge test and reading stories containing false facts, participants completed a second general knowledge test collaboratively or individually. By comparing rates of correct answers and production of misinformation answers of participants who worked collaboratively to those who worked individually, our goal was to determine whether errors were more likely to be pruned by group conversation, or conversely, whether errors would propagate. Alternatively, it was also possible that we would see some evidence of both processes. Examining performance on the third and final general knowledge test allowed us to explore the downstream effects of collaboration, particularly whether any observed benefits or decrements following collaboration persist when working alone.

In Experiment 2, story reading occurred after the second general knowledge test. Critically, this design change allowed us to explore whether collaborating prior to being exposed to misinformation would reduce participants' experience of illusions of knowledge. Since past research has shown that participants are more confident in their answers and performance on general knowledge tests is more accurate when participants complete them collaboratively (Pociask, et al., 2014), we hypothesized that this experience would protect participants against producing misinformation as answers on the final general knowledge test. In other words, the goal of Experiment 2 was to test whether collaborative retrieval of general knowledge can provide inoculation against memory errors under specific learning conditions.

II. Experiment 1

Method

Participants

The present sample consists of data collected from 96 Stony Brook University undergraduates who participated for course credit. All participants were fluent in English.

Materials

Stories. Four fictional short stories were revised for use in the present study (modified from Barber, et al., 2008; Marsh, 2004). These stories were on average 1900 words in length. Across the four stories, there were references to 72 critical facts (i.e. 36 misinformation items and 36 neutral items, described later). Critical facts were selected to represent a wide range of difficulty levels, and across the four stories, the hardest fact was answered by 9% out of 94 Stony Brook students, and the easiest fact answered by 91%. Each story contained an equal dispersion of facts within this difficulty range, and fact difficulty was matched across stories so that no story was any more or less difficult than any other.

Each story included 18 critical general knowledge facts, half of which were framed neutrally (i.e. the fact was mentioned, but the answer was not referenced), and the other half were framed with misinformation (i.e. an incorrect, but believable answer was referenced). For instance, a fact framed with misinformation would state, “I was supposed to be John Glenn, the first man on the moon” whereas in the neutral framing the fact would simply say, “I was supposed to be the first man on the moon” without directly mentioning the name of the astronaut. So that accurate information was also conveyed by the stories, six references to correct answers were included in each story, although these were not included either as critical facts or as fillers on the general knowledge tests, and were not the focus of any experimental hypotheses or analyses. Two versions of each story were created such that the facts presented with misinformation in one version of the story were presented neutrally in the other. The sequence in which the stories were presented was counterbalanced across subjects.

General knowledge test. The general knowledge test consisted of 144 general knowledge questions pertaining to a variety of topics (e.g. art, health, entertainment, geography, games/sports, history, literature, and science), and of these facts, 72 were critical facts referenced in the stories, 72 were filler items. The majority of the questions were selected from a previously normed fact list consisting of 300 general knowledge facts (Nelson & Narens, 1980). However, since many of the facts from this list were dated or too obscure to represent general knowledge for current college students (Tauber, et al., 2013), additional topical facts were also normed and included in the present experiments. All of the facts selected for use in the present experiments were normed with students at Stony Brook University to ensure that questions spanned an appropriate range of difficulty for our subject sample (the hardest question included in the general knowledge test was answered correctly by 5% of Stony Brook students and the easiest was answered by 98%).

Procedure

Throughout the course of the experiment, all participants completed three general knowledge tests in an I-C-I (individual, collaborative, individual) or I-I-I (individual, individual, individual) sequence. General knowledge tests were separated by distractor periods of about approximately five minutes during which participants completed an unrelated activity (i.e. a visuospatial puzzle or visual search task) that was designed to be engaging, but not overly taxing.

The first test was completed individually for all participants as a baseline measure of their general knowledge at the start of the experiment. Participants began the experiment by completing a 144-item general knowledge test on the computer. The questions were presented one by one on a computer screen, in a different random order for each participant. All questions could be answered with one or two words, which participants typed into the answer space displayed on the screen. Participants were instructed not to guess, and if they did know the answer, they were told to type the word “NEXT” to move on to the next question. Immediately following each question, all participants rated their confidence in the answers they provided on a scale of 1 (not at all confident) to 5 (absolutely confident).

After completing the initial general knowledge test, participants were given a short break (at least 5 minutes) during which they played a visuospatial puzzle game on the computer. During all breaks in the experiment, cell phone use was prohibited and computer use was restricted to ensure that participants could not look up any answers to the questions at any point.

Prior to the second general knowledge test, participants completed the reading comprehension task where they were exposed to the four fictional stories containing critical references to general knowledge facts. To control for differences in participant reading speeds, the stories were pre-recorded and presented visually and aurally via computer. Following each story, participants were required to answer basic comprehension questions to ensure that they were paying attention and adequately understanding the story. After the presentation of the stories, participants were given another five minute break during which they played a visual search game where they had to locate a list of objects in a natural scene.

Following the break, participants completed a second general knowledge test. Participants completed this test collaboratively in a triad or individually. This test included the same items that were on the initial test, but test items were presented in a different random order. If working collaboratively, participants were told that they should discuss each question and their thoughts on the answer as a group, but consensus was not required, and they should each write down their own individual answers and confidence ratings on separate answer sheets[2]. Participants working individually typed their answers and confidence ratings into the computer. Immediately after this test, participants were given another short five minute break where they again played the visual search game. Participants then completed one final general knowledge test, consisting of the same 144 questions that were once again presented in a new random order. All participants completed this test individually.

Finally, participants completed an end-of-experiment questionnaire which included demographics questions (e.g. major, GPA, ethnicity, etc.), questions about the experiment (e.g. about the difficulty and enjoyment of tasks, etc.), and questions about preferences for collaboration. As part of the debriefing procedures, all participants were shown the correct answers to the general knowledge questions and were told not to discuss the experiment with others.

Results

Data Analysis Approach

The primary goal of the present experiment involves the analysis of our two key experimental manipulations: collaboration and exposure to misinformation via story reading. Consistent with previous work in this literature (Barber, et al., 2008; Eslick, et al., 2011; Fazio, et al., 2013; Marsh & Fazio, 2006; Marsh, et al., 2003), in order to test the effects of exposure to

[2] Collaborative procedures requiring consensus and a single group response were found to yield similar results (Pociask, et al., 2014).

misinformation our primary focus was on comparing the rate at which correct answers and misinformation answers were produced on the general knowledge test for questions that were referenced with a misleading frame versus a neutral frame in the story reading phase. To assess whether the effects of misinformation exposure (i.e. on correct answer and misinformation production) are different when participants have the opportunity to engage in collaborative discussion, we compared the response patterns of participants who worked collaboratively relative to participants who worked individually.

Because the experimental design involves collaboration on Test 2 for participants in the I-C-I condition, as a first step in the data analysis process it was important to assess how group membership affected response variation, or in other words, whether group members influenced each other's performance by collaborating. The extent to which group members influenced each other has implications for how the data are analyzed. Specifically, if group members' scores are not significantly correlated with one another (i.e. there is little to no variation in performance due to group membership), this suggests that the observations from each participant can be treated as independent, and the standard analysis of variance (ANOVA) or OLS regression procedures would be justifiable. However, if there are significant positive correlations among group members' scores, an alternative data analysis approach that takes the grouped data structure into account is warranted (e.g. Harlow, 2014; Peugh, 2010).

We assessed whether dependencies existed within the data structure as a result of collaboration by first calculating intraclass correlation coefficients (ICCs) using group as a clustering variable at Test 2 (when collaboration occurred) and Test 3 (immediately following collaboration) for several key test performance variables. The ICC values can be interpreted as the proportion of performance variation that occurs across groups, and also as the expected correlation between the scores of participants working in the same group (Peugh, 2010). For triads, the upper limit of the ICC is +1 and the lower limit is -.5 (Kenny, Mannetti, Pierro, Livi, & Kashy, 2002). Significant positive ICC values indicate that group members' scores are positively correlated with each other, while independence is indicated by an ICC of zero[3]. As previously discussed, if significant positive ICC values are observed, this suggests observations are not independent, which can lead to biased results, particularly an increase in the Type I error rate if the group structure is ignored (Kenny & La Voie, 1985; Peugh, 2010; Tabachnick & Fidell, 2007; West, Welch, & Galecki, 2015). Evaluation of the intraclass coefficients in the present experiment revealed, as expected, that scores from participants working collaboratively in a triad were related to one another; ICCs were positive and significantly greater than zero, see Table 1 for a summary of the ICCs and significance values. Thus, when working collaboratively in the present experiment, even though participants wrote down their own answers, the scores of participants in a group were correlated with one another.

As a second step in understanding the data structure and determining the appropriate data analysis approach we calculated the design effect (DE) which estimates the value that one would need to multiply the obtained standard error by in order to correct for the dependency due to correlated data (Harlow, 2014; Peugh, 2010). There is no strict significance test for DE values, but when DE values exceed 2.0, it is recommended that the grouped data structure is accounted for with multilevel modeling or other comparable analysis procedure (Muthén, 1994). As seen in

[3] Because even small, nonzero ICCs can have significant effects on the accuracy of statistical inferences, it is recommended that researchers choose a liberal criterion ($\alpha = .25$) when determining whether group scores are correlated (Kenny & La Voie, 1985; Myers, 1972). In the present experiment, ICCs were significant at the .01 or .001 level.

Table 2, the DE values for test performance variables in this experiment either exceed or closely approach this criterion, ranging from 1.94 to 2.95, see Table 1 for a summary of these values.

To summarize, the diagnostic information provided by the ICC and DE values show that the collaboration manipulation in the present experiment was effective; group members influenced each other's performance on the general knowledge tests. Therefore, given the experimental design and our research questions of interest, we adopted a multi-faceted data analysis approach. When only data from participants in the I-I-I condition were being analyzed or when analyzing Test 1 data (prior to collaboration), we conducted ANOVAs, *t* tests, or their non-parametric equivalents as needed. When analyzing data from Tests 2 and 3 (i.e. any comparisons across experimental conditions during or following collaboration), we used multilevel linear modeling (MLM) adapted for partially clustered data structures when data were normally distributed (Baldwin, Bauer, Stice, & Rohde, 2011; Bauer, Sterba, & Hallfors, 2008), and the generalized estimating equations (GEE) approach when data were significantly positively skewed (i.e. misinformation responses).

When analyzing the data using MLM, model specification was conducted using a maximum likelihood (ML) parameter estimator. We began with a model that did not take the group clustering into account, and then gradually increased the complexity of each successive model by adding the group clustering variable as a random factor. When considering the group clustering variable, we first tested the fit of a random intercept model, which takes into account the dependencies in the data by estimating variance associated with group differences in average performance. As a next step, we tested the fit of a model that also included a parameter for random slopes, which allows for the estimation of the variance associated with group differences in the relationship between the predictor variable (e.g. fact framing) and the dependent variable (e.g. test performance). The relative fit of competing models was evaluated by comparing the deviance values (-2 Log Likelihood), which reflect how well the parameter estimates and variance structures fit the sample data. Throughout the model fitting process, we retained more complex models only when they significantly improved model fit as determined by the comparison of the deviance values using likelihood ratio tests. Additional details regarding the model fitting process for each analysis is provided within the sections that follow and in supporting tables.

The GEE approach is an extension of the generalized linear modeling approach, which is a modeling procedure that allows for the dependent variable to be non-normally distributed (Agresti, 2002; Hanley, Negassa, & Forrester, 2003; Zeger, Liang, & Albert, 1988). In the present experiment, misinformation responses were relatively rare occurrences and the distribution of this response variable was positively skewed. Rather than assuming normality, we fit the misinformation response data with a negative binomial distribution using the GEE procedure. The GEE approach accounts for the correlations among participant responses due to repeated measures and group clustering, but unlike MLM, it does not estimate parameters for random effects and is thus referred to as a population average model (Coxe, West, & Aiken, 2009). Because GEE involves a quasi-likelihood estimation procedure as opposed to maximum likelihood, the model fitting procedure described above for MLM does not apply for these analyses, and inferences about the significance of the fixed effects of interest are reported using the Wald chi square tests (Agresti, 2002). For this analysis, data must be integers, so the dependent measure for misinformation production was the total number of times each participant responded to the critical questions with misinformation (i.e. the possible range of total misinformation responses being 0-36 for the subset of misleading and neutral critical items).

Test 1 Performance

At this phase of the experiment, all participants completed the general knowledge test individually as a baseline measure of knowledge at the start of the experiment. The general knowledge test consisted of three subsets of questions: critical questions to be referenced with a misleading frame in the story, critical questions to be referenced with a neutral frame in the story, and filler questions. Participant responses on the general knowledge tests were coded as correct responses, omissions, misinformation errors (i.e. the critical misinformation referenced in the story), and other errors (i.e. other erroneous responses not referenced in the story). Even though participants had not yet read the stories containing misinformation at Test 1, we still coded misinformation errors that occurred on this test in order document the rate at which participants already believed that the misinformation was indeed the correct answer. These served as the baseline level of misinformation errors. Misinformation production on Test 1 was rare, as these responses were observed for only about 5% of the critical questions.

We examined the performance of participants assigned to the I-C-I and I-I-I test sequences to determine whether there were a priori differences in knowledge between participants in each experimental condition, see Table 2 for a summary of means and standard errors for all test performance variables on Test 1. As expected, there were no differences in performance between the I-C-I condition and the I-I-I condition at this stage in the experiment, across all questions and within the critical and filler item subsets[4]. Importantly, performance was the same for critical items to be presented in a misleading frame and neutral frame.

[4] When comparing performance on the critical items to the filler items, participants answered fewer filler items correctly, $F(1,95) = 721.47$, $p<.001$, omitted more filler items as compared to critical items, $F(1,95) = 529.29$, $p<.001$, and answered more filler questions with errors, $F(1,95) = 7.86$, $p = .006$ (because error rates were generally low and positively skewed, we also ran this analysis with the non-parametric equivalent, the Wilcoxon Signed Rank Test, and the same

Test 2 Performance

Test 2 was completed either collaboratively in a triad or individually, and prior to this phase of the experiment, all participants were exposed to the stories which contained critical misleading and neutral fact references. The analyses that follow focus on the rates of correct responding and misinformation production for the critical questions that were referenced with misleading or neutral framings in the stories.

Correct Responses. We used MLM to assess how fact framing (misleading or neutral) and collaboration affected the retrieval of correct information following exposure to the fictional stories. The dependent measure was the proportion of misleading and neutrally framed items answered correctly, and since all participants viewed half of the critical items in a neutral frame and the other half with the misleading frame, fact framing was a repeated observation across participants. As a first step, we created a basic model (similar to a standard repeated measures analysis of variance) as a point of comparison where fact framing, condition, and their interaction were entered as fixed effects. As a next step, we included a random intercept in order to model the variation in means due to the fact that some participants worked in groups while others worked individually. A comparison of the deviance values revealed that the model taking group clustering into account with random intercepts was a significantly better fit, $\chi^2(1, N = 96) = 43.96$, $p < .001$. Finally, we compared the random intercept model to a more complex model that also included random slopes to allow the effect of fact framing to vary across groups and individual participants, which again yielded significantly better fit, $\chi^2(1, N = 96) = 41.02$, $p < .001$.

---

conclusion was reached, $z = 2.86$, $p<.01$). These results are consistent with prior norming indicating that the questions selected to be filler items were on average more difficult than those selected to be critical items.

Looking at the fixed effects estimates derived from the random intercepts and slopes model, we see that there was a significant effect of fact framing, $F(1, 57.09) = 11.18$, $p = .001$, where items referenced with a misleading frame in the story ($M = .56$ $SE = .02$) were less likely to be answered correctly on the general knowledge test compared to those referenced with a neutral frame ($M = .62$ $SE = .02$). There was also a significant effect of collaboration, $F(1, 57.84) = 27.30$, $p < .001$, where participants working in groups ($M = .71$ $SE = .02$) answered more questions correctly than their peers working individually ($M = .47$ $SE = .02$). Although numerically the difference between the correct response rate for misleading and neutral critical items is greater for participants working individually, the interaction between fact framing and collaboration was not significant, $F(1, 57.09) = .90$, $p = .35$, indicating that the effects of fact framing on the rate of correct responding were consistent regardless of whether participants collaborated or worked individually, see Figure 1. A summary of the regression coefficients for all three fixed effects can be found in Table 3.

We were also interested in whether the effects of fact framing and collaboration on correct responses depend on prior knowledge. More specifically, if a participant demonstrated that they possessed correct prior knowledge by answering a question correctly on Test 1, would they still be susceptible to misinformation on these questions? To investigate this possibility, we tracked participants' responses to the critical items from Test 1 to Test 2. We identified the questions where participants initially possessed knowledge of the correct answer and documented the rate at which these correct responses persisted on Test 2. Thus, the analyses that follow focus on the proportion of critical questions that a given participant consistently answered correctly on both tests, or put another way, the questions for which participants were resistant to the negative impact of misinformation exposure. The ICC value for this variable was not significant, ICC = .09, $F(15, 32) = 1.31$, $p = .25$, so these data were analyzed with a 2 (Fact framing: misleading versus neutral) × 2 (Condition: collaborative versus individual) repeated measures ANOVA. The main effect of condition was not significant, $F(1,94) = .60$, $p = .44$, but there was a main effect of fact framing, $F(1,94) = 17.27$, $p < .001$, and a significant interaction between fact framing and condition, $F(1,94) = 5.50$, $p = 02$. To understand this interaction, we ran a series of contrasts to isolate the effect of fact framing for those participants who worked collaboratively on Test 2 versus those who worked individually. When participants worked individually, there was reduction in consistent correct responses for the questions that were referenced with a misleading frame in the story ($M = .41$ $SE =.02$) as compared to a neutral frame ($M = .49$ $SE =.03$), $F(1,94) = 16.95$, $p < .001$. This outcome is particularly striking because it demonstrates the negative impact of misinformation exposure even on questions where participants possessed prior knowledge (see Fazio et al., 2013). However, fact framing did not affect the rate of consistent correct answers for participants who worked collaboratively on Test 2, $F(1,94) = 2.18$, $p = .15$, as the rate of consistent correct responses was no different when facts were referenced with a misleading frame ($M = .47$ $SE =.03$) or a neutral frame ($M = .49$ $SE =.03$). This outcome indicates the protective effects of collaboration in mitigating the effects of exposure to misinformation observed in the I-I-I condition.

Misinformation Responses. We used the GEE approach to assess how fact framing (misleading or neutral) and collaboration affected misinformation production on Test 2 following exposure to the fictional stories. The dependent measure for this analysis was the total number of times each participant responded to the critical questions with misinformation (i.e. the possible range of total misinformation responses being 0-36 for the subset of misleading and neutral critical items). Since all participants viewed half of the critical items in a neutral frame and the other half with the misleading frame, fact framing was specified as a repeated observation across participants. We fit the misinformation response data with a negative binomial distribution and used a log link function for this analysis.

The interaction between fact framing and condition was not significant, Wald $\chi^2(1, N = 96) = .002$, $p = .96$, but there was a significant effect of fact framing, Wald $\chi^2(1, N = 96) = 219.80$, $p < .001$, and a significant effect of condition, Wald $\chi^2(1, N = 96) = 25.24$, $p < .001$. When questions were referenced with misleading frames in the stories, participants were more likely to respond to those questions with misinformation on the general knowledge test ($M = 6.88$ $SE = .55$) as compared to those referenced with a neutral frame ($M = 1.34$ $SE = .12$). When participants worked collaboratively to answer general knowledge questions they were less likely to respond with misinformation ($M = 2.95$ $SE = .38$) than their peers who worked individually ($M = 5.27$ $SE = .56$), see Figure 2.

We were also interested in whether the effects of fact framing and collaboration on misinformation production depended on participants' prior knowledge. In other words, if participants knew the correct answer on Test 1, would this prior knowledge affect whether misinformation would be produced on Test 2 following fictional story reading? To investigate this question, we tracked participants' responses to the critical items from Test 1 to Test 2, focusing on shifts from a correct response on Test 1 to a misinformation response on Test 2. We identified the questions where participants initially possessed knowledge of the correct answer and documented the instances where these initial correct responses were replaced with misinformation responses on Test 2. Thus, the dependent measure for this analysis was the total number of times each participant responded to the critical questions (which they previously answered correctly) with misinformation on Test 2. As with the previous analysis of misinformation response data, we again used a negative binomial distribution and a log link function for this analysis. The interaction between fact framing and condition was not significant, Wald $\chi^2(1, N = 96) = 1.25$, $p = .26$, nor was the effect of condition, Wald $\chi^2(1, N = 96) = 2.89$, $p = .09$, but there was a significant effect of fact framing, Wald $\chi^2(1, N = 96) = 41.77$, $p < .001$, indicating that on questions where participants demonstrated prior knowledge, participants were more likely to respond with misinformation after viewing misleading references to these questions in the story ($M = 1.22$ $SE = .03$) as compared to neutral references ($M = .09$ $SE = .03$).

Test 3 Performance

Test 3 was completed individually for all participants with the goal of assessing the downstream effects of collaboration relative to repeated individual performance. The analyses that follow focus on the rates of correct responding and misinformation production for the critical questions that were referenced with misleading or neutral frames in the stories.

Correct Responses. We followed the same model fitting procedure as described earlier for the analysis of correct responses on Test 2. Like before, the random intercepts and slopes model provided the best fit to the data, as compared to the random intercepts only model, and the model that did not take the hierarchical data structure into account; for a summary of the model fitting process, see Table 4.

Looking at the fixed effects estimates in the random intercepts and slopes model of correct responses on Test 3, we see the same pattern of results as observed on Test 2. There was a significant effect of fact framing, $F(1, 52.21) = 9.77$, $p = .003$, where items presented in a misleading frame ($M = .55$ SE =.02) were less likely to be answered correctly as compared to those facts referenced with a neutral frame in the story ($M = .60$ SE =.02). There was also a significant effect of collaboration, $F(1, 46.96) = 17.67$, $p < .001$, where participants who previously worked in groups ($M = .67$ $SE = .02$) answered more questions correctly than their peers who worked individually throughout all three test phases ($M = .49$ $SE = .02$). Again, the difference between the correct response rate for misleading and neutral critical items was numerically greater for participants working individually, but the interaction between fact framing and collaboration was not significant, $F(1, 52.21) = .99$, $p = .32$, indicating that the effects of fact framing on the rate of correct responding were again consistent regardless of whether participants collaborated or worked individually, see Figure 3. A summary of the regression coefficients for all three fixed effects can be found in Table 6.

We also examined the within subjects comparisons of correct response rates from Test 2 to Test 3 to assess how performance changed across these two tests. In our previous experiments (Pociask, et al., 2014), a drop in accurate knowledge retrieval was observed as participants transitioned from the collaborative test to the final individual test, whereas consistent or slightly improved performance was observed for participants who worked individually. We expected to see a similar pattern of results here, such that there would be a decrease in accurate retrieval for those in the I-C-I condition from Test 2 to Test 3, but either no change or a slight improvement for those in the I-I-I condition.

The dependent measure for this analysis was the proportion of critical questions answered correctly. We used MLM to compare performance across tests for participants in the I-C-I condition and a repeated measures ANOVA for participants in the I-I-I condition. Beginning with the I-C-I condition, we first created a model which included fact framing and test as repeated observations across participants since all participants were exposed to both types of fact references and all completed Tests 2 and 3. As a first step, we created a basic model (similar to a two factor repeated measures ANOVA) as a point of comparison where fact framing (misleading versus neutral), test (Test 2 versus Test 3), and their interaction were included as fixed effects. We then compared this model to one that included a random intercept to take into account the variations in means due to group clustering, and finally a random intercepts and slopes model which also allowed the effect of fact framing and test to vary. Again, the random intercepts and slopes model yielded the best fit to the data, as compared to the random intercepts only model, and the model that did not take the group structuring into account; for a summary of the model fitting process, see Table 5.

Looking at the fixed effects estimates in the random intercepts and slopes model, there was a significant main effect of fact framing, $F(1, 53.57) = 5.83$, $p = .02$, a significant main effect of test, $F(1, 53.57) = 8.84$, $p = .004$, but the interaction between test and fact framing was not significant, $F(1, 53.57) = .02$, $p = .90$. These main effects indicate that together across both tests, items presented in a misleading frame ($M = .56$ $SE = .02$) were less likely to be answered correctly as compared to those facts referenced with a neutral frame in the story ($M = .61$ $SE = .01$), and as expected, as participants in the I-C-I condition transitioned from the collaborative test (Test 2, $M = .59$ $SE = .01$) to the final individual test (Test 3, $M = .57$ $SE = .01$) they were less likely to answer questions correctly.

When comparing performance across Tests 2 and 3 for participants in the I-I-I condition, there was a significant main effect of fact framing, $F(1, 47) = 17.00$, $p < .001$, a significant main effect of test, $F(1, 47) = 5.83$, $p = .02$, and a significant interaction between test and fact framing, $F(1, 47) = 6.24$, $p = .02$. While there was no change in the rate of correct responding for neutral items from Test 2 to Test 3 ($M = .52$ $SE = .02$), $F(1, 47) = .34$, $p = .55$, there was a change for misleading items such that participants were less likely to answer misleading questions correctly on Test 2 (M = .44 SE =.03) immediately following misinformation exposure, as compared to Test 3 (M = .45 SE =.03) the final individual test, F(1, 47) = 9.47, p = .003.

Finally, to document changes from the start of experiment to the end of the experiment, we compared the rate of correct responding from Test 1 to Test 3. Using MLM, the model taking into account group clustering for participants in the I-C-I condition failed to converge, so we instead used a repeated measures ANOVA with test (Test 1 vs. Test 2) and fact framing (misleading versus neutral) as repeated measures factors. There was a significant main effect of test, F(1, 47) = 139.51, p < .001, which indicated that participants in the I-C-I condition answered more questions correctly at the end of the experiment on Test 3 (M = .66 SE =.02) than they did on Test 1 (M = .50 SE =.02). The main effect of fact framing was only marginally significant, F(1, 47) = 3.56, p = .07, and the interaction between test and fact framing was not significant, F(1, 47) = 2.78, p = .10. This suggests that the improved performance across tests was relatively consistent and observed for both misleading and neutral critical items, which again illustrates the benefits of collaboration in reducing the negative impact of misinformation exposure.

When looking at how performance changed from Test 1 to Test 3 for those participants who completed repeated individual tests, we found that in contrast, there was no significant effect of test, F(1, 47) = 1.88, p = .18, a significant main effect of fact framing, F(1, 47) = 10.10, p = .003, and a significant interaction between test and fact framing, F(1, 47) = 9.84, p = .003. Focusing on the significant interaction, for misleading critical items, participants answered significantly fewer questions correctly on Test 3 as compared to Test 1, F(1, 47) = 6.20, p = .02. This was not the case for critical neutral items, where a small, nearly significant improvement was observed from Test 1 (M = .51 SE =.02) to Test 3 (M = .52 SE =.02), F(1, 47) = 3.99, p =

.05. Again, this finding demonstrates the negative impact of misinformation exposure in that correct responses were significantly reduced for critical misleading items. And, when contrasting the pattern of results observed with I-I-I participants relative to their peers in the I-C-I condition, we can see the benefits of collaboration in reducing the negative impact of misinformation exposure.

Misinformation production. We again used the GEE approach to assess how fact framing (misleading or neutral) and collaboration affected misinformation production on Test 3 where all participants worked individually to complete the general knowledge questions. Results were consistent with the previously observed pattern on Test 2. The interaction between fact framing and condition was not significant, Wald $\chi^2(1, N = 96) = 1.08$, $p = .30$, but there was a significant effect of fact framing, Wald $\chi^2(1, N = 96) = 99.37$, $p < .001$, and a significant effect of condition, Wald $\chi^2(1, N = 96) = 12.09$, $p = .001$. When questions were referenced with misleading frames in the stories, participants were more likely to respond to those questions with misinformation on the general knowledge test ($M = 6.40$ $SE = .52$) as compared to those referenced with a neutral frame ($M = 1.44$ $SE = .16$), see Figure 4. When participants worked collaboratively to answer general knowledge questions they were less likely to respond with misinformation ($M = 2.83$ $SE = .33$) than their peers who worked individually ($M = 5.00$ $SE = .54$).

We also investigated whether there were any changes in misinformation production from Test 2 to Test 3 in either the I-C-I or I-I-I conditions. We included test (Test 2 versus Test 3), fact framing (misleading versus neutral), condition (I-C-I versus I-I-I), and their interactions in the GEE model, and again used a negative binomial distribution and log link function for the analysis. The total number of misinformation responses was the dependent measure. There was a significant main effect of fact framing, Wald $\chi^2(1, N = 96) = 178.30$, $p < .001$, indicating that overall, across Test 2 and 3, participants were more likely to respond with misinformation when facts were referenced with a misleading frame in the story ($M = 6.63$ $SE = .38$) compared to a neutral frame ($M = 1.39$ $SE = .10$). There was also a main effect of condition, Wald $\chi^2(1, N = 96) = 22.48$, $p < .001$, indicating that overall participants in the I-C-I condition produced fewer misinformation items ($M = 2.89$ $SE = .25$) than those in the I-I-I condition ($M = 5.14$ $SE = .39$). Because there was no main effect of test or any significant interactions (all p values greater than .13), misinformation production was consistent across Tests 2 and 3.

Finally, to document changes from the start of experiment to the end of the experiment, we compared the rate of misinformation production from Test 1 to Test 3. We included test (Test 1 versus Test 3), fact framing (misleading versus neutral), condition (I-C-I versus I-I-I), and their interactions in the GEE model, and again used a negative binomial distribution and log link function for the analysis. There was a significant main effect of test, Wald $\chi^2(1, N = 96) = 37.99$, $p < .001$, a significant main effect of fact framing, Wald $\chi^2(1, N = 96) = 66.73$, $p < .001$, and a significant main effect of condition, Wald $\chi^2(1, N = 96) = 5.08$, $p = .02$. These main effects were qualified by a significant interaction between test and fact framing, Wald $\chi^2(1, N = 96) = 60.42$, $p < .001$, and a significant interaction between test and condition, Wald $\chi^2(1, N = 96) = 7.88$, $p = .005$, but the fact framing by condition interaction was not significant, Wald $\chi^2(1, N = 96) = .61$, $p = .43$, nor was the three way interaction, Wald $\chi^2(1, N = 96) = .80$, $p = .37$. Explaining the two significant interactions, the test by fact framing interaction illustrates that while there were no differences in misinformation production across Tests 1 and 3 for neutral items, Wald $\chi^2(1, N = 96) = .60$, $p = .44$, misinformation production significantly increased from Test 1 to Test 3 for misleading items, Wald $\chi^2(1, N = 96) = 123.28$, $p <.001$. The test by condition interaction shows that there were no differences in misinformation production between conditions on Test 1, Wald $\chi^2(1, N = 96) = .002$, $p = .96$; misinformation was significantly greater for the I-I-I condition on Test 3 relative to the I-C-I condition, Wald $\chi^2(1, N = 96) = 19.82$, $p < .001$.

Summary of Test Performance Results in Experiment 1

To briefly summarize, results from Test 2 demonstrated that after being exposed to misleading fact references in the stories, participants were less likely to answer questions correctly and more likely to respond with the misinformation that they read. This reduction in accurate retrieval was observed for all participants, regardless of whether they engaged in collaborative discussion or worked individually on this test. However, even though collaboration did not completely eliminate the negative impact of misinformation exposure, it was nonetheless associated with key benefits in test performance. When participants engaged in collaborative discussion the rate of correct responding increased relative to repeated individual testing; and for questions where accurate prior knowledge was demonstrated, collaboration eliminated the negative impact of misinformation exposure such that a reduction in consistent correct responses was only participants who worked individually. On Test 3, when all participants completed one final individual test, we observed that the benefits of collaboration and the costs of misinformation exposure persisted. All participants were negatively impacted by exposure to misleading fact references in the stories earlier in the experiment; however, this negative impact was reduced for those who had collaborated earlier compared to those who did not.

Confidence Ratings

On every test, after answering each question participants also rated their confidence in their answer on a scale of 1 (not at all confident) to 5 (extremely confident). In order to understand how fact framing and collaboration affected how confident participants were in their answers, we calculated the change in confidence ratings from Test 1 to Test 2. For all questions that were answered on both tests, we subtracted the confidence rating provided on Test 1 from the value provided on Test 2, and in doing so, we were able to note the changes in confidence ratings across these two tests. The possible range for change scores was -4 to +4, reflecting a decrease from 5 to 1, and an increase from 1 to 5, respectively. The analyses that follow focus on the average change in confidence ratings for the critical items presented with a misleading or neutral frame in the stories. The ICC value for this variable was not significant, ICC = .02, $F(15, 32) = 1.05$, $p = .44$, so these data were analyzed with a 2 (Fact framing: misleading versus neutral) × 2 (Condition: collaborative versus individual) repeated measures ANOVA. The main effect of fact framing was not significant, $F(1, 94) = .07$, $p = .80$, but there was a main effect of condition, $F(1, 94) = 27.23$, $p < .001$, and most critically, a significant interaction between fact framing and condition, $F(1, 94) = 5.48$, $p = .02$. Focusing on the significant interaction, a series of *t* tests revealed that while those in the I-C-I condition reported a significantly greater increase in confidence from Test 1 to Test 2 relative to those in the I-I-I condition for both misleading and neutral items, [$t(94) = 3.89$, $p < .001$ and $t(94) = 5.69$, $p < .001$, respectively], there was a greater difference for the neutral items ($d = 1.16$) compared to misleading items ($d = .79$). This indicates that collaborative discussion was associated with increased confidence for all items, but exposure to misleading information in the stories reduced the magnitude of this confidence boost for the misleading critical items.

We also tracked the changes in confidence ratings from Test 2 to Test 3 (subtracting the confidence value reported on Test 2 from Test 3), to assess any shifts in confidence, from the collaborative test to the final individual test for those in the I-C-I condition and from the second to third individual test for those in the I-I-I condition. The ICC value for this variable was significant, ICC = .25, $F(15, 32) = 2.00$, $p < .05$, so we used multilevel modeling to analyze these data. We began by fitting a model similar to a repeated measures ANOVA (with fact framing, condition, and their interaction included as fixed effects), and compared the fit of this model with one that included group as a random intercept. However, taking into account the group structure of the data with a random intercept did not significantly improve the model fit ($\chi^2(1, N = 96) = 2.43$, $p = .12$), so the simpler model was retained for our analyses. No main effects were significant, nor was the interaction (all p values were greater than .19). While a small, yet significant increase in confidence was observed from Test 2 to Test 3 when averaging across all participants ($M = .08$ $SE = .02$), $t(191) = 3.81$, $p < .001$, the average change in confidence across the final two tests did not vary depending on fact framing or experimental condition.

Finally, in order to capture the shift in confidence from the very beginning of the experiment to the end, we tracked confidence changes from Test 1 to Test 3, subtracting the confidence rating reported on Test 1 from the value reported on Test 3. The ICC value for this variable was not significant, ICC = -.09, $F(15, 32) = .76$, $p = .71$, so these data were analyzed with a 2 (Fact framing: misleading versus neutral) × 2 (Condition: collaborative versus individual) repeated measures ANOVA. There was no main effect of fact framing, $F(1, 94) = 1.38$, $p = .24$, and no interaction, $F(1, 94) = 2.57$, $p = .11$, but a significant main effect of condition, $F(1, 94) = 22.62$, $p < .001$. The mean difference in confidence ratings from Test 1 to Test 3 for participants in the I-C-I condition ($M = .81$ $SE = .08$) was significantly greater than the mean difference observed for the I-I-I participants ($M = .36$, $SE = .05$), indicating that the increase in confidence over the course of the experiment was greater for those who had the opportunity to collaborate.

Discussion of Experiment 1

The goal of Experiment 1 was to test the error pruning mechanisms of collaboration under conditions where participants were systematically exposed to misinformation prior to retrieval of general knowledge. Exposure to misinformation has been previously shown to decrease accurate knowledge retrieval (Eslick, et al., 2011; Fazio, et al., 2013; Fazio & Marsh, 2008; Marsh & Fazio, 2006; Marsh, et al., 2003), and collaborative discussion has been shown to increase accurate knowledge retrieval (Pociask, et al., 2014). By including both manipulations—misinformation exposure and collaboration—the novel question tested by Experiment 1 was whether the benefits of collaboration (i.e. improved retrieval accuracy) could be maintained in this impoverished retrieval context.

Results from Experiment 1 indicated that the benefits of collaboration did indeed persist despite exposure to misinformation. Collaborative discussion improved the accuracy of general knowledge retrieval and increased participants' confidence in their responses. While participants who engaged in collaborative discussion significantly outperformed those that worked individually, they were not completely immune to the negative effects of misinformation exposure. Consistent with related work on this topic (e.g., Marsh, et al., 2003), we found that misinformation exposure via fictional stories had a negative impact on retrieval of accurate general knowledge. Specifically, after reading misleading information in the stories, participants were less likely to answer general knowledge questions correctly and more likely to respond with the misinformation that they had just read. Within the context of this important replication, a novel finding from Experiment 1 was that all participants experienced the negative impact of misinformation exposure, regardless of whether they completed repeated individual tests or had the opportunity to engage in collaborative discussion. But again, demonstrating yet another key benefit of collaboration, with regard to questions where participants demonstrated accurate prior knowledge of the answer, only participants who worked individually were negatively affected by misinformation exposure, as a reduction in the consistent correct responses was observed for participants who worked individually, but not for those who worked collaboratively.

In sum, Experiment 1 demonstrated that when participants were systematically exposed to misinformation via fictional story reading prior to the retrieval of general knowledge, the benefits of collaborative discussion were still observed, and although collaboration did not completely eliminate the transmission of misinformation, it significantly reduced its impact. But what would happen if participants collaborated before being exposed to misleading information instead of afterwards? Would collaboration confer a protective benefit? To our knowledge, there are no studies to date that have investigated the effects of collaboration in terms of protection against future errors. Thus, to continue our investigation of the effects of collaboration and misinformation exposure on retrieval of general knowledge we investigated these novel questions in Experiment 2.

III. Experiment 2

Following from and building upon the results of Experiment 1, in Experiment 2 we investigated yet another way that collaboration might reduce the impact of misinformation exposure—through protection or inoculation against errors. Specifically, the primary question of Experiment 2 was whether collaboration occurring prior to misinformation exposure (relative to repeated individual test-taking) would protect participants against the negative impact of misinformation. In order to investigate this question, Experiment 2 followed a similar design as Experiment 1, but the story reading phase (which included misleading fact references) occurred after the completion of the second general knowledge test. This subtle yet critical design change allowed us to investigate whether improved performance as a result of collaboration—increased accuracy and increased confidence—would serve to reinforce participants' correct knowledge prior to misinformation exposure, and consequently, mitigate the negative effects.

Based on the results of Experiment 1 and previous research highlighting the benefits of collaborative discussion with respect to general knowledge retrieval, we hypothesized that that collaborative discussion would again lead to improved performance and increased confidence, and as a result of this, participants would be more resilient to misinformation presented in the stories as compared to their peers who worked individually. But, just because a person is highly confident in an answer, this does not necessarily mean that this answer will be more resistant to change. For example, as evidenced by the hypercorrection effect, after being provided with correct feedback, errors endorsed with high levels of confidence were more likely to be corrected on a subsequent test as compared to lower confidence errors (Butterfield & Metcalfe, 2001, 2006). Extending this idea to the present experiment, if participants see the fact references in the story as corrective feedback, they might be inclined to change their answers, and thus it is possible that the confidence boost from collaboration would not serve to protect against misinformation exposure. Moreover, given that the effects of misinformation exposure are quite robust and difficult to eliminate (e.g., Marsh & Fazio, 2006), and were observed in Experiment 1 for all participants, it is possible that the benefits of collaboration might not extend to this novel situation.

Method

Participants

A new group of 96 Stony Brook University undergraduates participated in this study for course credit. All participants were fluent in English.

Materials/Procedure

The same tests and stories used in Experiment 1 were used in Experiment 2. As described above, a similar experimental procedure was followed for Experiment 2. Throughout the course of the experiment, all participants completed three general knowledge tests in an I-C-I (individual, collaborative, individual) or I-I-I (individual, individual, individual) sequence. The critical difference between Experiments 1 and 2 was in the timing of the story reading phase where participants were exposed to misinformation. In Experiment 2, this occurred after the second general knowledge test was completed, and before the final individual general knowledge test. All other aspects of the procedure purposefully remained the same across experiments.

Experiment 2 Results

Data Analysis Approach

The novel aspect of Experiment 2 was the timing of collaboration relative to exposure to the stories. Whereas participants in Experiment 1 were exposed to stories prior to collaboration, in Experiment 2 participants were exposed to the stories following collaboration. With this design, the primary goal of Experiment 2 was to investigate whether collaboration occurring prior to reading misleading information in the stories has any protective influence against the negative impact of misinformation exposure.

As described with Experiment 1, given our research questions and research design, it was necessary to assess the dependencies in our data structure in order to determine the appropriate data analysis approach. The results of these analyses confirm that participants influenced each other's responding as a result of collaboration, see Table 7 for a summary of the ICCs and DE values for all test performance variables in Experiment 2. For this reason, when analyzing normally distributed data obtained from Tests 2 and 3 (i.e. correct responses) we again used multilevel modeling adapted for partially clustered data structures to take into account the dependencies due to collaboration as well as the independence of the participants in the individual condition. We used the same model fitting process described in Experiment 1 to select the appropriate model from which to evaluate the fixed effects of interest. For misinformation responses, which again were positively skewed, we used the generalized estimating equations (GEE) approach, and fit the data with a negative binomial distribution and log link function.

Test 1 Performance

At this phase of the experiment, all participants completed the general knowledge test individually as a baseline measure of knowledge at the start of the experiment. We used the same general knowledge test as in Experiment 1, which consisted of the same three subsets of questions: critical questions to be referenced with a misleading frame in the story, critical questions to be referenced with a neutral frame in the story, and filler questions. Participant responses on the general knowledge tests were coded as correct responses, omissions, misinformation errors (i.e. the critical misinformation referenced in the story), and other errors (i.e. other erroneous responses not referenced in the story). Even though participants had not yet read the stories containing misinformation at Test 1, we still coded misinformation errors that occurred on this test in order document the baseline rate at which participants already believed that the misinformation was indeed the correct answer. Again, as in Experiment 1, misinformation was rare on Test 1, as these responses were observed in response to only about 5% of critical questions.

We examined the performance of participants assigned to the I-C-I and I-I-I test sequences to determine there were a priori differences in knowledge between participants in each experimental condition, see Table 8 for a summary of means and standard errors for all test performance variables on Test 1. As expected, there were no differences in performance between the I-C-I condition and the I-I-I condition at this stage in the experiment, across all questions and within the critical and filler item subsets[5]. Importantly, performance was the same for critical items to be presented in a misleading frame and neutral frame.

Test 2 Performance

Test 2 was completed either collaboratively in a triad or individually. Analyzing test performance at this stage of the experiment allows us to evaluate how collaboration alone affected performance on the general knowledge test relative to repeated individual performance. Because participants were not exposed to the fictional stories at this point in the experiment, in the sections that follow we focus on overall test performance across all questions, rather than on the subsets of critical questions.

We used multilevel modeling to assess how collaboration affected the rate of correct responding. The overall proportion of correct responses on the general knowledge test was used as the dependent measure in the model, and collaboration versus individual test taking was the independent measure. Compared to the basic model (similar to a one-way ANOVA), the model with a random intercept for group clustering was a significantly better fit, $\chi^2(1, N = 96) = 77.22$, $p < .001$. Relative to those participants working individually, collaboration significantly improved performance on the general knowledge test, $F(1, 62.21) = 36.97$, $p < .001$, as

[5] As observed in Experiment 1, when comparing performance on the critical items to the filler items, participants answered fewer filler items correctly, $F(1,95) = 703.50$, $p<.001$, omitted more filler items as compared to critical items, $F(1,95) = 574.17$, $p<.001$, and answered more filler questions with errors, $F(1,95) = 13.80$, $p < .001$ (because error rates were generally low and positively skewed, we also ran this analysis with the non-parametric equivalent, the Wilcoxon Signed Rank Test, and the same conclusion was reached, $z = 3.42$, $p = .001$). These results are consistent with prior norming indicating that the questions selected to be filler items were on average more difficult than those selected to be critical items.

participants working in groups answered more questions correctly (M = .60 SE =.02), compared to their peers who completed a second individual test (M = .39 SE =.02).

We also examined the within subjects comparisons from Test 1 to Test 2 to assess how performance changed across tests depending on collaboration. Again, the dependent measure was the overall proportion of questions answered correctly, and test was a repeated observation across participants since all participants completed both Tests 1 and 2. As a first step, we created a basic model (similar to a standard repeated measures analysis of variance) as a point of comparison where test (Test 1 versus Test 2), condition (I-C-I versus I-I-I), and their interaction were included as fixed effects. As a next step, we included a random intercept in order to model the variation in means due to the fact that some participants worked in groups while others worked individually. A comparison of the deviance values revealed that the model taking group clustering into account with random intercepts was a significantly better fit, $\chi^2(1, N = 96) = 9.01$, $p < .01$. We then compared the random intercept model to a more complex model which also included random slopes for the effect of test, but this model failed to converge, so we retained the random intercept model as the final model.

There was a significant main effect of condition, $F(1,41.29) = 11.45$, $p<.01$, and a main effect of test, $F(1, 96) = 208.87$, $p<.001$, and a significant interaction, $F(1, 96) = 152.14$, $p<.001$. Examining the change in performance from Test 1 to Test 2 separately for those in the I-C-I and I-I-I conditions, we see that in both conditions there was statistically significant improvement across tests; however, the difference in performance across tests was greater for those who had the opportunity to collaborate, $b = .23$, $t(48) = 14.06$, $p <.001$, as compared to those who completed repeated individual tests, $b = .02$, $t(48) = 3.58$, $p <.01$, see Figure 5.

Test 3 Performance

Prior to completing Test 3—which was completed individually for all participants—everyone was exposed to the stories which contained critical misleading and neutral fact references. This was the critical novel feature of Experiment 2 and it was introduced to test the question whether prior collaboration will protect participants from incorporating misinformation to which they are exposed subsequently. The analyses that follow focus on investigating the effect of prior collaboration on the rates of correct responding and misinformation production for the critical questions that were referenced with misleading or neutral frames in the stories.

Correct Responses. We used multilevel modeling to assess how fact framing (misleading or neutral) and prior collaboration (relative to repeated individual testing) affected the retrieval of correct information following exposure to the fictional stories. The dependent measure was the proportion of misleading and neutrally framed items answered correctly, and since all participants viewed half of the critical items in a neutral frame and the other half with the misleading frame, fact framing was a repeated observation across participants. We followed the same basic model fitting procedure as described in previous sections: we started with a basic model (similar to a standard repeated measures analysis of variance) as a point of comparison, then included a random intercept to account for variations in means due to group clustering, and finally we included a random slope to allow for the effect of fact framing to vary across groups and participants. Successive comparisons of these models revealed that the random slopes and intercept model provided the best fit for the data; see Table 9 for a summary of the model fitting process.

Turning our attention to the fixed effects of interest, there was a significant effect of fact framing, $F(1, 34.43) = 232.42$, $p < .001$, where items referenced with a misleading frame in the story ($M = .45$ $SE = .02$) were less likely to be answered correctly on the general knowledge test compared to those referenced with a neutral frame (*M* = .58 *SE* =.02). There was also a significant effect of collaboration, $F(1, 40.28) = 18.87$, $p <.001$, where participants working in groups (*M* = .60 *SE* =.01) answered more questions correctly than their peers working individually (*M* = .43 *SE* =.02). Although numerically the difference between the correct response rate for misleading and neutral critical items is greater for participants working individually, the interaction between fact framing and collaboration was not significant, $F(1, 34.43) = .12$, $p = .73$, indicating that the effects of fact framing on the rate of correct responding were relatively consistent regardless of whether participants collaborated or worked individually, see Figure 6. A summary of the regression coefficients for all three fixed effects can be found in Table 10.

As in Experiment 1, we were again interested in whether the effects of collaboration and misinformation exposure depend on prior knowledge. To investigate this possibility, we tracked participants' correct responses to the critical items from Test 1 to Test 3. We identified the questions where participants initially possessed knowledge of the correct answer and documented the rate at which these correct responses persisted on Test 3 (following misinformation exposure via story reading). Thus, the analyses that follow focus on the proportion of critical questions participants consistently answered correctly, both on the initial general knowledge test and on the final general knowledge test. The ICC value for this variable was not significant, ICC = -.08, $F(15, 32) = .77$, $p = .70$, so these data were analyzed with a 2 (Fact framing: misleading versus neutral) × 2 (Condition: collaborative versus individual) repeated measures ANOVA. The main effect of condition was not significant, $F(1,94) = .62$, $p = .43$, but there was a main effect of fact framing, $F(1,94) = 16.77$, $p < .001$, and a significant interaction between fact framing and condition, $F(1,94) = 5.82$, $p = 02$. To understand this interaction, we ran a series of contrasts to isolate the effect of fact framing for those participants in the I-C-I condition versus those in the I-I-I condition. When participants worked individually, there was reduction in consistent correct responses for the questions that were referenced with a misleading frame in the story (*M* = .38 *SE* =.03) as compared to a neutral frame (*M* = .46 *SE* =.03), $F(1,47) = 15.14$, $p < .001$, demonstrating the negative impact of misinformation exposure, even on questions where participants possessed prior knowledge. However, fact framing did not affect the rate of consistent correct answers for participants who worked collaboratively on Test 2, $F(1,47) = 2.36$, $p = .13$, as the rate of consistent correct responses was no different when facts were referenced with a misleading frame (*M* = .44 *SE* =.03) or a neutral frame (*M* = .46 *SE* =.02).

We also examined a more stringent measure of prior knowledge—consistent correct answers provided on both Test1 and Test 2. We identified the questions where participants consistently provided the correct answer on both Test 1 and Test 2, and then documented the rate at which these correct responses persisted on Test 3 (following misinformation exposure via story reading). The ICC value for this variable was not significant, ICC = -.01, $F(15, 32) = .96$, $p = .52$, so these data were analyzed with a 2 (Fact framing: misleading versus neutral) × 2 (Condition: collaborative versus individual) repeated measures ANOVA.

The main effect of condition was not significant, $F(1,94) = .76$, $p = .39$, but there was a main effect of fact framing, $F(1,94) = 18.56$, $p < .001$, and a significant interaction between fact framing and condition, $F(1,94) = 5.78$, $p = 02$. To understand this interaction, we ran a series of contrasts to isolate the effect of fact framing for those participants in the I-C-I condition versus those in the I-I-I condition, and we observed the same pattern of results as described above. When participants worked individually throughout the experiment, there was reduction in consistent correct responses for the questions that were referenced with a misleading frame in the story (*M* = .37 *SE* =.03) as compared to a neutral frame (*M* = .45 *SE* =.02), $F(1,47) = 16.30$, $p < .001$, demonstrating the negative impact of misinformation exposure, even on questions where participants possessed prior knowledge. However, fact framing did not affect the rate of consistent correct answers for participants in the I-C-I condition, $F(1,47) = 2.93$, $p = .09$, as the rate of consistent correct responses was no different when facts were referenced with a misleading frame (*M* = .43 *SE* =.03) or a neutral frame (*M* = .45 *SE* =.02).

Finally, in order to document how the overall rate of correct responses changed from the very beginning of the experiment to the final test, we compared correct response rates to the critical items on Test 1 to Test 3 with a repeated measures ANOVA with test (Test 1 vs. Test 2) and fact framing (misleading versus neutral) as repeated measures factors. There was a significant main effect of test, $F(1,47) = 202.66$, $p <.001$, a marginal main effect of fact framing, $F(1,47) = 3.73$, $p =.06$, and a significant interaction between test and fact framing, $F(1,47) = 8.08$, $p = .007$. To understand this interaction, we compared the change in correct response rates across tests for neutral and misleading critical items separately, finding significant improvement across tests for items presented with a misleading frame [$t(47) = -10.38$, $p <.001$] and a neutral frame [$t(47) = -114.03$, $p <.001$]; however, the improvement across tests was greater for neutral items ($d = .92$), as compared to misleading items, ($d = 1.31$).

We ran the same analyses for participants working individually, and there was a significant main effect of test, $F(1,47) = 4.95$, $p = .03$, a significant main effect of fact framing, $F(1,47) = 6.27$, $p = .02$, and a significant interaction between test and fact framing, $F(1,47) = 21.96$, $p = .001$. To understand this interaction, we compared the change in correct response rates across tests for neutral and misleading critical items separately, finding significant improvement across tests for the critical neutral items, $t(47) = 2.60$, $p = .01$, but a significant reduction in correct responses for the critical misleading items, $t(47) = -3.53$, $p = .001$. This comparison shows the negative impact of misinformation over the course of the experiment.

Misinformation Responses. We used the GEE approach to assess how fact framing (misleading or neutral) and collaboration affected misinformation production following exposure to the fictional stories. The dependent measure for this analysis was the total number of times each participant responded to the critical questions with misinformation (i.e. the possible range of total misinformation responses being 0-36 for the subset of misleading and neutral critical items), and since all participants viewed half of the critical items in a neutral frame and the other half with the misleading frame, fact framing was specified as a repeated observation across participants. We fit the misinformation response data with a negative binomial distribution and used a log link function for this analysis.

There was a significant effect of fact framing, Wald $\chi^2(1, N = 96) = 13.02$, $p < .001$, a significant effect of condition, Wald $\chi^2(1, N = 96) = 133.52$, $p < .001$, and a significant interaction between fact framing and condition, Wald $\chi^2(1, N = 96) = 13.23$, $p < .001$. In order to understand the interaction, we looked at the effect of condition on misinformation production separately for the critical items presented with a misleading frame and neutral frame. Participants who previously collaborated (I-C-I, $M = 1.85$ $SE = .19$) and participants who repeatedly worked individually (I-I-I, $M = 1.81$ $SE = .25$) did not differ in misinformation production for items presented with a neutral frame in the stories, Wald $\chi^2(1, N = 96) = .02$, $p = .90$, however, for items presented with a misleading frame, Wald $\chi^2(1, N = 96) = 17.10$, $p < .001$, participants who worked collaboratively answered fewer questions with misinformation ($M = 4.75$ $SE = .60$) than their peers who repeatedly worked individually ($M = 11.02$ $SE = 1.12$).

We were also interested in whether the effects of fact framing and collaboration depended on participants' prior knowledge. As a first step in examining this question, we tracked participants' responses to the critical items from Test 1 to Test 3. We identified the questions where participants initially possessed knowledge of the correct answer and documented the instances where these initial correct responses were replaced with misinformation responses on Test 3. The dependent measure for this analysis was the total number of times each participant responded to the critical questions (which were previously answered correctly) with misinformation on Test 3. As with the previous analysis, we fit the misinformation response data with a negative binomial distribution and used a log link function for this analysis.

The interaction between fact framing and condition was not significant, Wald $\chi^2(1, N = 96) = .98$, $p = .32$, but there was a significant effect of condition, Wald $\chi^2(1, N = 96) = 5.01$, $p = .03$, and a significant effect of fact framing, Wald $\chi^2(1, N = 96) = 44.87$, $p < .001$. On questions where participants demonstrated prior knowledge, participants who previously worked collaboratively were less likely to respond with misinformation ($M = .45$ $SE = .11$) compared to those who repeatedly worked individually ($M = 1.64$ $SE = .30$). Collapsing across conditions, participants were more likely to respond with misinformation after viewing misleading references to these questions in the story ($M = 1.97$ $SE = .30$) as compared to neutral references ($M = .11$ $SE = .05$), see Figure 7.

We also examined a more stringent measure of prior knowledge—consistent correct answers provided on both Test1 and Test 2. We identified the questions where participants consistently provided the correct answer on both Test 1 and Test 2, and then documented the rate at which these correct responses were replaced with misinformation responses on Test 3. The same basic pattern emerged such that the interaction between fact framing and condition was not significant, Wald $\chi^2(1, N = 96) = .07$, $p = .79$, but there was a significant effect of condition, Wald $\chi^2(1, N = 96) = 8.34$, $p = .004$, and a significant effect of fact framing, Wald $\chi^2(1, N = 96) = 34.41$, $p < .001$. On questions where participants demonstrated prior knowledge, participants who previously worked collaboratively were less likely to respond with misinformation (*M* = .35 *SE* = .10) compared to those who repeatedly worked individually (*M* = 1.59 *SE* = .30). Collapsing across conditions, participants were more likely to respond with misinformation after viewing misleading references to these questions in the story (*M* = 1.85 *SE* = .30) as compared to neutral references (*M* = .09 *SE* = .05).

Finally, in order to document how misinformation responses changed from the very beginning of the experiment to the final test, we compared misinformation production on Test 1 to Test 3. We included test (Test 1 versus Test 3), fact framing (misleading versus neutral), condition (I-C-I versus I-I-I), and their interactions in the GEE model, and again used a negative binomial distribution and log link function for the analysis. There was a significant main effect of test, Wald $\chi^2(1, N = 96) = 124.55$, $p < .001$, a significant main effect of fact framing, Wald $\chi^2(1, N = 96) = 59.21$, $p < .001$, and a significant main effect of condition, Wald $\chi^2(1, N = 96) = 7.30$, $p = .007$. These main effects were qualified by a significant interaction between test and fact framing, Wald $\chi^2(1, N = 96) = 119.56$, $p < .001$, and a significant interaction between test and condition, Wald $\chi^2(1, N = 96) = 7.40$, $p = .007$, and a significant interaction between fact framing and condition, Wald $\chi^2(1, N = 96) = 13.98$, $p <.001$. The three way interaction was not significant, Wald $\chi^2(1, N = 96) = .80$, $p = .37$.

Explaining the significant interactions, the test by fact framing interaction illustrates that while there were no differences in misinformation production across Tests 1 and 3 for neutral items, Wald $\chi^2(1, N = 96) = .58$, $p = .49$, misinformation production significantly increased from

Test 1 to Test 3 for misleading items, Wald $\chi^2(1, N = 96) = 223.28$, $p <.001$. The test by condition interaction shows that there were no differences in misinformation production between conditions on Test 1, Wald $\chi^2(1, N = 96) = .13$ $p = .71$; misinformation was significantly greater for the I-I-I condition on Test 3 relative to the I-C-I condition, Wald $\chi^2(1, N = 96) = 26.98$, $p <.001$. Finally, the fact framing by condition interaction shows that while that overall across Tests 1 and 3 there were no differences between conditions for neutral items, Wald $\chi^2(1, N = 96) = .67$ $p = .41$, there was a difference for misleading items, Wald $\chi^2(1, N = 96) = 29.00$, $p <.001$, such that those in the I-I-I condition produced greater levels of misinformation.

Confidence Ratings

On every test, after answering each question participants also rated their confidence in their answer on a scale of 1 (not at all confident) to 5 (extremely confident). In order to understand how fact framing and collaboration affected how confident participants were in their answers, we calculated the change in confidence ratings from Test 1 to Test 2. For all questions that were answered on both tests, we subtracted the confidence rating provided on Test 1 from the value provided on Test 2, and in doing so, we were able to note the changes in confidence ratings across these two tests. The possible range for change scores was -4 to +4, reflecting a decrease from 5 to 1, and an increase from 1 to 5, respectively.

The analyses that follow focus on the average change in confidence ratings from Test 1 to Test 2 for the critical items presented with a misleading or neutral frame in the stories, although, it is worth pointing out that at this stage in the experiment participants had not yet been exposed to the stories, so we did not expect any differences with respect to fact framing. The ICC value for this variable was not significant, ICC = -.03, $F(15, 32) = .90$, $p = .57$, so these data were analyzed with a 2 (Fact framing: misleading versus neutral) × 2 (Condition: collaborative versus individual) repeated measures ANOVA. As expected, the main effect of fact framing was not significant, $F(1, 94) = .19$, $p = .67$, nor was the fact framing by condition interaction, $F(1, 94) = 2.82$, $p = .10$, but there was a main effect of condition, $F(1, 94) = 47.04$, $p < .001$, indicating that participants who worked collaboratively ($M = .94$ $SE = .06$) reported a significantly greater boost in confidence ratings across the tests compared to participants who worked individually ($M = .25$ $SE = .04$). Using a single sample *t* test, we compared the average change in confidence ratings against a value of zero (i.e. no change), and results indicated that there was a significant increase in confidence for participants who worked collaboratively, $t(47) = 10.90$, $p < .001$, as well as those who completed both tests individually, $t(47) = 4.31$, $p < .001$.

We also tracked the changes in confidence ratings from Test 2 to Test 3, to assess any final shifts in confidence, from the collaborative test to the final individual test for those in the I-C-I condition and from the second to third individual test for those in the I-I-I condition. The ICC value for this variable was not significant, ICC = -.09, $F(15, 32) = .76$, $p = .71$, so these data were analyzed with a 2 (Fact framing: misleading versus neutral) × 2 (Condition: collaborative versus individual) repeated measures ANOVA. The main effect of fact framing was not significant, $F(1, 94) = .77$, $p = .38$, nor was the interaction, $F(1, 94) = 3.06$, $p = .08$, but there was a significant main effect of condition, $F(1, 94) = 17.75$, $p < .001$. Participants who had previously worked collaboratively reported a significant decrease in confidence from Test 2 to Test 3 ($M = -.13$ $SE = .04$), $t(47) = -2.92$, $p = .005$, while participants who worked individually reported a significant increase in confidence ($M = .15$ $SE = .05$), $t(47) = 3.02$, $p = .004$.

Finally, in order to capture the shift in confidence from the very beginning of the experiment to the end, we tracked confidence changes from Test 1 to Test 3.The ICC value for this variable was not significant, ICC = -.03, $F(15, 32) = .93$, $p = .55$, so these data were analyzed with a 2 (Fact framing: misleading versus neutral) × 2 (Condition: collaborative versus individual) repeated measures ANOVA. There was no main effect of fact framing, $F(1, 94) = .01$, $p = .94$, but a significant main effect of condition, $F(1, 94) = 18.57$, $p < .001$, and a significant interaction between fact framing and condition, $F(1, 94) = 11.51$, $p = .001$. To understand the interaction, we ran separate one-way repeated measures ANOVAs for participants in the I-C-I and I-I-I conditions. Participants in the I-C-I condition, there was a significant main effect of fact framing, $F(1, 47) = 6.41$, $p = .02$, such that there was a greater increase in confidence for neutral items ($M = .88$ $SE = .08$) as opposed to misleading items ($M = .75$ $SE = .09$). For participants in the I-I-I condition, there was also a significant main effect of fact framing, $F(1, 47) = 5.27$, $p = .03$, but in the opposite direction; there was a greater increase in confidence for items presented with a misleading frame ($M = .44$ $SE = .07$) as compared to a neutral frame ($M = .32$ $SE = .06$).

Summary of Experiment 2 Results

The results of Experiment 2 yielded a number of key findings. First, not only did collaboration significantly improve performance relative to repeated individual retrieval, but also, working collaboratively prior to misinformation exposure provided some protection against the negative impact of misinformation exposure. Specifically, collaboration reduced the rate at which participants produced misinformation, and helped participants maintain correct prior knowledge, even when this prior knowledge was contradicted with misinformation in the stories. Consistent with Experiment 1, the reduction in consistent correct responses was again observed only for participants who repeatedly worked individually. Participants who worked collaboratively also showed greater boosts in confidence throughout the experiment.

While significant benefits were observed as a result of collaboration, the negative impact of misinformation exposure was not completely eliminated. Following misinformation exposure, all participants exhibited a reduction in correct responses and an increase in misinformation production, although as mentioned earlier, misinformation production was lower for participants who previously collaborated.

Comparisons across Experiments 1 and 2

Individual Performance. In Experiment 1, participants in the I-I-I condition read the fictional stories after the baseline general knowledge test, and in Experiment 2, those in the I-I-I condition completed two general knowledge tests before reading the stories. In order to investigate whether the effects of misinformation exposure varied depending on prior test taking (i.e. one or two individual tests), we compared the performance of participants on Test 2 in Experiment 1 with Test 3 in Experiment 2. A 2 (Fact framing: misleading versus neutral) × 2 (Number of prior tests: one versus two) repeated measures ANOVA revealed that there was a significant effect of fact framing, $F(1, 94) = 33.33$, $p < .001$, but there was no difference in the proportion of correct responses depending on whether participants had taken one or two tests prior to misinformation exposure $F(1, 94) = .31$, $p = .58$, and no interaction between fact framing and number of prior tests, $F(1, 94) = .03$, $p = .86$. In other words, participants who worked individually in both Experiments 1 and 2 performed similarly; they were less likely to answer questions correctly when the fact was referenced with a misleading frame ($M = .42$ SE =.02) relative to a neutral frame ($M = .51$ SE =.02), and this effect did not vary depending on the number of prior tests the participants completed.

We also examined whether misinformation production differed depending on the number of prior tests. Using the GEE approach, fact framing was specified as a repeated observation across participants, and we fit the misinformation response data with a negative binomial distribution and used a log link function for this analysis. Again, there was a significant effect of item type, Wald $\chi^2(1, N = 96) = 218.00$, $p < .001$, but there were no differences depending on the number of prior tests, Wald $\chi^2(1, N = 96) = 1.71$, $p = .19$, and there was no interaction, Wald $\chi^2(1, N = 96) = .58$, $p = .45$. Participants were more likely to answer questions with misinformation when the fact was referenced with a misleading frame (*M* = 9.92 *SE* =.70) relative to a neutral frame (*M* = 1.77 *SE* =.15), and this effect did not vary depending on the number of prior tests the participants completed.

Collaborative Performance. In Experiment 1 participants in the I-C-I condition read the fictional stories prior to the second test where they engaged in collaborative discussion to answer the questions, but in Experiment 2, participants in the I-C-I condition collaborated first and then read the stories. In order to assess whether the effects of collaboration change when participants are exposed to misleading information before they engage in group discussion we compare performance on Test 2 across these experiments.

We used MLM to assess whether exposure to misinformation affected the retrieval of correct information during collaborative discussion. The dependent measure was the proportion of critical items answered correctly. Fact framing (misleading versus neutral) was a repeated observation across participants, and misinformation exposure (i.e. Experiment 1, exposure and Experiment 2, no exposure) was a between subjects factor. As a first step, we created a basic model (similar to a standard repeated measures analysis of variance) as a point of comparison where fact framing, misinformation exposure, and their interaction were entered as fixed effects. As a next step, we included a random intercept in order to model the variation in means due to group membership. A comparison of the deviance values revealed that the model taking group clustering into account with random intercepts was a significantly better fit, $\chi^2(1, N = 96) = 94.05$, $p < .001$. Finally, we compared the random intercept model to a more complex model that also included random slopes to allow the effect of fact framing to vary across groups, which again yielded significantly better fit, $\chi^2(1, N = 96) = 79.17$, $p < .001$.

Looking at the fixed effects estimates derived from the random intercepts and slopes model, there was no effect of fact framing, $F(1, 32) = 1.16$, $p = .29$, no effect of misinformation exposure, $F(1, 32) = .01$, $p = .91$, and no significant interaction, $F(1, 32) = 2.22$, $p = .15$. Even though participants in Experiment 1 were exposed to misinformation prior to group discussion, there were no reliable differences in their correct responses as compared to participants who had not been exposed to misinformation at this stage.

We used the GEE approach to assess whether misinformation exposure affected misinformation production during collaborative discussion. Fact framing was specified as a repeated observation across participants, and we fit the misinformation response data with a negative binomial distribution and used a log link function for this analysis. We found that there was a significant effect of fact framing, Wald $\chi^2(1, N = 96) = 32.07$, $p < .001$, a significant effect of misinformation exposure, Wald $\chi^2(1, N = 96) = 11.98$, $p = .001$, and a significant interaction, Wald $\chi^2(1, N = 96) = 107.54$, $p < .001$. Focusing on the significant interaction, we see that for the critical misleading items, Wald $\chi^2(1, N = 96) = 90.52$, $p < .001$, participants who were exposed to the misleading information prior to collaborative discussion (I-C-I, Experiment 1: $M = 4.94$ SE =.62) were more likely to produce this misinformation on the general knowledge test as compared to those who had not yet been exposed to this information (I-C-I, Experiment 2: $M = 1.15$ SE =.10). This cross-experimental comparison illustrates that misinformation errors are more likely to emerge when participants have been systematically exposed to the misleading information prior to collaboration, and importantly that misinformation errors are rarely spontaneously generated when participants have not been exposed to misleading information.

Oddly, there were also differences in misinformation production across experiments for the neutral items, as participants in Experiment 2 were more likely to respond to these questions with misinformation, Wald $\chi^2(1, N = 96) = 18.07$, $p < .001$. It is unclear what to make of this difference in misinformation production for neutral items across experiments, especially considering that initial base rates for misinformation production were similar. Recall, as presented earlier in the results of Test 1 for both experiments, that baseline misinformation responses were observed for only about 5% of critical items, or on average fewer than 4 items per participant, and there were no differences between items that would ultimately be presented with a neutral or misleading framing in the stories later in either experiment. That being said, for unknown reasons, participants in the I-C-I condition in Experiment 2 were slightly more likely to produce misinformation to these neutral items on Test 2 as compared to their peers in Experiment 1.

In order to assess whether there were any differences depending on whether collaborative discussion occurred following or prior to misinformation exposure we compared performance across experiments on Test 3. We used MLM to assess whether there were any differences in retrieval of correct information. The dependent measure was the proportion of critical items answered correctly. Fact framing (misleading versus neutral) was a repeated observation across participants, and timing misinformation exposure (i.e. Experiment 1, prior to collaboration and Experiment 2, following collaboration) was a between subjects factor. As a first step, we created a basic model (similar to a standard repeated measures analysis of variance) as a point of comparison where fact framing, timing of misinformation exposure, and their interaction were entered as fixed effects. As a next step, we included a random intercept in order to model the variation in means due to group membership. A comparison of the deviance values revealed that the model taking group clustering into account with random intercepts was a significantly better fit, $\chi^2(1, N = 96) = 22.75$, $p < .001$. Finally, we compared the random intercept model to a more complex model that also included random slopes to allow the effect of fact framing to vary across groups, which again yielded significantly better fit, $\chi^2(1, N = 96) = 22.29$, $p < .001$.

Looking at the fixed effects estimates derived from the random intercepts and slopes model, there was a significant effect of fact framing, $F(1, 32) = 6.84$, $p = .01$, but no effect of timing of misinformation exposure, $F(1, 32) = .47$, $p = .50$, and no significant interaction, $F(1, 32) = .06$, $p = .81$. On Test 3, participants in the I-C-I condition across both experiments were less likely to respond correctly to misleading critical items ($M = .63$ $SE = .02$) as compared to neutral items ($M = .67$ $SE = .01$), but there were no differences depending on when the misinformation exposure occurred.

Finally, we used the GEE approach to assess whether misinformation production differed across experiments on Test 3, the final test that all participants completed individually. Fact framing was specified as a repeated observation across participants, and timing of misinformation exposure (i.e. prior to collaboration, Experiment 1 or following collaboration, Experiment 2) was a between subjects factor. We fit the misinformation response data with a negative binomial distribution and used a log link function for this analysis. There was a significant effect of fact framing, Wald $\chi^2(1, N = 96) = 59.40$, $p < .001$, but the effect of timing of misinformation exposure was not significant, Wald $\chi^2(1, N = 96) = 2.97$, $p = .09$, and neither was the interaction, Wald $\chi^2(1, N = 96) = 1.57$, $p = .21$. These results indicate that participants were more likely to respond to the misleading questions with misinformation ($M = 4.60$ $SE = .40$) than the neutral questions ($M = 1.53$ $SE = .16$), and effect of misinformation exposure was consistent across experiments.

Discussion of Experiment 2

The primary goal of Experiment 2 was to investigate whether collaboration occurring prior to reading misleading information in fictional stories would have any protective influence against the negative impact of misinformation exposure. While the effects of collaboration have seen a resurgence of interest within the cognitive literature in recent years (for a review, see Rajaram & Pereira-Pasarin, 2010), to our knowledge, the question of whether collaboration is associated with any protective benefits has not yet been explored. Given the findings from Experiment 1, as well as from our previous research (Pociask, et al., 2014) which point to the benefits of collaborative discussion with respect to general knowledge retrieval, we hypothesized that that collaborative discussion would again lead to improved performance and increased confidence, and as a result of this, participants would be more resilient to misinformation presented in the stories as compared to their peers who worked individually.

With respect to accurate responding, misinformation exposure had a negative impact on retrieval of accurate general knowledge, and this was the case for participants who completed repeated individual tests and also for those that had the opportunity to work collaboratively. When considering misinformation responses, the negative impact of misinformation was greater for participants who worked individually, as they were more likely to respond with the misinformation that they read in the stories as compared to their peers who had the opportunity to engage in collaborative discussion. When considering questions where participants demonstrated accurate prior knowledge of the answer (in this case, either on Test 1 only, or on both Tests 1 and 2), consistent with findings observed in Experiment 1, only participants who worked individually were negatively affected by misinformation exposure, as a reduction in the consistent correct responses was observed for participants who worked individually, but not for those who worked collaboratively.

The shifts in confidence throughout the experiment as a result of collaborative discussion versus repeated individual testing and misinformation exposure were quite interesting. From Test 1 to Test 2, all participants reported significant increases in confidence, although this increase was greater for those who worked collaboratively on Test 2, which is consistent with the pattern of results observed in Experiment 1. Then, from Test 2 to Test 3 (following story reading), participants who had previously worked collaboratively became less confident in their responses, while participants who previously worked individually became more confident. This illustrates that seeing the misleading information in the stories affected participants differently depending on if they had previously collaborated or completed repeated individual tests. Participants in the I-C-I condition adjusted confidence ratings downward following misinformation exposure, showing that became less sure of the answer, while those in the I-I-I condition adjusted confidence ratings upward, showing that they became even surer of the answer. Finally, when considering the shifts in confidence from the start of the experiment (Test 1) to the conclusion of the experiment (Test 3), there was a greater increase in confidence for neutral items as opposed to misleading items for participants in the I-C-I condition, but for participants in the I-I-I condition, the opposite was true; there was a greater increase in confidence for items presented with a misleading frame as compared to a neutral frame. Coupled with the test performance findings which illustrate that those in the I-I-I condition also responded with higher levels of misinformation, this finding suggests that participants who worked individually may have been more likely to view the misleading fact references as correct feedback.

Largely consistent with Experiment 1, findings from Experiment 2 suggest that relative to repeated individual testing, collaboration reduced but did not completely eliminate the negative impact of exposure to misleading information in fictional stories. Direct comparisons across experiments revealed that participants in the I-I-I conditions performed similarly, and the effects of misinformation exposure did not vary depending on whether participants took one or two individual tests prior to the story reading. Rates of correct responding were similar across experiments for participants in the I-C-I condition on Tests 2 and 3. However, participants in Experiment 1 who were exposed to misinformation prior to collaboration were more likely to produce misinformation for the misleading items on Test 2, and participants who were exposed to misinformation following collaboration were more likely to respond with misinformation on Test 3.

IV. General Discussion

In the present experiments, we explored the malleability of general knowledge as a result of our two key experimental manipulations: collaboration and exposure to misinformation via fictional story reading. With the present experiments we sought to build upon and extend our previous findings which emphasized the benefits of collaboration for knowledge retrieval (e.g. error pruning, relearning through retrieval) leading to improved accuracy and confidence. In our previous experiments, we did not manipulate participants' prior knowledge in any way. But in contrast, by including the misinformation manipulation in the present experiments—before collaborative discussion in Experiment 1 and after collaborative discussion in Experiment 2—we were able to investigate whether the previously observed benefits of collaboration would again emerge in this new context where misinformation was systematically introduced, and thus counteract the powerfully disruptive effects of exposure to misinformation.

Importantly, across both experiments we found that although participants who worked collaboratively were not completely resistant to the negative effects of reading misinformation, relative to repeated individual test taking, collaborative discussion was associated with several key benefits. First, in both experiments collaboration increased correct responding and led to greater increases in confidence in the accuracy of one's responses throughout the experiment. Furthermore, collaborative discussion eliminated the reduction in consistent correct responses following misinformation exposure that was observed for participants who repeatedly worked individually, suggesting that correct knowledge was more likely to persist for participants who worked collaboratively. Finally, as observed in Experiment 2, participants were less likely to produce misinformation on the general knowledge test if they had previously collaborated, which shows that collaboration prior to misinformation exposure can make participants more resilient to making these types of errors. This benefit is particularly striking, as it is evidence of the protective benefits of collaboration, a benefit that to date has not been explored within the collaborative memory literature. In sum, as alluded to earlier, the emergence of collaborative benefits in spite of the impoverished retrieval context represents a critical new finding that builds upon our previous work investigating the effects of collaboration on knowledge retrieval.

The observation of these new collaborative benefits are promising, and as a follow-up, future experiments should focus on directly investigating the mechanisms that underlie the benefits of collaboration observed in the present studies. For example, by incorporating an error monitoring task into the story reading phase following collaboration (Fazio & Marsh, 2008; Marsh & Fazio, 2006), it would be possible to directly test whether participants who engaged in collaborative discussion more successfully identify errors as compared to others who completed the previous general knowledge test individually.

Another aim of the present experiments was to assess how collaboration affects the transmission of misinformation, and consistent with previous findings (Fazio, et al., 2013; Marsh & Fazio, 2006; Marsh, et al., 2003), the present experiments further demonstrate the disruptive effects of misinformation exposure on general knowledge retrieval. After reading fictional stories, participants were less likely to respond correctly to critical questions that were referenced with a misleading frame (relative to a neutral frame) in the stories and more likely to produce misinformation as answers in response to these questions. Within the context of this replication, a novel finding from the present experiments was that these negative effects were observed regardless of whether participants had the opportunity to engage in collaborative discussion following misinformation exposure (Experiment 1), or prior to misinformation exposure (Experiment 2). The effects of exposure to misinformation have been shown to be quite robust and difficult to completely eliminate (Eslick, et al., 2011; Marsh & Fazio, 2006) even in situations where participants are encouraged to directly draw upon correct information that invalidates incorrect information read in stories (Rapp, Hinze, Kohlhepp, & Ryskin, 2014). Thus, with these findings in mind, it is not surprising that collaborative discussion did not eliminate the negative impact of exposure to misinformation.

Although our central aim concerned the effects of collaborative retrieval, cross-experimental comparisons of individual performance on Test 2 in Experiment 1 and Test 3 in Experiment 2 allowed us to test another potential source of inoculation against misinformation—repeated individual testing. That is, in Experiment 1 participants completed only one test prior to reading the fictional stories and in Experiment 2 participants completed two tests, and by comparing misinformation production across experiments, we were able to investigate whether repeatedly retrieving answers across two general knowledge tests might reduce the negative impact of misinformation exposure, relative to participants who were only tested once before. This turned out not to be the case, as no differences were observed in correct or misinformation responses regardless of whether participants completed one or two general knowledge tests prior to reading the stories. The finding that participants in the I-I-I conditions across both experiments were equally susceptible to misinformation extends findings from a recent experiment that demonstrated prior study but not prior retrieval was effective in reducing the negative impact of misinformation exposure (Mullet, Umanath, & Marsh, 2014).

Cross-experimental comparisons also allowed us to test whether misinformation exposure prior to collaboration would alter the effects of collaborative discussion. As previously discussed, a key benefit of collaboration that has been repeatedly observed within the context of knowledge retrieval is the increase in correct responses, and interestingly, correct response rates were similar across experiments for participants in the I-C-I condition, suggesting that this benefit of collaboration was not reduced by misinformation exposure. This finding once again highlights the error pruning mechanisms of collaborative discussion; regardless of whether participants were exposed to misinformation prior to collaboration or following collaboration, rates of correct responding were unaffected.

Differences were observed across experiments with respect to misinformation production. Participants in Experiment 1 who were exposed to misinformation prior to collaboration were more likely to produce misinformation for the misleading items on Test 2, which makes sense given that participants in Experiment 2 had not yet read the stories and would have no reason to produce the misinformation at this stage. On Test 3, comparisons revealed that misinformation production was similar across experiments, that is, participants in Experiment 2 who were exposed to misinformation following collaboration were just as likely to respond with misinformation on Test 3 as compared to their peers who had been exposed to the misinformation earlier in the experiment.

The findings from the present experiments have important implications for education, given that group work is widely advocated for in nearly all educational settings (Johnson & Johnson, 2009), and fictional materials are commonly used to supplement the teaching of factual information (Marsh, et al., 2012). The present studies demonstrate both that collaborative retrieval is more beneficial than individual retrieval when dealing with exposure to misinformation, but also that students can still be susceptible to the effects of misinformation exposure when given the opportunity to discuss the answers as a group. To ensure that fictional materials and group work serve to enhance student education rather than contribute to its detriment, when implementing these teaching methods in their classrooms, educators should be aware ways to counteract the acceptance of inaccurate information. While simply warning participants that they may encounter inaccuracies has not been shown to be effective in reducing the negative impact of misinformation (Marsh & Fazio, 2006), providing much more explicit warnings about misinformation as well as feedback about correct information prior to exposure has been shown to be beneficial in reducing the negative misinformation (Butler, et al., 2009; Umanath, Butler, & Marsh, 2012). This is consistent with one of the benefits of collaboration that we discussed earlier, where in Experiment 2 there was a reduction in misinformation production when participants had engaged in collaborative discussion prior to reading the stories. Here, the collaborative discussion prior to story reading served as a source of feedback. Taken together, these findings suggests that it may be beneficial for teachers to have students discuss material with peers prior to reading a work of historical fiction or viewing a historical reenactment on film and then provide them with accurate feedback following this discussion in order to maximize their abilities to reject inaccuracies they encounter from these fictional sources. When employed in this way, educators can take advantage of the benefits associated with using fictional materials in the classroom, as well as the benefits of group work (Roseth, Johnson, & Johnson, 2008), while also reducing the costs of exposing students to misinformation.

On a daily basis we are all consumers of information from a plethora of sources, some more accurate than others (Lewandowsky, et al., 2012). As such, it is important to be aware of how being exposed to these inaccuracies, and discussing them with others, can affect what we know, or rather, what we think we know to be true. The present experiments illustrate some key benefits of discussing general knowledge with others, but at the same time, also demonstrate that group discussion is not a panacea. In conclusion, discussing information with groups can lead to improved accuracy in retrieval of general knowledge, but we should still be wary of the powerful effects of reading misleading information, even in the context of discussion with others.

V. Tables and Figures

Table 1. Assessment of Data Dependencies: Summary of Intraclass Correlations and Design Effect Values For Test Performance Variables in Experiment 1

Variables	Test 2 ICC[a]	Test 3 ICC	Test 2 DE[b]	Test 3 DE
Proportion Correct	.88***	.58***	2.75	2.16
Proportion Omitted	.85***	.47**	2.70	1.94
Proportion Error	.82***	.64***	2.64	2.28
Proportion Misinformation Error	.98***	.68***	2.95	2.37
Proportion Other Error	.59***	.60***	2.17	2.21

***$p < .001$, **$p < .01$; [a] Intraclass Correlation Coefficient; [b] Design Effect

Table 2. Performance on Test 1 as a Function of Experimental Condition

	I-C-I		I-I-I	
	M	*SE*	*M*	*SE*
Overall, Across All Items				
Correct Responses	0.40	0.02	0.40	0.02
Omissions	0.46	0.02	0.45	0.02
Errors	0.14	0.01	0.15	0.01
Misinformation Errors	0.03	0.002	0.03	0.003
Other Errors	0.11	0.01	0.12	0.01
Misleading Items				
Correct Responses	0.50	0.02	0.49	0.02
Omissions	0.36	0.02	0.37	0.02
Errors	0.14	0.01	0.14	0.01
Misinformation Errors	0.05	0.01	0.05	0.01
Other Errors	0.09	0.01	0.09	0.01
Neutral Items				
Correct Responses	0.51	0.03	0.51	0.02
Omissions	0.37	0.02	0.36	0.02
Errors	0.13	0.01	0.13	0.01
Misinformation Errors	0.04	0.01	0.04	0.01
Other Errors	0.08	0.01	0.09	0.01
Filler Items				
Correct Responses	0.28	0.02	0.30	0.02
Omissions	0.57	0.02	0.54	0.02
Errors	0.15	0.01	0.16	0.01
Misinformation Errors	--	--	--	--
Other Errors	--	--	--	--

Table 3. Effects of Fact Framing and Collaboration on the Proportion of Correct Responses on Test 2: Summary of Regression Coefficients

	Regression Coefficients (*b*)	Standard Error (*SE*)	95% Confidence Interval
Intercept [a]	.51***	.03	.46, .56
Fact Framing [b]	-.07***	.02	-.10, -.03
Condition [c]	.23***	.05	.13, .32
Fact Framing × Condition	.03	.03	-.03, .10

[a] Reflects the average proportion correct for individual participants on neutral items.

[b] Reflects the reduction in proportion correct for misleading items

[c] Reflects the increase in proportion correct when participants collaborate

*** $p < .001$

Table 4. Model Fitting Summary for the Effects of Fact Framing and Collaboration on the Proportion of Correct Responses on Test 3

	Model 1: Basic Model	Model 2: Random Intercepts	Model 3: Random Slopes and Intercepts
Model Summary			
Deviance statistic (-2 Log Likelihood)	-212.02	-227.05	-252.04
Number of estimated parameters	6	7	8
Likelihood Ratio Chi Square Tests	Difference in Deviance Values		
Model 1 vs. Model 2	15.03***		
Model 2 vs. Model 3	24.99***		

*** $\chi^2(1, N = 96), p < .001$

Table 5. Model Fitting Summary for the Effects of Test and Fact Framing on the Proportion of Correct Responses across Tests 2 and 3

	Model 1: Basic Model	Model 2: Random Intercepts	Model 3: Random Slopes and Intercepts
Model Summary			
Deviance statistic (-2 Log Likelihood)	-328.3	-358.04	-402.79
Number of estimated parameters	9	10	11
Likelihood Ratio Chi Square Tests	Difference in Deviance Values		
Model 1 vs. Model 2	29.74***		
Model 2 vs. Model 3	44.75***		

*** $\chi^2(1, N = 96)$, $p < .001$

Table 6. Effects of Fact Framing and Collaboration on the Proportion of Correct Responses on Test 3: Summary of Regression Coefficients

	Regression Coefficients (*b*)	Standard Error (*SE*)	95% Confidence Interval
Intercept [a]	.52***	.03	.47, .57
Fact Framing [b]	-.07***	.02	-.10, -.03
Condition [c]	.16***	.05	.07, .26
Fact Framing × Condition	.03	.03	-.03, .09

[a] Reflects the average proportion correct for individual participants on neutral items.

[b] Reflects the reduction in proportion correct for misleading items

[c] Reflects the increase in proportion correct when participants collaborate

*** $p < .001$

Table 7. Assessment of Data Dependencies: Summary of Intraclass Correlations and Design Effect Values For Test Performance Variables in Experiment 2

Variables	Test 2 ICC[a]	Test 3 ICC	Test 2 DE[b]	Test 3 DE
Proportion Correct	.95***	.41**	2.90	1.81
Proportion Omitted	.96***	.36**	2.92	1.72
Proportion Error	.76***	.38**	2.52	1.75
Proportion Misinformation Error	.54***	.56***	2.09	2.12
Proportion Other Error	.79***	.48**	2.58	1.95

***$p < .001$, **$p < .01$; [a] Intraclass Correlation Coefficient; [b] Design Effect

Table 8. Performance on Test 1 as a Function of Experimental Condition

	I-C-I		I-I-I	
	M	*SE*	*M*	*SE*
Overall, Across All Items				
Correct Responses	0.37	0.02	0.38	0.02
Omissions	0.47	0.02	0.46	0.02
Errors	0.16	0.01	0.17	0.01
Misinformation Errors	0.03	0.002	0.03	0.002
Other Errors	0.14	0.01	0.14	0.01
Misleading Items				
Correct Responses	0.47	0.02	0.48	0.02
Omissions	0.38	0.02	0.36	0.02
Errors	0.15	0.01	0.15	0.02
Misinformation Errors	0.04	0.004	0.05	0.005
Other Errors	0.11	0.01	0.10	0.01
Neutral Items				
Correct Responses	0.47	0.02	0.48	0.02
Omissions	0.38	0.02	0.37	0.02
Errors	0.15	0.01	0.15	0.01
Misinformation Errors	0.05	0.01	0.04	0.00
Other Errors	0.10	0.01	0.11	0.01
Filler Items				
Correct Responses	0.25	0.02	0.25	0.02
Omissions	0.57	0.02	0.56	0.03
Errors	0.18	0.01	0.18	0.02
Misinformation Errors	--	--	--	--
Other Errors	--	--	--	--

Table 9. Model Fitting Summary for the Effects of Fact Framing and Collaboration on the Proportion of Correct Responses on Test 3

	Model 1: Basic Model	Model 2: Random Intercepts	Model 3: Random Slopes and Intercepts
Model Summary			
Deviance statistic (-2 Log Likelihood)	-352.52	-361.05	-365.69
Number of estimated parameters	6	7	8
Likelihood Ratio Chi Square Tests	Difference in Deviance Values		
Model 1 vs. Model 2	8.53**		
Model 2 vs. Model 3	4.64*		

** $\chi^2(1, N = 96)$, $p < .01$
* $\chi^2(1, N = 96)$, $p < .05$

Table 10. Effects of Fact Framing and Collaboration on the Proportion of Correct Responses on Test 3: Summary of Regression Coefficients

	Regression Coefficients (*b*)	Standard Error (*SE*)	95% Confidence Interval
Intercept [a]	.50***	.03	.46, .54
Fact Framing [b]	-.12***	.01	-.15, -.11
Condition [c]	.16***	.04	.08, .23
Fact Framing × Condition	.03	.02	-.03, .04

[a] Reflects the average proportion correct for individual participants on neutral items.

[b] Reflects the reduction in proportion correct for misleading items

[c] Reflects the increase in proportion correct when participants collaborate

*** p <.001

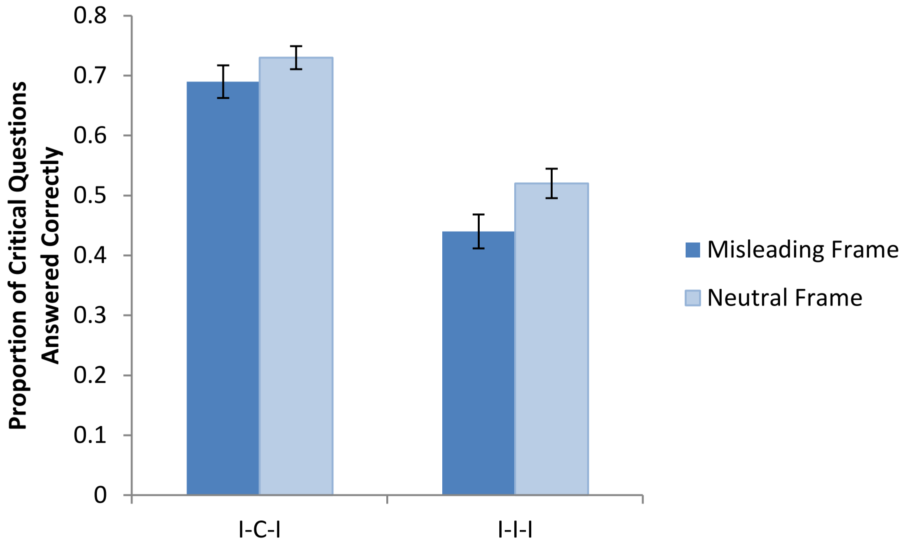

Figure 1. Proportion of critical questions answered correctly on Test 2 as a function of fact framing and experimental condition. There was a reduction in correct responses when questions were referenced with a misleading frame in the story relative to a neutral frame, and participants who worked collaboratively (I-C-I) answered more questions correctly compared to participants who worked individually (I-I-I).

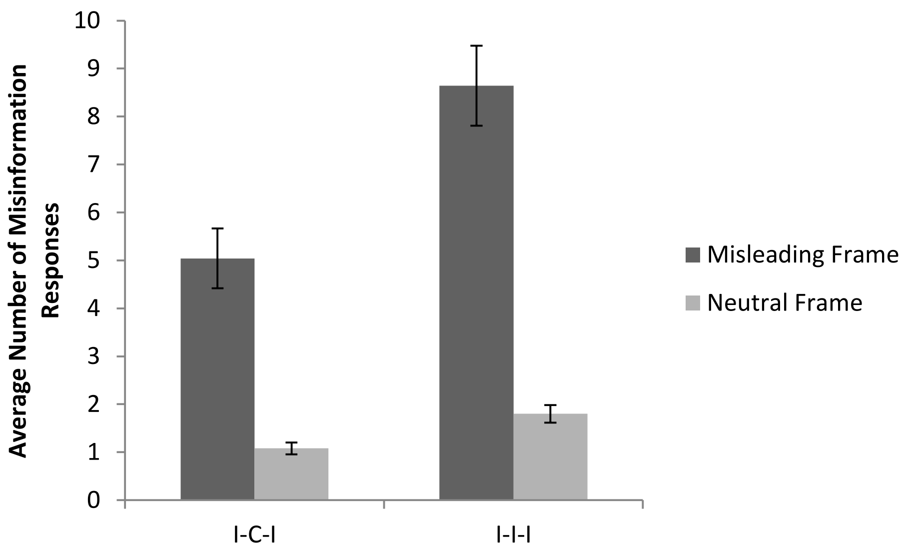

Figure 2. Average number of critical questions answered with misinformation on Test 2 as a function of fact framing and experimental condition. When questions were referenced with misleading frames in the stories relative to neutral frames, participants were more likely to respond with misinformation.

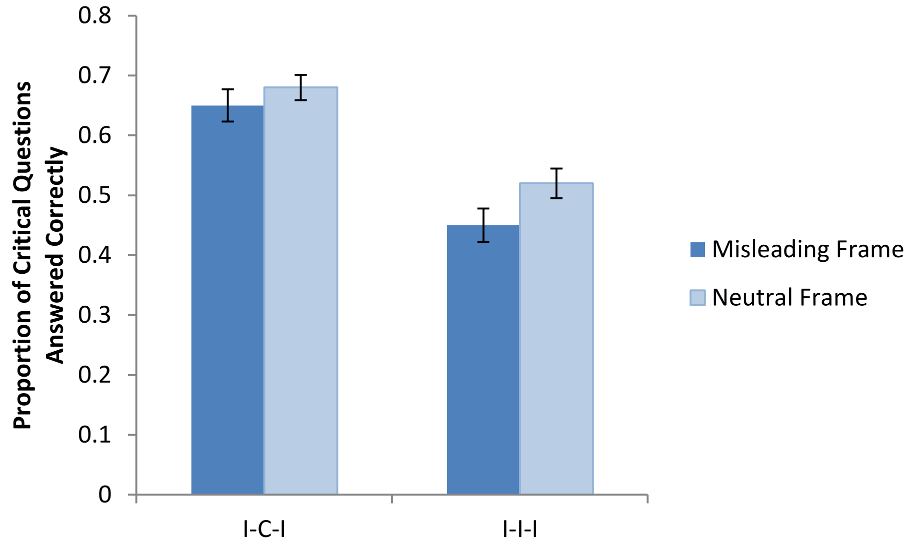

Figure 3. Proportion of critical questions answered correctly on Test 3 as a function of fact framing and experimental condition. Consistent with the pattern of results observed on Test 2, participants who previously worked collaboratively still outperformed those who repeatedly worked individually, and overall, questions referenced with a misleading frame in the story were less likely to be answered correctly.

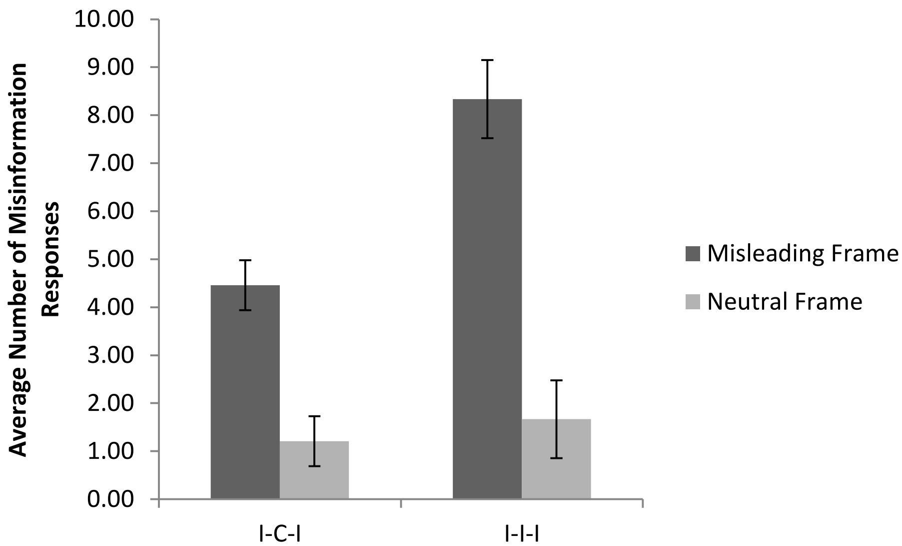

Figure 4. Average number of critical questions answered with misinformation on Test 3 as a function of fact framing and experimental condition. The negative impact of misinformation exposure persisted on the final individual test of the experiment; consistent with the results from Test 2, when questions were referenced with misleading frames in the stories relative to neutral frames, participants were more likely to respond with misinformation.

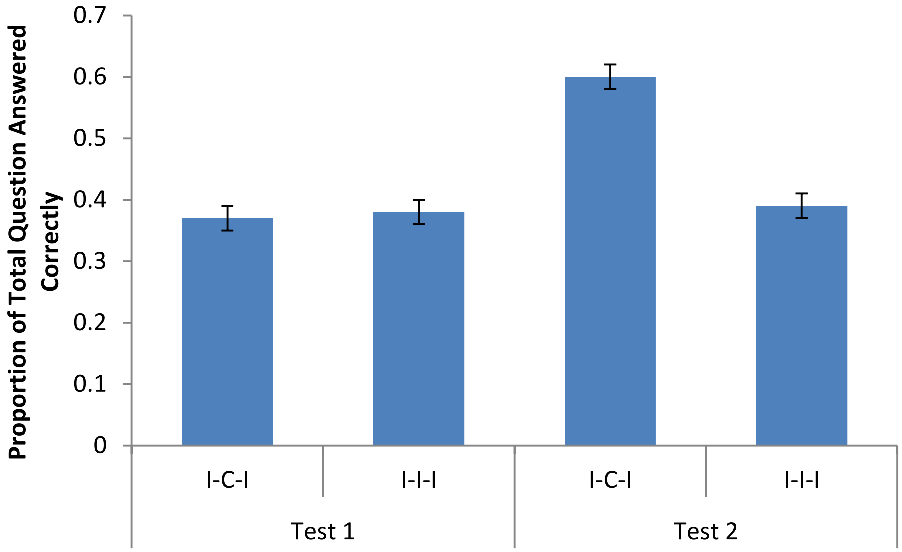

Figure 5. Proportion of test questions answered correctly on Tests 1 and 2 as a function of experimental condition. On Test 1, all participants worked individually, and regardless of condition assignment, all participants performed similarly. On Test 2, participants who engaged in collaborative discussion (I-C-I) answered more questions correctly than those who completed the test on their own (I-I-I).

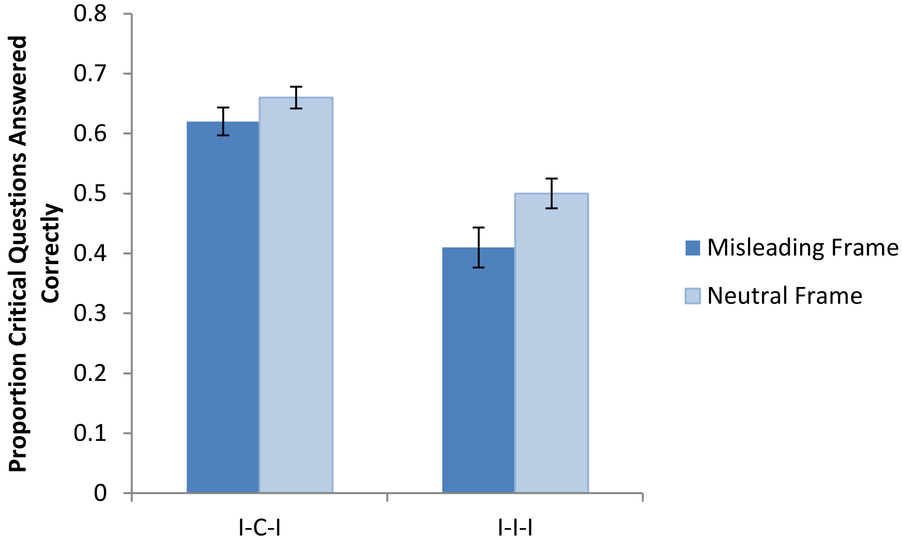

Figure 6. Proportion of critical questions answered correctly on Test 3 as a function of fact framing and experimental condition. Participants were less likely to answer questions correctly when it was referenced with a misleading frame relative to a neutral frame, and participants were more likely to answer questions correctly if they worked collaboratively (I-C-I).

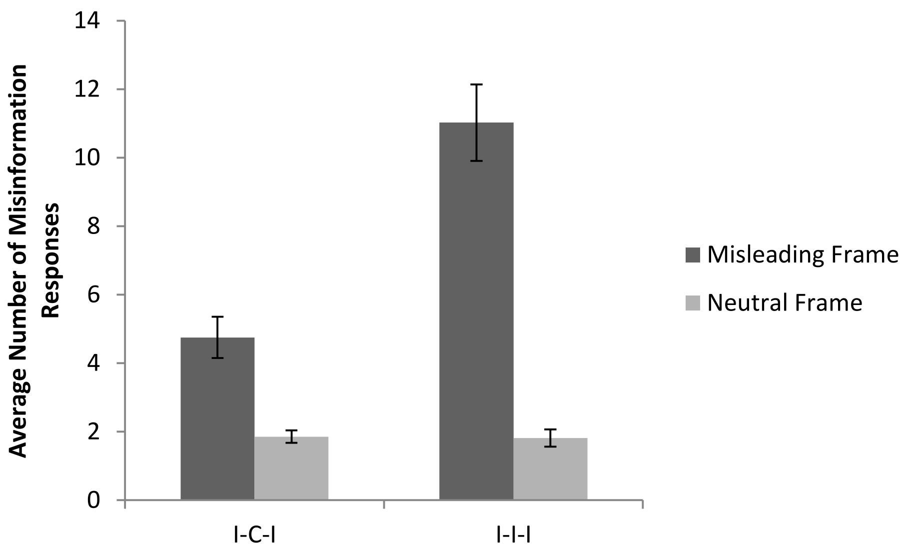

Figure 7. Average number of critical questions answered with misinformation on Test 3 as a function of fact framing and experimental condition. Participants were more likely to produce misinformation in response to questions that were referenced with a misleading frame in the story relative to a neutral frame, and this discrepancy was greater for those participants who completed repeated individual tests (I-I-I) as compared to those who previously worked collaboratively (I-C-I).

VI. References

Agresti, A. (2002). *Categorical data analysis*: John Wiley & Sons.

Andersson, J., Helstrup, T., & Rönnberg, J. (2007). Collaborative memory: How is our ability to remember affected by interaction with others? In S. Magnussen & T. Helstrup (Eds.), *Everyday memory* (pp. 135). Hove and New York: Psychology Press.

Andersson, J., & Rönnberg, J. (1996). Collaboration and memory: Effects of dyadic retrieval on different memory tasks. *Applied Cognitive Psychology, 10*(2), 171-181.

Assefi, S. L., & Garry, M. (2003). Absolut® Memory Distortions Alcohol Placebos Influence the Misinformation Effect. *Psychological Science*, *14*(1), 77-80.

Baldwin, S. A., Bauer, D. J., Stice, E., & Rohde, P. (2011). Evaluating models for partially clustered designs. *Psychological methods, 16*(2), 149.

Barber, S. J., Harris, C. B., & Rajaram, S. (2014). Why Two Heads Apart Are Better Than Two Heads Together: Multiple Mechanisms Underlie the Collaborative Inhibition Effect in Memory.

Barber, S. J., Rajaram, S., & Aron, A. (2010). When two is too many: Collaborative encoding impairs memory. *Memory & Cognition, 38*(3), 255-264.

Barber, S. J., Rajaram, S., & Marsh, E. J. (2008). Fact learning: How information accuracy, delay, and repeated testing change retention and retrieval experience. *Memory*, *16*(8), 934-946.

Basden, B. H., Basden, D. R., Bryner, S., & Thomas, R. L. (1997). A comparison of group and individual remembering: Does collaboration disrupt retrieval strategies? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*(5), 1176.

Basden, B. H., Basden, D. R., & Henry, S. (2000). Costs and benefits of collaborative remembering. *Applied Cognitive Psychology, 14*(6), 497-507.

Basden, D. R., Basden, B. H., & Galloway, B. C. (1977). Inhibition with part-list cuing: Some tests of the item strength hypothesis. *Journal of experimental psychology: Human learning and memory, 3*(1), 100.

Bauer, D. J., Sterba, S. K., & Hallfors, D. D. (2008). Evaluating group-based interventions when control participants are ungrouped. *Multivariate Behavioral Research, 43*(2), 210-236.

Blumen, H. M., & Rajaram, S. (2008). Influence of re-exposure and retrieval disruption during group collaboration on later individual recall. *Memory, 16*(3), 231-244.

Bottoms, H. C., Eslick, A. N., & Marsh, E. J. (2010). Memory and the Moses illusion: Failures to detect contradictions with stored knowledge yield negative memorial consequences. *Memory, 18*(6), 670-678.

Butler, A. C., Zaromb, F. M., Lyle, K. B., & Roediger, H. L. (2009). Using popular films to enhance classroom learning the good, the bad, and the interesting. *Psychological Science, 20*(9), 1161-1168.

Butterfield, B., & Metcalfe, J. (2001). Errors committed with high confidence are hypercorrected. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*(6), 1491.

Butterfield, B., & Metcalfe, J. (2006). The correction of errors committed with high confidence. *Metacognition and Learning, 1*(1), 69-84.

Choi, H.-Y., Blumen, H. M., Congleton, A. R., & Rajaram, S. (2014). The role of group configuration in the social transmission of memory: Evidence from identical and reconfigured groups. *Journal of Cognitive Psychology, 26*(1), 65-80.

Clark, S. E., Hori, A., Putnam, A., & Martin, T. P. (2000). Group collaboration in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*(6), 1578.

Clifasefi, S. L., Garry, M., Harper, D. N., Sharman, S. J., & Sutherland, R. (2007). Psychotropic placebos create resistance to the misinformation effect. *Psychonomic Bulletin & Review, 14*(1), 112-117.

Coman, A., Manier, D., & Hirst, W. (2009). Forgetting the Unforgettable Through Conversation Socially Shared Retrieval-Induced Forgetting of September 11 Memories. *Psychological Science, 20*(5), 627-633.

Congleton, A. R., & Rajaram, S. (2011). The influence of learning methods on collaboration: Prior repeated retrieval enhances retrieval organization, abolishes collaborative inhibition, and promotes post-collaborative memory. *Journal of Experimental Psychology: General*, *140*(4), 535.

Congleton, A. R., & Rajaram, S. (2014). Collaboration Changes Both the Content and the Structure of Memory: Building the Architecture of Shared Representations.

Coxe, S., West, S. G., & Aiken, L. S. (2009). The analysis of count data: A gentle introduction to Poisson regression and its alternatives. *Journal of personality assessment, 91*(2), 121-136.

Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of experimental psychology, 58*(1), 17.

DePrince, A. P., Allard, C. B., Oh, H., & Freyd, J. J. (2004). What's in a name for memory errors? Implications and ethical issues arising from the use of the term" false memory" for errors in memory for details. *Ethics & Behavior, 14*(3), 201-233.

Eakin, D. K., Schreiber, T. A., & Sergent-Marshall, S. (2003). Misinformation effects in eyewitness memory: the presence and absence of memory impairment as a function of warning and misinformation accessibility. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*(5), 813.

Eslick, A. N., Fazio, L. K., & Marsh, E. J. (2011). Ironic effects of drawing attention to story errors. *Memory, 19*(2), 184-191. doi: 10.1080/09658211.2010.543908

Fazio, L. K., Barber, S. J., Rajaram, S., Ornstein, P. A., & Marsh, E. J. (2013). Creating illusions of knowledge: Learning errors that contradict prior knowledge. *Journal of Experimental Psychology: General, 142*(1), 1.

Fazio, L. K., & Marsh, E. J. (2008). Slowing presentation speed increases illusions of knowledge. *Psychonomic Bulletin & Review, 15*(1), 180-185.

Finlay, F., Hitch, G. J., & Meudell, P. R. (2000). Mutual inhibition in collaborative recall: evidence for a retrieval-based account. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*(6), 1556.

Frenda, S. J., Nichols, R. M., & Loftus, E. F. (2011). Current issues and advances in misinformation research. *Current directions in Psychological science, 20*(1), 20-23.

Gallo, D. A. (2010). False memories and fantastic beliefs: 15 years of the DRM illusion. *Memory & Cognition, 38*(7), 833-848.

Gallo, D. A., Roediger, H. L., & McDermott, K. B. (2001). Associative false recognition occurs without strategic criterion shifts. *Psychonomic Bulletin & Review, 8*(3), 579-586.

Gerrig, R. J., & Prentice, D. A. (1991). The representation of fictional information. *Psychological Science, 2*(5), 336-340.

Hanley, J. A., Negassa, A., & Forrester, J. E. (2003). Statistical analysis of correlated data using generalized estimating equations: an orientation. *American journal of epidemiology, 157*(4), 364-375.

Harlow, L. L. (2014). The Essence of Multivariate Thinking: Basic Themes and Methods (Second ed.). New York: Routledge.

Harris, C. B., Barnier, A. J., & Sutton, J. (2012). Consensus collaboration enhances group and individual recall accuracy. *Quarterly Journal of Experimental Psychology, 65*(1), 179-194. doi: 10.1080/17470218.2011.608590

Harris, C. B., Keil, P. G., Sutton, J., Barnier, A. J., & McIlwain, D. J. (2011). We remember, we forget: Collaborative remembering in older couples. *Discourse Processes, 48*(4), 267-303.

Hirst, W., & Echterhoff, G. (2012). Remembering in conversations: The social sharing and reshaping of memories. *Psychology, 63*(1), 55.

Hirst, W., & Manier, D. (2008). Towards a psychology of collective memory. *Memory, 16*(3), 183-200.

Hyman, I. E., Jr., Cardwell, B. A., & Roy, R. A. (2013). Multiple causes of collaborative inhibition in memory for categorised word lists. *Memory, 21*(7), 875-890.

Hyman, I. E. J., & Loftus, E. F. (1998). Errors in autobiographical memory. *Clinical psychology review, 18*(8), 933-947.

Johansson, N., Andersson, J., & Rönnberg, J. (2005). Compensating strategies in collaborative remembering in very old couples. *Scandinavian Journal of Psychology, 46*(4), 349-359.

Johnson, D. W., & Johnson, R. T. (2009). An Educational Psychology Success Story: Social Interdependence Theory and Cooperative Learning. *Educational Researcher, 38*(5), 365-379. doi: 10.3102/0013189x09339057

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

Kenny, D. A., & La Voie, L. (1985). Separating individual and group effects. *Journal of Personality and Social Psychology, 48*(2), 339.

Kenny, D. A., Mannetti, L., Pierro, A., Livi, S., & Kashy, D. A. (2002). The statistical analysis of data from small groups. *Journal of Personality and Social Psychology, 83*(1), 126.

Lewandowsky, S., Ecker, U. K., 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.

Lewis, C. H., & Anderson, J. R. (1976). Interference with real world knowledge. *Cognitive Psychology, 8*(3), 311-335.

Loftus, E. F. (1975). Leading questions and the eyewitness report. *Cognitive Psychology, 7*, 560-572.

Loftus, E. F. (1979). The malleability of human memory: Information introduced after we view an incident can transform memory. *American Scientist*, 312-320.

Loftus, E. F. (2005). Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning & Memory, 12*(4), 361-366.

Loftus, E. F., Miller, D. G., & Burns, H. J. (1978). Semantic integration of verbal information into a visual memory. *Journal of experimental psychology: Human learning and memory, 4*(1), 19.

MacLeod, M. D., & Saunders, J. (2005). The role of inhibitory control in the production of misinformation effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*(5), 964.

Marsh, E. J. (2004). Story stimuli for creating false beliefs about the world. *Behavior Research Methods, Instruments, & Computers, 36*(4), 650-655.

Marsh, E. J., Butler, A. C., & Umanath, S. (2012). Using fictional sources in the classroom: Applications from cognitive psychology. *Educational Psychology Review, 24*(3), 449-469.

Marsh, E. J., & Fazio, L. K. (2006). Learning errors from fiction: Difficulties in reducing reliance on fictional stories. *Memory & Cognition, 34*(5), 1140-1149.

Marsh, E. J., Meade, M. L., & Roediger, H. L., III. (2003). Learning facts from fiction. *Journal of Memory and Language, 49*(4), 519-536. doi: 10.1016/s0749-596x(03)00092-5

McCabe, D. P., & Smith, A. D. (2002). The effect of warnings on false memories in young and older adults. *Memory & Cognition, 30*(7), 1065-1077.

Meade, M. L., Nokes, T. J., & Morrow, D. G. (2009). Expertise promotes facilitation on a collaborative memory task. *Memory, 17*(1), 39-48.

Meade, M. L., & Roediger, H. L. (2002). Explorations in the social contagion of memory. *Memory & Cognition, 30*(7), 995-1009.

Michael, R. B., Garry, M., & Kirsch, I. (2012). Suggestion, cognition, and behavior. *Current directions in Psychological science, 21*(3), 151-156.

Mullet, H. G., Umanath, S., & Marsh, E. J. (2014). Recent study, but not retrieval, of knowledge protects against learning errors. *Memory & Cognition*, 1-11.

Multhaup, K. S., & Conner, C. A. (2002). The effects of considering nonlist sources on the Deese–Roediger–McDermott memory illusion. *Journal of Memory and Language*, *47*(2), 214-228.

Muthén, B. O. (1994). Multilevel Covariance Structure Analysis. *Sociological Methods & Research, 22*(3), 376-398.

Myers, J. L. (1972). *Fundamentals of experimental design*. Boston: Allyn and Bacon, Inc.

Nelson, T. O., & Narens, L. (1980). Norms of 300 general-information questions: Accuracy of recall, latency of recall, and feeling-of-knowing ratings. *Journal of Verbal Learning and Verbal Behavior*, *19*(3), 338-368. doi: http://dx.doi.org/10.1016/S0022-5371(80)90266-2

Nickerson, R. S. (1984). Retrieval inhibition from part-set cuing: A persisting enigma in memory research. *Memory & Cognition, 12*(6), 531-552.

Parker, S., Garry, M., Engle, R. W., Harper, D. N., & Clifasefi, S. L. (2008). Psychotropic placebos reduce the misinformation effect by increasing monitoring at test. *Memory, 16*(4), 410-419.

Peugh, J. L. (2010). A practical guide to multilevel modeling. *Journal of School Psychology, 48*(1), 85-112. doi: http://dx.doi.org/10.1016/j.jsp.2009.09.002

Pezdek, K., & Lam, S. (2007). What research paradigms have cognitive psychologists used to study “false memory,” and what are the implications of these choices? *Consciousness and Cognition, 16*(1), 2-17.

Pociask, S., Rajaram, S., & Marsh, E. J. (2014). Collaboration in an Educational Context: Malleability of Knowledge as a Result of Discussion. *Manuscript in preparation*.

Rajaram, S. (2011). Collaboration Both Hurts and Helps Memory A Cognitive Perspective. *Current directions in Psychological science, 20*(2), 76-81.

Rajaram, S., & Pereira-Pasarin, L. P. (2010). Collaborative memory: Cognitive research and theory. *Perspectives on Psychological Science*, *5*(6), 649-663.

Rapp, D. N., Hinze, S. R., Kohlhepp, K., & Ryskin, R. A. (2014). Reducing reliance on inaccurate information. *Memory & Cognition*, *42*(1), 11-26.

Roediger, H., III, McDermott, K., Pisoni, D., & Gallo, D. (2004). Illusory recollection of voices. *Memory*, *12*(5), 586-602.

Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*(4), 803.

Roediger, H. L., & McDermott, K. B. (2000). Distortions of memory. *The Oxford handbook of memory*, 149-162.

Roediger, H. L., & Neely, J. H. (1982). Retrieval blocks in episodic and semantic memory. *Canadian Journal of Psychology/Revue canadienne de psychologie, 36*(2), 213.

Roseth, C. J., Johnson, D. W., & Johnson, R. T. (2008). Promoting early adolescents' achievement and peer relationships: The effects of cooperative, competitive, and individualistic goal structures. *Psychological Bulletin*, *134*(2), 223-246. doi: 10.1037/0033-2909.134.2.223

Ross, M., Spencer, S. J., Blatz, C. W., & Restorick, E. (2008). Collaboration reduces the frequency of false memories in older and younger adults. *Psychology and Aging, 23*(1), 85.

Schacter, D. L. (1999). The cognitive neuropsychology of false memories: Introduction. *Cognitive neuropsychology, 16*(3-5), 193-195.

Szary, J., & Dale, R. (2013). *Dyadic cooperation enhances retrieval and recall of crossword solutions.* Paper presented at the Thirty Fifth Annual Meeting of the Cognitive Science Society, Austin, TX.

Szary, J., & Dale, R. (2014). *Familiarity Modulates the Dynamics of Collaborative Inhibition in a Trivia Game*. Paper presented at the Thirty Sixth Annual Meeting of the Cognitive Science Society, Austin, TX.

Tabachnick, B. G., & Fidell, L. S. (2007). *Using multivariate statistics* (Sixth ed.). Boston: Pearson/Allyn & Bacon.

Takahashi, M. (2007). Does collaborative remembering reduce false memories? *British Journal of Psychology, 98*(1), 1-13. doi: 10.1348/000712606x101628

Tauber, S. K., Dunlosky, J., Rawson, K. A., Rhodes, M. G., & Sitzman, D. M. (2013). General knowledge norms: Updated and expanded from the Nelson and Narens (1980) norms. *Behavior research methods, 45*(4), 1115-1143.

Thorley, C., & Dewhurst, S. A. (2007). Collaborative false recall in the DRM procedure: Effects of group size and group pressure. *European Journal of Cognitive Psychology, 19*(6), 867-881. doi: 10.1080/09541440600872068

Thorley, C., & Dewhurst, S. A. (2009). False and veridical collaborative recognition. *Memory, 17*(1), 17-25.

Tousignant, J. P., Hall, D., & Loftus, E. F. (1986). Discrepancy detection and vulnerability to misleading postevent information. *Memory & Cognition, 14*(4), 329-338.

Umanath, S., Butler, A. C., & Marsh, E. J. (2012). Positive and negative effects of monitoring popular films for historical inaccuracies. *Applied Cognitive Psychology, 26*(4), 556-567.

Wade, K. A., Sharman, S. J., Garry, M., Memon, A., Mazzoni, G., Merckelbach, H., & Loftus, E. F. (2007). False claims about false memory research. *Consciousness and Cognition, 16*(1), 18-28. doi: http://dx.doi.org/10.1016/j.concog.2006.07.001

Weigold, A., Russell, E. J., & Natera, S. N. (2014). Correction of False Memory for Associated Word Lists by Collaborating Groups. *The American journal of psychology, 127*(2), 183-190. doi: 10.5406/amerjpsyc.127.2.0183

Weldon, M. S., & Bellinger, K. D. (1997). Collective memory: collaborative and individual processes in remembering. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*(5), 1160.

West, B. T., Welch, K. B., & Galecki, A. T. (2015). *Linear Mixed Models: A Practical Guide Using Statistical Software* (Second ed.). Boca Raton: Taylor and Francis Group, LLC.

Westerberg, C. E., & Marsolek, C. J. (2006). Do instructional warnings reduce false recognition? *Applied Cognitive Psychology, 20*(1), 97-114.

Zaragoza, M. S., Belli, R. F., & Payment, K. E. (2006). Misinformation effects and the suggestibility of eyewitness memory. *Do justice and let the sky fall: Elizabeth F. Loftus and her contributions to science, law, and academic freedom*, 35-63.

Zaragoza, M. S., McCloskey, M., & Jamis, M. (1987). Misleading postevent information and recall of the original event: Further evidence against the memory impairment hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*(1), 36.

Zeger, S. L., Liang, K.-Y., & Albert, P. S. (1988). Models for longitudinal data: a generalized estimating equation approach. *Biometrics*, 1049-1060.

Zhu, B., Chen, C., Loftus, E. F., Lin, C., He, Q., Chen, C., . . . Dong, Q. (2010). Individual differences in false memory from misinformation: Cognitive factors. *Memory*, *18*(5), 543-555.