

The Halo Effect: Evidence for Unconscious Alteration of Judgments

Richard E. Nisbett and Timothy DeCamp Wilson
University of Michigan

Two different videotaped interviews were staged with the same individual—a college instructor who spoke English with a European accent. In one of the interviews the instructor was warm and friendly, in the other, cold and distant. The subjects who saw the warm instructor rated his appearance, mannerisms, and accent as appealing, whereas those who saw the cold instructor rated these attributes as irritating. These results indicate that global evaluations of a person can induce altered evaluations of the person's attributes, even when there is sufficient information to allow for independent assessments of them. Furthermore, the subjects were unaware of this influence of global evaluations on ratings of attributes. In fact, the subjects who saw the cold instructor actually believed that the direction of influence was opposite to the true direction. They reported that their dislike of the instructor had no effect on their ratings of his attributes but that their dislike of his attributes had lowered their global evaluations of him.

Although the halo effect is one of the oldest and most widely known of psychological phenomena, surprisingly little is known about its nature. The halo effect is generally defined as the influence of a global evaluation on evaluations of individual attributes of a person, but this definition is imprecise with respect to the strength and character of the influence. At one extreme, the halo effect might be due simply to an extrapolation from a general impression to unknown attributes. Global evaluations might color presumptions about specific traits or influence interpretation of the meaning or affective value of ambiguous trait information. Thus, if we like a person, we often assume that those attributes of the person about which we know little are also favorable. (Politicians often seem to capitalize on this tendency by appearing warm and friendly but saying little about the issues.) Such a phenomenon could best be described as a de-

duction from an implicit personality theory holding that "nice people tend to have nice attributes and less nice people have less nice attributes." Similarly, the global evaluation may alter the interpretation of the meaning or evaluation of ambiguous attributes. Thus, if one is told that a warm and friendly person is impetuous, a quite different set of behaviors come to mind from those that occur when one is told that the impetuous person is angry and hostile. The evaluation of the attribute impetuous may differ in the two cases, but the difference could be a mere interpretive one arising from the ambiguity inherent in the term *impetuous*.

It is possible, however, that the halo effect is a much more powerful phenomenon. Global evaluations may be capable of altering perceptions of even relatively unambiguous stimuli, about which the individual has sufficient information to render a confident judgment. For example, a person's appearance may be perceived as more attractive if we like the person than if we do not. If so, then the halo effect involves much more than influence on presumptions about attributes or interpretations of the meaning of attributes, but extends to alteration of judgments about attributes for which we generally assume we are capable of rendering independent assessments.

The present research was supported by National Science Foundation Grant BNS75-23191.

We are indebted to Eugene Borgida and Hersz Zukier for their assistance and advice and to Lee Ross for criticism of an early draft of the manuscript.

Requests for reprints should be sent to Richard E. Nisbett, Research Center for Group Dynamics, Institute for Social Research, University of Michigan, Ann Arbor, Michigan 48106.

The literature to date on the halo effect does not provide a clear picture of either the strength or the nature of the phenomenon. Although correlational studies are plentiful, they do not serve to demonstrate even that the phenomenon exists. In any array of trait ratings, correlations tend to be positive, that is, affectively consistent, but it is always possible that the correlations are based on veridical social perception. For example, people with more pleasant personalities may actually be better looking on the average than people with less pleasant personalities. Alternatively, people's criteria for likability and good looks may differ in such a way as to produce positive correlations in the absence of any influence of one judgment or the other. For example, a college football player and a drama major may have different criteria for both likability and good looks in other males. The football player may prize a hail-fellow-well-met style and may regard brawn as essential to masculine good looks. The drama major may favor verbal wit and a lean, ascetic appearance. If an open personal style and brawniness are associated and if a witty style and lean, ascetic looks are associated, then both the football player and the drama major might show positive correlations between their judgments of likability and looks, even though they give high scores on both dimensions to utterly different people and even though neither judgment has in any way influenced the other.

To demonstrate that global evaluations alter evaluations of particular attributes requires an experimental approach. There have been two sorts of experimental studies in this area. One type employs adjectives rather than persons as the stimuli (e.g., Anderson, 1974; Asch, 1946). This research supports the generalization that the evaluative rating given to a trait adjective is a function of the total adjective set in which it is embedded. In a sense, then, the global evaluation alters the evaluation of a particular attribute. But this work suffices only to show that the total adjective set may influence the meaning or evaluation of a particular adjective whose meaning or value is somewhat ambiguous when considered in isolation, as in the example above of the interpretation of the adjective *impetuous*.

Another type of experimental study has employed actual persons as the stimuli (Clifford & Walster, 1973; Dion, 1972; Dion, Berscheid, & Walster, 1972; Miller, 1970), but these studies suffice to show only the weak "altered presumption" version of the halo effect because they provide extremely impoverished information about the attributes to be rated. For example, Miller (1970) gave photographs of attractive and unattractive target persons to subjects and asked them to rate the personalities of the target persons. Attractive people were ascribed more favorable traits than unattractive people, but this merely constitutes a demonstration that people's implicit personality theories include the assumption that physically attractive people have more attractive personalities. A convincing demonstration of the strong version of the halo effect phenomenon would require giving subjects sufficient information on the attributes to be rated to allow for an independent assessment of the attributes.

One study in the above tradition went substantially further. Landy and Sigall (1974) showed that evaluations of an essay given by male college students were substantially higher when the alleged writer of the essay was an attractive woman than when the alleged writer was an unattractive woman. The effect was very pronounced, especially when the essay was of relatively poor quality. This study supports the stronger interpretation of the halo phenomenon because, unlike the attributes in the above studies, the subjects had information about the woman's essay that was entirely sufficient to allow for an independent assessment of its quality. The Landy and Sigall study, however, was not precisely a study of the halo effect because it was a single attribute rather than a global evaluation that was manipulated and a person's product rather than an attribute that was measured.

The present experiment was designed to address an additional issue—the question of people's awareness of the halo effect. Nisbett and Wilson (1977) have recently argued that people have little awareness of the nature or even the existence of the cognitive processes that mediate judgments, inferences, and the production of complex social behavior. The

halo effect would appear particularly likely to be such a subterranean, unrecognized process. If people were capable of rendering independent assessments of attributes, they would surely do so. The halo effect, particularly if it is the strong, distorting influence postulated by some theorists, would appear to depend upon lack of awareness for its existence. On the other hand, it is not inconceivable that subjects might actually be able to report correctly on its operation if they were directly questioned about it (e.g., Oh yes, now that you mention it, maybe part of the reason I liked his looks was because I liked his personality). The present experiment directly assessed people's reports about the nature and direction of the influence of global evaluations on evaluations of attributes.

Method

Overview

College student subjects were asked to evaluate a psychology instructor, whom they viewed in a videotaped interview, on several dimensions. The subjects were told that the investigators were studying the possibility that ratings of an instructor presented in such a brief fashion might resemble ratings by students who had taken an entire course with the instructor. The subjects were shown one of two different interviews with the same instructor, a native French-speaking Belgian who spoke English with a fairly pronounced accent. In one interview, the instructor presented himself as a likable teacher, respectful of his students' intelligence and motives, flexible in his approach to teaching, and enthusiastic about his subject matter. In the other interview, he appeared to be quite unlikable, cold and distrustful toward his students, rigid and doctrinaire in his teaching style. After viewing the videotaped interview, the subjects rated the instructor's likability, as well as the attractiveness of his physical appearance, his mannerisms, and his accent. In order to explore the question of subject awareness, some subjects were then asked whether the instructor's likability had affected their ratings of his appearance, mannerisms, and accent. Others were asked the reverse question, namely, whether their evaluation of his appearance, mannerisms, or accent had affected their ratings of his likability. It was anticipated that the subjects would rate the instructor as having a more attractive physical appearance, more attractive mannerisms, and a more attractive accent when he was likable than when he was unlikable. It was also anticipated that the subjects would be unaware of these effects on their judgments and would deny that their global evaluation of the instructor had affected their other ratings.

Procedure

The subjects were 118 University of Michigan students enrolled in introductory psychology; 62 were males and 56 were females. The results for the two sexes were entirely similar and therefore were combined for presentation. The subjects participated in the experiment in groups of 6-17. Sessions were conducted in such a way as to insure that no overt communication between subjects took place. (But in any case, conclusions are not affected when the group rather than the individual was used as the unit of analysis.) On arrival, the subjects were ushered into a well-appointed conference room with chairs oriented toward a 19-inch (.48-meter) television monitor.

The subjects were told that the psychology department had recently conducted a large-scale evaluation of teachers of introductory psychology:

At the end of the term teachers were evaluated by their students on a number of dimensions, ranging from traditional questions of teaching effectiveness to questions about matters such as physical appearance and mannerisms. We are interested in finding out whether evaluations of teachers who are seen in brief videotaped interviews at all resemble the evaluations of these teachers by students who were exposed to them for a whole semester. We will ask you to look at these videotapes and then rate the teachers on some of the same dimensions employed in the original study.

The subjects were then shown two videotaped interviews. All subjects were first shown a 7-minute filler interview with a bright and pleasant instructor who answered such questions as, What is your general educational philosophy? Do you encourage student discussion? How do you handle evaluations of student work? Then half the subjects saw the Belgian instructor answer the same questions in his warm guise and half saw him answer the questions in his cold guise. Each of these interviews also lasted about 7 minutes.

The warmth or coldness of the instructor was apparent in his answers to virtually all the questions. For example, in his warm guise the instructor answered the question about student discussion by saying that, yes, he encouraged discussion and that he was stimulated by the give and take it provided and felt that it enhanced student interest in the lecture material. In his cold guise he answered the question by saying that he didn't allow much discussion because "there's a time to be a student and a time to be a professor" and he wouldn't be the professor if he didn't know more than they did. In his warm guise the instructor answered the question about evaluation of student work by saying that he gave paper assignments that the students seemed to like and profit from and gave exams that were a mixture of objective questions and thought questions. In his cold guise he said he gave weekly multiple-choice quizzes because otherwise "you can't trust them to do the reading."

Dependent Measures

After seeing each interview, all subjects were asked, "How much do you think you would like this teacher?" and were asked to rate him on an 8-point scale ranging from "like extremely" to "dislike extremely." The questionnaire then read: "Listed below are some characteristics of the teacher you saw. Please indicate how you think you would feel about each of the characteristics if you were to take a course from the teacher" The characteristics included physical appearance, mannerisms, and for the Belgian instructor, accent. The subjects rated each characteristic on the scale below:

- extremely appealing
- very appealing
- somewhat appealing
- slightly appealing
- slightly irritating
- somewhat irritating
- very irritating
- extremely irritating

In addition, 34 subjects were told that the investigators were interested in knowing "how much, if at all, your liking for the teacher influenced the ratings you just made." The subjects indicated their answers on scales of the following type:

My liking for the teacher caused me to rate his physical appearance:

- much higher
- higher
- slightly higher
- had no effect
- slightly lower
- lower
- much lower

Another 56 subjects were asked to indicate "how much, if at all, the characteristics you just rated influenced your liking of the teacher." The subjects indicated their answers on scales of the following type:

Physical appearance made me like him:

- much more
- more
- slightly more
- had no effect
- slightly less
- less
- much less

After responding to the questionnaire, the subjects were quizzed intensively about their reactions to the videotapes and to the questionnaire items, and then were debriefed.

Results

Effect of Global Evaluation on Evaluations of Attributes

The manipulation had a very pronounced effect on liking of the teacher. The *t* value

contrasting the warm condition mean of 5.48 with the cold condition mean of 3.18 is 8.62 ($df = 1/117, p < .00001$).¹ Figure 1 presents ratings of physical appearance, mannerisms, and accent as a function of the liking manipulation. For ease of presentation, all of the positive (*appealing*) ratings are grouped together and contrasted with the negative (*irritating*) ratings. A small number of subjects checked the point intermediate between *appealing* and *irritating*, and these subjects are included with the *irritating* ratings. It may be seen that the effect of the manipulation on attribute ratings was very pronounced. A substantial majority of the subjects who saw the teacher in his warm guise rated his physical appearance as appealing, whereas a substantial majority of those who saw him in his cold guise rated his appearance as irritating ($t = 4.71, df = 1/117, p < .00001$). A majority of warm condition subjects rated the teacher's mannerisms as appealing, whereas a majority of cold condition subjects rated his mannerisms as irritating ($t = 4.06, df = 1/117, p < .0001$). About half the warm condition subjects rated the teacher's accent as appealing

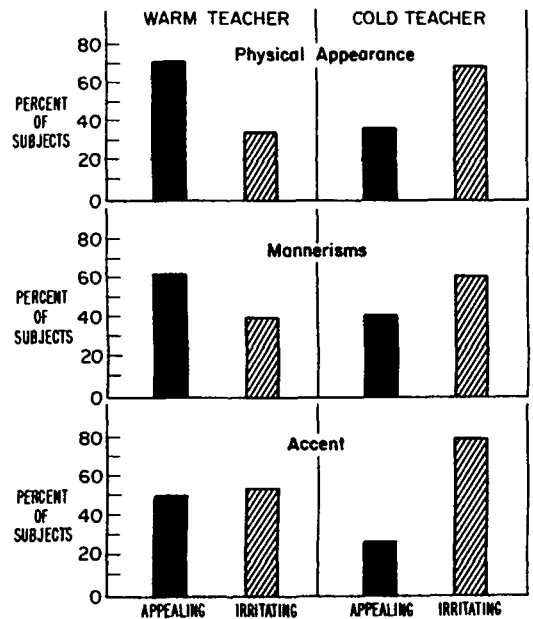


Figure 1. Ratings of Teacher's Physical Appearance, Mannerisms, and Accent as a Function of Manipulated Warmth versus Coldness.

¹ All *p* values are based on two-tailed tests.

and half as irritating, whereas the overwhelming majority of subjects in the cold condition rated his accent as irritating ($t = 3.78, df = 1/117, p < .0002$).

The results demonstrate that global assessment of a person can powerfully alter evaluations of particular attributes. That is, this is the case if one is willing to make the assumption that the attributes manifested by the confederate were in fact constant across experimental conditions. Although this is a reasonable assumption, it is by no means indisputable. It might have been the case, for example, that the teacher smiled a lot in his warm guise, thus making his appearance more appealing and/or frowned a lot in his cold guise, thus making it less appealing. Similarly, his mannerisms, both physical and verbal, might have differed when he was playing the two roles. It is even conceivable that his accent might have been less pronounced and harsh when he was playing the warm role than when playing the cold role.

In order to demonstrate that the teacher's physical appearance did not in fact differ across the two conditions, a follow-up study was conducted with 34 subjects from the same introductory psychology pool. These subjects




Figure 2. Reported Effect of Teacher Likability on Ratings of Physical Appearance, Mannerisms, and Accent.

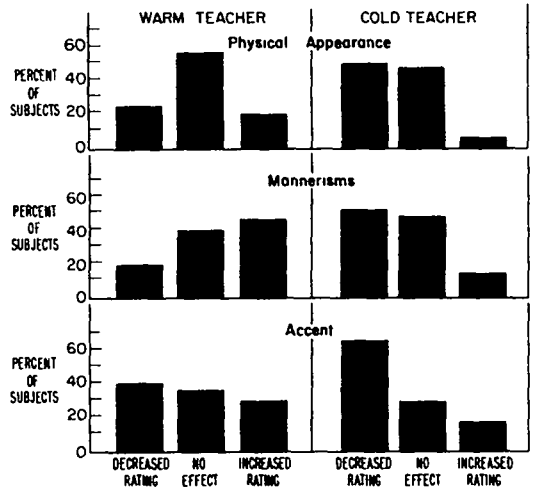


Figure 3. Reported Effect of Reactions to Physical Appearance, Mannerisms, and Accent on Ratings of Likability of Teacher.

were shown the interviews minus the audio portions (purportedly for a study of nonverbal communication) and were then asked to rate the physical appearance of the teacher, as well as his mannerisms. The subjects shown the warm version rated the teacher's physical appearance and mannerisms only trivially higher than the subjects shown the cold version (for both, $t < 1$). This indicates that the physical appearance and nonverbal mannerisms of the teacher did not in fact differ across conditions. Thus, it can be assumed that the ratings of physical appearance for the original subjects differed solely because of the global affective difference induced by the audio portion of the interview.

Awareness of the Influence of Global Evaluation

The subjects were apparently unaware of the nature of the influence of global evaluation on their ratings. Thirty-four of the subjects were asked if their liking for the teacher had affected their ratings of his attributes. Figure 2 presents these subjects' answers to this question. For ease of presentation, the data are divided into three categories: (a) decreased rating (representing 3 scale points), (b) no effect (representing 1 scale point), and (c) increased rating (representing 3 scale points). It may be seen that a majority of

subjects, in both the warm and cold conditions, believed that their overall liking for the teacher had not influenced their ratings. Moreover, for all ratings in both conditions, as many subjects felt that their overall evaluation had increased their ratings as believed it had decreased their ratings. None of the distributions differs significantly from the zero, no effect point (all six one-sample *t*s not significant).

Fifty-six subjects were asked if their evaluations of the teacher's appearance, mannerisms, or accent had affected their liking of the teacher. Figure 3 presents the results. It may be seen that the subjects who saw the warm teacher were divided on the question of whether the ratings of particular attributes affected liking. Only for the physical appearance question did the majority of subjects express the opinion that appearance had had no effect on liking. For the other two questions, approximately equal numbers of subjects believed that the ratings had increased, decreased, or had no effect on liking. None of the three distributions differs significantly from a mean of zero.

For the subjects who saw the cold teacher, the results were quite different. On balance, the subjects believed that their evaluations of each of the three attributes had decreased their liking for the teacher. The mean of all three distributions differs from zero (*t* for physical appearance = 3.28, $p < .01$; *t* for mannerisms = 2.80, $p < .01$; *t* for accent = 2.74, $p < .01$).

Thus, the subjects appear to have gotten matters precisely backwards. Their liking for the teacher was manipulated, and this affected their ratings of particular attributes. Yet the subjects did not acknowledge this effect, and the subjects who saw the cold teacher actually reported the opposite effect. These subjects reported that their evaluations of the teacher's attributes lowered their liking for him, although they denied that their liking for the teacher had affected their ratings of his attributes. (Logically, of course, it is possible that the subjects who saw the cold teacher are partly right. It is conceivable that the teacher's appearance, mannerisms, and accent were all intrinsically unattractive. In that event, the subjects who saw the cold teacher

might have been accurate when they reported that their liking for the teacher was decreased because of their evaluations of these attributes. If that were the case, then the subjects in the warm condition erred in believing that the teacher's appearance, mannerisms, and accent were neutral and had had no effect on their liking of him.)

Conversations with the subjects following the experiment served to reinforce the implications of Figures 2 and 3. The subjects were asked if they were quite sure that their global evaluations had not influenced their ratings of the teacher's attributes. Most subjects persisted in their denial—some warily, perhaps because they suspected the experimenter knew something they didn't, and some heatedly. One young woman said, in a somewhat exasperated tone, "Before he even began talking, I made my judgment about his appearance and then I stuck with it." The conversations gave reason to suspect, in fact, that some of the subjects in Figure 2 who reported that their liking for the teacher had influenced their ratings did not actually believe this. One of the subjects who had indicated this said, on being probed about his answer: "Actually, I just turned the question around I disliked his accent a little, so that made me dislike him more."

Discussion

The present results support the strong interpretation of the halo effect phenomenon. They indicate that global evaluations alter evaluations of attributes about which the individual has information fully sufficient to allow for an independent assessment. These results, it should be noted, are consistent with the very earliest theorizing about the phenomenon. Thorndike (1920), who gave the phenomenon its name, clearly believed that it represented far more than an effect on presumptions about or interpretations of the evaluative meaning of attributes, but rather that it represented a fundamental inability to resist the affective influence of global evaluation on evaluation of specific attributes. Subsequent theorists have tended to share this view, though in the absence of strong evidence to support it.

The layman has probably never doubted the existence of an inability to separate evaluations of attributes from global evaluations, at least at the affective extremes (and so long as he was referring to the foibles of others). The strong interpretation of the halo effect is embedded in such sayings as "love is blind" and "a face only a mother could love." The present study, however, suggests that global evaluations may have a pronounced impact on evaluations of specific attributes even when the global evaluations in question are less extreme than love and hate.

The present results, and those of Landy and Sigall (1974), also provide prima facie evidence that people lack awareness of the influence of one evaluation on another. Indeed, it would appear that the altered judgments require the absence of awareness. The subjects in the present study would surely have been disconcerted to know that they were incapable of rendering an independent evaluation of such characteristics as physical appearance, mannerisms, and accent. They would have been disconcerted, that is, if they could have been persuaded that they had been influenced by their global evaluation. The subjects were quite unfazed by the debriefing, probably because each subject exempted himself from the accusation that he had been prone to such an influence. The subjects in the Landy and Sigall (1974) study would undoubtedly have been even more horrified if they could have been convinced that they had upgraded their evaluations of the essay of an attractive woman or downgraded their evaluations of the essay of a plain woman.

But the prima facie case is not the only one that can be made. The subjects in the present study were directly quizzed about their beliefs concerning the nature and direction of influence. They overwhelmingly reported that their ratings of particular attributes had not been affected by their global

evaluations. The subjects who saw the cold instructor actually believed that the influence ran in the opposite direction—that their dislike of the teacher's attributes prompted their unfavorable global evaluation.

When considering the establishment of rules concerning blind review, conflict of interest, nepotism, and the like, it would therefore seem advisable to consider more than the possibility that some individuals in the system may be venal and corrupt. The protestations of even the most virtuous and disinterested participants that they are capable of independent judgments should be considered suspect.

References

- Anderson, N. H. Information integration theory: A brief survey. In D. H. Krantz, R. C. Atkinson, R. D. Luce, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology* (Vol. 2). San Francisco: Freeman, 1974.
- Asch, S. E. Forming impressions of personality. *Journal of Abnormal and Social Psychology*, 1946, 41, 258-290.
- Clifford, M. M., & Walster, E. The effect of physical attractiveness on teacher expectations. *Sociology of Education*, 1973, 46, 248-258.
- Dion, K. K. Physical attractiveness and evaluation of children's transgressions. *Journal of Personality and Social Psychology*, 1972, 24, 207-213.
- Dion, K. K., Berscheid, E., & Walster, E. What is beautiful is good. *Journal of Personality and Social Psychology*, 1972, 24, 285-290.
- Landy, D., & Sigall, H. Beauty is talent: Task evaluation as a function of the performer's physical attractiveness. *Journal of Personality and Social Psychology*, 1974, 29, 299-304.
- Miller, A. G. Role of physical attractiveness in impression formation. *Psychonomic Science*, 1970, 19, 241-242.
- Nisbett, R. E., & Wilson, T. D. Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 1977, 7, 231-259.
- Thorndike, E. L. A constant error in psychological ratings. *Journal of Applied Psychology*, 1920, 4, 25-29.

Received August 30, 1976 ■