JOURNAL OF EXPERIMENTAL SOCIAL PSYCHOLOGY **28,** 186–206 (1992)

How Attitudes Guide Memory-Based Judgments: A Two-Process Model

KARL CHRISTOPH KLAUER

Freie Universität Berlin

AND

ELSBETH STERN

Max-Planck-Institut für psychologische Forschung

Received February 13, 1991

Two experiments investigate the process whereby attitudes bias memory-based judgments such as trait ratings in the direction of affective consistency. It is assumed that both the attitude toward the attitude object to be judged as well as the affective connotation of the trait, event, or quality with respect to which the attitude object is to be evaluated are activated automatically and in addition, automatically compared for affective consistency with the result of a first affirmative or negative response tendency. In a grammatical judgment task involving pairs of persons' names and trait adjectives, responding "correct" was facilitated for affectively consistent pairs relative to inconsistent pairs, whereas the opposite effect was observed when the grammatical judgment required subjects of a second group to respond "false" to these same pairs. In the second experiment, we compared ratings obtained by two groups, one that worked under time pressure, and one where subjects were allowed as much time per judgment as desired. Ratings under time pressure exhibited stronger bias in the direction of affective consistency, but did not contain larger amounts of unsystematic errors. The results suggest that an affirmative or negative first response tendency is automatically extracted on the basis of the affective aspects of the judgmental situation and is spontaneously used as an a priori hypothesis for the adequate response.

People's judgments about persons and, more generally, attitude objects, often exhibit a tendency toward affective consistency. Thus, a trait, event, or quality with positive affective connotation is often more strongly ascribed to a positively evaluated attitude object than to a negatively eval-

We thank Heiner Welchert for his help in conducting Experiment I.

0022-1031/92 $3.00

uated attitude object and vice versa for traits, events, and qualities with negative affective connotations. Tendencies of this kind are found even when an effort is made to ensure that the relevant material presented to subjects does not justify differences in judgments as a function of the attitude and the affective connotation of the trait (for example, Nisbett & Wilson, 1977; Wetzel, Wilson, & Kort, 1981; Zillig, 1928). Several factors likely to contribute to this phenomenon, sometimes termed the halo effect, have been identified, including biased assimilation of new information (Lord, Ross, & Lepper, 1979; Lord, Lepper, & Preston, 1984) and schema-directed thought about the bits of information stored in memory (Tesser, 1977; Millar & Tesser, 1986), among others. The present paper is concerned with the factors involved in the cognitive process of producing a judgment and thus with the processes between presentation of the judgmental question and the subject's response.

There is evidence for an automatic component (Shiffrin & Schneider, 1977; Shiffrin & Dumais, 1981) in the processes whereby attitudes bias judgments. Indirect evidence is provided by studies demonstrating that subjects can exert little control over the tendency toward affective consistency (Nisbett & Wilson, 1977; Wetzel et al., 1981; Zillig, 1928). Furthermore, a number of studies have shown effects of the global affective tone of stimuli in tasks where the affective tone is irrelevant (Posner & Snyder, 1975; Zajonc, 1980), and the automatic activation of the global emotional tone of prime words has also been shown to contribute to the effects of so-called category priming (Higgins, Rholes, & Jones, 1977; Bargh, 1984; Bargh, Bond, Lombardi, & Tota, 1986; Erdley & D'Agostino, 1988). Thus, the affective aspects of stimuli are often quickly and spontaneously accessed even when they are irrelevant to the task subjects are required to perform.

Fazio and his co-workers (Fazio, 1986; Fazio, Sanbonmatsu, Powell, & Kardes, 1986) have demonstrated that attitudes can be automatically activated by the mere presence of an attitude object and that the activation spreads to affectively related stimuli. According to Fazio, the automatic activation of the attitude leads to some selective processing of the information available in the immediate situation, so that the individual's definition of the situation is more likely to be congruent with his or her attitude toward the object. The automatically activated attitudes influence subsequent information processing (Fazio & Williams, 1986; Houston & Fazio, 1989) as well as behavior (Fazio, Chen, McDonel, & Sherman, 1982; Fazio, Powell, & Herr, 1983; Fazio & Williams, 1986) in the direction of affective consistency. Similar results and arguments have also been forwarded by Dovidio, Evans, & Tyler (1986) and Perdue, Dovidio, Gurtman, & Tyler (1990) in the context of racial stereotypes and intergroup bias, respectively. Perdue et al., for example, argue that the positive affective tone of in-group designators such as *us* and *we* may be activated

automatically upon exposure to such words, and through the spread of that activation in a semantic network, "may bias the retrieval of evaluatively congruent material from semantic memory, in an automatic process apparently outside the awareness of the perceiver." (Perdue et al., 1990, p. 483). According to this *accessibility* hypothesis, judgments can differ as a function of the attitude, since retrieval processes are biased in favor of affectively consistent material and categories.

The present paper proposes a model that postulates a different mechanism whereby attitudes and the affective connotation of traits bias judgments. We consider social judgments that ask subjects to evaluate the appropriateness, strength, typicality, or truth of a proposed relation between an attitude object and a trait, quality, or event; an example being the question "How appropriate is the characterization 'competent' for the politician 'Genscher'?" It is proposed that two processes are instigated upon the presentation of this question, one automatic, the second controlled (Shiffrin & Schneider, 1977; Shiffrin & Dumais, 1981). It is assumed that the affective components of both the attitude object "Genscher" and the trait "competent" are activated quickly and automatically and, *in addition,* automatically compared with respect to affective consistency, resulting in a first response tendency or a priori hypothesis for the adequate response. In the case of affective *consistency* (both affective components positive or both negative), the first response tendency is to *affirm* the proposed relation of attitude object and trait. In the case of *inconsistency* (one affective component positive, the other negative), the first response tendency is to *reject* that relation. The first response tendency does not however give rise to an immediate response. Rather, it functions as an a priori hypothesis against which relevant information activated by means of a *controlled* memory search is weighted in an attempt to critically assess the a priori hypothesis. The judgment itself is thus based on an a posteriori hypothesis that integrates a first response tendency based on affective consistency and relevant descriptive information accessed in a controlled process. Particularly when the relevant information is sparse, little correction of the a priori hypothesis occurs, and the judgments are expected to exhibit strong tendencies toward affective consistency. A related two-process model has recently been proposed by Devine (1989) in the context of racial stereotypes and will be considered more fully below.

According to the present model, subjects spontaneously evaluate the affective consistency of the attitude object and the trait in question. The degree of affective consistency then determines, or at least contributes to, a first feeling of plausibility or implausibility that we have termed the a priori hypothesis or first response tendency for the appropriate response. Thus, prior to reflecting on the judgmental question itself in a controlled process, subjects perform a spontaneous plausibility check. In doing so,

they primarily use the affective consistency of the components of the question although, as discussed below, they may sometimes also be influenced by descriptive aspects. The accessibility hypothesis (Fazio et al., 1986; Dovidio et al., 1986; Perdue et al., 1990), on the other hand, claims that bias is introduced later on, in the course of reflecting on the question, through biased retrieval and subsequent use of affectively consistent material and categories. According to the accessibility hypothesis, certain attitude objects can function as prime stimuli that facilitate retrieval of affectively consistent material in memory through an "affective" priming effect (Fazio et al., 1986; Erdley & D'Agostino, 1988; Perdue et al., 1990) similar to the well-known semantic priming effect (for example, Neely, Keefe, & Kent, 1989). In contrast, the present model does not rely on a possible automatic facilitation for affectively consistent material in semantic memory to account for halo bias. Rather, it is assumed that bias stems from a spontaneous plausibility check that is guided by the automatically evaluated degree of affective consistency of the components of the judgmental question. Although the goal of the present research was not to compare the relative importance of both potential sources of bias, some of our findings bear on this issue, as discussed below.

Two experiments were performed to investigate the assumptions and consequences of the present two-process model. That model assumes that the affective components of the attitude and the trait are activated automatically and, in addition, automatically compared with respect to affective consistency, resulting in a first affirmative or negative response tendency. This assumption is clearly stronger than the "mere" activation and accessibility hypothesis (Fazio et al., 1986; Dovidio et al., 1986; Perdue et al., 1990), and the goal of the *first* experiment is to test the assumption of automatic extraction of affective consistency. The *second* experiment investigates whether or not the resulting first response tendency is spontaneously used in judgments in the manner posited by the two-process model. That model assumes that ratings are based on an a posteriori hypothesis. The a posteriori hypothesis integrates relevant descriptive information accessed during a controlled memory search with an automatically derived a priori hypothesis. If the controlled memory search is abbreviated through appropriate manipulations, the a priori hypothesis undergoes less correction by relevant descriptive information and should have greater weight in the subjects' judgments.

EXPERIMENT I: AUTOMATIC EXTRACTION OF AFFECTIVE CONSISTENCY

Subjects are presented pairs of words. The words are either trait adjectives (T) or persons' names (P), and pairs consist of two traits (TT), two persons' names (PP), or one of each (PT). The subject's task is to respond "correct" if a given pair falls into the grammatical category PT

and to respond "false" otherwise. Subjects' attention is thereby focused on the grammatical aspects of the stimuli, and subjects are not asked to consider their attitudes.

The persons' names denote well-known contemporary European politicians and the traits have either positive or negative affective connotations. Thus, depending on the subject's attitudes, some of the pairs in the PT categories will be affectively consistent, others inconsistent. In the case of an affectively consistent PT pair, the automatically extracted response tendency is affirmative according to the two-process model, and thus congruent with the response tendency induced by the grammatical judgment task (that is, to respond "correct" for PT pairs). For affectively inconsistent PT pairs, both response tendencies are incongruent. Thus, as in a Stroop effect (Stroop, 1938; Marcel, 1983), one might expect shorter response latencies for affectively consistent relative to inconsistent PT pairs, that is, facilitation when both response tendencies are congruent relative to when both are incongruent.

Results of this kind could still be accounted for by the "mere" activation and accessibility hypothesis (Fazio et al., 1986; Dovidio et al., 1986; Perdue et al., 1990). Affectively consistent PT pairs consist of two words with similar global affective tone, whereas the elements of inconsistent pairs have opposite affective connotations. Thus, if the global affective tone of one of the words of a PT pair is activated automatically, and if that activation also spreads to affectively related words (Fazio et al., 1986; Dovidio et al., 1986; Erdley & D'Agostino, 1988; Perdue et al., 1990), processing the second word of the PT pair would be facilitated for affectively consistent pairs relative to inconsistent pairs, quite independent of the response labeled "correct" attached to the PT category. The same *affective priming* would, for example, also be expected if the response label for PT pairs had been "false", according to the accessibility hypothesis.

A *second* group is therefore given the same task with reversed response labels, that is, they are asked to respond "false" to pairs in the PT category and "correct" to pairs in the grammatical categories PP and TT. The automatically extracted response tendency posited by the two-process model is then congruent with the required response "false" in the case of *affectively inconsistent* PT pairs. Hence, shorter response latencies are now expected for affectively *inconsistent* PT pairs relative to affectively consistent PT pairs. In sum, depending on the response label, affective consistency should have opposite effects on the response latencies for PT pairs under the two-process model, whereas the accessibility hypothesis cannot account for an interaction of the response label and affective consistency.

The predicted "second-order" effect of affective consistency of PT pairs has in fact been obtained in a previous study by Klauer and Stern (1989).

The present experiment is an attempt to replicate this finding with different subjects and stimuli and additional controls. As discussed below, the present study also controls for possible confounding effects of *descriptive* consistency (Klauer & Stern, 1989) through the use of Peabody quadruples (Peabody, 1967; Felipe, 1970) of traits. Peabody quadruples consist of four traits that all refer to one underlying bipolar dimension; for example, bold, cautious, timid, and reckless all refer to the dimension "fearfulness." "Cautious" and "timid" mark one end of the dimension and carry opposite affective connotations. The same is true of the second pair, "bold" and "reckless," that marks the descriptively opposite end of the underlying dimension. A second difference between that study and the present one consists in the construction of affectively consistent and inconsistent word pairs. While Klauer and Stern (1989) used preselected stimuli with known affective connotations, the present study employed subjects' subsequent evaluative judgments to select positively and negatively evaluated attitude objects.

Method

Subjects. Forty-five undergraduate students at a college of education in Munich, Germany participated and received DM 10 for their participation. They were randomly assigned to the two groups, the first group consisted of 23 subjects and the second, 22.

Material. Sixteen contemporary and well-known European politicians were selected who had recently figured prominently in German newspapers and television, and their family names were used as persons' names. Sixteen Peabody quadruples were taken from a list of 20 such quadruples constructed and normed for the German speech community by Borkenau and Ostendorf (1987). For each subject, four different blocks of 32 pairs of words were constructed. To each person's name, a Peabody quadruple was randomly assigned, and the four resulting pairs of words, each consisting of the person's name and one of the four traits of the Peabody quadruple, were distributed over the four blocks in random order. Thus, the same Peabody quadruple was used for each politician in each block.

For each block, a total of 16 pairs was thus obtained in the PT category. For each block, the 16 persons' names were also randomly paired into eight pairs of two different names each, yielding eight PP pairs. Finally, for each block, eight TT pairs were randomly selected from the traits not used in the PT pairs of that block. Thus, each block consisted of 16 pairs of the PT category, eight pairs of the PP category and eight pairs of the TT category. Per block, a subject saw each name twice as well as 32 of the total of 64 traits. The first block was used as a practice block, the other three blocks as experimental blocks.

Procedure. The subjects underwent the experimental sessions individually. Each was first shown a list of the 16 politicians' full names (first

and last names) and was asked whether he/she was familiar with these persons (all subjects reported knowing the politicians). Then, the four blocks of word pairs were presented. Finally, subjects were asked to classify the politicians into six categories according to their overall evaluation of each politician. The categories ranged from "very negative," "negative," "mildly negative," "mildly positive," "positive," to "very positive." Each session lasted approximately three quarters of an hour.

With respect to the four blocks of word pairs, subjects were told that they were to perform a simple reaction time task involving a grammatical judgment. Their task was to press a key labeled "correct" or a key labeled "false" as quickly as possible to indicate their grammatical judgment of a pair of words. Subjects of the first (second) group were asked to respond "correct" ("false") to pairs in the grammatical category PT, and "false" ("correct") to pairs from the grammatical categories PP and TT. Subjects were instructed to maximize both the speed and accuracy of their responses.

The presentation was controlled by a personal computer. The word elements of a pair were presented in capital letters within a 6 cm by 5 cm frame in the center of a CRT monochrome monitor. Trials were self-initiated. A subject initiated a trial by pushing a button and then saw numbers in the center of the screen counting down 3 s. A fixation point followed for a period of 300 ms and vanished 200 ms before the stimulus onset. The two words were then presented simultaneously, one word taking the position of the fixation point and the other written either above or below the first word, as determined by a random number generator. A given pair remained visible on the screen until the subject responded. The subject's response was recorded along with the latency of response (from stimulus onset to response) to the nearest millisecond using a software timer and video synchronization by Heathcote (1988). The subject was provided feedback as to the type of key pressed. For example, if the "correct" key had been pressed, the sentence "You responded with 'correct' " appeared for one second on the center of the screen. Then, the subject was invited to initiate the next trial. Two blocks were separated by a short break of approximately 2 min.

After the four blocks of word pairs had been presented, subjects were asked to classify the 16 politicians into the six evaluative categories described above. The subjects' evaluations of the politicians were used to identify affectively consistent and inconsistent PT pairs. For each subject, politicians in the categories "very negative" and "negative" were classified as negative, politicians in the categories "very positive" and "positive" as positive and politicians in the categories "mildly positive" and "mildly negative" as neutral. Traits were classified as positive or negative according to the norms provided by Borkenau & Ostendorf (1987). PT pairs comprising two positive or two negative words were used as affectively consistent pairs, whereas PT pairs with one positive and one negative word were used as affectively inconsistent pairs. Since each politician had been paired with the two positive and two negative traits forming a Peabody quadruple, each positive or negative politician gave rise to two affectively consistent and two inconsistent PT pairs. Three of these four pairs appeared in the three experimental blocks. Note that the selected politicians were thus equally likely to appear as members of consistent pairs as of inconsistent pairs, so that differences in familiarity and word length were controlled for. Similarly, by pairing each person's name with all four traits of a Peabody quadruple, each person is associated with descriptively similar as well as descriptively opposite traits of both positive and negative affective connotation. Again, since three of the resulting four PT pairs were randomly selected to appear in the three experimental blocks, a possible confound of descriptive consistency with affective consistency is removed through randomization. A neutral condition was also formed and consisted of the PT pairs involving a "neutral" politician.

TABLE 1
AVERAGE FREQUENCIES OF EVALUATIVE CATEGORIES

Very negative	Negative	Mildly negative	Mildly positive	Positive	Very positive
2.8	2.0	4.0	4.1	2.0	1.1

Results

Table 1 shows the average frequencies with which the evaluative categories were assigned to the 16 politicians. As can be seen, an average of 3.1 positively evaluated politicians (categories "positive" and "very positive"), 4.8 negatively evaluated politicians (categories "negative," and "very negative"), and 8.1 "neutral" politicians could be selected per subject.

Subjects committed few errors in making judgments of the grammatical category. The average error rate across subjects was 2.4%. Latencies pertaining to errors as well as latencies larger than 2.5 s were excluded, thereby discarding a total of 5.8% of the data. For each subject, the response latencies were log-transformed, and the means for affectively consistent, inconsistent, and neutral PT pairs were then computed. Table 2 presents (antilogs of the) average response latencies for each group and consistency condition as well as the error data.

Using the data for affectively consistent and inconsistent PT pairs as shown in Fig. 1, a MANOVA with between-subjects factor "response label" and within-subjects factor "consistency" was performed to test whether the effects of affective consistency differed between the two

TABLE 2
AVERAGE RESPONSE LATENCIES (ms) AND PERCENTAGES OF ERRORS

Index	Response label	Consistency condition		
		Consistent	Neutral	Inconsistent
Latency	False	1259	1285	1213
	Correct	1111	1191	1199
Errors	False	1.2	3.5	1.9
	Correct	1.8	2.3	2.7

groups that employ different response labels. As predicted, a significant interaction of consistency and response label emerged ($F(1, 43) = 14.83$, $p < .001$). Neither consistency nor response label gave rise to a significant main effect ($F(1, 43) = 1.74$, $p = .19$, and $F(1, 43) = 1.20$, $p = .28$, respectively). Furthermore, one-sided t tests performed in each group reveal, as predicted by the two-process model, that responding "correct" is facilitated by affective consistency relative to affective inconsistency ($t(22) = -3.38$, $p < .01$, whereas the opposite effect is obtained, as predicted, in the group responding "false" ($t(21) = 1.99$, $p < .05$).[1]

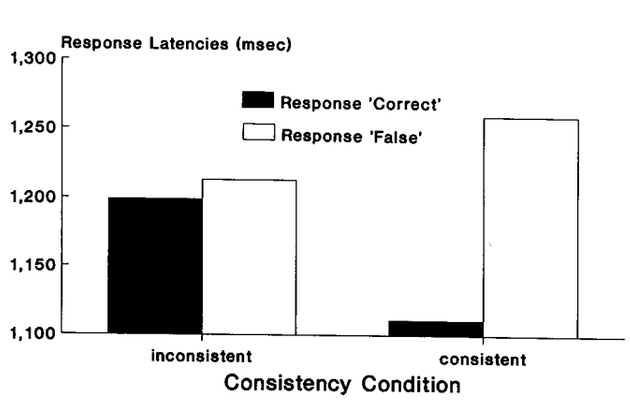

FIG. 1. Mean response latencies in Experiment I as a function of affective consistency and response label.

[1] As pointed out by the editor and an anonymous reviewer, it is also interesting to consider the data for PP and TT pairs. Since these pairs have little claim for signalling either true or false statements, subjects may not spontaneously check their plausibility as postulated for PT pairs, so that one would not expect an interaction of response label and affective consistency. The data for PP and TT pairs revealed no such interaction, nor any other reliable effects. It should be noted, however, that for methodological reasons, these data do not provide a strong test of the hypothesis. Since the study had been planned in terms of PT pairs, PP and TT pairs were constructed less carefully. For example, subjects usually saw substantially different numbers of affectively consistent and inconsistent PP and TT pairs, and several subjects saw either no affectively consistent PP pairs or no affectively inconsistent PP pairs.

Fazio et al. (1986) and Perdue et al. (1990) noted that there is considerable discussion in the cognitive literature about the appropriate use of stimuli for baseline trials and that it is difficult to define satisfactory baseline conditions. Problems of this kind also arise in the present study with respect to the neutral PT pairs as candidates for baseline stimuli. For example, it is plausible that less well-known politicians receive less extreme evaluations and therefore are likely to figure as members of "neutral" pairs. On the other hand, less familiar words give rise to larger response latencies, so that this baseline condition is possibly biased. With this reservation in mind, we also performed two-sided t tests against the baseline condition to see whether the effects of affective consistency tend to be of facilitatory or inhibitory nature. In both groups, *facilitation* relative to the baseline was observed when the postulated a priori hypothesis and the required response were congruent, that is, for affective consistency in the group responding "correct" ($t(22) = -3.59$, $p < .01$) and for affective inconsistency in the group responding "false" ($t(21) = 4.14$, $p < .01$). No reliable difference between baseline condition and the experimental conditions was found when the a priori hypothesis and the required response were incongruent, that is, for affective inconsistency in the group responding "correct" ($t(22) = .32$, $p = .75$) and for affective consistency in the group responding "false" ($t(21) = -1.24$, $p = .23$).

Discussion

The effects of affective consistency are found to differ in the predicted directions as a function of the response label used in the grammatical judgment task. As discussed above, the pattern of results is consistent with the assumption that affective consistency is automatically extracted and contributes to a first affirmative or negative response tendency and cannot be accounted for by the accessibility hypothesis (Fazio et al., 1986; Dovidio et al., 1986; Erdley & D'Agostino, 1988; Perdue et al., 1990) that predicts facilitation for affectively consistent material irrespective of the response label. Although there may have been a small amount of overall facilitation for affectively consistent PT pairs, the possible facilitation was not strong enough, in the present case, to overcome the pronounced second-order effect of affective consistency as a function of the response label.

The primary objective of the first experiment was to demonstrate that the affective components are not only automatically activated, but also automatically compared with respect to affective consistency, resulting in a first affirmative or negative response hypothesis. When asked, subjects often consider affectively consistent pairs as descriptively more appro-

priate than affectively inconsistent pairs, so that affective consistency may to some extent be confounded with descriptive consistency. It is possible, furthermore, that the spontaneous plausibility check that is postulated by the present two-process model also has a descriptive component. For example, when there are strong learned associations between a person and a given trait (such as Einstein and "intelligent"), presenting both words simultaneously could give rise to a spontaneous feeling of plausibility (Neely, Keefe, and Kent, 1989). Since in accounting for bias in the direction of affective consistency, we were primarily interested in a possible independent effect of *affective* consistency, a secondary objective of the present experiment was to assess whether the effects of affective consistency could also be found when descriptive consistency is controlled through the use of Peabody quadruples. The results replicate the predicted second-order effect of affective consistency, also obtained previously by Klauer & Stern (1989), in a situation in which descriptive consistency is controlled for.

EXPERIMENT II: SNAP JUDGMENTS

While the first experiment provides evidence for the automatic extraction of a first affirmative or negative response tendency as a function of affective consistency, the *second* experiment investigates whether the automatically activated a priori hypothesis is also used in judgments in the manner posited by the two-process model. According to that model, the a priori hypothesis is critically assessed using relevant descriptive information accessed during a controlled memory search. If that memory search is abbreviated through time pressure, the first response tendency undergoes less correction by relevant descriptive information and should have greater relative weight in the subjects' actual judgments. Thus, it is predicted that memory-based judgments obtained under time pressure exhibit stronger biases in the direction of affective consistency.

On the other hand, it has been argued that reflection upon a judgment or attitude increases the affective consistency of the relevant information stored in memory (Tesser & Leone, 1977; Tesser, 1978; Millar & Tesser, 1986; Sherman, Judd, & Park, 1989), a point of view that would lead one to expect the opposite effect, namely, an increased halo bias if the subjects are allowed time for thought. Concerning the amount of *unsystematic* errors in judgments also assessed in this study, a plausible first expectation is to predict an increase of unsystematic errors in judgments under time pressure as a consequence of a speed accuracy trade-off (Lohman, 1989), so that the reliability of the subjects' responses decreases. Under the two-process model, however, subjects' responses are assumed to be based to a large extent on the automatically extracted a priori hypothesis when working under time pressure. Since the a priori hypothesis itself relies on comparatively *stable* affective information, subjects' responses should then remain reliable even under time pressure. In sum, under time pressure, the amount of *systematic* errors should increase, not, however, the amount of *unsystematic* errors.

The second experiment compares ratings of two groups; the first group works under time pressure, and the second group is allowed as much time per judgment as desired. Subjects were asked to rate politicians on several traits. To separate systematic tendencies going back to affective consistency from those relying on descriptive information, traits forming Peabody quadruples were used. To evalute the reliability of the ratings, subjects in each group were asked to repeat their judgments after a short break.

Method

Subjects. Forty FU Berlin undergraduates participated in partial fulfillment of an introductory psychology course requirement. Subjects were randomly assigned to two groups of 20 subjects each.

Material. Four Peabody quadruples of trait adjectives were selected on the basis of their relevance for politicians from a list of 20 such quadruples constructed and normed by Borkenau and Ostendorf (1987). Thus, a total of 16 traits was used. Four contemporary German politicians were selected on the basis of a pilot study. In the pilot study, 20 members of the student population provided overall evaluations of 20 such politicians. The four politicians selected were the two most consistently positively evaluated and the two most often negatively evaluated.

Procedure. Subjects were asked to provide ratings for questions such as "How typical is trait XYZ for politician ABC?" on a scale from 1 = "not at all typical" to 8 = "very typical." Per block, subjects responded to 64 (= 4 politicians × 16 traits) such questions. Each subject worked through two *identical* blocks separated by a short break of 2 min. Within each block, judgments were grouped by politician, and the sequence of traits was randomized with the restriction that two traits pertaining to the same Peabody quadruple did not immediately follow one another.

The questions were presented on a CRT monitor of a personal computer, and subjects indicated their rating by pressing one of eight keys labeled "1" to "8." After a subject had responded, there was a pause of 2 s followed by the next question.

Subjects of the group working without time pressure were instructed to provide impartial judgments, and subjects working under time pressure were instructed to provide impartial judgments as quickly as possible. The latter subjects were also set a time limit of 12 s per judgment. If a subject working under time pressure did not respond within 12 s, the subject heard a beep and was shown the message "time limit exceeded." In this case, the judgment was left out, and the next question was presented.

Subjects were run in groups of five. They began with 10 practice ratings

using material not appearing in the experimental blocks. The group working under time pressure was then shown one trial that demonstrated the events following exceedance of the time limit as described above. Next, subjects worked through the experimental blocks. Each session lasted approximately 45 min.

Results

In the group working under time pressure, five ratings were lost due to exceeding the time limit, amounting to a .173% loss of the ratings obtained in this group.

Affective versus descriptive consistency: Indices. Recall that a Peabody quadruple consists of four traits all referring to one underlying bipolar dimension, and that two traits mark one end of this dimension, referred to as a in the following and the other two the opposite end, referred to as b. Two traits marking the same end have opposite affective connotations, one trait is generally evaluated positively and the other negatively. Thus, each Peabody quadruple crosses two factors, the first factor being the end of the bipolar dimension with values a and b and the second factor being the affective connotation of the trait with values + and −. Hence, the four traits of a Peabody quadruple may be mapped onto the four cells, $a+$, $a-$, $b+$, $b-$, of the resulting two by two design, and a subject's ratings of a given politician with respect to these traits may be denoted z_{a+}, z_{a-}, z_{b+}, and z_{b-}, respectively.

If a subject uses only descriptive information and disregards affective connotations, the subject's ratings with respect to the Peabody quadruple mentioned above (not used in this study) should approximately exhibit the pattern shown in Fig. 2. If the subject considers the person rated fearless, the traits "bold" and "reckless" should obtain relatively high ratings, whereas the traits "cautious" and "timid" should receive comparatively smaller ratings. Thus, most of the variance in the subject's ratings obtained for one politician and a given Peabody quadruple should be accounted for by the descriptive factor "end of bipolar dimension" and thus, the absolute size of the main effect *DES* for the descriptive factor should be comparatively large:

$$DES = |z_{a+} + z_{a-} - z_{b+} - z_{b-}|/2.$$

If on the other hand, a subject primarily considers the affective aspects at the expense of descriptive information, the traits "bold" and "cautious" should receive high ratings, whereas the traits "reckless" and "timid" should receive comparatively smaller ratings, as shown in Fig. 3, if the subject's attitude toward the person rated is positive. Thus, most of the variance should be accounted for by the affective factor, so that the second main effect *EVAL* should dominate the ratings:

$$EVAL = |z_{a+} + z_{b+} - z_{a-} - z_{b-}|/2.$$

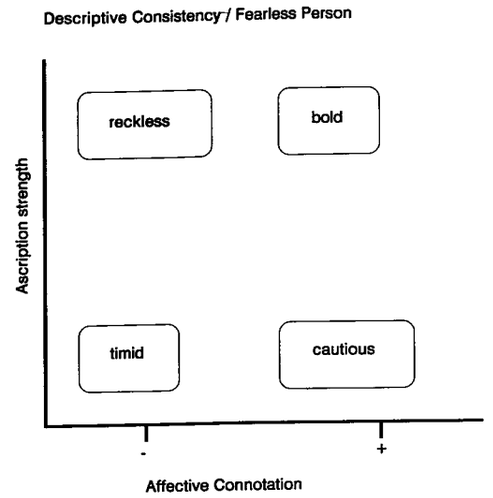

FIG. 2. Ratings dominated by descriptive consistency.

The two main effects are based on orthogonal contrasts and can vary independently. Since both indices as well as their difference also vary as a function of the overall variance of a subject's ratings, each subject's ratings were z-transformed before computing the subjectwise means of the indices *DES* and *EVAL*.

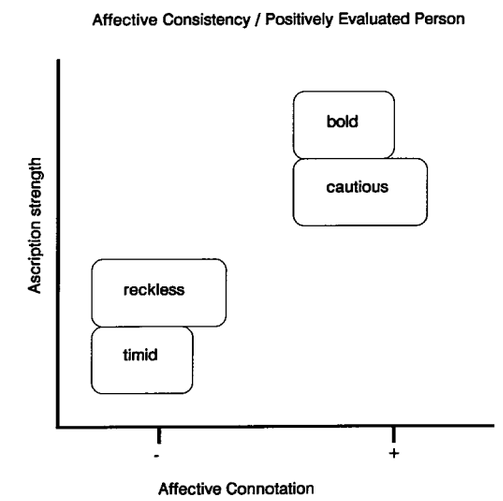

FIG. 3. Ratings dominated by affective consistency.

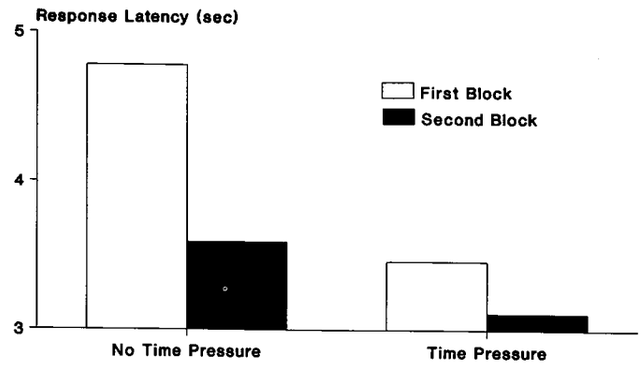

FIG. 4. Mean response latencies in Experiment II as a function of block and time pressure.

Manipulation check. Response latencies were log-transformed, and Fig. 4 shows the (antilogs of the) average response latencies per group and block. A MANOVA with between-subjects factor "time pressure" and within-subjects factor "block" reveals a main effect for time pressure ($F(1, 38) = 8.78, p < .01$) as well as a main effect for block ($F(1, 38) = 68.82, p < .001$). There is also a substantial interaction of the two factors ($F(1, 38) = 14.37, p = .001$). As shown in Fig. 4, subjects worked faster under time pressure than without time pressure and also worked faster in the second block than in the first. In addition, the difference between groups is larger in the first block than in the second, possibly reflecting a ceiling effect such that subjects working under time pressure in the second block approached the upper limit of the speed with which judgments of this kind can be delivered. All in all, the manipulation was successful. One-sided tests within blocks show that, under time pressure, subjects, as expected, worked faster during the first block of ratings ($t(38) = 3.8, p < .001$) as well as during the second block ($t(38) = 1.81, p < .05$).

Affective versus descriptive consistency: Results. The indices *EVAL* and *DES* were submitted to a MANOVA with between-subjects factor "time pressure" and within-subjects factors "block" and "consistency" (index *EVAL* vs index *DES*). The overall contribution of affective consistency was found to be larger than that of descriptive consistency: $F(1, 38) = 8.78, p < .01$. This main effect was moderated by two significant two-way interactions. As expected, the consistency factor interacted with time pressure ($F(1, 38) = 10.11, p < .01$) and also, unexpectedly, with the block factor ($F(1, 38) = 4.98, p < .05$). No other main effect or interaction reached significance. Figure 5 shows the interaction of the consistency factor with time pressure. One-sided t tests reveal, as predicted, that the contribution of affective consistency increases under time pressure

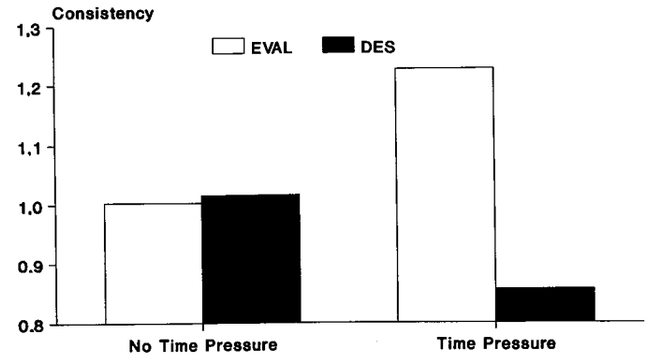

FIG. 5. Averages of indices *EVAL* and *DES* in Experiment II as a function of time pressure.

($t(38) = -3.28, p < .01$), while the contribution of descriptive factors decreases ($t(38) = 2.71, p < .01$).

Figure 6 shows the interaction of the consistency and the block factor. Two-sided paired-sample t tests reveal a tendency for the contribution of affective consistency to increase from the first to the second block ($t(39) = -1.81, p = .08$) and a simultaneous decrease in the contribution of descriptive consistency ($t(39) = 2.39, p < .05$).

The amount of unsystematic errors. To assess the amount of unsystematic errors, the correlation of the 64 ratings of the first and the second block were computed for each subject. In the groups working with and without time pressure, the average reliability across subjects amounts to .73 and .75, respectively. The difference is not significant ($t(38) = .30, p = .768$).

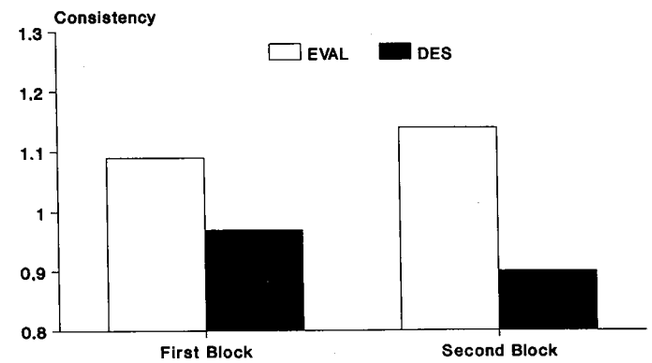

FIG. 6. Averages of indices *EVAL* and *DES* in Experiment II as a function of block.

Discussion

Ratings under time pressure are more strongly biased in the direction of affective consistency than ratings without time pressure. It is possible that ratings under time pressure also differ from ratings without time pressure in other aspects not related to affective or descriptive consistency. For example, ratings under time pressure might be more extreme and polarized than ratings without time pressure. Such tendencies should not, however, affect the indices *DES* and *EVAL* differentially and therefore cannot account for the observed shifts in the relative weights of affective and descriptive consistency. Furthermore, it is interesting to note that both groups worked considerably faster for the second block of ratings, as if they had to some degree imposed time pressure on themselves. Consistent with the two-process model, the relative contribution of affective consistency also increases from first to second block. The latter finding indicates that the effect may not be limited to time pressure induced by explicit instructions, but generalizes to self-imposed time pressure. It is also interesting to note that Jamieson and Zanna (1989) report similar effects of time pressure in a situation requiring subjects to appraise new information, suggesting that similar mechanisms may be at work when the relevant information stems from external rather than internal (memory) sources.

While the amount of systematic halo bias was found to increase under time pressure, a comparable increase of *unsystematic* errors could not be shown. Rather, ratings under time pressure were found to be as reliable as ratings obtained in a situation where there was time for thought, indicating that in both situations, subjects base their judgments on a relatively stable, if different basis. When subjects work under time pressure, this stable basis is determined to a greater extent by a first response tendency extracted from the affective components of the judgmental situation according to the two-process model, a point of view that is supported by the finding that ratings under time pressure are *systematically* biased in the direction of affective consistency.

GENERAL DISCUSSION

The first experiment as well as similar findings by Klauer and Stern (1989) suggest that the affective component of attitudes as well as the affective connotation of traits may be activated automatically and that the activated affective aspects may undergo additional automatic processing, specifically that they may be automatically compared with respect to affective consistency. The degree of affective consistency forms the basis for a spontaneous feeling of plausibility that we have termed the a priori hypothesis or first response tendency for the appropriate response. The extraction of affective consistency appears to be both spontaneous and inescapable. Even though affective as well as descriptive aspects of the stimulus words were irrelevant to the task that subjects were required to perform, evidence was found in the first experiment that affective consistency is extracted upon exposure to appropriate pairs of attitude objects and traits, resulting in a positive or negative response tendency. The second experiment suggests that the response tendency based on affective consistency functions as an a priori hypothesis for the adequate response in trait ratings. According to the two-process model, ratings are based on an a posteriori hypothesis. The a posteriori hypothesis integrates relevant descriptive information that is accessed during a controlled memory search with the automatically produced a priori hypothesis. If the controlled memory search is abbreviated through time pressure, the a priori hypothesis undergoes less correction by relevant descriptive information and has greater relative weight in the subjects' judgments. Thus, judgments under time pressure should exhibit more bias in the direction of affective consistency. Furthermore, since the a priori hypothesis itself is based on relatively stable evaluative information, the reliability of ratings obtained under time pressure should not suffer. Both predictions were upheld, that is, memory-based judgments obtained under time pressure exhibited more pronounced *systematic* bias in the direction of affective consistency, while no increase of *unsystematic* errors was found.

In accounting for bias toward affective consistency, we were primarily interested in a possible independent contribution of affective consistency and have accordingly presented the model with an emphasis on the affective aspects of stimuli, whereas we attempted to control for possible contributions of the descriptive aspects of stimulus words through the use of Peabody quadruples. It is quite likely, however, that the spontaneous plausibility check postulated here can also integrate descriptive aspects. Devine (1989), for example, has proposed a two-process model in the context of racial stereotypes that argues that specific stereotype-congruent beliefs acquired through a lifetime of socialization experiences can be activated automatically even by nonprejudiced persons and must be inhibited in a subsequent, controlled process. Although specific beliefs, rather than an overall feeling of plausibility, are activated according to Devine's model, it seems likely that a spontaneous plausibility check as postulated by the present model would also integrate strong stereotype-congruent, descriptive associations. Similarly, it is plausible that strong learned associations between a person and a given trait (such as Einstein and "intelligent") could mediate a spontaneous feeling of plausibility upon exposure to such juxtapositions.

The present model differs from related models (Fazio, 1986; Fazio et al., 1986, Dovidio et al., 1986; Devine, 1989; Perdue et al., 1990) in that it does not postulate a "spreading-activation" or similar mechanism. According to spreading-activation accounts, associated material and cate-

gories can be activated upon exposure to certain attitude objects and subsequently are likely to be accessed and used in the course of one's reflections. In contrast, the present model claims that the affective components of both the attitude object and the trait in question are activated separately and used to perform a spontaneous plausibility check, prior to reflecting on the judgmental question itself in a controlled process. The latter point of view received strong support in the first experiment, where the facilitation for affectively consistent material—as predicted by spreading-activation accounts and observed in the group responding "correct"—turned into inhibition in the group responding "false." Spreading-activation accounts and the present approach do not contradict each other, however, and may both capture important independent aspects of the processes whereby attitudes bias judgments.

The present account is also consistent with the relatively small success of manipulations aimed at reducing biases in the direction of affective consistency such as instructing subjects to provide neutral judgments, to disregard their attitudes (Nisbett & Wilson, 1977; Wetzel et al., 1981), or to consider themselves members of a neutral jury (Lord et al., 1979) because the extraction of affective consistency that is responsible for the biases is not under subjects' control. On the other hand, if subjects are explicitly instructed to extract and consider an a priori hypothesis that proceeds from the *opposite* attitude, this a priori hypothesis and the one automatically derived should approximately neutralize each other. That idea is realized in Lord et al.'s (1984) *consider-the-opposite* instruction for reducing biases in the direction of affective consistency. The consider-the-opposite strategy, adapted to the present situation, would ask subjects to consider in the course of their reflections how their judgment would be if their global evaluation of the person to be judged were the opposite. The consider-the-opposite strategy has in fact been found successful in eliminating biases in the direction of affective consistency in the context of biased assimilation of information (Lord et al., 1984) as well as for the memory-based judgments considered here (Klauer, 1988).

If the present two-process model continues to receive empirical support, a number of interesting questions remains to be addressed. For example, the role of attitude strength or accessibility (Fazio, 1986) as well as the strength of the affective connotation of the trait have not yet been investigated systematically. It would also be interesting to see whether first response tendencies are automatically extracted from the affective aspects of social situations to guide more general attitude-related behavior than the memory-based judgments considered here.

REFERENCES

Bargh, J. A. (1984). Automatic and conscious processing of social information. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition.* Hillsdale, NJ: Erlbaum.

Bargh, J. A., Bond, R. N., Lombardi, W. J., & Tota, M. E. (1986). Automatic information processing and social perception: The influence of trait information presented outside of conscious awareness on impression formation. *Journal of Personality and Social Psychology,* **50,** 869–878.

Borkenau, P., & Ostendorf, F. (1987). Ein Versuch der Trennung deskriptiver und evaluativer Bedeutungskomponenten in Eigenschafts-beschreibenden Begriffen [An attempted separation of descriptive and evaluative components in trait concepts]. *Archiv für Psychologie,* **139,** 189–207.

Devine, P. G. (1989). Stereotypes and prejudice: their automatic and controlled components. *Journal of Personality and Social Psychology,* **56,** 5–18.

Dovidio, J. F., Evans, N., & Tyler, R. B. (1986). Racial stereotypes: The contents of their cognitive representations. *Journal of Experimental Social Psychology,* **22,** 22–37.

Erdley, C. A. & D'Agostino, P. R. (1988). Cognitive and affective components of automatic priming effects. *Journal of Personality and Social Psychology,* **54,** 741–747.

Fazio, R. H. (1986). How do attitudes guide behavior? In R. M. Sorrentino & E. T. Higgins (Eds.), *The handbook of motivation and cognition: Foundations of social behavior.* New York: Guilford Press.

Fazio, R. H., Chen, J., McDonel, E. C., & Sherman, S. J. (1982). Attitude accessibility, attitude-behavior consistency, and the strength of the object-evaluation association. *Journal of Experimental Social Psychology,* **18,** 339–357.

Fazio, R. H., Powell, M. C., & Herr, P. M. (1983). Toward a process model of the attitude-behavior relation: Accessing one's attitude upon mere observation of the attitude object. *Journal of Personality and Social Psychology,* **44,** 723–735.

Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology,* **50,** 229–238.

Fazio, R. H. & Williams, C. J. (1986). Attitude accessibility as a moderator of the attitude-perception and attitude-behavior relations: An investigation of the 1984 presidential election. *Journal of Personality and Social Psychology,* **51,** 505–514.

Felipe, A. I. (1970). Evaluative vs. descriptive consistency in trait inferences. *Journal of Personality and Social Psychology,* **16,** 627–638.

Heathcote, A. (1988). Screen control and timing routines for the IBM microcomputer family using a high-level-language. *Behavior Research Methods, Instruments, and Computers,* **20,** 289–297.

Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology,* **13,** 141–154.

Houston, D. A., & Fazio, R. H. (1989). Biased processing as a function of attitude accessibility: Making objective judgments subjectively. *Social Cognition,* **7,** 51–66.

Jamieson, D. W., & Zanna, M. P. (1989). Need for structure in attitude formation and expression. In A. R. Pratkanis, S. J. Breckler, & A. G. Greenwald (Eds.), *Attitude structure and function.* Hillsdale, NJ: Erlbaum.

Klauer, K. C. (1988). Die Consider-the-opposite Strategie für Korrektur und Analyse von Halo-Fehlern [The consider-the-opposite strategy for correcting and analysing halo errors]. *Zeitschrift für experimentelle und angewandte Psychologie,* **35,** 597–627.

Klauer, K. C. & Stern, R. (1989). Hinweise für die automatische Auswertung affektiver Konsistenz [Evidence for the automatic extraction of affective consistency]. *Psychologische Beiträge,* in press.

Lohman, D. F. (1989). Estimating individual differences in information processing using speed-accuracy models. In R. Kanfer, P. L. Ackerman, & R. Cudeck (Eds.), *The Minnesota symposium on learning and individual differences: Abilities, motivation, and methodology* (pp. 119–163). Hillsdale, NJ: Erlbaum.

Lord, C. G., Lepper, M. R., & Preston, E. (1984). Considering the opposite: A corrective

strategy for social judgment. *Journal of Personality and Social Psychology,* **47,** 1231–1243.

Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology,* **37,** 2098–2109.

Marcel, A. J. (1983). Conscious and unconscious perception. *Cognitive Psychology,* **15,** 197–300.

Millar, M. G. & Tesser, A. (1986). Thought-induced attitude change: The effects of schema structure and commitment. *Journal of Personality and Social Psychology,* **51,** 259–269.

Neely, J. H., Keefe, D. E., & Kent, L. R. (1989). Semantic priming in the lexical decision task: roles of prospective prime-generated expectancies and retrospective semantic matching. *Journal of Experimental Psychology: Memory, Learning, and Cognition,* **15,** 1003–1019.

Nisbett, R. E. & Wilson, T. (1977). The halo effect: Evidence for unconscious alteration of judgments. *Journal of Personality and Social Psychology,* **35,** 250–256.

Peabody, D. (1967). Trait inferences: evaluative and descriptive aspects. *Journal of Personality and Social Psychology Monograph,* **2,** 1–18.

Perdue, C. W., Dovidio, J. F., Gurtman, M. B., & Tyler, R. B. (1990). Us and them: Social categorization and the process of intergroup bias. *Journal of Personality and Social Psychology,* **59,** 475–486.

Posner, M. I. & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola symposium.* Hillsdale, NJ: Erlbaum.

Sherman, S. J., Judd, C. M., & Park, B. (1989). Social cognition. *Annual Review of Psychology,* **40,** 281–326.

Shiffrin, R. M. & Dumais, S. T. (1981). The development of automatism. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition.* Hillsdale, NJ: Erlbaum.

Shiffrin, R. M. & Schneider, W. (1977). Controlled and automatic human information processing. II. Perceptual learning, automatic attending, and a general theory. *Psychological Review,* **84,** 127–190.

Stroop, J. R. (1938). Factors affecting speed in serial verbal reactions. *Psychological Monographs,* **50,** 38–48.

Tesser, A. (1978). Self-generated attitude change. *Advances in Experimental Social Psychology,* **11,** 289–338.

Tesser, A. & Leone, C. (1977). Cognitive schemas and thought as determinants of attitude change. *Journal of Experimental Social Psychology,* **13,** 340–356.

Wetzel, C. G., Wilson, T. D., & Kort, J. (1981). The halo effect revisited: Forewarned is not forearmed. *Journal of Experimental Social Psychology,* **17,** 427–439.

Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist,* **35,** 151–175.

Zillig, M. (1928). Einstellung und Aussage [Attitude and statement]. *Zeitschrift für Psychologie,* **106,** 58–106.