
Attention Allocation and Impression Formation

James L. Hilton

Jill G. Klein

University of Michigan

William von Hippel

Ohio State University

An experiment examined the attention allocation strategies that perceivers employ when they are confronted with expectancy-relevant information and the mediating roles that these strategies play in impression formation. A dichotic listening task was used to measure subjects' attention allocation strategies. Subjects monitored a control child in one ear, while in the other ear they heard a target child perform in a manner that was either initially consistent or inconsistent with their expectations. They then evaluated the target on a number of dimensions. Subjects who received initially inconsistent information from the target allocated more attention to him and became more cognitively complex in their final evaluations than subjects who received initially consistent information. Furthermore, the greater attention paid by these subjects led them to become more moderate in their evaluations of the target. The results are discussed in light of the person perception literature.

Many perception theorists assume that people allocate their attention according to the consistency between their expectations and the events they observe. Bruner (1957), for example, proposed that when perceivers first encounter a target, they make an initial categorization based on the fit between the target cues and the specifications of the category. Following this initial categorization, perceivers engage in a cue search in which they are open to relevant information (i.e., they allocate attention to the target). Once a sufficient amount of consistent information is encountered, perceivers perform a confirmation check in which they become increasingly less open to relevant information. At some point, confirmation is completed and irrelevant or incongruent information is "gated out" (i.e., attention is allocated elsewhere). Similarly, Neisser (1976) has argued that perception consists of an interaction between a perceiver's schemata and the information available in

the environment. According to Neisser's view, schemata are both plans for collecting information and cognitive structures that respond to the information once it has been collected. As such, they first determine where attention will be directed and then accommodate to the information once it has been encoded.

More recently, Fiske and Neuberg (1990) have proposed a continuum model of impression formation in which attention allocation plays a central role. According to their model, impression formation begins when perceivers assign the target to an initial category (e.g., athlete, woman). Assuming that the target is of at least minimal interest, perceivers then allocate attention to the target and examine the consistency between their category-based expectations and the attribute information that is available from the target. If the target information is consistent with their expectations, perceivers then allocate attention elsewhere, and the impressions that they subsequently form of the target reflect their knowledge of the target's category rather than their knowledge of the target's specific characteristics. If, however, the target information is inconsistent with their

Authors' Note: Preparation of this article was facilitated by a National Science Foundation grant to the first author (#BNS-8717784). Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation. We gratefully acknowledge the assistance of John Darley, Hazel Markus, Richard Nisbett, and Charles Perdue, who commented on an earlier draft of the article. Jill G. Klein is now with the J. L. Kellogg Graduate School of Management, Northwestern University. Correspondence regarding the article should be sent to James L. Hilton, Research Center for Group Dynamics, Institute for Social Research, University of Michigan, Ann Arbor, MI 48106.

PSPB, Vol. 17 No. 5, October 1991 548-559

© 1991 by the Society for Personality and Social Psychology, Inc.

category-based expectations, they continue to attend to the target in an attempt to recategorize the target. If they are successful in recategorizing the target, they will then allocate attention elsewhere, and the impressions that they subsequently form will reflect their knowledge of the target's new category. If they are unsuccessful in recategorizing the target, however, they will continue to attend to the target and will begin to form impressions of the target on a piecemeal basis.

Evidence consistent with the assumption that perceivers allocate their attention according to the consistency between their expectations and the information that is available about the target has emerged in a variety of forms. In the attribution literature, for example, Pyszczynski and Greenberg (1981) have shown that when perceivers receive expectancy-consistent information, they engage in relatively little attributional processing. When they receive expectancy-inconsistent information, however, they engage in relatively elaborate attributional processing. Similarly, in the person memory literature, a variety of investigators have shown that expectancy-inconsistent information is better recalled than expectancy-consistent information (e.g., Belmore & Hubbard, 1987; Hastie & Kumar, 1979; Hastie, Park, & Weber, 1984; Srull, 1981), possibly indicating that inconsistent information receives more elaborate processing (for reviews see Hastie, 1980; Srull & Wyer, 1989). And in the impression formation literature, a number of investigators have shown that perceivers spend more time examining expectancy-inconsistent information than expectancy-consistent information (e.g., Bargh & Thein, 1985; Belmore, 1987; Erber & Fiske, 1984; Hemsley & Marmurek, 1982; White & Carlston, 1983).

Although the strategy of allocating attention according to the consistency between expectations and observations is often a reasonable response to the pragmatics of situations that heavily tax attentional resources (White & Carlston, 1983), it may also lead to an asymmetry that favors expectancy confirmation over expectancy disconfirmation. Consider the following example. Imagine that you are on a graduate admissions committee and that you have dedicated part of a busy afternoon to interviewing two students who have applied to your program. Having reviewed their academic records, you expect the first student to be intelligent and well prepared and the second student to be dull and academically lacking. Consistent with your expectations, the first student begins by saying something bright and witty whereas the second begins with something foolish and wrong. Given these responses, you quickly conclude the interview on the assumption that your expectations were correct; you decide to admit the first student and reject the second. But what if the interviews had begun differ-

ently? What if the first student had begun the interview with a foolish statement and the second student had begun with an insightful observation? Under these circumstances, you might be hesitant to conclude that your expectations were correct, and you might extend the interviews in an attempt to learn more about the two students. Moreover, depending on how the interviews progress, you might even end up rejecting the first student and accepting the second.

More formally, when perceivers receive information that is initially consistent with their expectations, they often assume that their expectations are correct and quickly allocate their attention elsewhere. In allocating their attention elsewhere, however, they should find themselves poorly situated to detect subsequent information that could disconfirm their original expectations. When perceivers receive information that is initially inconsistent with their expectations, however, they cannot assume that their expectations are correct, and they do not allocate their attention elsewhere. Given that they do not allocate their attention elsewhere, perceivers who are confronted with initially inconsistent information should be well situated to detect subsequent information that could either confirm or disconfirm their original expectations. Although confirmation should be a relatively quick process, disconfirmation should be a relatively slow process in which the accumulation of individuating information pulls perceivers away from their expectancy-based impressions. Consistent with this logic, Darley, Fleming, Hilton, and Swann (1988) found that perceivers required considerably more evidence to disconfirm an expectation than to confirm an expectation. Disconfirmation appeared to occur in a stepwise fashion in which each additional disconfirming response moved perceivers away from their initial expectations. In contrast, confirmation appeared to occur in a threshold fashion in which relatively few confirming responses were sufficient to confirm the expectation.

The current study was designed with two goals in mind. The first goal was to examine, in a relatively direct manner, the attention allocation strategies that perceivers employ when they are confronted with expectancy-consistent and expectancy-inconsistent information. The second goal was then to examine the roles that these allocation strategies play in mediating expectancy confirmation and disconfirmation. To these ends, subjects were led to believe that they were participating in a study aimed at developing student evaluation forms for use in the classroom. They were told that their task would be to role-play the part of a teacher evaluating two children. All subjects learned that one child (i.e., the control) was from a middle-class background and had been an aver-

age student during the previous year. To induce a negative expectation of intellectual performance, half the subjects were told that the other child (i.e., the target) was from a low-socioeconomic-status (SES) background and had done poorly in school during the previous year. To induce a positive expectation of intellectual performance, the other half of the subjects were told that the target was from a high-SES background and had done well in school during the previous year. All subjects then heard the control child and the target child perform simultaneously on different oral exams. As the subjects listened to the two children take the two exams, we measured the extent to which they attended to each child's performance by asking them to indicate whether each child answered each question correctly. By using a dichotic listening task (Cherry, 1953) of this sort, we were able to measure where the subjects were allocating their attention on a moment-to-moment basis (cf. Moray, 1969; Neisser, 1976; Treisman, 1969).

Independent of the expectancy manipulation, we also manipulated the initial consistency between the expectations that the subjects were given and the performances that they subsequently heard. Although all subjects heard *the same overall performance*, half the subjects heard a performance in which the target child started out doing well, and half heard a performance in which he started out doing poorly. (For discussions of this manipulation see Jones & Goethals, 1971, and Jones, Rock, Shaver, Goethals, & Ward, 1968.) Thus, half the subjects heard a performance that should have been initially consistent with their expectations, and half heard a performance that should have been initially inconsistent with their expectations. The performance of the control child was constant across conditions. After the exam, the subjects rated the target on a number of expectancy-relevant and expectancy-irrelevant dimensions. By manipulating the order of the target's performance, we were able to examine the subjects' attention allocation strategies as a function of the initial consistency between their expectancies and the target's behaviors while holding the target's overall behavior constant. We measured attention by calculating the proportion of times the subjects indicated whether the target or control answered each question correctly.

Our predictions are based on the analysis presented earlier and reflect an integration of much of the impression formation literature. Specifically, when perceivers are confronted with expectancy-inconsistent information, they should allocate more attentional resources to the target than when they are confronted with expectancy-consistent information. In the current study, this means that when the positive-expectancy target initially does

poorly or when the negative-expectancy target initially does well, subjects should pay greater attention to his performance than when the positive-expectancy target initially does well or when the negative-expectancy target initially does poorly. Moreover, as subjects pay greater attention to the target's performance, their ratings of the target should increasingly reflect this performance rather than their initial expectations. Given that the overall performance of the target is moderate, greater attention to the target's performance should lead both positive- and negative-expectancy subjects to become less extreme in their evaluations of the target. Finally, to the extent that subjects who receive information that is initially inconsistent with their expectations begin to think about the target in a more individuated manner (Fiske & Neuberg, 1990), their views of the target should become increasingly differentiated and complex.

METHOD

Subjects

Seventy-one introductory psychology students participated in the study in partial fulfillment of their course requirements.

Procedure

Subjects were contacted by phone and scheduled to arrive in the lab individually. Upon their arrival, they were shown a letter, ostensibly from the National Education Foundation (NEF), in which the goal of the study was explained. According to the letter, the purpose of the study was to develop student evaluation forms that would be both easier for teachers to use and easier for parents to understand than the forms currently in use. The letter indicated that during the course of the study the subjects would be asked to put themselves in the position of a teacher and evaluate the performances of several children who had been randomly selected in a nationwide survey of school districts.

Subjects were then shown a letter that was ostensibly from the Department of Educational Psychology. This letter stated that the department would be conducting the study for NEF. The letter explained that the main goal in the experiment would be to simulate certain key aspects of the classroom environment and to have the subjects evaluate the performances of several students. At this point the subjects were reminded that teachers often find themselves in the position of having to respond to simultaneous demands on their attention. The letter then stated that in order to simulate this situation, the subjects would be asked to put on a set of headphones and listen to two children simultaneously taking

two different oral exams. One child would be heard taking one test in one ear, and another child would be heard taking a different test in the other ear. As they listened to the children take the two tests, they would be asked to indicate whether they thought each child answered each question correctly. The letter went on to state that this task was designed to be quite difficult and that they would probably find themselves having to switch their attention between the two ears rather than trying to follow both ears at once.

When subjects finished reading both letters, the experimenter highlighted the key points covered in each of the letters and explained that the study was still in the pilot phase. She told the subjects that they would be receiving background information about two children that she hoped would include a background information sheet describing each child's home environment, an end-of-year evaluation from each child's fifth-grade teacher, a photograph of each child, and a short videotape of each child in his or her home environment. She then told the subjects that they would be evaluating a single pair of sixth-grade children and asked them to pick a number between 1 and 25 to determine which pair they would hear. After the subjects chose a number, she rummaged through the filing cabinet, apparently searching for the relevant folders. As she was searching for the folders, she pointed out that the folders were sometimes incomplete because NEF had been particularly slow in sending the photographs and videotapes. She then pulled two folders out of the cabinet that she said corresponded to the numbers the subjects had picked. She then noted that although the folder for one child (i.e., the target) was complete, the folder for the other child (i.e., the control) lacked both the photograph and the videotape. This was done in order to minimize the chances that subjects would subsequently confuse the target and the control during the evaluation phase of the experiment. At this point, the expectancy manipulation was introduced.

The information that was available about the target (a.k.a. Brian) was designed to create either a positive or a negative expectation of his pending performance and consisted of a blend of category-based and target-based information (Jones & McGillis, 1976). In the positive-expectancy conditions the target was described as having come from an upper-class background. His father was said to be an attorney and his mother was described as a free-lance writer. Together they were said to have an annual income of \$98,000, and the videotape showed the target playing basketball next to a large suburban home with two new cars parked in the driveway. The teacher evaluation read, "Overall, Brian is an exceptional stu-

dent. He does well on exams and participates in classroom activities. He is a model student. His social skills continue to develop normally and he gets along very well with his classmates."

In the negative-expectancy conditions, the target was described as having come from a lower-class background. The background information indicated that he lived with his mother, who was described as a divorced seamstress with an annual income of \$9,800. The videotape showed the target playing basketball in a gravel driveway next to a dilapidated mobile home with a broken-down, rusted truck parked in the driveway. The teacher evaluation read, "Brian is not an exceptional student. He does poorly on exams and in his participation in classroom activities. This year he has had a hard time in most of his subjects. His social skills continue to develop normally and he gets along well with his classmates."

The information about the other child was held constant across all conditions. The control (a.k.a. Matt) was always described as having come from a middle-class background. His father was said to be a department manager in a large retail store and his mother was described as a medical records technician. Together they were said to have an annual income of \$36,000. His teacher evaluation stated, "Matt fits in well with his peers and his school work is exactly where you would expect a fifth-grade graduate to be performing."

Once subjects finished reviewing these materials, the experimenter reminded them that they would be asked to listen to the two children simultaneously perform on two different tests. They then put on the headphones and were given an opportunity to practice dichotic listening. After the practice session, subjects were told that they would hear Brian answer questions in their right ear and Matt answer questions in the left ear. In order to disguise the actual length of the tests, they were also informed that each test would be approximately 40 questions in length. They were then seated in front of a computer keyboard that had keys labeled *correct* and *incorrect* on both the right and left sides. They were told that they should indicate whether Brian and Matt had answered each of their respective questions correctly by pressing the appropriately labeled key on each side of the keyboard. To help the subjects keep track of who was speaking in which ear, Brian's folder was placed on the right side of the keyboard and Matt's folder on the left side. The experimenter then stated that the subjects were not expected to respond to all the answers given by both children. Instead, she stated that they would probably find themselves needing to switch between the two performances. She then indicated that, for the sake of convenience, they should start the session by attending

to the child on the right and that once they felt they had enough information about this child to make their evaluation, they should switch their attention to the child on the left. She further indicated that once they felt comfortable in evaluating that child, they should feel free to switch back and forth between the children as they pleased. After answering any questions the subjects had, the experimenter started the tape containing the two children's performances.

All subjects heard the target and control take separate oral examinations that consisted of approximately 20 questions.¹ The questions that made up each exam were drawn from the Otis-Lennon Mental Abilities Test and ranged in difficulty from the fourth-grade through the ninth-grade level. They consisted of analogies, mathematical word problems, vocabulary questions, and analytical reasoning questions. Although the target's overall performance was the same across all conditions (60% correct), the order in which he answered questions varied by condition. In the conditions in which the target initially performed well, the target answered 9 of the first 10 questions correctly but then missed 7 of the next 10 questions. In the conditions in which the target's initial performance was poor, the target missed 7 of the first 10 questions but then answered 9 of the next 10 questions correctly. These differences in performance were accomplished by varying the order in which the individual questions were presented rather than varying the answers given to questions. Across all conditions the control answered 72% of his questions correctly.

When the dichotic listening task was completed, the remaining dependent measures were collected by having the subjects complete evaluation forms for each child, first for the target and then for the control, that were similar to the evaluation forms used by Darley and Gross (1983). These forms required the subjects to rate each child on his mental ability in nine curriculum areas: reading comprehension, reading ability, verbal ability, abstract reasoning, mathematical concepts, mathematical computation, science, general knowledge, and social studies. Responses were given on a grade-level scale ranging from 3rd grade—0 months to 8th grade—9 months. In addition, each child was evaluated on 19 skills or traits representing five dimensions (i.e., work habits, motivation, sociability, emotional maturity, and cognitive skills). Responses to these items were given on 9-point bipolar scales. Subjects were also asked to evaluate each child's test performance. More specifically, they were asked how many questions each child answered correctly and how many of the first and last 10 questions each child answered correctly. They were also asked how many of the test questions that each child received were easy,

moderate, or difficult and how many of each of these types of questions each child answered correctly. Finally, subjects were asked to indicate the grade level of the child, the grade level of the test, and the socioeconomic status of the child. When both evaluation forms were completed, subjects were probed for suspicion, debriefed, thanked, and dismissed.

RESULTS

Impressions of the Target

To examine subjects' final impressions of the target, we computed two indexes from the ratings they provided on the evaluation form that was completed at the end of the experiment.² The first index was designed to measure subjects' impressions of the target's general mental ability. It was constructed by averaging across the 9 items that assessed the subjects' impressions of the target's mental ability (Cronbach's alpha = .94). The second index was designed to measure subjects' impressions of the target on the other five dimensions (i.e., work habits, motivation, sociability, emotional maturity, and cognitive skills). It was constructed by averaging across the 19 items that assessed the subjects' impressions of the target's supplementary skills and traits (Cronbach's alpha = .95). Once the two indexes had been created, they were subjected to separate 2 (Expectancy) \times 2 (Order) analyses of variance (ANOVAs).

Given the strong manner in which the expectancy was manipulated (i.e., subjects were given information about family educational background, SES, and prior academic performance) and the robustness of the primacy effect (Anderson, 1965; Anderson & Norman, 1964), the subjects' ratings of the target on both of these indexes are perhaps best conceived of as manipulation checks. The means for the two indexes are presented in Figure 1.

Consistent with our expectations, both the mental ability index and the supplementary skills and traits index showed a main effect for the expectancy manipulation, indicating that the target was evaluated more positively when the expectation was positive than when it was negative, $F(1, 60) = 8.84, p < .01$; $F(1, 60) = 16.46, p < .001$, respectively. Similarly, both indexes showed a main effect for the order manipulation, indicating that the target was rated more positively when he performed well initially than when he performed well later, $F(1, 60) = 17.94, p < .001$; $F(1, 60) = 8.52, p < .01$, respectively. An unpredicted interaction emerged on the mental ability index, $F(1, 60) = 7.02, p < .02$, indicating that the order manipulation had a greater impact in the negative expectancy conditions than in the positive expectancy conditions. This interaction did not emerge on the supple-

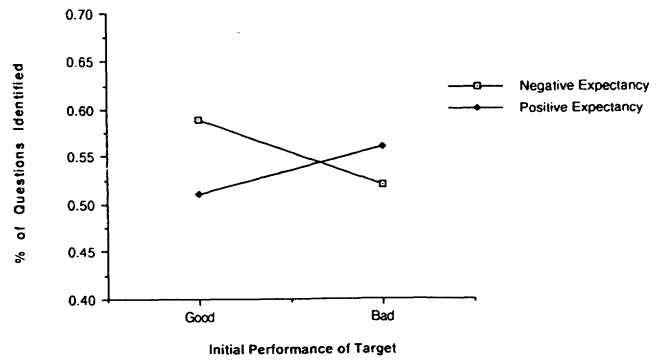
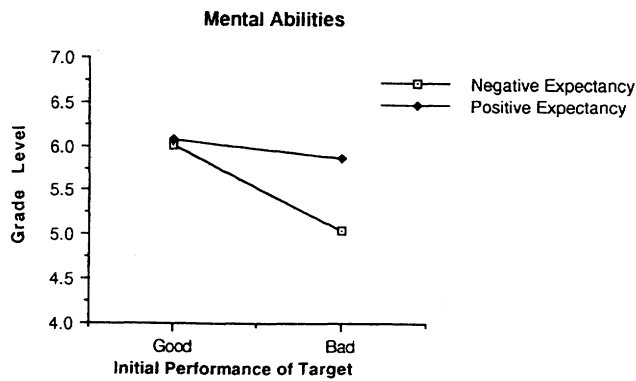


Figure 2 Proportion of the target's questions that were identified as a function of the subject's expectancy and the target's initial performance.

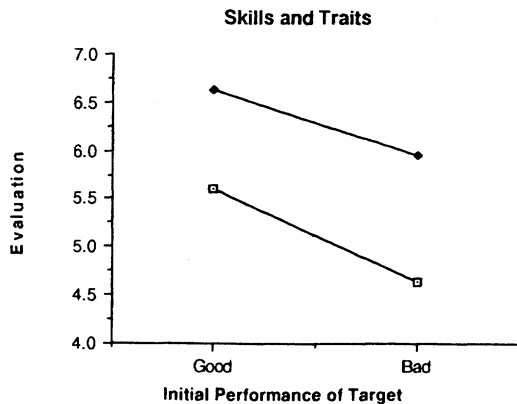


Figure 1 Mean ratings of the target's mental abilities and supplementary skills and traits as a function of the subject's expectancy and the target's initial performance.

mentary skills and traits index, $F(1, 60) < 1$, n.s. The interaction for the ability items was unexpected but makes sense in light of research suggesting that whereas the cause of failure in an ability domain may be ambiguous, success is probably due to true possession of ability (Darley & Goethals, 1980; Heider, 1958; Skowronski & Carlston, 1989). Thus, a child who initially performs well on an exam is seen as high in mental ability.

Attention Allocation

To examine the attention allocation strategies that subjects employed during the experiment, we calculated the proportion of times that the subjects indicated whether the target answered each of the 20 questions correctly and used this as our measure of attention. According to our predictions, subjects should have spent more time

attending to the target's performance when they were confronted with information that was initially inconsistent with their expectations. That is, we predicted that subjects would pay greater attention in the positive expectancy/poor initial performance condition and the negative expectancy/good initial performance condition. To test this prediction, we submitted the attention measure (i.e., the proportion of times that the subjects responded to the target's performance) to a 2 (Expectancy) \times 2 (Order) ANOVA. Consistent with predictions, the two-way interaction emerged with all four means in the predicted directions, $F(1, 60) = 4.46$, $p < .05$ (see Figure 2).³ Subjects spent a larger proportion of time attending to the target's performance when they received information that was initially inconsistent with their expectations than when they received information that was initially consistent with their expectations. Neither of the main effects emerged (both $F_s < 1$).

Effect of Attention on Evaluative Extremity

According to our analysis, as perceivers pay greater attention to a target's performance, their ratings of the target should increasingly reflect the target's actual performance rather than their initial expectations. Given the design of the current study, this means that greater attention to the target's mediocre performance should have led both high- and low-expectancy subjects to become less extreme in their evaluations of the target. To test this prediction, we conducted a path analysis in which we examined the effect that attention to the target's performance had on the extremity of subjects' evaluations of the target.⁴ In this analysis, a performance index was formed by averaging across the subjects' judgments concerning the number of exam questions the target answered correctly, the proportion of easy, mod-

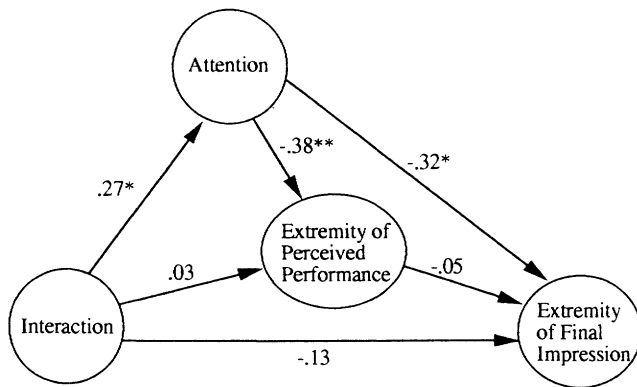


Figure 3 Path analysis of the effects of initial consistency on attention and perceived extremity of performance and final impressions. * $p < .05$; ** $p < .01$.

erate, and difficult questions answered correctly, the number of questions correctly answered out of the first 10, and the number of questions correctly answered out of the last 10 (Cronbach's $\alpha = .81$). The extremity of the perceived performance was then computed by subtracting the grand mean of this index from each subject's individual responses on the index and then taking the absolute value of the resulting score. This score was then used as an indicator of the extremity of the subjects' impressions of the target's performance, higher scores indicating more extreme positive or negative perceptions of the performance. Similarly, to compute the extremity of the subjects' final impressions of the target, we subtracted the grand mean of the mental ability index combined with the skills and traits index from each subject's responses on these indexes and then took the absolute value of the resulting score. This score was then used as an indicator of the extremity of the subjects' final impressions of the target, higher scores again indicating more extreme positive or negative evaluations. Finally, the subjects' responses concerning the correctness of the target's answers were used as indicators of attention.

Consistent with predictions, the path analysis (Figure 3) revealed that the interaction between expectancy and order had a positive effect on attention, which, in turn, had a negative effect on the extremity of subjects' evaluations of the target's performance and their final impressions of him. Moreover, there was no direct effect of the interaction on the extremity of their evaluations of his performance or their final impressions. In other words, the more subjects attended to the target, the less extreme were their evaluations of his performance and their final impressions. Interestingly, although attention had an effect on both the extremity of subjects' evaluations of the target's performance and their final impressions of

him, there was no direct relationship between performance and final impressions.

Effect of Consistency on Complexity of Impressions

Consistent with the notion that disconfirmation is a process in which perceivers reevaluate their initial expectations, subjects who were exposed to initially disconfirming evidence should have come to realize that some elements of their expectation were accurate and some elements needed revision. Consequently, these subjects should have added new elements of information to their impressions. In contrast, subjects who encountered initially confirming information should have continued to evaluate the target on the one or two dimensions that their initial expectancy included. The final result of this process should have been that subjects who encountered initially disconfirming information became more cognitively complex in their evaluations of the target than subjects who were exposed to information that was initially consistent with their expectations.

To examine the complexity of the subjects' impressions, we examined the intercorrelations between the subjects' ratings of the target on all 28 of the evaluative dimensions (cf. Scott, Osgood, & Peterson, 1979). Specifically, a correlation matrix was calculated for subjects in the two conditions where the manipulations resulted in the subjects' receiving information that was initially consistent with their expectations, and another was calculated for subjects in the two conditions where the manipulations resulted in information that was initially inconsistent. To the extent that subjects in the initially consistent conditions formed impressions of the target that were based primarily on the positive or negative expectancy information they had received, the correlations between the traits included in the matrix should be relatively high. Conversely, to the extent that subjects in the initially inconsistent conditions formed more complex impressions of the target that reflected an integration of their expectations and the individuating information they received, the correlations in the matrix should be relatively low.

To determine whether the two correlation matrices were significantly different from each other, they were subjected to a multigroup LISREL analysis. Poor goodness of fit in this analysis would indicate that the matrices are unlikely to have arisen from the same population and that there are significant differences between the two matrices. Consistent with our predictions, all three measures of goodness of fit indicated that a model treating the two matrices as sampling variations from a single population matrix provides a poor fit with the data, $\chi^2(406) = 1244.8$, $p < .001$, $GFI = .67$, and root mean square residual = .128.

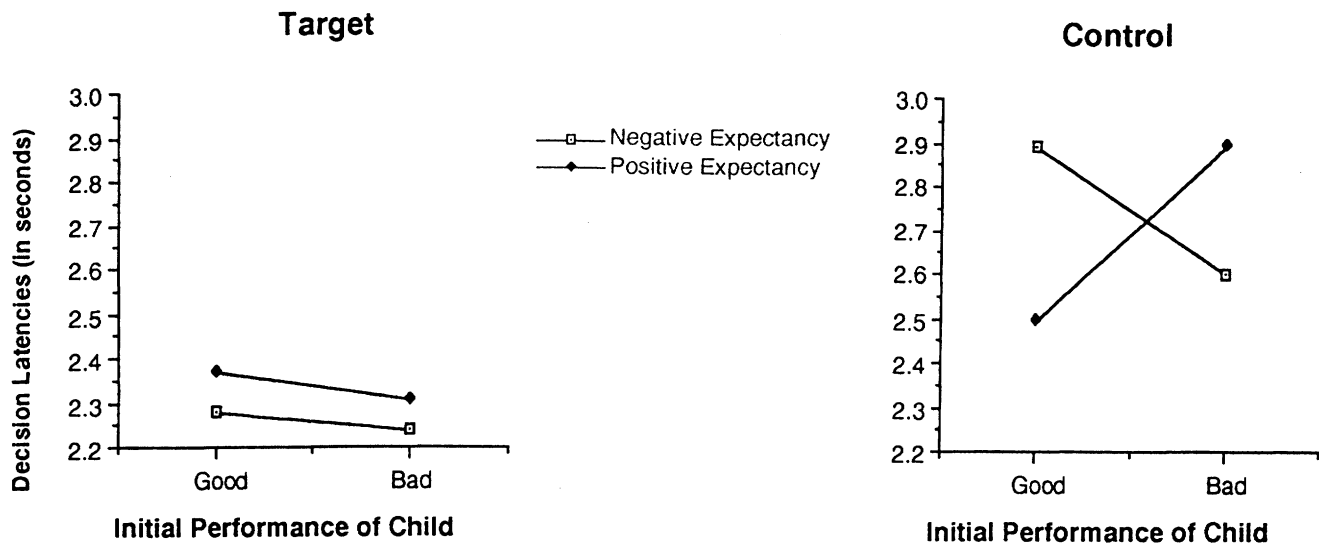


Figure 4 Decision latencies in identifying correctness of target child's or control child's answers as a function of the subject's expectancy and the child's initial performance.

Although the LISREL analysis reveals that the two matrices differ significantly from one another, it does not indicate the direction of this difference. In fact, this analysis could result in a significant difference if the average correlations in the two matrices were identical but the patterns within the matrices were different. To test the prediction that the correlations in the initially consistent conditions were consistently larger than the correlations in the initially inconsistent conditions, a sign test was conducted. In the sign test we compared the number of times that a correlation was larger in the consistent conditions than in the inconsistent conditions with the number of times that a correlation was larger in the inconsistent conditions than in the consistent conditions. Because the correlation matrices were based on the subjects' ratings of the target across 28 dimensions, this allowed for a total of 378 comparisons. Consistent with our predictions, the correlations were larger in the initially consistent conditions on 313 of the comparisons and larger in the initially inconsistent conditions on only 65 of the comparisons, $z = 12.6$, $p < .001$.

Decision Latencies

Finally, although we had no a priori predictions concerning the latencies that would be associated with subjects' decisions concerning the correctness of the target's and control's answers on the oral exam, when these latencies were submitted to a 2 (Expectancy) \times 2 (Order) \times 2 (Student) mixed-model ANOVA,⁵ an interesting effect emerged (see Figure 4). When we calculated the average time that it took each subject to make a decision about

the correctness of the target's or control's answer and examined only those instances in which the subject actually made such a decision, a three-way interaction emerged, $F(1, 58) = 4.52$, $p < .04$. No other effects emerged (all other F s < 1 except for the Expectancy \times Order interaction, $F[1, 58] = 1.59$, n.s.).⁶

Simple effects analyses revealed that in those instances in which subjects were attending to the target's performance, there were no effects for the manipulations, all F s < 1 , n.s. In those instances in which they were attending to the control's performance, however, the manipulations had an interactive effect, $F(1, 58) = 5.16$, $p < .03$, indicating that when subjects received information about the target that was initially inconsistent with their expectations, they were slower to respond to the control than when they received information about the target that was initially consistent with their expectations. To the extent that longer latencies on this measure can be taken as an indicator of cognitive load, this analysis suggests that subjects in the initially inconsistent conditions continued to operate under a relatively heavy cognitive load even after they shifted their attention to the control.

DISCUSSION

Taken together, the results of the current study provide support for the notion that when perceivers' attentional resources are heavily taxed, they allocate attention according to the fit between their expectations and the target's initial behaviors and that these allocation strategies, in turn, play an important role in determining the

impressions that perceivers subsequently form. Subjects who received information that was initially consistent with their expectations devoted relatively little attention to the target. In contrast, subjects who received information that was initially inconsistent with their expectations continued to devote attentional resources to the target. Moreover, this asymmetry appears to have played a critical role in the impressions that subjects subsequently formed. Subjects who received information that was initially consistent with their expectations tended to form relatively extreme impressions that were based primarily on the expectations they had been given rather than on the target's behavior. Conversely, subjects who received information that was initially inconsistent with their expectations tended to form more moderate impressions of the target that reflected a more complicated integration of their expectations and the target's behavior. In the paragraphs that follow, we discuss some of the broader implications that these results have for our understanding of the processes underlying the confirmation bias in particular and person perception more generally.

Expectancy Confirmation and Disconfirmation

Previous investigations of the confirmation bias have made it clear that the confirmation bias is a multiply mediated phenomenon. Sometimes the bias emerges because the evidence that perceivers encounter is ambiguous and perceivers use their expectations to make the meaning of this evidence less ambiguous. The expectation that someone is happy, for example, will lead perceivers to see crying as evidence of extreme joy, while the expectation that someone is sad will lead perceivers to see crying as evidence of extreme sorrow (Trope, 1986). At other times, the bias emerges because perceivers rely on flawed hypothesis-testing strategies. Both Wason (1960) and Snyder and Swann (1978), for example, have shown that when perceivers engage in hypothesis testing, they tend to examine cases that could confirm the hypothesis and fail to examine cases that could disconfirm it. At still other times, the bias emerges because perceivers overlook inconsistent information when they search their memories for relevant evidence (e.g., Cohen, 1981; Fischhoff & Beyth, 1975; Snyder & Cantor, 1979) or because they discount information that is inconsistent with their expectations (Kulik, 1983; Pyszczynski & Greenberg, 1981).

In addition to these factors, the current research suggests that the confirmation bias also emerges because perceivers resort to a kind of attentional triage when confronted with situations that tax their attentional resources. When the information that perceivers first receive confirms their expectations, they shift their at-

tention away from the target and are less likely to notice subsequent expectancy-relevant information. When the information that perceivers first receive disconfirms their expectations, however, they continue to allocate attention to the target and continue to notice expectancy-relevant information. This asymmetry leads to a general bias toward confirmation, because perceivers who receive initially confirming information quit attending to the target and thus can no longer notice subsequent information that could disconfirm the expectation.

Although the results of the current study suggest that the confirmation bias emerges in part as a function of perceivers' pragmatic attempts to deal with situations that swamp their attentional resources, several issues remain to be investigated. First, it is possible that before perceivers will attend to expectancy-inconsistent information, they must be motivated to form accurate impressions. Neuberg (1989) and Neuberg and Fiske (1987), for example, have demonstrated that in the absence of a motivation to be accurate, perceivers tend to ignore expectancy-inconsistent information. In light of their work, it is important to recognize that the evaluative context created by our cover story may have helped motivate the subjects to be accurate, and this may have encouraged them to attend to expectancy-inconsistent information.

Second, it remains to be seen whether attention overload is necessary to trigger a strategy of attention allocation that relies on the consistency between target information and the perceivers' expectations. It is possible that this strategy emerges only in response to attention overload. If so, then greater attentional loads should lead to a greater likelihood that the confirmation bias will emerge. The alternative possibility, consistent with the idea that attentional resources are frequently exceeded (Fiske & Taylor, 1984), is that perceivers habitually rely on this strategy. That is, perceivers may allocate their attention according to the consistency between their expectations and the behaviors they observe, even when the situation does not place a heavy demand on their attentional resources. If this second possibility is correct, then one explanation for the success that manipulations of perceiver accountability (Tetlock, 1983), impression formation goals (Neuberg, 1989; Neuberg & Fiske, 1987), and interaction goals (Darley et al., 1988; Shapiro & Hilton, 1988) have had at reducing the confirmation bias is that they place the perceiver in a more deliberative mind set (Gollwitzer & Kinney, 1989) in which he or she attends more carefully to all information.

Impression Formation

One of the important results to emerge from the path analysis was the finding that attention to the target led

to more moderate ratings of the target. This finding stands in apparent contradiction to a variety of findings indicating that attention to the stimulus will lead to extremization (e.g., Darley & Gross, 1983; Lord, Ross, & Lepper, 1979; Tesser & Leone, 1977). But it is important to examine the nature of the information that subjects in these studies receive. In both the Darley and Gross study and the Lord et al. study, for example, the evidence that subjects received was ambiguous. In the current study, the evidence that subjects heard was relatively unambiguous. The questions themselves were selected to be highly diagnostic of the target's ability, in that they were all close to the target's reported grade level in terms of difficulty (for a discussion of diagnosticity see Trope, 1986), and subjects were able to identify successfully the accuracy of the child's response 84% of the time.

As noted previously, when subjects are confronted with ambiguous evidence, they frequently use their expectations to disambiguate the evidence and then incorporate the disambiguated evidence into their final impressions. This suggests that when the behavioral evidence confronting subjects is ambiguous, greater attention to the evidence should lead perceivers with different expectations to form increasingly divergent impressions of the target. When the evidence is unambiguous, in contrast, greater attention should lead perceivers with different expectations to form increasingly convergent impressions of the target. It is important to note, however, that even when the nature of the evidence is held constant, it is unlikely that a simple relationship between attention and impression formation will emerge. Although our results suggest that attention to inconsistent information will lead perceivers to modify their expectancy-based impressions of a target, it is easy to imagine a case in which a perceiver who is committed to a certain impression will engage in an elaborate construction of excuses to explain away any expectancy-inconsistent information. Under such circumstances, initial impressions are unlikely to be modified. Stated more generally, to the extent that a perceiver is motivated to explain away inconsistent information, the relationship between attention to inconsistent information and the final impressions that are formed will weaken.

Cognitive Complexity

Another important set of findings to emerge were those concerning cognitive complexity. Previous research on cognitive complexity (e.g., Mayo & Crockett, 1964; Nidorff & Crockett, 1965; Rosenkrantz & Crockett, 1965) suggests that individuals who are cognitively complex are better able to accommodate to new information than individuals who are not. Our results extend this research in that they provide evidence that causality can

flow in the opposite direction. Attention to inconsistent information can drive accommodative changes that result in greater cognitive differentiation. Recall that subjects who were exposed to information that was initially inconsistent with their expectations became more cognitively complex in their evaluations of the target than subjects who were exposed to information that was initially consistent with their expectations. It is worth noting in this regard that these subjects were also less extreme in their final evaluations of the target—a finding that is consistent with the notion that cognitive complexity leads perceivers to greater moderation in their evaluations (Linville, 1982; Linville & Jones, 1980).

The cognitive complexity results that emerged are also consistent with recent research on person memory. In their review of the literature, Srull and Wyer (1989) suggest that the presence of inconsistent information results in several processing effects. Among these effects, Srull and Wyer note that inconsistent information causes perceivers to engage in greater processing than consistent information does. This processing is thought to occur as perceivers attempt to reconcile inconsistent behaviors with their expectancies. One possibility is that the process of reconciliation forces perceivers to become more complex as they attempt to form a coherent impression. Srull and Wyer go on to argue that, to the extent that inconsistent information leads to uncertainty in impression formation, perceivers tend to review consistent information in an attempt to bolster the validity of their initial impressions. Some evidence for this possibility can be found in the reaction time data from the current study. Recall that subjects were slower in their decision latencies in their evaluations of the control when the target behaved in a way that was initially inconsistent with their expectations. It is possible that these subjects were slowed in their evaluations because, even after they switched their attention away from the target, they continued to review his behavior in an attempt to resolve the inconsistencies.

Srull and Wyer also note that the presence of inconsistent information should lead to better memory for both consistent and inconsistent information. This enhanced memory is a product of the increased processing that subjects engage in when they are trying to resolve inconsistencies in their impressions. The data from the current study suggest an additional factor that could lead to enhanced recall of consistent and inconsistent information—namely, greater attention to all of the target's behaviors following exposure to initially inconsistent information. Subjects in the current experiment paid more attention to the target and consequently gathered more information about him when his behaviors were initially inconsistent with their expectancies

than when his behaviors were initially consistent. Although no memory measures were collected, subjects who paid more attention to the target should have been in a better position to recall his behaviors.

CONCLUSIONS

At the broadest levels, the current study demonstrates some of the benefits that can accrue from the study of expectancy disconfirmation. Too often, expectancy confirmation is treated as the psychologically interesting outcome while expectancy disconfirmation is treated as the obvious and uninteresting outcome. Yet, as a field, we have a substantial understanding of the psychological factors that lead to expectancy confirmation and cognitive stability. What we lack is an understanding of the factors that lead to expectancy disconfirmation and cognitive change. Research on expectancy disconfirmation may go a considerable way toward rectifying this situation.

NOTES

1. The control's test actually contained two extra questions in order to equate the amount of time required to complete the two exams.
2. Throughout the analyses reported, the responses from five subjects are missing because of equipment failure and the responses from two subjects are missing because they confused the target and the control.
3. An alternative attention measure was also computed by counting only the times that subjects were correct in their indication of whether the target answered each of the 20 questions correctly. Analyses on this alternative measure of attention revealed the same, significant cross-over interaction.
4. The responses from five subjects are missing from this analysis because these subjects failed to complete all the measures needed to construct the perceived performance scale.
5. The reader may have noticed that it would also be possible to submit the earlier-reported attention data to the same 2 (Expectancy) \times 2 (Order) \times 2 (Student) mixed-model ANOVA. Because submitting the attention data to this analysis would raise serious questions about the dependency between attention to the target and attention to the control that are not raised by the latency data, in the earlier analysis we decided to report the results from a design that looked only at attention to the target. It is worth noting, however, that when the attention data were submitted to the three-way ANOVA, the three-way interaction was significant, and more important, the two-way simple interaction involving attention to the target was identical to the interaction reported in the text.
6. The latencies from two subjects are missing because of equipment failure.

REFERENCES

- Anderson, N. H. (1965). Averaging versus adding as a stimulus combination rule in impression formation. *Journal of Experimental Psychology*, *70*, 394-400.
- Anderson, N. H., & Norman, A. (1964). Order effects in impression formation in four classes of stimuli. *Journal of Abnormal and Social Psychology*, *69*, 467-471.
- Bargh, J. A., & Thein, R. D. (1985). Individual construct accessibility, person memory, and the recall-judgment link: The case of information overload. *Journal of Personality and Social Psychology*, *49*, 1129-1146.
- Belmore, S. M. (1987). Determinants of attention during impression formation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 480-489.
- Belmore, S. M., & Hubbard, M. L. (1987). The role of advance expectancies in person memory. *Journal of Personality and Social Psychology*, *53*, 61-70.
- Bruner, J. S. (1957). Going beyond the information given. In J. S. Bruner, E. Brunswik, L. Festinger, F. Heider, K. F. Muenzinger, C. E. Osgood, & D. Rapaport (Eds.), *Contemporary approaches to cognition* (pp. 41-69). Cambridge, MA: Harvard University Press.
- Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and with two ears. *Journal of the Acoustical Society of America*, *25*, 554-559.
- Cohen, C. E. (1981). Person categories and social perception: Testing some boundaries of the processing effects of prior knowledge. *Journal of Personality and Social Psychology*, *40*, 441-452.
- Darley, J. M., Fleming, J. H., Hilton, J. L., & Swann, W. B., Jr. (1988). Dispelling negative expectancies: The impact of interaction goals and target characteristics on the expectancy confirmation process. *Journal of Experimental Social Psychology*, *24*, 19-36.
- Darley, J. M., & Goethals, G. R. (1980). People's analyses of the causes of ability-linked performance. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 13, pp. 1-37). Orlando, FL: Academic Press.
- Darley, J. M., & Gross, P. H. (1983). A hypothesis-confirming bias in labeling effects. *Journal of Personality and Social Psychology*, *44*, 20-33.
- Erber, R., & Fiske, S. T. (1984). Outcome dependency and attention to inconsistent information. *Journal of Personality and Social Psychology*, *47*, 709-726.
- Fischhoff, B., & Beyth, R. (1975). "I knew it would happen": Remembered probabilities of once-future things. *Organizational Behavior and Human Performance*, *13*, 1-16.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum model of impression formation from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23, pp. 1-74). Orlando, FL: Academic Press.
- Fiske, S. T., & Taylor, S. E. (1984). *Social cognition*. Reading, MA: Addison-Wesley.
- Gollwitzer, P. M., & Kinney, R. F. (1989). Effects of deliberative and implemental mind-sets on illusion of control. *Journal of Personality and Social Psychology*, *56*, 531-542.
- Hastie, R. (1980). Memory for behavioral information that confirms or contradicts a personality impression. In R. Hastie, T. M. Ostrom, E. B. Ebbesen, R. S. Wyer, D. L. Hamilton, & D. E. Carlston (Eds.), *Person memory: The cognitive basis of social perception* (pp. 155-177). Hillsdale, NJ: Lawrence Erlbaum.
- Hastie, R., & Kumar, P. A. (1979). Person memory: Personality traits as organizing principles in memory for behaviors. *Journal of Personality and Social Psychology*, *37*, 25-38.
- Hastie, R., Park, B., & Weber, R. (1984). Social memory. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 2, pp. 151-212). Hillsdale, NJ: Lawrence Erlbaum.
- Heider, F. (1958). *The psychology of interpersonal relations*. New York: Wiley.
- Hemsley, G. D., & Marmurek, H.H.C. (1982). Person memory: The processing of consistent and inconsistent person information. *Personality and Social Psychology Bulletin*, *8*, 433-438.
- Jones, E. E., & Goethals, G. R. (1971). Order effects in impression formation: Attribution context and the nature of the entity. In E. E. Jones, D. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior*. Morristown, NJ: General Learning Press.
- Jones, E. E., & McGillis, D. (1976). Correspondent inferences and the attribution cube: A comparative reappraisal. In J. H. Harvey, W. J. Ickes, & R. F. Kidd (Eds.), *New directions in attribution research* (Vol. 1). Hillsdale, NJ: Lawrence Erlbaum.
- Jones, E. E., Rock, L., Shaver, K. G., Goethals, G. R., & Ward, L. M. (1968). Pattern of performance and ability attribution: An unexpected primacy effect. *Journal of Personality and Social Psychology*, *10*, 317-340.
- Kulik, J. A. (1983). Confirmatory attribution and the perpetuation of social beliefs. *Journal of Personality and Social Psychology*, *44*, 1171-1181.
- Linville, P. W. (1982). The complexity-extremity effect and age-based stereotyping. *Journal of Personality and Social Psychology*, *42*, 193-211.

- Linville, P. W., & Jones, E. E. (1980). Polarized appraisals of outgroup members. *Journal of Personality and Social Psychology, 38*, 689-703.
- Lord, C. G., Ross, L., & Lepper, M. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology, 37*, 2098-2109.
- Mayo, C. W., & Crockett, W. H. (1964). Cognitive complexity and primacy-recency effects in impression formation. *Journal of Abnormal and Social Psychology, 68*, 335-388.
- Moray, N. (1969). *Listening and attention*. Baltimore: Penguin Books.
- Neisser, U. (1976). *Cognition and reality*. New York: W. H. Freeman.
- Nidorff, L. J., & Crockett, W. H. (1965). Cognitive complexity and the integration of conflicting information in written impressions. *Journal of Social Psychology, 66*, 165-169.
- Neuberg, S. L. (1989). The goal of forming accurate impressions during social interactions: Attenuating the impact of negative expectancies. *Journal of Personality and Social Psychology, 56*, 374-386.
- Neuberg, S. L., & Fiske, S. T. (1987). Motivational influences on impression formation: Outcome dependency, accuracy-driven attention, and individuating processes. *Journal of Personality and Social Psychology, 53*, 431-441.
- Pyszczynski, T. A., & Greenberg, J. (1981). Role of disconfirmed expectancies in the instigation of attributional processing. *Journal of Personality and Social Psychology, 40*, 31-38.
- Rosenkrantz, P. S., & Crockett, W. H. (1965). Some factors influencing the assimilation of disparate information in impression formation. *Journal of Personality and Social Psychology, 2*, 397-402.
- Scott, W. A., Osgood, D. W., & Peterson, C. (1979). *Cognitive structure: Theory and measurement of individual differences*. Washington, DC: Winston.
- Shapiro, J., & Hilton, J. L. (1988, May). *The impact of interaction goals on expectancy confirmation*. Paper presented at the 60th Annual Meeting of the Midwestern Psychological Association, Chicago.
- Skowronski, J. J., & Carlston, D. E. (1989). Negativity and extremity biases in impression formation: A review of explanations. *Journal of Personality and Social Psychology, 105*, 131-142.
- Snyder, M., & Cantor, N. (1979). Testing hypotheses about other people: The use of historical knowledge. *Journal of Experimental Social Psychology, 15*, 330-342.
- Snyder, M., & Swann, W. B., Jr. (1978). Hypothesis-testing processes in social interaction. *Journal of Personality and Social Psychology, 36*, 1202-1212.
- Srull, T. K. (1981). Person memory: Some tests of associative storage and retrieval models. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 440-463.
- Srull, T. K., & Wyer, R. S., Jr. (1989). Person memory and judgment. *Psychological Review, 96*, 58-83.
- Tesser, A., & Leone, C. (1977). Cognitive schemas and thought as determinants of attitude change. *Journal of Experimental Social Psychology, 13*, 340-356.
- Tetlock, P. E. (1983). Accountability and complexity of thought. *Journal of Personality and Social Psychology, 45*, 74-83.
- Treisman, A. M. (1969). Strategies and models of selective attention. *Psychological Review, 76*, 282-299.
- Trope, Y. (1986). Identification and inferential processes in dispositional attribution. *Psychological Review, 93*, 239-257.
- Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology, 12*, 129-140.
- White, J. D., & Carlston, D. E. (1983). Consequences of schemata for attention, impressions, and recall in complex social interactions. *Journal of Personality and Social Psychology, 45*, 538-549.