

The Implications of Corrections: Then Why Did You Mention It?

Julie G. Bush, Hollyn M. Johnson, and Colleen M. Seifert

University of Michigan

330 Packard Rd.

Ann Arbor, Michigan 48104-2994

jgbush@umich.edu, hollyn.johnson@um.cc.umich.edu, seifert@umich.edu

Abstract

How can misreported information be effectively corrected? Wilkes and Leatherbarrow (1988) found that people relied upon invalidated information to answer questions despite their awareness of its inaccuracy, a phenomenon called the "continued influence effect" (Johnson & Seifert, in press). But corrections in which an assertion is made and then denied (e.g., "X is true ... actually, X is untrue") may violate important conversational assumptions. Grice (1967/1989) and others have argued that people expect speakers to offer only information that is both truthful and conversationally relevant; thus, people may seek interpretations for corrections that will incorporate both the literal meaning and the conversational implications of the contradictory statements. Our hypothesis was that corrections would be more successful when they *explained why* the original information was asserted. An empirical study showed that corrections that accounted for conversational implications (e.g., "X, which had originally been believed because of Y, is actually untrue") could more effectively reduce the continued use of discredited information. Additionally, the results show that reiterating the literal content of a correction may actually be perceived as implying that the correction statement should be disbelieved. Since the conversational implications of corrections critically shape comprehension, their examination is crucial in domains (such as courtrooms, newspapers, and classrooms) where informational updates frequently occur.

The Implications of Corrections

Successful comprehension involves more than interpreting the literal meanings of words we speak and hear; in addition, it involves conversational assumptions that help us assign meaning based on the context surrounding the occurrence of an utterance. For example, speakers are assumed to offer only information that they believe is true, and that they intend to be informative and conversationally relevant to their audiences (Grice, 1967/1989). Given these conventions, the problem of correcting prior information becomes much more complex than simply negating the literal content of the information. Instead, an attempt to correct misinformation may give rise to conflicts given the conventions of conversation: If a retraction has been uttered truthfully, why would a well-intentioned speaker have presented the earlier information in the first place?

Consider a hypothetical example of a reporter covering a crisis situation as it unfolds. A person with a gun is observed atop a tall building, threatening to shoot people below. The gunman is believed to be holding two building workers as hostages. Later, after it turns out that the supposed hostages had actually never come to work that day, the reporter must announce that "the suspect is *not* holding any hostages." Situations requiring the update of information that has since been shown to be invalid are common in everyday life, where access to information can change rapidly. In these cases and in the example, the speaker intends the audience to disregard the earlier reference; in the example, listeners should view the situation in the same way as would a person who tuned in later, and thus never heard any suggestion that there might have been hostages. In situations such as jury decision-making and tutoring sessions, instructing people to disregard incorrect information may happen frequently and may be critical for accurate understanding.

Unfortunately, audiences exposed to corrections often differ in important ways from those who never heard the incorrect information at all. In particular, people have been shown in many domains to continue to rely on subsequently corrected information when making further inferences, a phenomenon termed the *continued influence effect* (Johnson & Seifert, in press). It has been shown that discredited information about personality characteristics influences subjects' perceptions of people (e.g., Ross, Lepper, & Hubbard, 1975; Schul & Burnstein, 1985; Wyer & Budesheim, 1987), jurors can be influenced by testimony ruled inadmissible (e.g., Carretta & Moreland, 1983; Thompson, Fong, & Rosenhan, 1981), and readers respond differently to questions depending on whether they had been exposed to critical information that was subsequently retracted (Wilkes & Leatherbarrow, 1988). This continued influence effect is not merely due to a failure to notice or remember the discrediting episode, nor to any more general memory failures (Carretta & Moreland, 1983; Ross et al., 1975; Wilkes & Leatherbarrow, 1988).

Previous work on corrections has focused on the literal contradiction in meaning between the original and updated statements (Johnson & Seifert, in press). However, such analyses fail to address the context of the communication, in which the literal meaning of a discrediting statement must be reconciled with conversational conventions. Grice's (1967/1989) theory of conversational logic argues that communication depends upon shared assumptions about the pragmatic meaning of discourse, in which "implications and suggestions" are additionally conveyed "over and above the

strictly asserted content of the speaker" (p. 141). Perhaps discrediting is difficult to understand because they often seem to violate these conversational conventions; specifically, corrections may be difficult to interpret when the listener cannot believe that a coherent, good-faith speaker would have uttered both the initial assertion and its later correction. In attempting to resolve this conflict, the listener might assume the continued relevance of the discredited information despite being aware of the correction.

Given our conversational conventions, correction statements may pose problems in two important ways. First, listeners typically assume that conversational contributions will be relevant to the topic under discussion. In the case of discrediting, Grice's *maxim of relation* would imply that both the original assertion and its subsequent negation were relevant to the topic. Thus although the literal content of the two statements is directly contradictory, listeners may continue to treat the original information as relevant. Second, listeners conventionally assume that speakers only provide information when they have good reason to believe it is true (Grice's *maxim of quality*). In corrections, the speaker must be assumed to have had good reason for believing both the original assertion and its subsequent negation. Thus despite the contradictory content of the two statements, listeners may continue to treat the original information as accurate. Consequently, in either case, the discredited information could continue to influence listeners' later inferences.

The present research investigates ways of reducing the continued influence effect. Our design was modeled after Wilkes and Leatherbarrow's (1988) paradigm, in which subjects read a series of newspaper-style reports about the investigation of a fire in a stationery warehouse. Control subjects were given no explicit information about the fire's cause; subjects in the correction conditions were given causal information that was subsequently discredited. The critical messages in the series concerned the contents of a closet. Subjects in the correction conditions were initially told that the closet had contained volatile materials (paint cans and gas cylinders), and were later told that the closet had actually been empty; control subjects learned only that the closet had been empty. The continued influence effect was demonstrated in subjects' responses to a series of inferential questions about the reports (e.g., What could have caused the toxic fumes? Why did the fire spread so quickly?): Correction subjects mentioned volatile materials as the fire's cause significantly more than did control subjects, despite successfully remembering the story facts and the correction itself. In the present experiment, we examine the contribution of both conversational implications and literal content to the continued influence effect.

Conversational Implications

In the Wilkes and Leatherbarrow (1988) paradigm, the continued influence effect may have stemmed from subjects' inferring that the original information as well as the correction statement was relevant and of good informational quality. If so, then providing correction statements which release subjects from such inferences should reduce the continued influence effect. For example, in the rooftop

terrorist situation described above, the reporter could have delivered an explanatory correction by saying, "the suspect is *not* holding any hostages; although two missing workers had been initially believed to be held as hostages, we now know that in actuality those workers had never come to work today." This statement provides a reason why someone might have initially believed in good faith that there were two hostages.

In the present experiment, we included two conditions in which the correction statement explained why the original information might have been uttered in good faith. The *explain relevance* correction suggested that the original information was found to be no longer relevant to current conditions. The *explain quality* condition suggested that the original information had been found to be of poor informational quality. We predicted that subjects receiving explanatory corrections would show less continued influence than subjects receiving simple negations. The explanations were predicted to release subjects from assuming that the original assertion was accurate and relevant to subsequent understanding, thus leading them to respond in accordance with the correction statement.

Literal Content

Alternatively, one might argue that the continued influence effect had been obtained in the prior studies because the literal content of the correction message in the prior studies had been inadequate. Specifically, one might argue that the message failed to rule out all involvement of the discredited information (e.g., subjects reading about the warehouse fire might infer that volatile materials, although not in the closet, had been elsewhere on the premises at the time of the fire). If so, enhancing the strength of the correction statement should reduce the continued use of discredited information. Thus we included in the present experiment an *enhanced negation* condition to determine whether elaborating upon the literal content of a correction improves understanding. If inadequacy of literal content had produced the effect found in previous studies, then one would expect to see a reduction in the effect for subjects in the enhanced negation condition. On the other hand, an extension of the conversational logic perspective predicts that listeners will draw further inferences when speakers seem to repeat themselves without conveying any additional information. In particular, if the literal content is uninformative, listeners will infer that its reiteration implies a level of doubtfulness about the truth status of the statement (cf. Gruenfeld & Wyer, 1992; Wegner, Wenzlaff, Kerker, & Beattie, 1981). In line with the conversational logic perspective, we predicted that subjects receiving the enhanced negation correction would show as much or more continued influence as would subjects receiving the *simple negation* correction used in prior studies.

In the present experiment, subjects read one of six versions of the series of reports about the investigation of a warehouse fire (adapted with minor modifications from Wilkes & Leatherbarrow, 1988). There were two control conditions: a *no-reference control* in which subjects never heard any suggestion that gas and paint had been involved, and an *undisputed reference control* in which subjects heard

without contradiction that there had been paint and gas in the closet at the time of the fire. Together, these conditions were intended to provide a baseline and a ceiling, respectively, for subjects' likelihood of mentioning the critical materials. In addition, there were four correction conditions, in which subjects were initially told that gas and paint were in the closet before the fire. The *simple negation* correction told subjects that "the closet had actually been empty before the fire." The *enhanced negation* correction additionally asserted that no paint or gas had ever been stored on the premises. The *explain relevance* correction asserted that the closet had been empty because a trucker's strike had interfered with an anticipated delivery of paint and gas. The *explain quality* correction asserted that the closet contained cans of coffee and soda rather than cans of paint and gas.

We expected that the continued influence effect -- making more inferences involving the discredited information compared to when it had not been mentioned at all -- would be obtained for the simple negation condition, replicating Wilkes and Leatherbarrow (1988) and Johnson and Seifert (in press). We further expected that subjects receiving either of the explanatory corrections would make fewer inferences involving the volatile materials than would subjects receiving negation corrections, because the explanations

would release these subjects from expecting the original assertion to be relevant. In addition, we expected that the enhanced negation correction would result in at least as much influence of the discredited information as the simple negation correction, because subjects would attempt to view the enhanced information as informative and would infer that its utterance conveys doubtfulness about the validity of the correction statement.

Method

See Bush, Johnson, and Seifert (1994) for an unabridged description of all methods and results.

Subjects

One hundred and sixteen University of Michigan undergraduates participated in groups of 6 to 12 for a single session lasting approximately 50 minutes.

Materials

The materials were a series of 13 newspaper-style reports (2-4 sentences long) describing the investigation of a warehouse fire, adapted from the materials of Wilkes and

Table 1: Text of Critical Messages in All Conditions

Condition	Content of Target Message
Message 5 Text	
4:30 a.m. Message received from Police Investigator Lucas saying that they have reports that	
No-Reference Control	<i>the closet was empty before the fire.</i>
All Other Conditions	<i>cans of oil paint and pressurized gas cylinders had been present in the closet before the fire.</i>
Message 11 Critical Text	
10:40 a.m. A second message received from Police Investigator Lucas regarding the investigation into the fire.	
Control Conditions	
No-Reference Control	It stated that <i>the two firefighters taken to the hospital had been released.</i>
Undisputed Reference Control	It stated that <i>the two firefighters taken to the hospital had been released.</i>
Negation Correction Conditions	
Simple Negation	It stated that there were no cans of paint or gas cylinders in the closet that had reportedly contained them; <i>the closet had actually been empty before the fire.</i>
Enhanced Negation	It stated that <i>there was now clear evidence that no paint or gas had ever been stored anywhere on the premises. The closet had actually been empty before the fire.</i>
Explanatory Correction Conditions	
Explain Quality	It stated that there were no cans of paint or gas cylinders in the closet that had reportedly contained them; <i>the closet had actually been filled with cans of coffee and soda canisters.</i>
Explain Relevance	It stated that there were no cans of paint or gas cylinders in the closet that had reportedly contained them; <i>a delivery had been anticipated, but it had never arrived, due to a truckers' strike.</i>

Leatherbarrow (1988). Booklets for the six conditions were identical with the exception of two critical messages, which were varied in accordance with experimental condition.

For five of the conditions, the statement that the closet contained the volatile materials (cans of oil paint and pressurized gas cylinders) was presented as Message 5; however, in the *no-reference control* condition, Message 5 stated instead that a closet in the warehouse had been empty (see Table 1 for text).

For all conditions, Message 11 stated, "10:40 a.m. A second message received from Police Investigator Lucas regarding the investigation into the fire." The subsequent text of Message 11 was varied by condition, as shown in Table 1 above.

Two outcome measures were also prepared. The first, a "summary sheet," asked subjects to provide a free-recall report of the messages' contents, and to answer a question about the fire's cause. The second measure was an open-ended questionnaire which contained eleven questions on explicit facts (e.g., What was in the storage hall? Where on the premises was the fire located?), followed by nine other questions requiring subjects to make inferences about the event (e.g., What could have caused the toxic fumes? What do the police currently suspect was responsible for the fire?), and concluding with two questions assessing whether subjects had been aware of any correction or contradiction in the series (i.e., What was the point of the second message from the police? Were you aware of any corrections in the reports that you read?).

Procedure

Each subject received a booklet of reports and was instructed to read through it page by page at his or her own pace. Subjects were told that they would be asked to recall the information later. When each individual subject had finished reading, the experimenter collected his or her booklet of reports, and the subject received the summary sheet. When everyone had finished the summary task, all subjects worked on an unrelated distractor task for 10 minutes. Afterwards, subjects received the open-ended questionnaire and were instructed to answer each question based on their understanding of the reports.

Results and Discussion

Inferences

All scoring was done by a trained coder blind to experimental conditions. The *thematic-inference* measure consisted of "negligence" theme responses made on the nine inference questions. Such responses were consistent with believing that the warehouse had contained carelessly stored volatile materials that could have caused the fire; although logically consistent with the initial message about the volatile materials, they were inconsistent with the correction message (see Table 2 for sample responses consistent with a negligence theme). Included were responses using key words (e.g., oil, paint, gas), more general terms for the closet materials (e.g., flammables, combustibles), and responses that made qualified but uncontroverted use of the key words

Table 2: Sample Negligence Theme Responses

Sample Inference Questions and Responses

Q: What might have made the fire particularly intense?

A: Oil fires are hard to put out.

Q: What could have caused the toxic fumes?

A: They thought oil and paint cans.

Q: What could have caused the toxic fumes?

A: Burning paint.

(e.g., "if there was paint" [in the closet]). Excluded were responses expressing themes other than negligence, such as stored paper, faulty wiring, structural properties of the building, and intentional fire-setting, or directly controverted references to the volatile items.

Table 3 below shows the mean number of general negligence theme inferences for subjects within each condition (column 1). All results reported throughout the paper are significant at the .05 level unless otherwise noted. An analysis of variance revealed that the six groups differed significantly in average number of negligence inferences; $F(5, 110) = 7.67$, $MSe = 3.48$. A series of planned comparisons revealed that, on average, the undisputed reference condition contained significantly more negligence inferences than the two negation conditions ($F(1, 110) = 7.09$); the two negation conditions contained significantly more negligence inferences than the two explanatory correction conditions ($F(1, 110) = 4.14$); and the two explanatory correction conditions contained significantly more negligence inferences than the no-reference condition ($F(1, 110) = 4.51$). An additional planned comparison was performed between the two negation conditions, revealing that subjects in the enhanced negation condition produced marginally more negligence inferences than in the simple negation condition ($F(1, 110) = 2.73$, $p = .10$). Thus, as in the previous studies, a continued influence effect was found (Johnson & Seifert, in press; Wilkes & Leatherbarrow, 1988), and the amount of continued influence was reduced in the explanatory correction conditions and increased in the enhanced negation condition relative to the simple negation condition.

An additional *direct-reference* measure was computed, consisting of a count of all direct and uncontroverted references to the volatile materials themselves (paint cans or gas cylinders). It showed a statistically significant pattern that closely resembles that obtained for negligence theme inferences. Finally, a *correction-recall* measure, which indicated if subjects showed an awareness of having received a correction on either the summary sheet or responses to the two contradiction questions, was computed. It showed differences among conditions, $X^2 = 22.91$. As expected, subjects in all of the correction groups showed high levels of recall of the correction (96% of subjects), whereas subjects in the control conditions rarely made responses suggestive of having heard a correction (15% of subjects). These results provide support for the hypothesis that a correction is more effective if it contains an explanation addressing why the conflicting reports were given in the first

Table 3: Mean Number of Inferences by Condition

Condition	Negligence	Unhedged	Hedged
No-Reference Control	0.70	0.65	0.05
Explain Relevance	1.74	1.21	0.53
Explain Quality	1.85	1.40	0.45
Simple Negation	2.16	1.79	0.37
Enhanced Negation	3.16	2.42	0.74
Undisputed Reference Control	4.05	3.95	0.11

place, and more generally that the conversational implications of correction statements are no less influential than the literal implications.

Hedged Responses

The *thematic-inference* measure above was further broken down into two subcategories: *unhedged* negligence theme inferences, and *hedged* negligence theme inferences. The latter category encompassed references to the volatile materials that were couched in terms of uncertainty (e.g., "If there were paint in the closet [it could have caused the toxic fumes];" "supposedly there were explosives in the closet"), or that posited distinct instances of the volatile materials other than the controverted instance (e.g., "paint [that came] off the walls"), or that suggested an interpretation that integrated elements of both the original information and the correction (e.g., gas was "put back" into the closet before the fire). The remaining negligence theme inferences, all of which were stated directly and without qualification, fell into the unhedged negligence inferences category. A comparison of the level of unhedged negligence references for the various groups, shown in columns 2 and 3 above, indicates whether subjects' interpretations differed in their complexity or degree of uncertainty (cf. Grice (1967/1989) on "doubtfulness," pp. 140-141).

A series of planned comparisons revealed that the undisputed reference group made significantly more unhedged negligence inferences than the two negation groups ($F(1, 110) = 12.87$); the two negation groups made marginally more unhedged negligence inferences than the two explanatory correction groups ($F(1, 110) = 3.69, p < .06$); and the two explanatory correction groups failed to differ significantly from the no-reference condition ($F(1, 110) = 1.7$). For the hedged inferences, a comparison between the two control conditions and the four correction conditions planned as a manipulation check showed that the correction conditions contained significantly more ambiguous negligence inferences ($F(1, 110) = 8.11$). Thus, being exposed to a correction *does* change how people talk about the discredited information, in that they use more hedges than do subjects whose exposure to the information is undisputed. Furthermore, these results suggest that the impact of explanations is powerful enough to eliminate significant influence of the discredited information on how subjects think and speak about the events in question.

In the explanatory correction conditions, the level of unhedged inferences involving the discredited materials did not differ from the level in the no-reference control group. This is consistent with the idea that subjects, faced with corrections, try to come up with plausible explanations for

why the conflicting messages would have been given, and that the hedged references refer to the integrative explanations that subjects have generated. In the two explanatory negation conditions, subjects were provided with an explanation, and thus they were more successful than subjects in the negation conditions at avoiding the discredited information in their responses. To the extent that these subjects continued to incorporate the discredited information into their interpretations, they tended to develop interpretations that *integrated* this information with the explanation provided in the discrediting statement. Thus, when faced with a complex comprehension situation, the only difference between the inferential responses of these subjects and those of no-reference control subjects was that these subjects continued to speak of the discredited information in ways that did not literally conflict with the discrediting statement.

Conclusion

Earlier demonstrations of the continued influence effect have shown that people who are exposed to a discrediting statement often continue to make use of the information that has been discredited (e.g., Johnson & Seifert, in press). The present results further show that being exposed to a correction statement can mitigate (although not remove) reliance upon the discredited information when the correction statement provides an explanation for why the earlier misinformation might have been uttered. Furthermore, these results demonstrate that enhancing the strength of a discrediting statement without providing an explanation for its utterance can actually *decrease* the statement's effectiveness. These results suggest models of understanding informational updates must address the pragmatic issues of the manner in which the updated information is conveyed.

The present results suggest that one cause of the continued influence effect (Wilkes & Leatherbarrow, 1988; Johnson & Seifert, in press) was the failure of the discrediting statement to account for the conversational implications that favored believing in the continued relevance and informativeness of the initial information. Evidence for the importance of conversational implications in understanding corrections comes from two directions: (a) *Release from conversational implications*: the obtained reduction in the number of inferences based on the discredited information for the explanatory negation groups relative to the simple negation group; and (b) *Inadequacy of literal enhancement*: the obtained marginal *increase* in the number of inferences based on the discredited information for the enhanced negation group relative to the simple negation group. The present results do not suggest, however, that such conventions are

the only important factor in comprehension: although subjects in the explanatory correction groups showed less of the continued influence effect, they nevertheless continued to make more affected inferences than did control subjects who had never been exposed to the misinformation.

These findings imply that correcting the literal content of the misinformation is insufficient to produce an accurate understanding; instead, the conversational implications of the initial assertion of the information must also be addressed. In particular, this study shows that explicitly addressing the implications raised by corrections, -- that is, the implications that the initial assertion remains truthful and relevant -- can mitigate the further influence and use of the misinformation. Thus, if a speaker wants to make a correction well understood, the speaker should take pains to explain *why* he or she originally believed the now-discredited information to have been correct. The speaker should also be aware more generally that the conversational implications of correction statements are no less influential than their literal implications.

In news, courtroom, and classroom settings, communicators may convey corrected information without providing any explanation for their apparent violation of conversational conventions. For example, in tutoring systems (e.g., Anderson, 1990; Anderson, Boyd, Farrell, & Reiser, 1984; Brown & Burton, 1975), it may be especially important to explain the implications of the occurrence of misinformation in order to assist students in accepting and incorporating corrections in instruction. Similarly, in natural language dialog systems (e.g., Colby & Enea, 1973) and models of explanation (e.g., Kass, Leake, & Owens, 1986), resolving conflicting statements or viewpoints may lead to similar comprehension problems. Since it may be critical to make one's audience accurately understand updated information, a speaker must carefully address the complete meaning of utterances, including their conversational implications, when correcting misinformation.

References

- Anderson, J. R. (1990). Analysis of student performance with the LISP tutor. In Frederiksen, N., Glaser, R., Lesgold, A., & Shafto, M. G. (Eds.), *Diagnostic monitoring of skill and knowledge acquisition*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Anderson, J. R., Boyd, C. F., Farrell, R., and Reiser, B. J. (1984). Cognitive principles in the design of computer tutors. Office of Naval Research Technical Report #84-1.
- Brown, J. S., & Burton, R. R. (1975). Multiple representations of knowledge for tutorial reasoning. In D. G. Bobrow & A. M. Collins (Eds.), *Representation and Understanding*. New York: Academic Press.
- Bush, J. G., Johnson, H. M., & Seifert, C. M. (1994). Conventions of corrections: How to correct modified information within its conversational context. Unpublished manuscript. University of Michigan, Ann Arbor, MI.
- Carretta, T. R., & Moreland, R. L. (1983). The direct and indirect effects of inadmissible evidence. *Journal of Applied Social Psychology*, *13*, 291-309.
- Colby, K., & Enea, H. (1973). Heuristic methods for computer understanding of natural language in context-restricted on-line dialogues. *Mathematical Biosciences*, *1*.
- Grice, H. P. (1967/1987). Logic and conversation. In H. P. Grice (1989), *Studies in the Way of Words* (pp. 1-143). Cambridge, MA: Harvard University Press.
- Gruenfeld, D. H., & Wyer, R. S. (1992). Semantics and pragmatics of social influence: How affirmations and denials affect beliefs in referent propositions. *Journal of Personality and Social Psychology*, *62*, 38-49.
- Johnson, H. M., & Seifert, C. M. (in press). Sources of the continued influence effect: When discredited information in memory affects later inferences. *Journal of Experimental Psychology: Learning and Memory*.
- Kass, A., Leake, D., & Owens, C. (1986). Programming the theory: SWALE, a program that explains. In R. Schank (Ed.), *Explanation Patterns: Understanding Mechanically and Creatively*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Ross, L., Lepper, M. R., & Hubbard, M. (1975). Perseverance in self-perception and social perception: Biased attributional processes in the debriefing paradigm. *Journal of Personality and Social Psychology*, *32*, 880-892.
- Schul, Y., & Burnstein, E. (1985). When discounting fails: Conditions under which individuals use discredited information in making a judgment. *Journal of Personality & Social Psychology*, *49*, 894-903.
- Thompson, W. C., Fong, G. T., & Rosenhan, D. L. (1981). Inadmissible evidence and juror verdicts. *Journal of Personality and Social Psychology*, *40*, 453-463.
- Wegner, D. M., Wenzlaff, R., Kerker, R. M., & Beattie, A. E. (1981). Incrimination through innuendo: Can media questions become public answers? *Journal of Personality and Social Psychology*, *40*, 822-832.
- Wilkes, A. L., & Leatherbarrow, M. (1988). Editing episodic memory following the identification of error. *Quarterly Journal of Experimental Psychology*, *40A*, 361-387.
- Wyer, R. S., & Budesheim, T. L. (1987). Person memory and judgments: The impact of information that one is told to disregard. *Journal of Personality and Social Psychology*, *53*, 14-29.