Journal of Experimental Social Psychology **37,** 373–385 (2001)
doi:10.1006/jesp.2000.1458, available online at http://www.idealibrary.com on IDEAL®

Overconfidence as Dissonance Reduction

Hart Blanton

University at Albany, State University of New York

and

Brett W. Pelham, Tracy DeHart, and Mauricio Carvallo

University at Buffalo, State University of New York

Received April 10, 2000; revised August 25, 2000; accepted August 28, 2000; published online April 3, 2001

People are often more confident than accurate. Past accounts of this *overconfidence effect* have focused on social-cognitive mechanisms, such as the biasing effects of judgmental heuristics and the faulty integration of relevant information. The current studies tested the idea that overconfidence is also a product of dissonance reduction. Specifically, we argue that overconfidence can result from a desire to see the self as knowledgeable and competent. In two studies, the motive to view the self as an accurate perceiver elevated confidence, independent of its effect on accuracy. This effect was diminished by manipulations derived from cognitive dissonance theory. In Study 1, confidence ratings were debiased by an affirmation manipulation designed to boost feelings of self-worth. In Study 2, confidence ratings were debiased by a manipulation designed to lower the aversive implications of feeling uncertain. These findings argue for a motivational perspective on overconfidence.

Doubt is not a pleasant condition, but certainty is absurd.

—Voltaire [François Marie Arouet] (1768)

Accurate knowledge provides many rewards, chief among these being increased abilities to predict, control, and respond to the social world. Feeling confident in the accuracy of one's beliefs and judgments should thus foster a sense of security in the face of both mundane and important decisions. Unfortunately, the comfort gained through such confidence is often unwarranted. A great deal of research indicates that people are often more confident than they are correct. As Fischhoff, Slovic, and Lichtenstein (1977) put it, people are "wrong too often when they are certain that they are right" (p. 561).

This *overconfidence effect* has been demonstrated for a wide variety of judgments. The most common technique for assessing overconfidence involves asking people to answer a number of general knowledge questions and then having them estimate the probability that they have answered each question correctly. If respondents' mean confidence scores are higher than their mean accuracy scores, this is taken as evidence of overconfidence (see Fishhoff, 1982, for review). Another common method of assessing overconfidence involves asking people to evaluate their ability to solve problems in the laboratory. These studies show that people think that they can solve problems that they cannot, think that they have made progress toward correct solutions when they have not, and think that they have drawn correct conclusions when they have not (see Metcalfe, 1999, for a review). Systematic tendencies toward overconfidence occur with each of these methodologies. This is true for both trivial and consequential judgments and for judgments ranging from moderate to high levels of difficulty. Moreover, overconfidence holds up in the face of numerous debiasing techniques (see Fischhoff, 1982; Seiber, 1974).

Most theoretical accounts of overconfidence have focused on social cognitive mechanisms. It has been suggested, for instance, that the use of judgmental heuristics often results in an inability to discriminate inferences from

The order of the first two authors was determined alphabetically. We thank Jim Jaccard for his comments on earlier drafts.

Address correspondence and reprint requests to Hart Blanton, Social Science 112, Department of Psychology, University at Albany, Albany, NY 12222 or to Brett Pelham, Department of Psychology, Park Hall, University at Buffalo, Amherst, NY 14260. E-mail: hblanton@csc.albany.edu or brettpel@acsu.buffalo.edu.

0022-1031/01 $35.00

retrieved memories (e.g., Fischhoff, Slovic, & Lichtenstein, 1977). As an example, people who assess confidence on the basis of the availability heuristic (Tversky & Kahneman, 1974) may confuse easily drawn inferences for easily remembered facts (Block & Harper, 1991; Fischhoff, Slovic, & Lichtenstein, 1977; Fischhoff, 1982; see Nisbett & Wilson, 1977). Other cognitive accounts have implicated the weighing and sampling of information during judgment. For instance, overconfidence can occur when people selectively focus on evidence that is consistent with a conclusion and disregard or ignore evidence that is inconsistent with it (e.g., Koriat et al., 1980; and see Wason, 1968; Klayman & Ha, 1989, for related discussions). Similarly, feelings of confidence can occur when people give greater weight to the evidence for and against a conclusion than they give to the evidence for and against the alternatives to that conclusion (e.g., McKenzie, 1997).

Overconfidence as Dissonance Reduction

In the current study, we explore a motivated account of overconfidence. We argue that overconfidence results at times from a desire to see the self as a competent or accurate perceiver. According to this perspective, undue confidence often arises when uncertainty would challenge valued beliefs about the self as knowledgeable and competent. This perspective is not meant to invalidate or replace explanations based on social-cognitive mechanisms. Instead, it is meant to complement these accounts by showing that cases of overconfidence can result from motivated distortions.

The suggestion that people are motivated to believe that their judgments and beliefs are valid is not new. This idea has a long history in the cognitive dissonance literature on postdecisional regret. Cognitive dissonance theory predicts that, once a decision is made, people will change attitudes that are logically related to the decision to minimize the doubts they would otherwise experience. Most of work in this area has been conducted using variations on Brehm's (1956) "free choice paradigm." Brehm asked participants to rate the attractiveness of several household appliances. Afterward, he gave each participant an unexpected opportunity to choose a free gift from the items they had just rated. Specifically, he had them choose between one of two products they had previously rated as roughly equal in attractiveness. The main finding was that, once a choice between appliances had been made, participants tended to evaluate the chosen appliance more favorably than it had been rated initially.

The tendency for people to enhance their evaluations of chosen objects has been referred to as the "spread of alternatives." This apparently robust tendency has been taken as evidence for a motivated shift in judgment confidence. It is not altogether clear, however, whether changes such as this provide information about overconfidence. The desirability of a consumer item is an inherently subjective matter. It is thus not possible to say whether product reevaluation reflects a move toward more or less accurate evaluations. More direct evidence that postdecisional dissonance reduction can lead to inflated confidence estimates comes from a classic set of studies by Knox and Inkster (1968). Their participants were gamblers placing $2 bets on horses at a race track. Experimenters approached betters 30 s before or a few seconds after they had placed their bets and asked them to report their confidence that the horses they had chosen would win. Consistent with the idea of postdecisional dissonance reduction, gamblers were more confident of their bets after they had just placed them than before. Similar results have been shown with postdecisional increases in confidence by voters in national elections. Much like Knox and Inkster's gamblers, voters have been shown to be more confident that their chosen candidates will win after they have placed their votes than before (Frenkel & Doob, 1976; Regan & Kulduff, 1988).

Of course, many of the classic studies showing postdecisional changes in confidence may result from the effects of nonmotivated mechanisms. For instance, an increase in confidence following a purchase, a wager, or a vote may occur because acts such as these cause people to actively imagine the desired outcome, resulting in an increased perception that it might occur (Anderson, 1983; Sherman, Cialdini, Schwartzman, & Reynolds, 1985). Our point in reviewing classic dissonance studies is not to argue that any one of these makes a definitive case for motivated effects on confidence estimation. We instead point to them to highlight an area of active theorizing about the effects of motivation on confidence. This research tradition seems to have been ignored in the overconfidence literature. In fact, we know of no research that has made a link between dissonance theory and overconfidence for either one's judgments or one's knowledge. In the two studies that follow, we seek to establish such a link by demonstrating that overconfidence is influenced by much the same factors that motivate consumers, gamblers, and voters to feel confident about their choices. We argue that overconfidence reflects the motive to maintain a view of the self as a knowledgeable perceiver who makes sound judgments.

STUDY 1: SELF-AFFIRMATIONS AND OVERCONFIDENCE

If overconfidence derives, in part, from the motive to maintain a positive view of the self, overconfidence should be higher when a perceived lack of knowledge in a specific area would pose a greater threat to the self. In Study 1, we hypothesized that confidence would increase for judgments in which a judgmental error would threaten a preconceived notion that the self is knowledgeable about a specific consumer judgment. We further predicted that this tendency to increase confidence in the face of a potential threat would be

attenuated among people who had recently been made to feel secure in their global self-evaluations. This latter prediction derives from self-affirmation theory (Steele & Liu, 1983; Steele, 1988). This perspective assumes that the dissonance-reduction motive is oriented primarily toward maintaining global feelings of competence and worth—not toward dismissing every self-concept threat that might challenge these feelings. Steele, Spencer, & Lynch (1993) obtained support for this idea in a study using the free-choice paradigm. When participants were first given a chance to affirm a positive (though task irrelevant) view of the self, they did not show the typical spread of alternatives. Thus, to the extent that the spread of alternatives and judgmental overconfidence grow out of the same motive for self-enhancement, these findings suggest a simple way of reducing overconfidence. Overconfidence should be diminished when people are given a chance to affirm a valued aspect of their identity (cf. Blanton, Cooper, Skurnik, & Aronson, 1997). Moreover, the affirmation-driven reduction in overconfidence should occur primarily among individuals for whom feelings of uncertainty would challenge prior views of the self as knowledgeable.

Study 1 tested these predictions. Participants were exposed to an affirmation manipulation and then given a confidence-estimation task. The task used in this study differs from those typically used in free-choice studies in two ways: First, the accuracy of judgments could be objectively assessed. Second, participants were told that they would receive feedback from the experimenter on the accuracy of their judgments. The judgment task used was a "blind taste test," in which participants were asked to discriminate between two popular soft drinks, Coke and Pepsi. In the grand scheme of things, this is not a very important judgment. However, people often hold strong preferences for one of these two colas, and this strongly implies a perceived ability to discriminate between them. Pilot testing revealed, however, that this perception is often inaccurate. Many people are not able to distinguish between the two colas, and there appears to be little or no connection between the strength of people's stated preferences for one cola over the other and their ability to discriminate between them. We thus hypothesized that a felt inability to discriminate between the two colas during a blind taste test would arouse greater dissonance among those holding stronger brand preferences, leading to inflated confidence estimates. Dissonance should be reduced, however, by a self-affirmation manipulation, leading to a reduction in confidence estimates. These predictions were tested in Study 1.

Method

Procedure

Participants were 141 students in a research methods class who agreed to participate as part of a class demonstration. This demonstration was presented as two separate exercises. In the first exercise, ostensibly a pilot test, participants completed a thought-listing questionnaire for a visiting experimenter whom they had not previously met. For those in the affirmation condition, the thought listing provided them with an opportunity to affirm an important value or identity. These participants wrote a short description of an area of their lives that was both (1) important to them and (2) made them feel proud. They were told that this could be any aspect of their identity, a talent, a relationship, or a basic value. Those in the control condition were asked to list the names of the first 30 U.S. states that came to mind, presumably for the purpose of assessing which states come to mind most quickly. To ensure that the task was not threatening, participants were told to list only 30 states and were assured that it was not a problem if they could not list 30 states in the time provided.

After the affirmation manipulation, participants began the second "exercise," which was the blind cola taste test. Participants came to the front of the room in groups of 5–10 students to receive two paper cups, each of which held 2.5 ounces of chilled cola. Precautions were taken to ensure that neither the participants themselves nor the assistant who delivered the cola samples were aware of the contents of the cups. The cups were always labeled using two consecutive numbers (e.g., 1 and 2, 3 and 4, and 5 and 6). During the actual taste test, participants ate half a saltine cracker, drank the cola in the lower numbered cup, ate the remainder of the cracker, and drank the cola in the higher numbered cup. Each participant always received one cup of Coke and one cup of Pepsi. Based on their random assignment to a tasting order, half the participants tasted Coke first and half tasted Pepsi first. After receiving the two cups of cola, participants took their cups back to their seats, where they sampled them and then judged which cup held Coke and which held Pepsi. Participants reported their judgments using their cup numbers, beginning with the lower numbered cup. Participants were given a detailed debriefing in a class lecture approximately 3 weeks after the taste test. In addition, students wrote a brief research paper based on the findings of this study.

Measures

Cola liking. Prior to the affirmation manipulation, participants rated their liking for both Coke and Pepsi. These questions asked, "How much do you like Coke [Pepsi]?" Responses were made on 9-point scales, with endpoints labeled *not at all* and *very much.* A preference for one cola over the other was computed by taking the absolute difference between the two ratings. Thus, on the basis of this operational definition, those who had a strong preference for one cola over the other indicated liking one cola a great deal more than the other. This new variable, *preference,* is described in greater detail in the results section.

Postaffirmation mood. Following the thought-listing procedure, participants rated their *current mood.* They did so using two 7-point semantic differentials. These scales had endpoints ranging from *sad* to *happy* and from *extremely negative* to *extremely positive* (Cronbach's α = 0.78). We collected information on mood in an effort to ensure that any effects of the self-affirmation manipulation were not merely mood effects in disguise (see Steele, 1988). It is worth noting that we did not include a manipulation check to ensure that the affirmation manipulation enhanced self-evaluation. This approach is standard practice in self-affirmation studies, reflecting the "hydraulic" nature of dissonance effects. Manipulation checks can reduce cognitive dissonance motivation by providing an opportunity to change a relevant belief or perception and thereby alleviate dissonance. For instance, the most logical manipulation check for this study would be a state measure of self-esteem. However, if participants in the control condition are given a chance to report that they feel good about themselves, this itself could prove to be self-affirming (Steele, Spencer, & Lynch, 1993). Given the face validity of our self-affirmation manipulation, we do not consider this a serious methodological concern.

Accuracy. Following the affirmation manipulation, participants tasted and attempted to label the two colas. On the basis of these labels, participants' responses were dummy coded to reflect the accuracy of their judgments (0 for inaccurate versus 1 for accurate).

Confidence. Participants rated their confidence in their cola judgments by stating the likelihood that they had labeled the two colas correctly, with responses ranging between 50 and 100%. To aid the interpretation of this scale, they were told that a rating of 50% should indicate chance guessing, whereas a rating of 100% should be used only to indicate that they were absolutely certain of their judgments. To provide a response anchor that would reduce error variance and to ensure that participants appreciated that this was not an easy task, they were told (truthfully) that roughly 60–65% of the students in pilot tests were able to label the colas correctly.

Results

Mean Accuracy and Confidence

In all, 73% of the participants labeled the two colas accurately. Given that the sample demonstrated some degree of cola-discrimination skill, it might be expected that it would also show some degree of confidence. This was the case, as the average confidence rating was M = 78% (SD = 19.30; 95% confidence interval = 75 to 82%). In absolute percentage terms, the difference between confidence and accuracy is conventionally viewed as an overconfidence effect (in this case, an overconfidence effect of 5 percentage points). However, to state that this difference reflects true overconfidence, it would be necessary for the sample confidence ratings and the sample accuracy ratings to have ratio-level properties, which is untenable.[1] Rather than embarking on the difficult task of documenting absolute levels of overconfidence and trying to meet the strong psychometric assumptions that are inherent in such documentation, we chose a more appropriate focus for our analyses. Specifically, we chose to isolate factors that influence confidence ratings independent of the effects they have on accuracy (see Radecki & Jaccard, 1995). Factors that systematically inflate confidence ratings, while holding accuracy constant, can be said to be exerting influences that are "unwarranted." These factors can lead to levels of confidence that meet the psychometric requirements necessary for the label "overconfidence." However, there is no need to address this definition if this methodological approach is used. Thus, the remaining analyses do not attempt to assess the magnitude of overconfidence but instead identify factors that influence confidence ratings independent of their effects on accuracy.

Cola Preferences

Inspection of the mean preference ratings showed that liking for both Coke (M = 5.55, SD = 2.36) and Pepsi (M = 5.65, SD = 2.49) were a little above the scale midpoints. A repeated-measures analysis of variance revealed that cola liking did not differ between the two colas and that there were no premanipulation differences in these ratings. A more detailed picture of cola liking emerged from an examination of individual differences in the tendency to show a liking for one versus the other. Recall that a new variable, *preference*, was computed by taking the absolute value of the difference between liking for Coke and liking for Pepsi. Because the ratings for Coke and Pepsi ranged from 1 to 9, values for preference have a theoretical range of 0 (no difference) to 8 (maximal difference). The mean

[1] To determine the magnitude of overconfidence, one has to be able to compare the true underlying dimension of confidence with the true underlying dimension of accuracy. For these comparisons to be veridical, one would have to assume that the measured constructs met a set of rigid assumptions. First, one would have to assume that, if one true dimension was regressed onto the other, this would yield a solution with an intercept of 0 and a slope of 1.0 [e.g., True Confidence = 0 + 1.0 × (True Accuracy)]. Then, with the observed measure of accuracy and the observed measure of confidence, one would have to assume that the regression of both observed measures onto their underlying true dimensions would have intercepts of zero and slopes of 1.0 [i.e., Measured Confidence = 0 + 1.0 × (True Confidence) and Measured Accuracy = 0 + 1.0 × (True Accuracy)]. If these conditions are not met, it is possible that, even in cases where there is perfect judgment accuracy, the measures themselves would not reflect it. Alternatively, there can be cases where even though there is lack of accuracy, the measured values would suggest otherwise. These strong measurement assumptions simply are not tenable and so we avoid addressing the magnitude of overconfidence in this article. For a general discussion of measurement issues related to comparing true underlying constructs, see Anderson (1982).

TABLE 1
Confidence in Cola Discrimination

Block	Variable	*B*	Std. Error	β	*t* value	*p* value
1	Constant	72.93	4.16		17.55	0.01
	Accuracy	1.60	3.62	0.04	0.44	0.66
	Preference	1.82	0.68	0.23	2.68	0.01
	Condition	−0.64	3.27	−0.02	−0.2	0.85
2	Constant	67.91	4.77		14.25	0.01
	Affirmation	3.47	3.69	0.08	0.94	0.35
	Preference	2.91	0.85	0.37	3.41	0.01
	Affirmation	7.33	5.02	0.20	1.46	0.15
	Condition × Preference	−2.94	1.42	−0.29	−2.07	0.04

Note. Regression analyses predicting confidence in ability to discriminate colas (ranging from 50% confidence to 100% confidence). Accuracy relates participants' accuracy at the cola discrimination task (0 = *inaccurate;* 1 = *accurate*), Preference relates participants' preference for one cola over the other (0 = *no preference* to 8 = *extreme preference*), and Condition relates the participants' experimental condition (0 = *control group;* 1 = *affirmation group*).

value for this new variable indicated a modest degree of discrimination between the colas for the sample (M = 2.70), but there was also a fair amount of variability (SD = 2.37). Scores covered the entire range of possible values, with 54% of the participants showing between a 0-point and 3-point cola preference, 30% of the sample showing between a 3-point and 5-point cola preference, and 12% showing between a 6-point and 8-point cola preference.

Accuracy and Confidence

If participants in the sample had clear insights into their abilities to discriminate the colas, their confidence ratings should be associated with the accuracy of their judgments. As expected on the basis of pilot testing, this was not the case. The point–biserial correlation between judgmental confidence and judgmental accuracy in this sample was close to zero, r (138) = .05, ns. Thus, in the case of this particular judgment, high levels of confidence did not indicate high levels of accuracy. If participants were unable to detect their own ability to discriminate the colas, how could they state preferences? A taste preference would seem to imply an ability to discriminate the two colas. However, the correlation between strength of preference and accuracy for the entire sample was also small and nonsignificant, r (136) = .10, ns. More important, when accuracy was regressed onto liking for Coke, liking for Pepsi, and the multiplicative cross-product between the two, no significant effects were uncovered (ps > .45). In short, participants' ability to discriminate the two colas had little to do with how much they reported liking either or both of the colas.

Preference and Confidence

Although preference was unrelated to judgmental accuracy, it *was* related to judgmental confidence. Within the entire sample, strength of preference was positively correlated with confidence, r (138) = 0.26, $p < .01$. To determine if the association between preference and confidence remained when accuracy was held constant, we regressed confidence onto both preference and accuracy. This yielded only a significant effect of preference, B = 1.98, $p < .01$. Thus, every 1-point increase in the 9-point preference scale resulted in a B = 1.98 percentage-point increase in confidence, controlling for accuracy. We believed that this association between perceived preference and judgment confidence reflected a self-protective motive. As preference increased, the motivation to believe in the veracity of one's judgment also increased. As suggested above, however, this relation should be attenuated by a self-affirmation manipulation. To test this prediction, we regressed judgmental confidence on strength of preference, a dummy code for experimental condition (0 = *control;* 1 = *affirmation*), and the multiplicative cross-product of these two variables. To ensure that we were investigating effects that occurred independent of accuracy, we also entered judgmental accuracy as a predictor in the equation. To allow us to inspect the independent main effects, these were entered in the first block of a hierarchical regression equation, with the cross-product entered in the second block.

The full results of the regression analysis are reported in Table 1. The first block, which accounted for a significant proportion of variance in the confidence ratings, R^2 = .05, $F(3, 133) = 2.78$, $p < .05$, revealed only the previously documented main effect of preference, B = 1.82, $p < .01$. Results from the second block added significantly to the amount of variance accounted for, R^2_{change} = .03, $F(1, 132) = 4.30$, $p < .05$. These results also added an important qualification to the main effect of preference. The significant cross-product term in the second block, B = −2.94, $p < .05$, showed that the slope of preference on confidence varied as a function of experimental condition. Starting with those in the control condition (condition = 0),

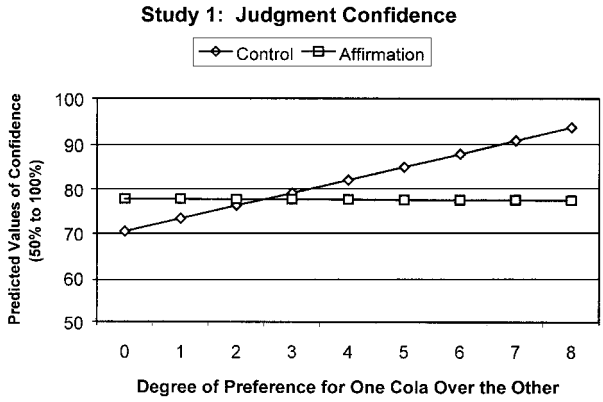

FIG. 1. The effects of cola preference (range from 0 to 9) and experimental condition (affirmation vs no affirmation) on judgment confidence (range from 50% confident to 100% confident) at the average level of accuracy for the sample ($M = .73$, $SD = .44$). High scores on preference indicate a greater discrimination between Coke and Pepsi (and thus a greater presumed motive to be confident).

the slope of preference on confidence is the value of the slope found in the second block, $B = 2.91$, $p < .001$ (or $B = 2.91 + 0 \times (-2.94) = 2.91$). Due to the significant cross-product, however, this slope was reduced by 2.94 units for those in the affirmation condition (condition = 1). Specifically, the slope of preference on confidence among those in the affirmation condition was, $B = -.03$, *ns*. [or $B = 2.91 + 1 \times (-2.94) = -.03$]. The two slopes of preference on confidence can be seen by inspecting Fig. 1. This maps out the expected value of confidence as a function of preference and condition, holding accuracy constant. As can be seen, independent of accuracy, preference led to a large and significant linear increase in confidence among those in the control condition. However, but it had little or no effect on confidence for those in the affirmation condition (see Jaccard, Turisi, & Wan, 1990, for interpretations of these types of interactions).

These results offer strong evidence for the dissonance-reducing effect of the affirmation on confidence ratings. As mentioned previously, however, there is a viable alternative interpretation. It may be that the effect of the affirmation manipulation can be attributed to its influence on mood rather than its influence on dissonance process. Arguing against this, affirmations did not influence mood, $t\ (136) = 1.40$, *ns*. Moreover, when the mood measure was entered into the previous regression equation, the interaction between experimental condition and preference remained significant and in the same direction. Thus, the affirmation effects reported here cannot easily be attributed to the induction of a positive mood state.

Factors Influencing the Relation between Accuracy and Confidence

The affirmation manipulation reduced the tendency for higher levels of preference to translate into higher levels of confidence. In a set of ancillary analyses, we sought to determine if affirmations accomplished this by helping participants become more attuned to their true abilities to discriminate. If this were the case, there should be a significant interaction between accuracy and the affirmation manipulation. Specifically, if the affirmation manipulation caused individuals to reflect more critically on their true abilities to discriminate when forming confidence judgments, and if the accuracy score was a reliable indicator of this true ability, then the accuracy score should be more strongly associated with confidence ratings in the affirmation condition as opposed to the control condition. A similar issue can be addressed with respect to the preference variable. If the higher preference for one cola over the other caused individuals to reflect less on their true abilities when forming confidence judgments, and if the accuracy score was a reliable indicator of this true ability, then the accuracy score should be less strongly associated with confidence ratings as preference increased. To test for these two effects, confidence ratings were regressed on to accuracy, preference, experimental condition, and the product terms for accuracy $\times$ condition and accuracy $\times$ preference. To control for confounds with the remaining two-way interaction, the preference $\times$ affirmation term was also included in the equation.

This analysis uncovered the previously documented interaction between preference and affirmation, $B = -3.21$, $p < .03$. In addition, it uncovered a significant interaction between preference and accuracy, $B = -3.83$, $p < .02$. The nature of the interaction was such that, for every unit increase in preference, the mean difference in confidence ratings for those who were accurate minus those who were inaccurate changed by -3.83 units. As an example, consider those participants who stated no preference (i.e., a preference score = 0). After controlling for accuracy, these individuals were $B = 19.98$ percentage points more confident when they were accurate than when they were not ($p < .02$). Among participants stating a preference of 5 units, however, the mean difference was near zero, $B = 0.83$, *ns* [or $B = 19.98 + 5(-3.83) = 0.83$]. For someone with the maximum preference of 8 units, the difference was $B = -10.66$ units [or $B = 19.98 + 10(-3.83) = -10.66$], though this was not a statistically significant effect, $p > .20$ (see Jaccard, Turisi, & Wan, 1990). In sum, only those having little preference between the two colas showed any evidence that they were engaging in a critical or nondefensive analysis of their true abilities. This finding provides an important qualification regarding the overall tendency for confidence and accuracy to be independent of one another. A lack of association was found during our original pilot testing and it was replicated with the sample as a whole at the level of the zero-order correlation. These new analyses indicate, however, that the lack of an overall correlation for the sample as a whole reflected

the fact that many in the sample had some degree of preference for one cola over the other. Among participants stating very low preferences, however, confidence was influenced by true accuracy.

The moderating effect of preference on the accuracy–confidence relationship bolsters the hypothesized link between dissonance and overconfidence. According to the dissonance perspective, those who had little preference for one cola over the other should have had little motivation to feel certain that they could discriminate one cola from the next. Consistent with this view, these uninvested participants appear to have given a more critical analysis of their true abilities when making their confidence ratings (see Seiber, 1974, for a similar set of findings). Arguing against this, the affirmation × condition interaction term in the above regression was not significant. Thus, the affirmation manipulation did not influence the degree of association between accuracy and confidence. This means that even those who were affirmed showed little correspondence between accuracy and confidence if they had stated a strong preference for one of the two colas.

Discussion

Many participants indicated a preference for one cola over the other. This would seem to imply an ability to discriminate between colas. To the contrary, we found that preference was not associated with increased accuracy. Instead, it was associated with increased confidence, even after controlling for the (negligible) effects of preference on accuracy. This supports a dissonance account of overconfidence. It suggests that overconfidence is particularly likely among those who hold (unwarranted) views of the self as competent. In further support of this perspective, we found that the relation between preference and confidence was attenuated following an affirmation manipulation. By shoring up people's feelings of self-worth, this procedure appears to have lowered the threat involved in admitting a lack of confidence among those likely to experience such a threat (Steele, Spencer, & Lynch, 1993). In further support of a dissonance interpretation, there was evidence that the motive to see the self as competent lead to less critical analyses of true ability levels during confidence assessments. This was reflected in the finding that strength of preference for one cola over the other lead to weaker correspondence between accuracy and confidence. In short, those who were motivated to feel confident engaged in a less critical analyses of their true ability to discriminate. If this pattern indicates motivated processing among those with a high preference, there was no evidence that the affirmation condition reduced confidence through this mechanism. The lack of a significant affirmation × accuracy interaction suggests that affirmation did not make participants engage in more realistic or less guarded analyses of their true abilities, even though it seems to have helped them make more cautious estimates in the face of a default motivation to do otherwise.

Generalizability of Results

One difficulty of linking our findings with those in the overconfidence literature is that our taste-test procedure bore little resemblance to the procedure typically used in studies of overconfidence. This procedure was chosen because it bore some resemblance to the types of consumer preference tasks used in the cognitive dissonance literature. However, most of studies in the overconfidence literature investigate people's confidence in their knowledge, not their confidence in their judgments (Fishhoff, 1982). A second criticism of this study focuses on the appropriateness of the dissonance manipulation. It could be argued that the affirmation manipulation was based on an overly restricted model of cognitive dissonance. Self-affirmation procedures derive from Steele's (1988) self-affirmation model of cognitive dissonance, which emphasizes the motive to maintain global feelings of self-worth. In contrast, most of the research in the cognitive dissonance literature has not emphasized this need but has instead focused on the need to justify specific actions (e.g., Aronson, 1968, 1992; Cooper & Fazio, 1984; Cooper, 1999; Tedeschi & Rosenfeld, 1981). We addressed each of these concerns in Study 2 by determining if self-justification motives inflated confidence in one's knowledge.

STUDY 2: LOWERING THE AVERSIVENESS OF NOT KNOWING

Procedure

Participants in Study 2 were 194 students taking the first of four scheduled examinations in an introductory social psychology class. The experimental stimuli appeared on the last page of the 30-question exam. On this page, entitled "exam feedback," participants were told that their instructor was interested in receiving feedback on some of the new questions on the test. In the first question of this feedback sheet, participants reporting how important it was for them to do well on the exam. We predicted that, as the importance of doing well in the course increased, the motive to feel confident in one's answers would increase similarly. Thus, the importance variable served the same theoretical role as the preference variable in Study 1.

Following the importance rating, all participants were reminded of four of the questions that had appeared on the test. They were then asked to consider the answers they had given to each and to rate their confidence in their answers. The first three confidence ratings pertained to questions taken from the text. The fourth pertained to a question taken from lecture. We predicted that increased importance would lead to increased confidence among participants in their

answers to these 4 questions, holding their accuracy constant. The dissonance-reduction manipulation was placed between the 3rd and 4th question. Thus, the primary dependent variable was the confidence rating for the 4th question. The dissonance-reduction manipulation involved a subtle prompt that reminded those in the experimental condition of a (true) grading policy that allowed them to drop their lowest exam score. To ensure that participants read this prompt, they were asked to rate whether they felt the drop rule was a good policy by circling "yes," "no," or "not sure."[2] By reminding participants that their exam score would be dropped if it proved to be very low, we hoped temporarily to lower the aversiveness of feeling uncertain about a test answer. This manipulation was inspired by Cooper and Fazio's (1984) model of cognitive dissonance. According to this model, the need to justify occurs only under circumstances in which one's actions lead to aversive consequences (see also Cooper & Worchel, 1970; Scher & Cooper, 1989). We thus predicted that those exposed to this question would also become momentarily less motivated to inflate confidence estimates, even if it was important to them to do well in the course. Because those in the control condition were not exposed to this question, a lack of certainty should still be aversive to them, particularly if it was important to them to do well.

To state our predictions formally, we predicted that the importance of doing well in the course would lead to increased confidence among participants in their answers on the exam, independent of any effects of importance on actual exam performance. Next, we predicted that this effect of importance on confidence would be attenuated by the dissonance-reduction manipulation. These predictions mirror those from Study 1, except that in Study 2 we operationalized dissonance motivation in terms of class importance (and not taste preference) and in that we used a dissonance-reduction manipulation designed to lower the aversiveness of feeling uncertain (rather than to diminish the threat to global self-worth).

Measures

Importance. The importance of doing well in the class was measured with the question, "How important is it for you to do well in this course?" Responses were made on a 7-point scale, with anchors *not at all* and *extremely.*

Confidence. As we had done in Study 1, we reminded participants that they could be incorrect to ensure they appreciated that there was reason to question their performance. To anchor their scores and reduce error variance, we informed participants (truthfully) that about 60% of students taking this introductory social psychology course in the past had answered the crucial fourth question correctly. Immediately thereafter, participants were asked, "What is the likelihood that you answered this question correctly?" Responses were made on a 10-point scale, with endpoints of 10 to 100%.

Results

Analyses of Premanipulation Responses

Most students thought it was extremely important to do well in class ($M = 6.50$ of a possible 7, $SD = 4.0$). Further inspection revealed a skew in the distribution, such that 54% of the class gave the highest possible importance rating, a 7 of 7, and 28% gave a rating of 6. The remaining students gave ratings that ranged between 2 and 5. This caused the sample sizes within each value of importance to be small. Because the dichotomy between values of 6 and 7 dominate parameter estimation and because regression diagnostics suggested that the fundamental trend in the data was captured by a simple dichotomization of importance, we created a new importance variable. Roughly half of the sample was categorized as "low importance" (if they gave a rating of 1 through 6), and the remaining participants were categorized as "high importance" (if they gave a rating of 7). To ensure that this dichotomization captured enough variability in importance to lead to the predicted variability in overconfidence, we conducted a one-way analysis of variance (ANOVA) on the mean score for the three first three (premanipulation) confidence questions, with the mean score for accuracy score used as a covariate. This yielded a significant effect of the covariate, $F(1, 188) = 75.92$, $p < .001$, and a significant effect of the dichotomization on importance, $F(1, 188) = 4.89$, $p < .03$. Controlling for accuracy, those who considered it extremely important to do well in the course were more confident ($M_{\text{adjusted}} = 68.16$, $SE = 1.98$) than those who were not so highly invested ($M_{\text{adjusted}} = 61.67$, $SE = 2.17$). Thus, the two levels of importance did yield enough variability to replicate the dissonance effect in Study 1, allowing a test of the hypothesized dissonance-reduction process.[3]

[2] Ninety-four percent stated that this policy was a good idea, 2% said it was not, and 4% were not sure. We found no differences in the results if dummy codes for the responses were entered as covariates or if we removed the respondents who did not think this was a good policy.

[3] An alternative strategy that would have been consistent with our regression diagnostics would have been to model effects of importance on confidence using a quadratic term rather than the simple dichotomization. We feature the analyses using the dichotomy in the text for ease of presentation. However, all of the effects we report using the dichotomization replicate in analyses using the quadratic term. This was illustrated in the analyses of premanipulation confidence. We regressed premanipulation confidence on premanipulation accuracy, importance, and the quadratic term for importance, with the quadratic entered by itself in the second step of a hierarchical regression. Consistent with the findings using the dichotomy, the quadratic term added significantly to the variance accounted in the regression equation, $F_{\text{change}}(1, 188) = 4.36$, $p < .03$.

Postmanipulation Confidence

Based on the performance of students in past semesters, it was anticipated that around 60% of the students would answer the experimental question correctly. However, this estimate was based on the performance of classes with mostly 3rd- and 4th-year students, whereas the current study was conducted in a class with mostly 1st- and 2nd-year students (attending a different university). In this group, only 36% of the students answered the question correctly. Despite this poor performance, students were quite confident that they had answered the crucial question correctly ($M = 66\%$, $SD = 22\%$; 95% confidence interval on confidence ratings = 62 to 69%). For reasons outlined in Study 1, however, we did not attempt to characterize the magnitude of overconfidence. Rather, we directed our analyses at detecting the factors that influenced confidence ratings independent of accuracy.

In contrast to the findings in Study 1, confidence was associated with true accuracy. Those who answered the question correctly were more confident ($M = 72.03$, $SD = 22.13$) than were those who did not ($M = 62.38$, $SD = 21.44$), $t(191) = 2.96$, $p < .01$. However, consistent with predictions (and with the results on the premanipulation confidence ratings), the importance of doing well in the class had an effect on confidence over an above its effects on accuracy. After covarying out the effects of accuracy, those who considered it important to do well in class were significantly more confident ($M_{\text{adjusted}} = 69.53$, $SE = 2.27$) than were those who considered it less so ($M_{\text{adjusted}} = 61.52$, $SE = 2.10$), $F(1, 189) = 6.63$, $p < .01$.

To determine if the dissonance-reduction procedure reduced the effect of importance on confidence, we performed a 2 (importance) × 2 (condition) ANCOVA on confidence for the experimental question, with accuracy on this question treated as the covariate. This revealed a significant effect of the covariate, $F(1, 186) = 7.09$, $p < .01$, a significant effect of importance, $F(1, 186) = 6.66$, $p < .01$, and a significant importance × condition interaction, $F(1, 186) = 5.03$, $p < .03$. There were no main effects of experimental condition, $F(1, 186) = 1.04$, *ns*. The nature of these results can be understood by inspecting Fig. 2. As predicted, and consistent with the results prior to the manipulation, confidence ratings in the control condition were higher among those who reported it was extremely important to do well in the course ($M_{\text{adjusted}} = 74.70$, $SE = 2.98$) than it was among those who did not ($M_{\text{adjusted}} = 59.79$, $SE = 3.17$), $F(1, 186) = 11.72$, $p < .01$. Consistent with predictions, however, this effect was eliminated in the dissonance-reduction condition. Confidence ratings of those in the dissonance-reduction condition were no higher among those who reported it was extremely important to do well in the course ($M_{\text{adjusted}} = 64.62$, $SE = 2.62$) than they were among those who did not ($M_{\text{adjusted}} = 63.57$, $SE = 3.28$), $F < 1$. Thus, the dissonance-reduction manipulation eliminated the tendency for those who wanted to do well in the course to boost their confidence in their exam performance.[4]

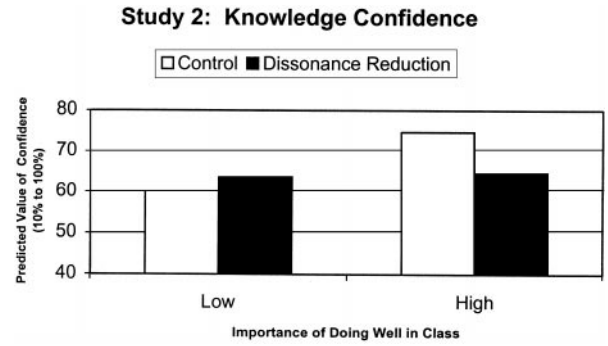

FIG. 2. The effects of the self-rated importance of doing well in class (range = 0–1) and experimental condition (control vs dissonance reduction) on confidence (range from 10 to 100%) at the average level of accuracy for the sample ($M = .36$, $SD = .48$). High scores on importance indicate a higher value placed on doing well in the class (and thus a greater presumed motive to be confident).

Factors Influencing the Relation between Accuracy and Confidence

To determine if either the importance variable or the dissonance-reduction manipulation influenced the relation between confidence and accuracy, the covariate was pulled out of the design and a 2 (Importance) × 2 (Condition) × 2 (Accuracy) ANOVA was run on the confidence ratings. In this design, the cell sample sizes ranged from $n = 15$ to $n = 38$ ($M = 24$, $SD = 7.56$). To the extent that importance, the experimental manipulation, or some combination of these factors altered the relationship between accuracy and confidence, significant interactions between these variables and accuracy should have been uncovered. No such interactions were found, though the first-order interaction between importance and condition remained significant, $F(1, 183) = 3.90$, $p < .05$. This result differs from the supplemental findings of Study 1. Recall that those who reported a stronger preference for one of the colas in Study 1 made a less critical analysis of their true abilities when rating their confidence.

[4] When the entire range of scores was used in a regression analysis, this revealed a significant curvilinear relationship between importance and confidence, as expressed in a significant quadratic term for importance, $F(1, 187) = 7.82$, $p < .01$. However, this quadratic effect was qualified by a significant interaction with condition, $F(1, 186) = 3.96$, $p < .05$. These results are consistent with those that were found using the simple dichotomization of importance. They indicate confidence increased as a (curvilinear) function of class importance but that this effect was reduced when students were reminded that they could drop their lowest grade.

GENERAL DISCUSSION

Overconfidence as Dissonance Reduction

The findings in this report suggest an important connection between cognitive dissonance and overconfidence. The findings in each study suggested that, as the cognitive dissonance associated with feeling uncertain increased, confidence increased in a way that was not warranted by increases in accuracy. More importantly, experimental manipulations designed to decrease the dissonance aroused by feelings of not knowing eliminated these effects. These findings are all the more compelling, given the fact that the two studies used such different methodologies. Study 1 focused on confidence for a seemingly trivial judgment task (a cola taste test). The motive to feel certain was operationalized in terms of a perceived preference between colas, and the dissonance-reduction procedure was designed to restore global feelings of self-worth. In contrast, Study 2 focused on confidence regarding one's knowledge in a consequential task (a graded exam). The motive to feel certain was operationlized in terms of a stated desire to perform well in class, and the dissonance-reduction manipulation was designed to lower the aversiveness of feeling uncertain. Probably as a result of these differences, there were minor differences in the factors predicting the relation between accuracy and confidence. The nature of these differences are spelled out in a section that follows. With respect to the focal hypotheses, however, the results of the two studies were surprisingly similar.

The specific factors that were hypothesized to influence confidence estimation in these studies were generated from two major models of cognitive dissonance, Steele's (1988) self-affirmation model (Study 1) and Cooper and Fazio's (1984) new-look model (Study 2). In future studies, these findings could be complemented by a consideration of factors suggested by other models of cognitive dissonance. For instance, overconfidence may be more common when there are social concerns to present the self as knowledgeable (Baumeister & Tice, 1984; Tedeschi & Rosenfeld, 1981) or when attention is focused on the self (Stone, 1997; Wicklund, 1975).[5] Also in future studies, dissonance theory may inspire new methods of debiasing confidence estimates. For instance, dissonance research suggests that the motive to boost confidence may be attenuated if a person is first given opportunities to lower the importance of feeling knowledgeable (Simon, Greenberg, & Brehm, 1995). In short, the literature on cognitive dissonance, with the different emphases of the different models, is a potentially rich source of insight into the factors that might increase or decrease overconfidence. It is our hope that these perspectives will be incorporated into future research on overconfidence.

Is Overconfidence "Just" Cognitive Dissonance?

In 1989, Berkowitz and Devine suggested that dissonance theory has been "overlooked" by researchers studying some of the most important effects it could have been used to predict. At about this same time, Aronson (1992) suggested that cognitive dissonance effects had largely been forgotten, only to be "rediscovered" by researchers in the social cognitive tradition. These tendencies, these researchers argued, have led to theoretical blind spots in which investigators have failed to see parallels between their own findings and those that that would be predicted by Festinger's (1957) original monograph on cognitive dissonance. Our primary conclusion is consistent with this critique. We suggest that the literature on overconfidence has largely ignored the seminal demonstrations of overconfidence that were generated within the cognitive dissonance tradition. Perhaps as a result, there has been little attention to the motivations that might drive people to feel certain of their knowledge when they should not.

More damaging still has been the explicit rejection of motivations as possible explanations for overconfidence. In a recent article, Metcalfe (1999) reviewed the literature on overconfidence and concluded that, although many real-life tendencies to maintain optimism in the face of threat may reflect a self-protective need to deceive the self, such mechanisms do not contribute to overconfidence in laboratory studies of confidence. In Metcalfe's defense, she made it clear that she was using a narrow definition of the self-deception motivation. This definition hinged on the desire to self-deceive even in the face of contrary evidence (after Fingarett, 1969). As she pointed out, this does not seem to occur in studies of overconfidence. Participants taking knowledge tests or performing judgement tasks seem to "want" to know the correct answers. Once these answers are revealed to them by the experimenters, they embrace them

[5] At first blush, the two studies appear to argue against a self-presentational account. We argued that participants would find it threatening to view the self as unknowledgeable or incompetent in valued domains. We thus predicted that they would inflate their confidence estimates as a way of minimizing this threat. In so doing, participants actually increased the likelihood that they would encounter public feedback from the experimenter stating that their confidence was higher than it should be. It thus seems that participants are inflating confidence in response to intrapsychic concerns over self-presentational ones. The self-presentational argument can be salvaged, however, by assuming that confidence ratings were elevated due to concerns for public presentation in the short-term, rather than of the long-term. In the short-term, participants may have wanted to give the appearance that they were certain of their abilities. In so doing, they may have ignored the increased likelihood that they would appear "overconfident" in the long-term. This suggests that confidence may be influenced by self-presentation manipulations that vary the salience of short-term versus long-term public social consequences of inflating confidence.

as valid. Our own studies share these features, and so they could not be construed as indicating a motive to deceive in the manner defined by Metcalfe. However, if we define self-deception more broadly, it seems clear that our participants were motivated to protect themselves from the implications of feeling uncertain. Our findings thus suggest that some form of self-protective motivation probably influences confidence estimations in the laboratory.

That said, we confess that we do not believe that cognitive dissonance theory provides an all-encompassing framework for understanding the voluminous literature on overconfidence. As Kunda (1992) and others have noted, the research tradition in the cognitive dissonance literature has not lent itself to identifying the specific mechanisms underlying dissonance-reduction effects. Instead, it has focused on identifying the conditions under which dissonance effects will or will not occur. To wit, if research on overconfidence had originally been subsumed under the "big tent" of cognitive dissonance theory, we doubt that investigators would have come to appreciate the role of judgmental heuristics or the importance of faulty information integration and review (e.g., Block & Harper, 1991; Fischhoff, Slovic, & Lichtenstein, 1977; Fischhoff, 1982; Koriat et al., 1980; McKenzie, 1997). More critical still, we doubt that a dissonance framework could have identified all of the conditions leading to overconfidence. For instance, Schwartz & Metcalfe (1992) demonstrated that the feeling of knowing on knowledge tasks increases when memory cues are made more accessible than usual. As an example, they demonstrated that subliminal presentation of the words "Prime Minister" increased feelings of knowing in response to the question "What was the name of Canada's first prime minister?" (see also Metcalfe & Shimamura, 1994). This result suggests that factors influencing memory accessibility are important triggers for overconfidence. This is an important insight, and we doubt that the cognitive dissonance literature would have been of much use pointing researchers in this direction.

Motivated Mechanisms

Future research should be brought to bear to understand the mechanisms underlying motivated confidence assessment. One possibility that was introduced in our ancillary analyses was the tendency to engage in a critical analysis of one's true accuracy. Despite the consistent results for the primary hypotheses, there were inconsistencies across the two studies with respect to this mechanism. In Study 1, the motive to feel certain reduced the association between accuracy and confidence. This suggests that those who stated a preference for one cola over the other engaged in a more defensive analysis of their true abilities than those who did not state a preference. This is an intriguing possibility but there was no evidence that the affirmation manipulation attenuated this effect. Moreover, this pattern of results did not replicate with the knowledge estimation task in Study 2. If the original finding in Study 1 replicates, future research will need to determine what factors determine when it will or will not occur. It would appear that the minimal condition for motivation to influence the degree of correspondence between accuracy and confidence is that the task itself must be one for which individuals *can* gain access to their judgment process. As Nisbett & Wilson (1977) have pointed out, however, people are often in the dark about these factors (see Kruger & Dunning, 1999).

Applications

Even when dissonance-reduction mechanisms do not increase the correspondence between accuracy and confidence at the individual level, dissonance-reduction mechanisms can bring confidence ratings for an entire group down to a level that is near the group's objective levels of accuracy. This may have important applied implications. As an example, Jaccard and Dodge (2001) recently found a link between a particular type of overconfidence and risky sexual behavior. In a longitudinal study using a large and representative sample of male and female high school students, they found that teens' confidence that they were knowledgeable about accurate use of birth control was associated with a subsequently *higher* probability of pregnancy. The reason for this appeared to be twofold. First, perceived knowledge about accurate use of birth control was only weakly correlated with actual knowledge (see also Radecki & Jaccard, 1995). Second, perceived knowledge was associated with subsequent increases in the frequency of sexual activity. This second tendency presumably existed because those teens who felt more confident about their birth control knowledge were also more likely to believe they could have sex without risk of pregnancy. Because the perception of knowledge was not a valid indication of true knowledge, any resulting increase in sexual frequency increased the likelihood of a poorly protected encounter.

The current studies suggest a mechanism that might account for this effect, as well as a possible strategy for addressing it. It seems reasonable that teens who become sexually active also become motivated to believe that they are knowledgeable enough to avoid pregnancy. In fact, a similar prediction was made by Festinger (1957) in his analysis of smoking. He argued that the act of smoking is inconsistent with the knowledge that smoking is dangerous. As a result, he predicted that smokers would generate cognitive distortions in order to minimize the perceived risks of their actions. Research has largely supported this prediction (Gibbons, Eggleston, & Benthin, 1997). If this dynamic also operates with sexual risk taking, then attempts to educate sexually active teens about the proper use of birth control

may encounter obstacles from people's overconfidence about their knowledge. From this analysis, a good first step in educating sexually active teens about birth control may be lowering the threat inherent in admitting a lack of knowledge. For instance, one might promote a learning context that affirms teens' self-worth, even while drawing attention to their ignorance about birth control. A strategy such as this might be of use in any number of applied contexts in which overconfidence causes resistance to information-based interventions. In short, one of the best ways to decrease overconfidence may be to decrease the threat inherent in admitting ignorance.

REFERENCES

Anderson, N. H. (1982). *Methods of information integration theory.* New York: Academic Press.

Anderson, C. A. (1983). Imagination and expectation: The effect of imagining behavioral scripts on personal intentions. *Journal of Personality and Social Psychology,* **45,** 293–305.

Aronson, E. (1968). Dissonance theory: Progress and problems. In R. Abelson, E. Aronson, W. McGuire, T. Newcomb, M. Rosenberg, & P. Tannenbaum (Eds.), *The cognitive consistency theories: A sourcebook* (pp. 5–27). Chicago: Rand McNally.

Aronson, E. (1992). The return of the repressed: Dissonance theory makes a comeback. *Psychological Inquiry,* **3,** 303–311.

Baumeister, R. F., & Tice, D. M. (1984). Role of self-presentation and choice in cognitive dissonance under forced compliance: Necessary or sufficient causes? *Journal of Personality and Social Psychology,* **46(1),** 5–13.

Blanton, H., Cooper, J., Skurnik, I., & Aronson, J. (1997). When bad things happen to good feedback: Exacerbating the need for self-justification through self-affirmation. *Personality and Social Psychology Bulletin,* **23(7),** 684–692.

Block, R. A., & Harper, D. R. (1991). Overconfidence in estimation: Testing the anchoring-and-adjustment hypothesis. *Organizational Behavior and Human Decision Processes,* **49(2),** 188–207.

Brehm, J. W. (1956). Post-decision changes in desirability of alternatives. *Journal of Abnormal and Social Psychology,* **52,** 384–389.

Cooper, J. (1999). Unwanted consequences and the self: In search of the motivation for dissonance-reduction. In E. Harmon-Jones & J. Mills (Eds.), *Cognitive dissonance: Progress on a pivotal theory in social psychology* (pp. 149–164). Washington, DC: American Psychological Association.

Cooper, J., & Fazio, R. H. (1984). A new look at dissonance theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 17, pp. 229–266). Orlando, FL: Academic Press.

Cooper, J., & Worchel, S. (1970). Role of undesired consequences in arousing cognitive dissonance. *Journal of Personality and Social Psychology,* **29,** 566–571.

Festinger, L. (1957). *A theory of cognitive dissonance.* Stanford, CA: Stanford Univ. Press.

Fingarett, H. (1969). *Self-deception.* London: Routledge & Kegan Paul.

Fischhoff, B., Slovic, P., & Lichtenstein, S. (1977). Knowing with certainty: The appropriateness of extreme confidence. *Journal of Experimental Psychology: Human Perception and Performance,* **3(4),** 552–564.

Frenkel, O. J., & Doob, A. N. (1976). Post-decision dissonance at the polling booth. *Canadian Journal of Behavioural Science,* **8(4),** 347–350.

Gibbons, F. X., Eggleston, T. J., & Benthin, A. C. (1997). Cognitive reactions to smoking relapse: The reciprocal relation between dissonance and self-esteem. *Journal of Personality and Social Psychology,* **72(1),** 184–195.

Gibbons, F. X., & Gerrard, M. (1997). Health images and their effects on health behavior. In B. P. Buunk & F. X. Gibbons (Eds.), *Health, coping, and social comparison,* (pp. 63–94). Hillsdale, NJ: Erlbaum.

Jaccard, J., & Dodge, T. (2000). Metacognition, risk behavior and risk outcomes: The role of perceived intelligence and perceived knowledge. Manuscript under editorial review.

Jaccard, J., Turrisi, R., & Wan, C. K. (1990). Interaction effects in multiple regression. In *Sage University Paper series on Quantitative Applications in the Social Sciences,* series no. 07-072. Newbury Park, CA: Sage.

Klayman, J., & Ha, Y. (1989). Hypothesis testing in rule discovery: Strategy, structure, and content. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **15(4),** 596–604.

Knox, R. E., & Inkster, J. A. (1968). Postdecision dissonance at post time. *Journal of Personality and Social Psychology,* **8(4/1),** 319–323.

Koriat, A., Lichtenstin, S., & Fischhoff, B. (1980). Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory,* **6,** 107–118.

Kunda, Z. (1992). Can dissonance theory do it all? *Psychological Inquiry,* **3(4),** 337–339.

Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it. *Journal of Personality and Social Psychology,* **77(6),** 1121–1134.

Metcalfe, J. (1999). Cognitive optimism: Self-deception or memory-based processing heuristics? *Personality and Social Psychology Review,* **2(2),** 100–110.

Metcalfe, J., & Shimamura, A. (1993). *Metacognition: Knowing about knowing.* Cambridge, MA: MIT Press.

McKenzi, C. R. M. (1997). Underweighting alternatives and overconfidence. *Organizational Behavior and Human Decision Processes,* **71(2),** 141–160.

Nisbett, R., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review,* **84,** 231–259.

Radecki, C. M., & Jaccard, J. (1995). Perceptions of knowledge, actual knowledge, and information search behavior. *Journal of Experimental Social Psychology,* **31(2),** 107–138.

Regan, D. T., & Kilduff, M. (1988). Optimism about elections: Dissonance reduction at the ballot box. *Political Psychology,* **9(1),** 101–107.

Scher, S. J., & Cooper, J. (1989). Motivational basis of dissonance: The singular role of behavioral consequences. *Journal of Personality and Social Psychology,* **56,** 899–906.

Schwartz, B. L., & Metcalfe, J. (1992). Cue familiarity but not target retrievability enhances feelings-of-knowing judgments. *Journal of Experimental Psychology Learning, Memory, and Cognition,* **18,** 1074–1083.

Sherman, S. J., Cialdini, R. B., Schwartzman, D. F., & Reynolds, K. D. (1985). Imagining can heighten or lower the perceived likelihood of contracting a disease: The mediating effect of ease of imagery. *Personality and Social Psychology Bulletin,* **11(1),** 118–127.

Simon, L., Greenberg, J., & Brehm, J. (1995). Trivialization: the forgotten mode of dissonance reduction. *Journal of Personality and Social Psychology,* **68(2),** 247–260.

Steele, C. M. (1988). The psychology of self-affirmation: Sustaining the integrity of the self. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (pp. 261–302). Hillsdale, NJ: Erlbaum.

Steele, C. M., & Liu, T. J. (1983). Dissonance processes as self-affirmation. *Journal of Personality and Social Psychology,* **41,** 831–846.

Steele, C. M., Spencer, S. J., & Lynch, M. (1993). Self-image resilience and dissonance: The role of affirmational resources. *Journal of Personality and Social Psychology*, **64(6),** 885–896.

Stone, J. (1999). What exactly have I done? The role of self-attribute accessibility in Dissonance. In E.Harmon-Jones & J. Mills (Eds.), *Cognitive dissonance: Progress on a pivotal theory in social psychology* (pp. 175–200). Washington, DC: American Psychological Association.

Tedeschi, J. T., & Rosenfeld, P. (1981). Impression management theory and the forced compliance situation. In J. T. Tedeschi (Ed.), *Impression management theory and social psychological research* (pp. 147–177). New York: Academic Press.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, **185(4157),** 1124–1131.

Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, **20(3),** 273–281.

Wicklund, R. A. (1975). Objective self-awareness. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 8, pp. 233–275). New York: Academic Press.