

On the Genesis of Confidence

Michael J. Gill
Lehigh University

William B. Swann, Jr.
University of Texas at Austin

David H. Silvera
Universitetet i Tromsø

Confidence in personality impressions is proposed to stem from the richness of people's mental representations of others. Representational richness produces confidence because it enhances the fluency with which people can make judgments, and it increases confidence even when it does not result in more accurate impressions. Results of 3 experiments support these propositions. A 4th experiment suggests that representational richness is increased by both pseudorelevant and relevant information, but not by irrelevant information. A 5th experiment suggests that representational richness has effects on confidence above and beyond the effects of meta-information (i.e., extraneous aspects of information). The implications of these findings for evaluating evidence of error in person perception and for reducing stereotyping and prejudice are discussed.

"We grew up together—I know her better than I know myself." "I've only talked to him twice, so I may be wrong about him." People often qualify reports of their impressions with statements such as these, suggesting that some of their beliefs about others are held with great conviction, whereas other beliefs are held only tentatively. In this article, we investigate the origin of these differences in confidence.

Understanding the origin of confidence is important because confidence seems to affect whether people translate their beliefs into behaviors. For example, a hiring committee may be intrigued by a job candidate on skimming her vita but will hire her only after boosting its confidence through conversations with the candidate herself. Likewise, one would entrust a large sum of money to a broker only after becoming confident that the broker will not take the money and run. This tendency for confidently held beliefs to be a more potent impetus for behavior than tentatively held beliefs has been confirmed in several empirical

studies (e.g., Berger & Mitchell, 1989; Fazio & Zanna, 1978; Pieters & Verplanken, 1995; Swann & Ely, 1984).

Presumably, confidently held beliefs motivate action because people believe confidence stems from an accurate perception of reality. But does it? The earliest social psychological research to address the relation of confidence and accuracy was conducted by Oskamp (1965). His work suggested that clinicians' confidence was affected by factors that did not influence the accuracy of their beliefs. Recent research using everyday person perceivers has shown similar results. For example, Dunning, Griffin, Milojkovic, and Ross (1990) had college students attempt to predict the behavior of their roommates, strangers whom they interviewed, or complete strangers. In all cases, confidence exceeded accuracy, and there was only a modest correlation between the two. Similarly, Swann and Gill (1997) found marked dissociations between the confidence and accuracy of dating partners' and roommates' global personality impressions. Moreover, Wells and Murray (1984) reported virtually no relation between the confidence and accuracy of eyewitness testimony. Finally, the work of cognitive psychologists on such phenomena as feelings of knowing and judgments of learning suggests that confidence-accuracy dissociations also occur in nonsocial domains (see Jacoby, Bjork, & Kelley, 1994, for a review). A consideration of the ubiquity of such confidence-accuracy dissociations raises the following question: If confidence does not arise from the possession of accurate beliefs, from where does it arise?

REPRESENTATIONAL RICHNESS AND THE GENESIS OF CONFIDENCE

We propose that the richness of mental representations plays an important role in producing confidence. Representational

Michael J. Gill, Department of Psychology, Lehigh University; William B. Swann, Jr., Department of Psychology, University of Texas at Austin; David H. Silvera, Department of Psychology, Universitetet i Tromsø, Tromsø, Norway.

This research was supported by a University of Texas continuing fellowship to Michael J. Gill. We thank Joel Blanchard, Cicely Bomar, Robert McCann, and Adriana Weisz for their assistance during the empirical phase of this work. Sincere thanks also go to Jennifer Bosson, Ryan Brown, Morris Goldsmith, Asher Koriat, Aric Kruglanski, and Vincent Yzerbyt for their thoughtful comments on an earlier version of this article.

Correspondence concerning this article should be addressed to Michael J. Gill, Department of Psychology, 17 Memorial Drive East, Lehigh University, Bethlehem, Pennsylvania 18015. Electronic mail may be sent to mjg6@lehigh.edu.

richness is defined in terms of the confluence of two factors: (a) representation integration and (b) relatively large amounts of information. When representations contain a relatively large amount of well-integrated information, they should be associated with relatively high levels of confidence. We emphasize the conjunction of integration and information because classic work in cognitive psychology (Anderson, 1976; Anderson & Bower, 1973) suggests that large amounts of unintegrated information can actually hamper people's ability to answer questions about the information (so-called "interference effects"). Although we define representational richness in terms of the confluence of integration and information, for the sake of clarity we elaborate separately on each component.

The term *representation integration* refers to the extent to which a representation is coherent and characterized by conceptual links among its elements (Myers, O'Brien, Balota, & Toyofuku, 1984; Smith, Adams, & Schorr, 1978). In the domain of impression formation, one of the primary means of achieving an integrated representation is to recognize that a group of behaviors implicate a particular personality trait (Park, 1989). Once this initial integration occurs, encoding processes may cause additional behaviors to be assimilated to the impression, increasing its degree of integration even more (Darley & Gross, 1983; Rosenhan, 1973). Given this conceptualization of integration, it seems likely that the consistency of the behavioral information on which people base their impressions will contribute to integration. The *objective* consistency of the information may not be critical, however, because high levels of motivation encourage people to reevaluate the meaning of inconsistent information to achieve well-integrated impressions (Murray & Holmes, 1993). Motivation may also influence representation integration in another manner. Specifically, Ostrom, Lingle, Pryor, and Geva's (1980) work suggests that people create coherent representations of a person on dimensions that are relevant to their impression formation goals. Thus, for example, a person whose goal is evaluating Hans's cleverness will strive to make sense of that aspect of Hans, thereby developing an integrated impression of his cleverness. Whether it arises from cognitive or motivational processes, representation integration is a critical ingredient of representational richness.

The amount of information composing a representation is also a critical ingredient of representational richness. One important issue to consider with respect to this dimension of representational richness is what types of information people will use to enrich their representations of others and what types of information they will ignore. For example, if one were interested in forming an impression of a woman's intelligence, one would presumably characterize her IQ as relevant information and her height as irrelevant information. Between these hypothetical extremes (where the degree of information relevance is obvious) lies *pseudorelevant information* that can be construed as being relevant to a given judgment but that actually has no necessary implication for the judgment. For example, learning that a man enjoys mystery movies or likes to read may *seem* relevant to a judgment of his intelligence, but in reality it may do little to enhance the accuracy of that judgment. In the spirit of research on social judgment, we assume that people have a tendency to use both pseudorelevant information and relevant information to enrich their impressions (e.g., DePaulo, Stone, & Lassiter, 1985; Dion, Berscheid, & Walster, 1972; Hilton & Fein, 1989;

Jones & Harris, 1967; Ross, Amabile, & Steinmetz, 1977; also see Nisbett & Ross, 1980, for a review). We thus predict that both types of information will foster confidence.

The Role of Judgment Fluency

We propose that judgment fluency mediates the impact of representational richness on confidence (see Benjamin & Bjork, 1996, for a discussion of the related constructs of perceptual fluency and retrieval fluency). The term *judgment fluency* refers to the speed and subjective ease with which judgments can be constructed or retrieved. Research suggests that information integration and amount (Myers et al., 1984; Park, 1989; Sherman & Klein, 1994; Smith et al., 1978) enhance fluency. This research, however, was not concerned with confidence, and participants were forced to rely on written information rather than on the types of information they would encounter in naturally occurring social settings (e.g., verbal statements from a person).

In our view, highly fluent judgments will be made with more confidence than judgments that are made slowly and with subjective difficulty. Indeed, fluency has been linked to confidence in nonsocial judgments (Benjamin & Bjork, 1996; Kelley & Lindsay, 1993; Koriat, 1993; Nelson & Narens, 1990; Shaw, 1996), although, to our knowledge, its impact on the confidence of social judgments has not been explored.

Representational Richness and Confidence— Accuracy Dissociation

Our interest in the genesis of confidence was motivated by a desire to understand confidence-accuracy dissociations. We propose that representational richness will sometimes be unrelated to accuracy for several reasons. First, rich representations can be constructed from nondiagnostic information. For example, one can have a rich representation suggesting that one's business partner is honest (e.g., "He attends my church," "He returned the office equipment that he borrowed," "He is a family man," etc.) despite the fact that the partner is stealing from the company. When nondiagnostic information enriches an impression, confidence will outstrip accuracy. (Conversely, accuracy may outstrip confidence when a representation that is not especially rich happens to contain highly diagnostic information.) Furthermore, even if the information one gains about others during social interaction is diagnostic of their underlying dispositions, one may draw incorrect inferences from that information and develop rich representations based on such faulty inferences. Thus, even the presence of diagnostic information does not guarantee that accuracy will increase as richness increases.

More subtle processes related to the elaboration of representations may also affect the extent to which richness and fluency are associated with accuracy (see Benjamin & Bjork, 1996, for a review). For example, frequent thought about a person may lead one to generate additional supporting reasons for one's beliefs about that person, thereby increasing the richness but not necessarily the veridicality of one's beliefs (e.g., Koriat, Lichtenstein, & Fischhoff, 1980). Moreover, the mere act of repeatedly accessing an impression may enhance the fluency of the impression, but it obviously will not enhance accuracy because the underlying representation has not changed (e.g., Shaw,

1996). In our view, any factors that enhance richness and fluency will contribute to confidence even when they do not affect accuracy.

We tested our ideas about representational richness, judgment fluency, and confidence in three experiments. In part, our experiments were a follow-up to Swann and Gill (1997, Study 3), who used structural equation modeling to examine these processes in the context of a single study. Although the results were consistent with our account of the genesis of confidence, they did not provide strong evidence about the causal relations among our theoretical variables. To remedy this shortcoming, we focus here on using experimental designs to test those causal relations. We sought support for our account of confidence by testing three hypotheses: (a) Representational richness fosters confidence (without necessarily increasing accuracy), (b) representational richness fosters judgment fluency, and (c) fluency fosters confidence. These three hypotheses were tested in Experiments 1, 2, and 3, respectively. Support for all three hypotheses would bolster our claims that representational richness produces confidence and that the impact of richness on confidence is mediated (at least in part) by fluency. In Experiment 4 we examined the moderating impact of information content on our confidence effects, and in Experiment 5 we attempted to reconcile our account of the genesis of confidence with the notion that meta-information gives rise to confidence.

EXPERIMENT 1: RICHNESS AND CONFIDENCE

In Experiment 1, participants formed impressions of a target person while viewing a videotaped interview. After watching the videotape, they attempted to predict the target's answers to a sexual history and a self-concept questionnaire and reported their confidence in those predictions. In accord with our definition of representational richness, we manipulated both the manner in which participants' impressions were integrated and the amount of information on which their impressions were based. Participants induced to integrate their impressions in terms of the target's sexual history and self-concept, and who received relatively large amounts of information about the target, were expected to report greater confidence than participants in the remaining three conditions combined. Furthermore, we reasoned that if representational richness affects confidence even when it does not affect accuracy, then our findings will provide at least a partial explanation of confidence-accuracy dissociations.

Our measure of accuracy required participants to "step into the target's shoes" and predict his or her self-ratings. Of course, it is possible that people are generally accustomed to reporting their impression of what a person is really like rather than their predictions of what a person would say he or she is like. To ensure that our effects were not an artifact of the nature of the predictions we asked participants to make, we added a pair of control conditions. Participants in these control conditions learned that they should form an impression of the target's sexual history and self-concept and then heard either a relatively small or a relatively large amount of information about the target before reporting their impression of the target's "true" sexual history and self-concept (rather than predicting what the target would say). We expected that our manipulation of information amount would produce confidence in these control conditions just as in our experimental conditions.

Method

Participants

Participants were 42 male and 45 female introductory psychology students at the University of Texas at Austin who received course credit. Data from 5 participants were deleted: Two deletions were due to experimenter error, and three were due to participants' failures to follow instructions.

Procedure

On arrival at the laboratory, participants met a same-gender experimenter who escorted them to a private cubicle equipped with a TV and VCR. After participants signed an informed consent statement, the experimenter introduced them to a study of "the accuracy of first impressions." Participants learned that they would view a videotaped interview with a target person and were asked to form an impression of the target person as they watched the videotape.

Independent Variables

Information integration: Undirected versus directed. Our integration manipulation was similar to manipulations used by Ostrom et al. (1980) to produce "thematically organized" impressions. Participants in the *undirected-integration condition* learned that they should merely "form an impression of the target person," whereas participants in the *directed-integration condition* learned that they should form an impression of the target person's sexual history, intellectual ability, social skills, artistic and musical ability, athleticism, and physical attractiveness. These dimensions were the ones on which participants would eventually make judgments of the target. The list of dimensions was repeated twice (once in writing and once orally) to facilitate retention, and participants were told that they would be attempting to predict how the target would rate himself or herself on these dimensions.

Information amount: Low versus high. After our integration manipulation, participants watched a videotaped interview with an opposite-gender target person. To increase the generality of our results, we used three female and two male target individuals. (The effect of target person did not interact with any of our manipulations, so we do not discuss it further.) Participants in the *low-information condition* listened to the target person describe his or her major and why he or she chose it, his or her career plans, and where he or she was from and what it was like growing up there. In contrast, participants in the *high-information condition* heard this information and then heard the target agree or disagree with 15 statements describing various behaviors and attitudes (e.g., "If my significant other criticizes me, I would yield to the criticism and change my behavior," "If I do poorly on a test, I feel miserable," "I believe in the old saying 'There is no place like home,'" etc.). These 15 items were read by an interviewer who was not visible on the videotape. Before being videotaped, targets were encouraged to answer the interviewer's questions truthfully. All participants viewed the target person for the same amount of time. However, for participants in the low-information condition, the audio portion of the videotape was silenced while the target discussed the 15 behavior and attitude items. To prevent these participants from feeling suspicious about the lack of sound on the videotape, all participants learned that we were interested in the impact of both verbal and nonverbal behavior on impressions and that both sources of information should be used while viewing the videotape.

Dependent Variables

Impression of the target. To create an index of the accuracy of participants' impressions, we compared the actual responses that targets made on two questionnaires to participants' predictions of targets' re-

sponses. One questionnaire was the Sexual History Questionnaire (SHQ), which is a modified version of Metzler, Noell, and Biglan's (1992) measure of high-risk sexual behavior. The SHQ consists of 10 items, answered on 5-point scales, concerning the respondent's number of sexual partners, frequency of condom use, number of one-night stands, frequency of discussing sexual history before engaging in intercourse with a new partner, and so on. The other questionnaire was the Self-Attributes Questionnaire (SAQ; Pelham & Swann, 1989). The SAQ consists of five items, answered on 10-point scales, asking respondents to rate their self-perceived intellectual ability, social skills, artistic and musical ability, athletic ability, and physical attractiveness. The SHQ and SAQ were reworded so that participants were predicting the target's responses on each questionnaire.

We collected these ratings because we were interested in the accuracy of participants' predictions of targets' responses. The computation of accuracy necessarily involves a decision on the researcher's part about what counts as an accurate judgment (Kruglanski, 1989). We reasoned that because we explicitly asked participants to predict how the target answered the SHQ and SAQ, the appropriate measure of accuracy would require some assessment of the extent to which participants succeeded in predicting targets' responses. We created an accuracy index for each participant for both the SHQ and the SAQ. Specifically, we used Pearson product-moment correlations to measure the extent to which the pattern of responses predicted by the participant correlated with the pattern of responses provided by the target person on each questionnaire. Accuracy correlations ranged from $-.32$ to $.86$ on the SHQ and from $-.84$ to 1.0 on the SAQ. The mean levels of accuracy were $.33$ and $.38$ for the SHQ and SAQ, respectively.

Confidence. After answering each item on the SHQ, participants responded to the following confidence item by filling in a blank:

What is the probability, ranging from 0% (meaning "I have no confidence that my answer is correct") to 100% (meaning "I have complete confidence that my answer is correct") that your answer to the question above is a correct prediction of the target person's response to that question?

After completing all five items on the SAQ, participants answered a single confidence item similar to the one above. The 10 confidence items from the SHQ ($\alpha = .90$) were averaged to form one index of confidence, whereas the single confidence item from the SAQ provided a second index.

Accuracy and confidence indexes for the SHQ were not computed for participants who rated the target as having had zero sexual partners. This was necessary because such participants were missing almost all the data needed to compute those indexes (i.e., they did not answer most of the SHQ items).

Control Conditions for Type of Target Rating

We added two control conditions to ensure that our effects were not an artifact of having participants predict the target's self-ratings. The procedures for these two conditions were identical to the procedures for the low- and high-information/directed-integrated conditions with one exception: Participants in these control conditions reported their impression of what the target was "really like" (on the SHQ and SAQ) rather than attempting to predict what the target would say he or she was like.

Results and Discussion

Confidence

Representational richness is defined in terms of the confluence of information integration and amount. Accordingly, we analyzed confidence by computing a planned comparison of the directed-integration/high-information condition with the aver-

age of the remaining three conditions. Because preliminary analyses revealed that the results for the SHQ and SAQ were highly similar, here and throughout this article we averaged the SHQ and SAQ data before conducting our analyses. We present results separately for each questionnaire when those results differ.¹ As can be seen in Figure 1, participants in the directed-integration/high-information condition were more confident of their judgments of the target than were participants in the remaining conditions, $F(1, 39) = 8.3, p < .007$. Moreover, the residual between-groups variability remaining after this comparison was not significant ($F < 1$), suggesting that our planned comparison described the pattern of means well.²

Because belief confidence may be correlated with belief extremity (e.g., Jones & Davis, 1965), we sought to test whether our confidence effects were independent of the extremity of participants' judgments of targets. First, we computed an extremity score for each participant. We did this by calculating the absolute value of participants' deviations from the midpoint of the rating scale on each item, averaging these within the SHQ and SAQ, and then averaging across questionnaires. When we compared response extremity in the directed-integration/high-information condition with response extremity in the remaining conditions, we found no difference ($F < 1$). Thus, representational richness seems to affect confidence even when it does not produce changes in the extremity of people's beliefs.

¹ Because participants who predicted that the target would report being a virgin lacked data for the SHQ, we excluded their data from analyses. In all relevant analyses, however, including these participants (by substituting their SAQ data for what would be an average of SHQ and SAQ data) did not alter our results.

² We included a single control condition to assess the extent to which metainformational cues (Yzerbyt, Schadron, Leyens, & Rocher, 1994) were responsible for the high levels of confidence expressed by participants in the high-information/directed-integration condition. Metainformational cues are content-free cues that affect people's beliefs about how informed they are. For example, in this study participants may have completely ignored the information on our videotapes and reported greater confidence simply because we told them that they would receive information about the target. We reasoned that because we had made clear that we were interested in the accuracy of their judgments, participants would attend to the actual information on our videotapes rather than being swayed by metainformational cues (see Fiske & Neuberg, 1990). To test our reasoning, participants in our control condition underwent the same procedure as participants in the directed-integration/high-information condition, except that they received an explicit metainformational cue suggesting that they should doubt the usefulness of the videotape for enabling them to judge the target: "Some people in this experiment receive information that helps them form an accurate impression of the target person's sexual history and self-concept, whereas other people receive information that is not helpful for forming an accurate impression. Your task is simply to form an impression with whatever information you happen to get." If participants were not attending to the actual information we presented but were being swayed by metainformational cues, the experimenter's suggestion that the information about the target might not be diagnostic should cause participants in this condition to report less confidence than participants in the directed-integration/high-information condition. They did not ($t < .43, ns$). This suggests that, rather than simply responding to metainformational cues, participants were attending to the information we provided and believed that the information helped them form accurate impressions. We return to the issue of the impact of metainformation on confidence in Experiment 5.

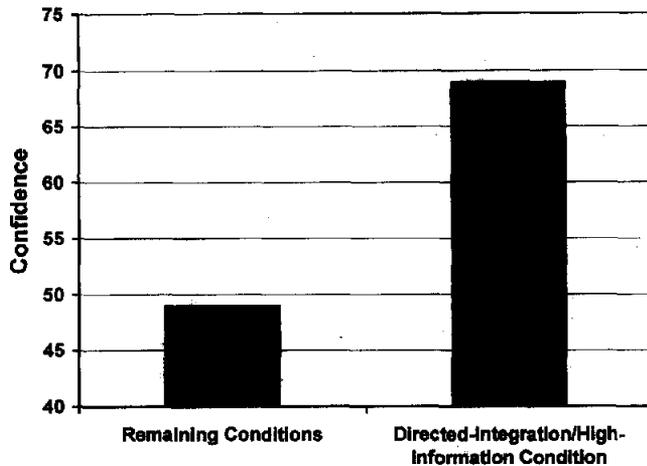


Figure 1. The impact of representational richness on confidence, assessed by comparing the confidence of participants who were directed to integrate information in judgment-relevant ways and who received relatively large amounts of information to all other participants. Confidence could range from 0% to 100%.

To ensure that our effects were not an artifact of having participants predict the target's responses rather than report their impression of the target, control participants offered their impression of what the target was really like. As predicted, participants in the low-information condition were significantly less confident ($M = 51\%$) than participants in the high-information condition ($M = 70\%$), $F(1, 20) = 7.5, p < .02$. This suggests that our confidence effects do not depend on having participants predict the target's self-ratings.

Accuracy

Did representational richness have the same impact on accuracy that it had on confidence? We transformed our accuracy correlations using Fisher's r -to- Z transformation and then averaged SHQ and SAQ accuracies. We submitted this accuracy index to the same planned comparison that was computed for confidence ratings and found that representational richness had no impact on the accuracy of participants' judgments ($F < 1$).

To confirm that our failure to find a relation between representational richness and accuracy was not an artifact of the particular manner in which we operationalized accuracy, we computed a second index of accuracy. Specifically, we computed the absolute value of the difference between perceiver predictions and target responses on each item, averaged these within the SHQ and SAQ, and then averaged across questionnaires. When we computed a planned comparison between the directed-integration/high-information condition and the remaining three conditions using this new index of accuracy, we found no difference, $F(1, 39) = 1.3, p = .26$.³

Within our control conditions, a one-way analysis of variance (ANOVA) revealed that our manipulation of information amount had no effect on our correlational measure of accuracy ($F < 1$). Our average deviation measure of accuracy (i.e., the average amount by which perceivers mispredicted target self-ratings) corroborated this finding ($F < 1$).

In summary, in both the experimental and control conditions, there was a tendency for confidence to increase with manipulated increases in the richness of participants' representations. Despite its effect on confidence, our manipulation of representational richness did not foster accuracy. Note that this failure to find an effect of representational richness on accuracy is not easily attributable to problems with our accuracy measure, as we did find gender differences in accuracy (see Footnote 3). Apparently, mental representations become rich (hence confidently held) even when they do not become faithful reflections of social reality.

EXPERIMENT 2: RICHNESS AND JUDGMENT FLUENCY

We believe that representational richness affects confidence by heightening the fluency with which people can make judgments. In Experiment 2 we tested the relation between richness and judgment fluency using a 2 (information integration: undirected and directed) \times 2 (information amount: low and high) factorial design. We predicted that participants in the directed-integration/high-information condition would show greater judgment fluency than participants in the remaining three conditions combined.

Method

Participants

Participants were 101 male introductory psychology students at the University of Texas at Austin who received course credit.

Procedure

The procedure for Experiment 2 was virtually identical to the procedure used in Experiment 1. One difference was that for participants in the low-information condition of this experiment, the experimenter turned off the VCR and TV at the appropriate time rather than merely turning down the volume (as was done in Experiment 1). This was done to determine whether our effects would generalize across procedures. Also, only one (female) target person was used. The accuracy of participants' judgments of her ranged from .09 to .99 on the SHQ and from -.76 to 1.0 on the SAQ, with mean levels of accuracy being .67 and .45 for the SHQ and SAQ, respectively. Finally, rather than providing open-ended confidence responses (ranging from 0% to 100%), participants answered the question "How confident are you of your previous response?" on a 5-point scale, with endpoints labeled *not at all* and *extremely* after each item of the SHQ and SAQ.

The most important change in this experiment was the inclusion of two measures of judgment fluency: response latencies and subjective

³ None of our planned comparisons interacted with participant gender ($F_s < 1$). We did find, however, that men were more confident than women on the SHQ, $F(1, 33) = 4.2, p = .05$ ($M_s = 60\%$ and 49% , respectively), whereas there was no gender difference in confidence on the SAQ, $F(1, 40) = 1.9, p = .17$. When examining our correlational measures of accuracy, we found that men were more accurate than women on the SHQ, $F(1, 33) = 28.7, p < .001$ (average $r_s = .53$ and $.17$ for men and women, respectively), whereas women were more accurate than men on the SAQ, $F(1, 40) = 19.6, p < .001$ (average $r_s = .15$ and $.74$ for men and women, respectively). Interpretation of these results is hampered because the genders of targets and perceivers are confounded.

ease judgments. Participants made their predictions of the target's self-ratings on a computer that recorded their response latencies for each prediction. Thus, for each participant, we recorded 10 latencies for the SHQ and 5 for the SAQ. Participants were encouraged (both orally and by means of computer-presented instructions) to make their predictions both quickly and accurately. Another measure of fluency was obtained after participants completed their SHQ and SAQ predictions. Specifically, each item of the SHQ and SAQ was presented again on the computer screen, and participants reported their subjective impression of the speed with which they were able to answer each item. These subjective ease judgments were made on 5-point scales, with endpoints labeled *relatively slowly* and *relatively quickly*.

An examination of our response latency data revealed that although the mean response latency for each of our 15 items fell between 6 and 13 s, there were several extreme values within the 40- to 60-s range. Rather than deleting all the data of individuals with these extreme values, we merely deleted the extreme values (Shoben, 1982). Extreme values were defined as any response latency greater than 2 *SDs* from the mean, and this procedure resulted in the deletion of approximately 4% of our response latency data. Neither the information integration manipulation nor the information amount manipulation affected the rate of data deletion (χ^2 's < 1.5, *ps* > .23). In addition, we deleted all the SHQ data from 3 participants (one from every condition except directed integration/low information) because more than 40% of their SHQ response latencies had been extreme values, suggesting that they had some difficulty answering the sexual history questions.

We obtained two indexes of judgment fluency (i.e., response latencies and subjective ease judgments) that we wished to combine into a single index. Subjective ease judgments were missing for 19 participants who predicted that the target would say she was a virgin (the bulk of SHQ items are irrelevant for virgins, so our computer program terminated on receiving a virgin rating). Thus, whereas 79 participants have response latency data for the SHQ and 101 participants have response latency data for the SAQ, only 78 participants have subjective ease data (data from 1 participant were not recorded because of computer error). We dealt with this disparity by analyzing judgment fluency only among the 78 participants who provided both response latency and subjective ease judgments.⁴

We computed a measure of judgment fluency that captured both the judgment speed and subjective ease aspects of the construct. First, we reverse-scored the 15 subjective ease items (so that low scores indicated fluency, as with response latencies) and then standardized response latencies and subjective ease judgments for each of our 15 items. Finally, we averaged these standardized scores to form an index of judgment fluency ($\alpha = .82$).

Results and Discussion

Judgment Fluency

We analyzed our index of judgment fluency by computing a planned comparison of the directed-integration/high-information condition with the average of the remaining three conditions. This contrast supported our prediction that fluency would be higher among participants who formed rich representations than among those who did not, $F(1, 76) = 10.2, p < .003$ (see Figure 2 for means). The residual between-groups variability remaining after this contrast was significant, $F(2, 75) = 3.2, p < .05$. Post hoc tests revealed that the directed-integration/low-information condition ($M = -0.09$) was associated with greater fluency than the undirected-integration/high-information condition ($M = 0.38$) but that neither of these conditions differed from the undirected-integration/low-information condition ($M = 0.20$). This pattern of means suggests that integration may

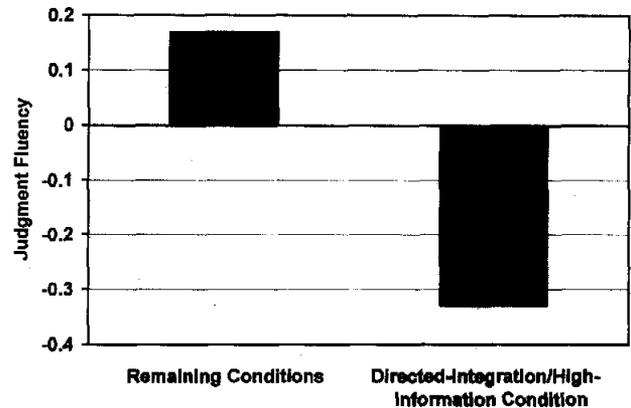


Figure 2. The impact of representational richness on judgment fluency, assessed by comparing the fluency among participants who were directed to integrate information in judgment-relevant ways and who received relatively large amounts of information to all other participants. Fluency scores were calculated by standardizing and then averaging response latencies and subjective ease judgments. Low scores indicate greater fluency.

enhance fluency even in the presence of small amounts of information and is consistent with the notion that relatively large amounts of unintegrated information hamper fluency (Anderson, 1976; Anderson & Bower, 1973). From our perspective, however, the most important point to be taken from these analyses is that fluency was greatest when people had been induced to form rich representations by integrating relatively large amounts of information.

Did Fluency Mediate the Impact of Richness on Confidence?

Our account of confidence suggests that the impact of representational richness on confidence should disappear when judgment fluency is statistically controlled. We conducted a correlational test of mediation to examine this possibility (see also Swann & Gill, 1997). First, we examined the impact of richness on confidence (averaged across questionnaires) by computing a planned comparison of the directed-integration/high-information condition with the average of the remaining conditions. This comparison revealed that confidence was higher when participants had been induced to form rich representations ($M = 3.7$) than when they had not ($M = 3.4$), $F(1, 76) = 4.4, p < .04$. The residual between-groups variability remaining after this

⁴ To make use of the additional SAQ response latencies available for participants who predicted that the target would say she was a virgin, we computed average SAQ response latencies for all 101 participants ($\alpha = .60$) and computed a planned comparison of the directed-integration/high-information condition with the average of the remaining three conditions. This comparison supported the notion that representational richness fostered judgment speed on the SAQ, $F(1, 99) = 6.8, p < .02$ (M s = 6.6 and 7.6 for the directed-integration/high-information condition and the remaining conditions, respectively). Incidentally, SHQ response latencies showed the same pattern, $F(1, 77) = 16.6, p < .001$ (M s = 7.9 and 9.7 for the directed-integration/high-information condition and the remaining conditions, respectively).

contrast was not significant, $F(2, 75) = 1.4, p = .26$. Next, we repeated this planned comparison after covarying fluency scores out of confidence judgments. This analysis of covariance (ANCOVA) revealed that the relation between fluency and confidence was strong and in the predicted direction, $F(1, 75) = 32.3, p < .001, \beta = -.59$, and that the previously significant relation between richness and confidence disappeared when fluency was statistically controlled ($F < 1$). This supports our idea that judgment fluency mediates the impact of representational richness on confidence.

Accuracy

Did representational richness affect accuracy in the same manner as it affected judgment fluency? After conducting a Fisher r -to- Z transformation, we averaged our correlational accuracy indexes from the SHQ and SAQ and submitted this accuracy index to the same planned comparison that was computed for judgment fluency. This comparison revealed that our manipulation of representational richness did not affect accuracy ($F < 1$).

As in Experiment 1, we computed the average absolute deviation of perceivers' predictions from targets' self-ratings on both the SHQ and SAQ to serve as a second index of accuracy. When we computed planned comparisons between participants in the directed-integration/high-information condition and participants in the remaining three conditions using these new indexes of accuracy, we uncovered a tendency for SHQ accuracy to be relatively low in the directed-integration/high-information condition ($M = 1.8$) compared with the average of the remaining conditions ($M = 1.6$), $F(1, 77) = 5.1, p < .03$. The opposite pattern was found on the SAQ, such that accuracy was relatively high in the directed-integration/high-information condition ($M = 0.6$) compared with the average of the remaining conditions ($M = 0.8$), $F(1, 99) = 6.6, p < .02$.

In summary, these results support the notion that representational richness increases the fluency of people's judgments. Furthermore, our correlational test of mediation was consistent with the notion that representational richness fosters confidence by increasing the fluency with which people can make judgments. Finally, richness was again found to have no necessary connection to accuracy: Our average deviation measure of accuracy suggested that richness enhanced accuracy on the SAQ but that richness decreased accuracy on the SHQ. Although an exploration of the conditions under which rich representations enhance accuracy is beyond the scope of this article, Funder's (1995) and Kenny's (1991) explications of the conditions under which diagnostic information is likely to be incorporated into impressions offer relevant insights.

Taken together, the results of Experiments 1 and 2 provide evidence for links between representational richness and confidence as well as between representational richness and judgment fluency. To clinch our argument that fluency *mediates* (at least in part) the relation between representational richness and confidence, we need to show that increases in fluency produce increased confidence. We sought such evidence in Experiment 3.

EXPERIMENT 3: JUDGMENT FLUENCY AND CONFIDENCE

Extensive research has suggested that when people are primed with a particular trait concept, they tend to interpret the behaviors of subsequently encountered people in terms of that concept (e.g., Bargh & Pietromonaco, 1982; Higgins, Rholes, & Jones, 1977; Srull & Wyer, 1980; see Higgins, 1989, for a review). In one well-known study (Higgins et al., 1977), the concept *adventurous* was primed for some participants, whereas the concept *reckless* was primed for others. Participants then took part in an ostensibly unrelated experiment in which they provided an open-ended personality description of a man whom they were told had climbed Mt. McKinley, gone white-water kayaking, driven in a demolition derby, and so on. An examination of participants' open-ended descriptions revealed that they spontaneously (and unwittingly) described the man in terms of the concept that had been primed earlier, suggesting that judgments pertaining to the primed dimension were especially fluent (i.e., came more quickly and easily to mind) compared with other judgments.

In Experiment 3, we capitalized on the judgment fluency engendered by priming manipulations to test whether fluency would cause confidence. Specifically, we primed participants with either the concept *ambition* or the concept *intelligence* before they formed impressions of a videotaped target. After watching the videotape, participants predicted the target's responses to a personality questionnaire. We predicted that participants who had been induced (via the priming manipulation) to form fluent impressions of the target's ambition would make more confident ratings of the target's ambition than of the target's intelligence, whereas the reverse pattern would hold for participants who had been induced to form fluent impressions of the target's intelligence.

Method

Participants

Participants were 16 male and 38 female introductory psychology students at the University of Texas at Austin who received course credit. Data from 1 participant were deleted because he failed to answer all of the questionnaire items, and data from another were deleted due to experimenter error.

Procedure

A female experimenter escorted each participant to a cubicle, obtained informed consent, and then explained that the participant would be taking part in "two separate studies." The rationale given for having participants complete two experiments was that both experiments took less than 15 min, and thus two experimenters had agreed to share participants who had signed up for 30 min of experimental credit. In reality, there was only one experiment.

Priming Manipulation

The experimenter explained that the first study was "an investigation of the meaning of psychological concepts to nonpsychologists." The cover story asserted that psychologists are often faulted for defining concepts in overly narrow ways. Participants learned that the purpose of the study was to collect a sample of definitions of psychological concepts from nonpsychologists, look for themes in those definitions,

and then explore the extent to which psychologists ignored those themes when defining the concepts. Participants in the *ambition-fluent condition* spent 5 min defining the concept *ambition*, whereas participants in the *intelligence-fluent condition* spent 5 min defining the concept *intelligence*. Participants were encouraged to write a lot, to use examples, to provide multiple definitions, and to keep writing and thinking until the experimenter asked them to stop. After 5 min of writing, participants were stopped, given a bogus oral and written debriefing that merely reiterated our cover story, and escorted out of the cubicle.

Measuring Confidence

The second experimenter was a man who led participants to a laboratory room on a different floor of the psychology building. This was done to maximize the apparent independence of the two studies. After obtaining informed consent, the second experimenter introduced participants to a study of the accuracy of first impressions. Participants learned that they would view a target person on videotape for 3 min and then attempt to predict his responses to some personality questionnaire items. All participants watched the same male target discuss his background, describe his hypothetical reactions to situations, and answer some attitude items (as in the high-information conditions of Experiments 1 and 2).

After watching the videotape, participants received a personality questionnaire that asked them to predict how the target person would rate himself on seven items. The items on the personality questionnaire included artistic ability, ambition, decisiveness, sociability, liberalism, patience, and intelligence. Each item was rated on a 7-point scale anchored with the phrases *far below average* and *far above average*. The primed concept (i.e., ambition or intelligence) was always in the second position on the questionnaire, whereas the unprimed concept (also ambition or intelligence) was in the seventh position. After making each prediction, participants filled in a blank in response to a confidence item that asked "On a scale ranging from 0% to 100%, how confident are you in the accuracy of the prediction you just made?"

After completing their ratings, participants were told that the experiment was over and were given a bogus debriefing that reiterated our interest in the accuracy of first impressions. Next, the experimenter offhandedly presented a final questionnaire and stated that "the psychology department is interested in whether or not people's responses in an experiment are affected when they have participated in another experiment immediately beforehand." The questionnaire asked participants whether they thought that the first study had any influence on their responses during the second study. Participants merely circled "yes" or "no" on the questionnaire. Participants who circled "yes" were asked to explain their answer. Because priming effects do not occur when participants are conscious of the link between a prime and a subsequently encountered stimulus (Fiske & Taylor, 1991), we deleted the data of the 6 participants who answered "yes" to this item and who explained their answer by noting that they rated the target in the second study on a dimension that they had been asked to define in the first study.

Results and Discussion

Was the target rated with greater confidence when judgment fluency had been increased by our priming manipulation? We conducted a 2 (fluent judgment: ambition and intelligence) \times 2 (confidence rating: ambition and intelligence) between-within ANOVA. This analysis revealed only the predicted interaction, $F(1, 44) = 15.1, p < .001$. As can be seen in Figure 3, participants in the ambition-fluent condition reported greater confidence when rating the target's ambition than when rating the target's intelligence, $F(1, 44) = 11.53, p < .002$, whereas participants in the intelligence-fluent condition reported greater

confidence when rating the target's intelligence than when rating the target's ambition, $F(1, 44) = 4.53, p < .04$.⁵

We then examined whether our confidence effects were independent of the extremity of participants' ratings of targets. First, we computed extremity scores for each participant by taking the absolute value of the extent to which his or her ambition and intelligence ratings deviated from the midpoint of the rating scale. We conducted a 2 (fluent concept: ambition and intelligence) \times 2 (response extremity: ambition and intelligence) between-within ANOVA that revealed no main effects and no interaction ($F_s < 1$). Thus, our manipulation of fluency did not affect response extremity, suggesting that the impact of fluency on confidence occurs independent of any changes in the extremity of participants' ratings of the target.

Our findings support the view that judgment fluency has a causal impact on confidence. Participants made relatively confident personality judgments when the fluency of those judgments had been heightened by a priming manipulation. Taken together, the results of Experiments 1, 2, and 3 support the idea that rich representations produce confidence (but not necessarily accuracy), and that this increased confidence is mediated (at least in part) by judgment fluency.

Our findings are also consistent, however, with the possibility that representational richness fosters confidence and that confidence enhances judgment fluency. If it is true that confidence enhances fluency, this does not rule out the possibility that fluency also fosters confidence (e.g., they might have reciprocal effects). Indeed, several previous studies support the notion that fluency produces confidence (e.g., Kelley & Lindsay, 1993; Koriat, 1993; Nelson & Narens, 1990; Shaw, 1996). Thus, even if confidence does enhance fluency (which has not been demonstrated), substantial data (including this experiment) attest to the fact that fluency fosters confidence. Having provided evidence supportive of the proposed relations among representational richness, judgment fluency, and confidence, we turn in Experiment 4 to the issue of what types of information people will use to enrich their impressions.

EXPERIMENT 4: THE MODERATING ROLE OF INFORMATION CONTENT

Earlier in this article, we described a continuum of information types ranging from irrelevant through pseudorelevant to relevant. We also theorized that whereas both pseudorelevant and relevant information would foster confidence, irrelevant information would not. In Experiment 4 we tested this prediction.

⁵ Because the primed concept was always in the second position and the unprimed concept was always in the seventh position, these results may have occurred simply because people are more confident on items early in a series than late in a series. To rule out this alternative explanation, we examined the linear trend of confidence across items. We found a weak tendency for confidence to increase across items in the ambition-prime condition, $F(1, 23) = 3.0, p = .10$, and no trend across items in the intelligence-prime condition, $F(1, 21) = 1.7, p = .21$. These results argue against the idea that our results are attributable to the tendency for confidence to decrease across items.

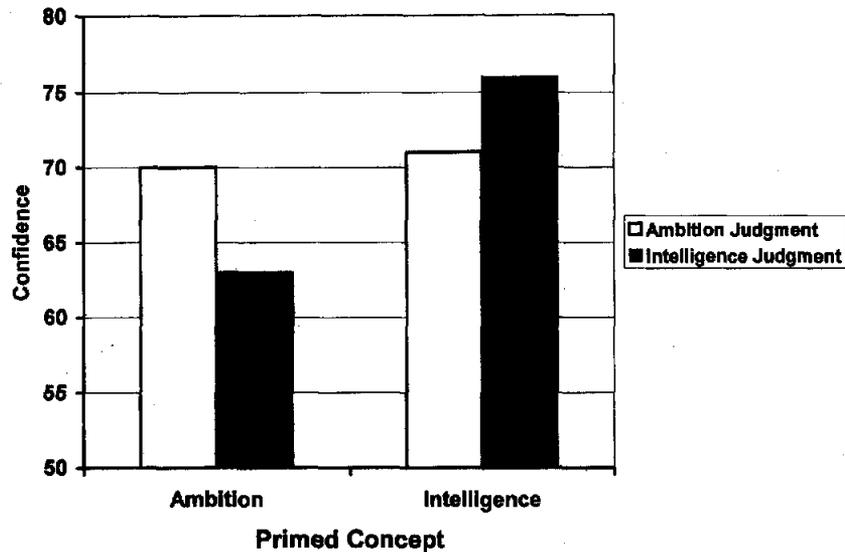


Figure 3. Confidence as a function of judgment dimension and priming manipulation. Confidence could range from 0% to 100%.

Method

Participants

Participants were 64 male and 52 female introductory psychology students at the University of Texas at Austin who received course credit.

Procedure

Participants arrived in groups of 4–8 and were escorted by a male experimenter to separate experimental cubicles. Once there, participants were seated in front of a computer and introduced to a study of “how people form impressions of one another.” All instructions were presented by computer while the experimenter waited in the hallway outside the cubicles. Participants learned that they would rate their impression of a woman named Lisa, who was ostensibly a student at the University of Texas. Next, all participants read that Lisa was an 18-year-old first-year student at the University of Texas who had not yet declared her major. After reading this information, participants made initial ratings of the extent to which they thought that Lisa was mature, artistic, kind, traditional, and intelligent (on 5-point scales with endpoints labeled *less than the average UT [University of Texas] student* and *more than the average UT student* and the midpoint labeled *same as the average UT student*). They also indicated how confident they were of their impressions on 9-point scales with scale points labeled from 20% through 100% in intervals of 10%.

The critical manipulation in this study was the type of information that participants received about Lisa before reporting their impressions and confidence for a second time. Participants were randomly assigned to one of four information conditions. Participants in the control condition were asked to “wait for 30 s while the computer records your ratings of Lisa.” In contrast, participants in the remaining three conditions were asked to examine presented information for 30 s. Participants in the *irrelevant-information condition* were shown six items of information about Lisa that objective judges had rated as being nondiagnostic with respect to Lisa’s kindness and intelligence. Some of this information was taken from Hilton and Fein’s (1989) “clearly irrelevant” category of information and included items such as “She found 20¢ in a pay phone at the Union when she went to make a phone call” and “Lisa lives within walking distance of the UT campus.” Participants in the

pseudorelevant-information condition were shown six items of information that objective judges had rated as nondiagnostic of intelligence and kindness but that provided some basis for personality inferences. Some of this information was based on Hilton and Fein’s pseudorelevant category of information and included items such as “Although she has not yet declared a major, Lisa is interested in psychology and biology” and “Lisa enjoys renting movies with her friends, especially old movies and mysteries.” Finally, participants in the *relevant-information condition* received six items of information, three of which objective judges had rated as being diagnostic of kindness and three of which the judges had rated as being diagnostic of intelligence. These items included, for example, “Over the summer, she volunteered for a group that constructed houses for disadvantaged people” and “Lisa is interested in physics and mathematics, having excelled at those subjects in high school.”⁶ After this manipulation, participants once again rated the extent to which they thought Lisa was mature, artistic, kind, traditional, and intelligent and rated their confidence in the accuracy of their impressions.

Results and Discussion

Our key prediction was that confidence would increase when participants received either pseudorelevant or relevant information, but not when they received no information or irrelevant

⁶ Our objective judges were seven graduate students at the University of Texas who rated the diagnosticity of all our stimulus information. Specifically, they were told that each item of information pertained to a hypothetical UT student. They then read each item of information and rated both the kindness and intelligence of the person described by the information on 5-point scales with endpoints labeled *less than the average UT student* and *more than the average UT student*, and the midpoint labeled *same as the average UT student*. The mean ratings for the irrelevant and pseudorelevant information were 3.0, suggesting that the information in those categories was perceived as nondiagnostic of kindness or intelligence by our raters. In contrast, the average rating of relevant information (on the appropriate trait) was 4.2. Ratings of relevant information were significantly higher than ratings of both irrelevant and pseudorelevant information, $F_s(1, 6) > 53.4$, $p_s < .001$.

information. Because participants in the relevant condition received information that was diagnostic of kindness and intelligence, we analyzed their confidence on those traits.

We submitted confidence ratings for kindness and intelligence to a 4 (information condition: control, irrelevant, pseudorelevant, and relevant) \times 2 (confidence rating: Time 1 and Time 2) between-within ANOVA and conducted planned comparisons of Time 1 versus Time 2 confidence ratings within each information condition. On the dimension of kindness, planned comparisons within the control and irrelevant conditions revealed no change in confidence from Time 1 to Time 2 ($F_s < 1.03$). Confidence did increase, however, both within the pseudorelevant condition, $F(1, 111) = 31.0, p < .001$ ($M_s = 2.3$ and 3.6 for Time 1 and Time 2 ratings, respectively), and within the relevant condition, $F(1, 111) = 129.7, p < .001$ ($M_s = 2.1$ and 5.0 for Time 1 and Time 2 ratings, respectively). On the dimension of intelligence, confidence did not increase from Time 1 to Time 2 in the control and irrelevant conditions ($F_s < 1$). Confidence did increase, however, both within the pseudorelevant condition, $F(1, 111) = 5.8, p < .02$ ($M_s = 3.4$ and 4.0 for Time 1 and Time 2, respectively), and within the relevant condition, $F(1, 111) = 126.5, p < .001$ ($M_s = 2.7$ and 5.3 for Time 1 and Time 2, respectively).

Another test of the impact of pseudorelevant information on confidence involved examining people's confidence in their ratings of the extent to which Lisa was mature, artistic, and traditional. Because participants did not receive information that bore directly on these aspects of Lisa's personality (even in the relevant-information condition), both the pseudorelevant and relevant conditions provided a test of the impact of pseudorelevant information on confidence. Notably, planned comparisons revealed that confidence always increased from Time 1 to Time 2 in the both the pseudorelevant and relevant conditions on all three traits, $F_s(1, 111 \text{ or } 112) > 5.5, p_s < .03$ (the degrees of freedom across traits fluctuated because of missing data). There were no statistically significant increases in confidence in the control and irrelevant conditions. This bolstered our claim that information need not have clear implications for a trait to boost confidence in judgments relating to the trait. Taken together, the results of this study support our prediction that relevant information is not the only information that fosters confidence: Confidence is also fostered by pseudorelevant information despite the fact that pseudorelevant information will not systematically foster accuracy.

We next examined whether our confidence effects were independent of the extremity of participants' ratings of targets. First, we computed extremity scores for each participant by finding the absolute value of the extent to which his or her trait ratings deviated from the midpoint of the rating scale at both Time 1 and Time 2. For each of our five traits, we conducted a 4 (information condition: control, irrelevant, pseudorelevant, and relevant) \times 2 (response extremity: Time 1 and Time 2) between-within ANOVA and conducted planned comparisons of Time 1 versus Time 2 response extremity within the pseudorelevant and relevant conditions. Within the pseudorelevant condition, there was not a statistically significant increase in response extremity from Time 1 to Time 2 on any of our five traits. Indeed, the largest change in response extremity was found when we examined participants' judgments of the extent to which Lisa was traditional and suggested that extremity decreased from Time 1

to Time 2 ($M_s = 0.80$ and 0.63 at Time 1 and Time 2, respectively), $F(1, 111) = 2.4, p = .13$ (all other $F_s < 1.9, p_s > .17$). This suggests that, within the pseudorelevant condition, our confidence effects occurred independent of any increase in the extremity of participants' ratings. Also, the fact that participants' ratings did not become more extreme when they received pseudorelevant information corroborates our claim that such information has no necessary implication for judgments and hence is unlikely to enhance accuracy. Within the relevant condition, we found no change in response extremity from Time 1 to Time 2 when we examined participants' ratings of the extent to which Lisa was artistic ($F < 1$). As in the pseudorelevant condition, we found a decrease in response extremity from Time 1 to Time 2 when we examined participants' ratings of the extent to which Lisa was traditional ($M_s = 0.66$ and 0.41 at Time 1 and Time 2, respectively), $F(1, 112) = 4.9, p < .03$. Finally, we found an increase in response extremity from Time 1 to Time 2 when we examined participants' ratings of the extent to which Lisa was kind, intelligent, and mature, $F_s(1, 111 \text{ or } 112) > 15.1, p_s < .001$. Having found an increase in extremity on these three dimensions, we conducted three ANCOVAs to determine whether relevant information affected confidence above and beyond its effect on response extremity. Specifically, we conducted 4 (information condition: none, irrelevant, pseudorelevant, and relevant) \times 2 (confidence rating: Time 1 and Time 2) between-within ANCOVAs with response extremity covaried out of confidence and conducted planned comparisons within the relevant information condition. These comparisons revealed that, on all three traits on which extremity increased from Time 1 to Time 2, the impact of relevant information on confidence was significant even when extremity was held constant, $F_s(1, 109 \text{ or } 110) > 40.1, p_s < .001$. These analyses suggested that, within the relevant information condition, changes in confidence were not an artifact of changes in the extremity of participants' ratings.

EXPERIMENT 5: RECONCILING DIVERGENT VIEWS OF THE GENESIS OF CONFIDENCE

We have presented evidence for an account in which confidence stems from the richness of people's mental representations. Of course, other researchers have presented different accounts of the genesis of confidence. One of the best known of these accounts is that of Leyens, Yzerbyt, and Schadron's (1994) *social judgability theory*. The theory holds that, when making social judgments, people pay attention to the extent to which their judgments follow cultural norms, preserve the integrity of the self, and are congruent with their personal theories about the world. For example, Leyens et al. argued that many people in Western cultures have a belief that stereotypes are not a valid basis on which to judge an individual, and thus such people report low levels of confidence when asked to make judgments in the absence of individuating information about a person (Yzerbyt, Schadron, Leyens, & Rocher, 1994).

At first blush, social judgability theory may seem to offer an alternative to our account of the genesis of confidence. We propose that, in many cases, however, social judgability and representational richness explanations can peacefully coexist. For example, had we asked people to explain why they reported high levels of confidence in our first experiment, they might

have reported a culturally acceptable response such as "Because I learned something from watching the interview." Finding that people both report and behave consistently with such cultural beliefs does not undermine our explanation of confidence; rather, it represents a view of confidence from a different level of analysis. Whereas the Leyens et al. (1994) analysis focuses on people's (presumably conscious) beliefs about when confidence is warranted, our work focuses on underlying cognitive processes of which people might not be especially aware. For example, although we believe that richness and fluency foster confidence, it seems perfectly reasonable to suppose that people may not think about confidence in this manner but may simply hold (and be able to report) the belief that confidence is warranted after one has "learned something" about a target. Of course, people who report having "learned something" probably have richer, more accessible impressions than those who report having learned nothing. These two perspectives on confidence seem perfectly harmonious to us.

There is one aspect of social judgment theory that suggests an alternative explanation for some of our findings, however. Specifically, Yzerbyt et al. (1994) suggested that metainformational cues are sufficient to produce confidence. Metainformational cues are content-free cues that cause people to believe that they have been informed, even when they have not received actual information. Yzerbyt et al. showed that people who were told that they received "subliminal information" about a target person (even though they had actually received no information) reported higher levels of confidence in their judgments of the target than did people who were not told that they had received subliminal information. A strong version of the metainformational view of confidence (which Yzerbyt et al. did not specifically endorse) would suggest that confidence never depends on the receipt of information that contributes to rich representations but depends only on people's belief that they have received information.

We question this strong version of the metainformational view and suggest that the metainformational and representational richness accounts of confidence can be combined to form a more comprehensive account of confidence. Specifically, we propose that both metainformation and representational richness contribute to confidence and that representational richness produces confidence in excess of that which would be produced by metainformation alone. (Because information necessarily contains metainformation, we do not make the complementary claim that metainformation produces confidence in excess of that which would be produced by representational richness alone.) Outside the laboratory, this contention would be difficult to examine because actual information about people (which contributes to representational richness) and metainformational cues are confounded—information contains metainformation and vice versa. Our claim can be tested in the laboratory, however, by presenting people with either no metainformation and no actual information about a target, only metainformation, or both metainformation and actual information.

We conducted such a test in an experiment billed as an investigation of the impact of auditory experiences on first impressions. Participants learned that, after listening to an audiotape, they would report their impression of a female volunteer between the ages of 18 and 22. In one condition, participants heard a tape containing no actual information about the target and

were given no reason to expect that the tape would enable them to form an accurate impression. In a second condition, participants heard a tape containing no actual information about the target but were told that the tape contained subliminally presented information about the target. Finally, in a third condition, participants heard a tape containing actual information about the target after being told that the tape contained information about the target. All participants predicted the target's responses to a sexual history and a self-concept questionnaire and reported their confidence in those predictions. We expected that confidence would increase incrementally across our three conditions.

Method

Participants

Participants were 37 male introductory psychology students at the University of Texas at Austin who received course credit.

Procedure

A male experimenter escorted individual participants to a cubicle, obtained informed consent, and explained that the participant would be taking part in a study of "the impact of auditory experiences on first impressions." Participants learned that they would listen to an audiotape and then attempt to predict how a female volunteer answered questions about her sexual history and self-concept. Participants also learned that the female volunteer was between the ages of 18 and 22.

Manipulating Metainformation and Actual Information

Participants were randomly assigned to one of three conditions: (a) a control condition, in which they received neither metainformation nor actual information about the target; (b) a *metainformation condition*, in which they received metainformation but no actual information; and (c) an *informed condition*, in which they received both metainformation and actual information about the target. Participants in the control condition listened to an audiotape comprising excerpts from *The Prophet* by Kahlil Gibran. Before hearing a tape of *The Prophet* that contained quiet, incomprehensible male and female voices in the background, participants in the metainformation condition learned that the tape contained a subliminally presented interview revealing "personality and background information about the target person." There was, in fact, no subliminally presented information on the tape (see Yzerbyt et al., 1994). Finally, participants in the informed condition listened to an actual interview with the target person after learning that they would hear an interview containing "personality and background information about the target person." The tape used in the informed condition was the audio portion of one of the videotapes used in Experiments 1 and 2. All audiotapes were the same length (slightly longer than 3 min).

Dependent Measures

Impression of the target. After hearing the audiotapes, participants predicted the target's responses to the SHQ and SAQ. As in Experiments 1 and 2, an index of accuracy was computed for each participant on both the SHQ and SAQ using Pearson product-moment correlations. Accuracy ranged from .47 to .95 on the SHQ and from -.64 to .89 on the SAQ. The mean levels of accuracy were .77 and .36 for the SHQ and SAQ, respectively.

Confidence. As in Experiment 1, participants indicated how confident they were of each of their predictions on the SHQ and provided one confidence rating for their predictions on the SAQ. The 10 confidence items from the SHQ ($\alpha = .91$) were combined to form an index

of confidence, whereas the confidence item from the SAQ was treated as a separate index.

Perception of metainformation. Participants answered two questions designed to tap the extent to which metainformation was perceived in each condition. Specifically, they rated how helpful the audiotape was in enabling them to form an accurate impression of both the target's sexual history and self-concept (on 7-point scales anchored with the phrases *not at all helpful* and *extremely helpful*).

Results and Discussion

Confidence

The mean confidence values for ratings of the target (collapsed across the SHQ and SAQ) were 38%, 51%, and 70% for participants in the control, metainformation, and informed conditions, respectively. As predicted, planned comparisons revealed statistically significant differences between confidence ratings in the control and metainformation conditions, $t(32) = 2.65, p < .02$, as well as between the metainformation and informed conditions, $t(32) = 3.6, p < .002$. These results support the idea that both metainformation and actual information contribute to confidence such that information has effects beyond those attributable to metainformation alone.

Next, we sought to test whether our confidence effects were independent of the extremity of participants' ratings of targets. First, we computed extremity scores for each participant on both the SHQ and the SAQ. We did this by calculating the absolute value of the extent to which the participant deviated from the midpoint of the rating scale on each item of the SHQ or SAQ and then taking the average of these deviations. When we computed one-way ANOVAs on these extremity scores comparing the control, metainformation, and informed conditions, we found that SHQ extremity was unrelated to our manipulations ($F < 1$), whereas SAQ extremity differed by condition, $F(2, 34) = 6.0, p < .007$. Planned comparisons indicated that there was no difference in extremity between the control and metainformation conditions ($t < 1$), whereas the metainformation condition ($M = 1.4$) was associated with less extremity than the informed condition ($M = 1.9$), $t(34) = 2.3, p < .03$. Importantly, when we reanalyzed SAQ confidence with SAQ extremity included as a covariate, we found that the differences between the control and metainformation conditions, $F(1, 33) = 3.2, p < .09$, and between the metainformation and informed conditions, $F(1, 33) = 7.9, p < .009$, were still evident. Thus, once again, our confidence effects seemed to be independent of any change in the extremity of participants' ratings.

Perceptions of Metainformational Cues

We reasoned that the strength of the metainformational cues in the metainformation and informed conditions would be equal. Our reasoning was based on the fact that both groups learned that they would hear "personality and background information about the target person" before listening to a 3-min audiotape. In contrast, we expected metainformational cues to be much weaker in the control condition because controls had no reason to believe that their audiotape contained information about the target. To test our reasoning, we averaged participants' ratings of the helpfulness of the audiotapes for forming accurate impressions of both sexual history and self-concept, and computed

planned comparisons on these averages. These comparisons revealed that participants in the control condition ($M = 1.3$) rated the audiotape as less helpful for forming an accurate impression of the target than did participants in both the metainformation condition ($M = 3.8$), $t(32) = 5.0, p < .001$, and the informed condition ($M = 3.7$), $t(32) = 4.9, p < .001$. Participants in the metainformation and informed conditions did not differ, $t(32) = -0.22, p = .83$. This suggests that the confidence difference between participants in the informed condition and participants in the metainformation condition was not caused by the presence of a stronger metainformational cue in the informed condition.

Accuracy

After conducting a Fisher r -to- Z transformation of our accuracy correlations, we found that SHQ accuracy did not differ by condition ($F < 1$). Although SAQ accuracy did differ by condition, $F(2, 34) = 4.3, p < .03$, the pattern of accuracy across our three conditions did not mirror the pattern of confidence. We compared the levels of accuracy in our three conditions using contrasts so that accuracy differences would be tested with the same degree of power as were confidence differences. These follow-up analyses revealed that participants in the metainformation condition (average $r = .10$) were significantly less accurate than participants in the informed condition (average $r = .59$), $t(34) = 2.9, p < .008$, but also slightly less accurate than participants in the control condition (average $r = .37$), $t(34) = 1.8, p = .08$. The difference between controls and participants in the informed condition did not approach significance, $t(34) = 1.2, p = .25$.

To confirm that the divergent impact of our manipulations on confidence and accuracy were not an artifact of the particular manner in which we operationalized accuracy, we computed a second index of accuracy. Specifically, we computed the difference between perceivers' predictions and targets' responses on each item of the SHQ and SAQ. We then averaged these numbers to create an index of the average deviation between perceiver and target responses on both the SHQ and SAQ. These deviation score measures of accuracy did not differ by condition when we examined either the SHQ or the SAQ ($F_s < 1$).

GENERAL DISCUSSION

The results of our experiments suggest that rich representations of others make people confident of those representations independent of any gains in accuracy. Furthermore, our evidence also suggests that the impact of richness on confidence is mediated by judgment fluency. (See Swann & Gill, 1997, Study 3, for an additional correlational test of these processes in the context of a single study.) The data also indicate that representational richness is increased by pseudorelevant as well as relevant information. Finally, representational richness was shown to affect confidence above and beyond the impact of metainformational cues.

One implication of our findings regarding the relations among richness, fluency, and confidence is that we should expect people in long-lasting or highly interdependent relationships (who presumably have relatively rich impressions of their partners) to have considerable confidence in their beliefs about their partners regardless of whether those beliefs are accurate. Indeed, Swann

and Gill (1997) have reported evidence that is consistent with this hypothesis: They found that relationship length and relationship involvement predict the confidence (but not the accuracy) of people's beliefs about their relationship partners. Moreover, consistent with the present account of confidence, Swann and Gill discovered that the effects of length and involvement on confidence were mediated by representational richness.

Our work also has important implications for the ongoing debate over the degree to which the person perception process is fraught with error. Some critics (e.g., May, 1989) contend that the errors and biases enumerated by psychologists are inconsequential because they are made with little confidence. The present work suggests that when errors are made, perceivers will not necessarily remain unconfident of them. Rather, errors may become associated a high degree of confidence if the representation underlying the error is especially rich. Thus, although people who receive little information about a target may make erroneous inferences with little confidence, those who receive a great deal of nondiagnostic evidence may become extremely confident of impressions that are highly erroneous.

Understanding the genesis of confidence may also be useful to people interested in reducing stereotyping and prejudice, as lowering people's confidence in stereotypes should decrease the prevalence of behaviors guided by those stereotypes. To date, research examining the impact of stereotype-inconsistent individuals on people's stereotypes has suggested that such individuals often fail to change the content of people's stereotypes but instead cause people to use "subcategories" to explain the inconsistent information (e.g., Weber & Crocker, 1983). We suggest that although stereotype *content* is often unaffected by inconsistent information, stereotype *confidence* might be lowered by such information. Presumably, stereotype-inconsistent information would reduce the integration of people's representations of a stereotyped group, thus reducing confidence in the stereotype. Alas, this prediction must remain for future researchers to test because prior research involving stereotype-inconsistent information has focused on stereotype content rather than stereotype confidence (see Hewstone, 1996, for a review).

Our work also fits nicely with much work in the literature of cognitive psychology. For example, several theorists in cognitive psychology have posited a link between fluency and confidence (Benjamin & Bjork, 1996; Jacoby et al., 1994; Kelley & Lindsay, 1993; Koriat, 1993; Nelson & Narens, 1990). Furthermore, the concept of representational richness ties together aspects of the social and cognitive literature on confidence. Manipulations that have produced confidence in that literature include incremental exposure to case history information (Oskamp, 1965), relationship length and relationship involvement (Swann & Gill, 1997), experience in an academic domain (Glenberg & Epstein, 1987), increments of information about a hypothetical baseball team (Peterson & Pitz, 1988), practice and effort (Paese & Sniezek, 1991), interdependence between individuals (Berscheid, Graziano, Monson, & Dermer, 1976), and the number of letters people retrieved on being asked to recall a letter string (Koriat, 1993). All these findings seem amenable to interpretation in terms of representational richness despite the fact that some of them were originally explained (implicitly) in terms of metainformational cues (e.g., Peterson & Pitz, 1988). Given our evidence about the importance of both metainformation and representational richness for producing confidence, we believe

that researchers interested in confidence should attempt to uncover which source of confidence is operative under various sets of circumstances and should test whether the confidence fostered by each process produces the same consequences. For example, does confidence affect the likelihood that people will translate beliefs into behaviors regardless of the source of the confidence?

CODA

Socrates defined *wisdom* as the ability to recognize what one knows and what one does not know. Our data suggest a potential limit on the extent to which people are capable of achieving the Socratic ideal of wisdom. Results of our studies suggest that people's judgments about the accuracy of their knowledge (i.e., their confidence judgments) are consistently influenced by a factor (i.e., representational richness) that is, at best, inconsistently related to accuracy. These results support the notion that people's judgments about the accuracy of their knowledge have no necessary connection to their true level of accuracy. Yet, confidence and accuracy are sometimes related. Indeed, they are modestly correlated in some instances (e.g., Dunning et al., 1990) and highly correlated in others (e.g., Koriat, 1993). Given this variability in people's ability to judge the state of their knowledge, the task of future researchers should be to identify the conditions under which people know what they know and the conditions under which they do not. We hope that the framework we have provided for understanding the genesis of confidence may assist researchers in this endeavor.

References

- Anderson, J. R. (1976). *Language, memory, and thought*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R., & Bower, G. H. (1973). *Human associative memory*. Washington, DC: Winston.
- Bargh, J. A., & Pietromonaco, P. (1982). Automatic information processing and social perception: The influence of trait information presented outside of conscious awareness on impression formation. *Journal of Personality and Social Psychology*, 43, 437-449.
- Benjamin, A. S., & Bjork, R. A. (1996). Retrieval fluency as a metacognitive index. In L. M. Reder (Ed.), *Implicit memory and metacognition* (pp. 309-338). Mahwah, NJ: Erlbaum.
- Berger, I. E., & Mitchell, A. A. (1989). The effect of advertising on attitude accessibility, attitude confidence, and the attitude-behavior relationship. *Journal of Consumer Research*, 16, 269-279.
- Berscheid, E., Graziano, W., Monson, T., & Dermer, M. (1976). Outcome dependency: Attention, attribution, and attraction. *Journal of Personality and Social Psychology*, 34, 978-989.
- Darley, J. M., & Gross, P. H. (1983). A hypothesis-confirming bias in labeling effects. *Journal of Personality and Social Psychology*, 44, 20-33.
- DePaulo, B. M., Stone, J. L., & Lassiter, G. D. (1985). Deceiving and detecting deceit. In B. R. Schlenker (Ed.), *Self and social life* (pp. 323-370). New York: McGraw-Hill.
- Dion, K. K., Berscheid, E., & Walster, E. (1972). What is beautiful is good. *Journal of Personality and Social Psychology*, 24, 285-290.
- Dunning, D., Griffin, D. W., Milojkovic, J. D., & Ross, L. (1990). The overconfidence effect in social prediction. *Journal of Personality and Social Psychology*, 58, 568-581.
- Fazio, R. H., & Zanna, M. P. (1978). Attitude qualities relating to the strength of the attitude-behavior relationship. *Journal of Experimental Social Psychology*, 14, 398-408.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression for-

- mation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. *Advances in Experimental Social Psychology*, 23, 1-74.
- Fiske, S. T., & Taylor, S. E. (1991). *Social cognition* (2nd ed.). New York: McGraw-Hill.
- Funder, D. C. (1995). On the accuracy of personality judgment: A realistic approach. *Psychological Review*, 102, 652-671.
- Glenberg, A. M., & Epstein, W. (1987). Inexpert calibration of comprehension. *Memory & Cognition*, 151, 84-93.
- Hewstone, M. (1996). Contact and categorization: Social psychological interventions to change intergroup relations. In C. N. Macrae, C. Stangor, & M. Hewstone (Eds.), *Stereotypes and stereotyping* (pp. 323-368). New York: Guilford Press.
- Higgins, E. T. (1989). Knowledge accessibility and activation: Subjectivity and suffering from unconscious sources. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 75-123). New York: Guilford Press.
- Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology*, 13, 141-154.
- Hilton, J. L., & Fein, S. (1989). The role of typical diagnosticity in stereotype-based judgments. *Journal of Personality and Social Psychology*, 57, 201-211.
- Jacoby, L. L., Bjork, R. A., & Kelley, C. M. (1994). Illusions of comprehension, competence, and remembering. In D. Druckman & R. A. Bjork (Eds.), *Learning, remembering, believing: Enhancing human performance* (pp. 57-80). Washington, DC: National Academy Press.
- Jones, E. E., & Davis, K. E. (1965). From acts to dispositions: The attribution process in person perception. *Advances in Experimental Social Psychology*, 2, 219-266.
- Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social Psychology*, 3, 1-24.
- Kelley, C. M., & Lindsay, D. S. (1993). Remembering mistaken as knowing: Ease of retrieval as a basis for confidence in answers to general knowledge questions. *Journal of Memory and Language*, 32, 1-24.
- Kenny, D. A. (1991). A general model of consensus and accuracy in interpersonal perception. *Psychological Review*, 98, 155-163.
- Koriat, A. (1993). How do we know that we know? The accessibility model of the feeling of knowing. *Psychological Review*, 100, 609-639.
- Koriat, A., Lichtenstein, S., & Fischhoff, B. (1980). Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 107-118.
- Kruglanski, A. W. (1989). The psychology of being "right": On the problem of accuracy in social perception and cognition. *Psychological Bulletin*, 106, 395-409.
- Leyens, J.-P., Yzerbyt, V. Y., & Schadron, G. (1994). *Stereotypes and social cognition*. London: Sage.
- May, R. S. (1989). Overconfidence in overconfidence. In A. Chikan, J. Kindler, & I. Kiss (Eds.), *Proceedings of the Fourth FUR Conference*. Dordrecht, the Netherlands: Kluwer Academic.
- Metzler, C. W., Noell, J., & Biglan, A. (1992). The validation of a construct of high-risk sexual behavior in heterosexual adolescents. *Journal of Adolescent Research*, 7, 233-249.
- Murray, S. L., & Holmes, J. G. (1993). Seeing virtues in faults: Negativity and the transformation of interpersonal narratives in close relationships. *Journal of Personality and Social Psychology*, 65, 707-722.
- Myers, J. L., O'Brien, E. J., Balota, D. A., & Toyofuku, M. L. (1984). Memory search without interference: The role of integration. *Cognitive Psychology*, 16, 217-242.
- Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. In G. Bower (Ed.), *The psychology of learning and motivation* (Vol. 26, pp. 125-141). San Diego, CA: Academic Press.
- Nisbett, R. E., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice Hall.
- Oskamp, S. (1965). Overconfidence in case-study judgments. *Journal of Consulting Psychology*, 29, 261-265.
- Ostrom, T. M., Lingle, J. H., Pryor, J. B., & Geva, N. (1980). Cognitive organization of person impressions. In R. Hastie, T. M. Ostrom, E. B. Ebbesen, R. S. Wyer, D. L. Hamilton, & D. E. Carlston (Eds.), *Person memory: The cognitive basis of social perception* (pp. 121-153). Hillsdale, NJ: Erlbaum.
- Paese, P. W., & Sniezek, J. A. (1991). Influences on the appropriateness of confidence in judgment: Practice, effort, information, and decision-making. *Organizational Behavior and Human Decision Processes*, 48, 100-130.
- Park, B. (1989). Traits attributes as on-line organizers in person impressions. In J. N. Bassili (Ed.), *On-line cognition in person perception* (pp. 39-59). Hillsdale, NJ: Erlbaum.
- Pelham, B. W., & Swann, W. B., Jr. (1989). From self-conceptions to self-worth: On the sources and structure of global self-esteem. *Journal of Personality and Social Psychology*, 57, 672-680.
- Peterson, D. K., & Pitz, G. F. (1988). Confidence, uncertainty, and the use of information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 85-92.
- Pieters, R., & Verplanken, B. (1995). Intention-behavior consistency: Effects of consideration set size, involvement, and need for cognition. *European Journal of Social Psychology*, 25, 531-543.
- Rosenhan, D. L. (1973). On being sane in insane places. *Science*, 179, 250-258.
- Ross, L., Amabile, T. M., & Steinmetz, J. L. (1977). Social roles, social control, and biases in social-perception processes. *Journal of Personality and Social Psychology*, 35, 485-494.
- Shaw, J. (1996). Increases in eyewitness confidence resulting from postevent questioning. *Journal of Experimental Psychology: Applied*, 2, 126-146.
- Sherman, J. W., & Klein, S. B. (1994). Development and representation of personality impressions. *Journal of Personality and Social Psychology*, 67, 972-983.
- Shoben, E. J. (1982). Semantic and lexical decisions. In C. Richard Puff (Ed.), *Handbook of research methods in human memory and cognition* (pp. 287-314). New York: Academic Press.
- Smith, E. E., Adams, N. A., & Schorr, D. (1978). Fact retrieval and the paradox of interference. *Cognitive Psychology*, 10, 438-464.
- Strull, T. K., & Wyer, R. S., Jr. (1980). Category accessibility and social perception: Some implications for the study of person memory and interpersonal judgments. *Journal of Personality and Social Psychology*, 38, 841-856.
- Swann, W. B., Jr., & Ely, R. M. (1984). The battle of wills: Self-verification vs. behavioral confirmation. *Journal of Personality and Social Psychology*, 46, 1287-1302.
- Swann, W. B., Jr., & Gill, M. J. (1997). Confidence and accuracy in person perception: Do we know what we think we know about our relationship partners? *Journal of Personality and Social Psychology*, 73, 747-757.
- Weber, R., & Crocker, J. (1983). Cognitive processes in the revision of stereotypic beliefs. *Journal of Personality and Social Psychology*, 45, 961-977.
- Wells, G. L., & Murray, D. M. (1984). Eyewitness confidence. In G. Wells & E. Loftus (Eds.), *Eyewitness testimony: Psychological perspectives* (pp. 155-170). New York: Cambridge University Press.
- Yzerbyt, V. Y., Schadron, G., Leyens, J. P., & Rocher, S. (1994). Social judgeability: The impact of meta-informational cues on the use of stereotypes. *Journal of Personality and Social Psychology*, 66, 48-55.

Received July 7, 1997

Revision received June 4, 1998

Accepted June 22, 1998 ■