

Requesting Information to Form an Impression: The Influence of Valence and Confirmatory Status

VINCENT Y. YZERBYT AND JACQUES-PHILIPPE LEYENS

Department of Psychology/EXSO, Catholic University of Louvain at Louvain-la-Neuve, 20 Voie du Roman Pays, B-1348 Louvain-la-Neuve, Belgium

Received November 1, 1989

The present study examines the impact of the valence and the confirmatory status of trait information on impression formation by giving subjects an opportunity to play an active role in their reception of information. In two experiments subjects indicated who, out of a series of candidates, corresponded to an a priori personality profile. These profiles were either positive or negative in tone. Also, whereas some profiles comprised univalent evidence (all personality traits were either positive or negative), others included mixed evidence (most traits were either positive or negative together with a few that were opposite). On the basis of earlier research on dispositional inference using behavioral instances as the cues for judgment, we expected the positive traits to carry less weight than the negative ones. Thus, subjects were predicted to require a greater number of positive traits than negative ones before making a decision. Subjects were also expected to request more traits when these confirmed rather than disconfirmed their expectations. The obtained pattern of data strongly supports our hypotheses. As a whole, the data indicate the existence of a "necessity-oriented" decision rule, as well as the presence of strong constraints in people's implicit theories of personality. Implications are drawn for further research on stereotypes and on lay epistemology. © 1991 Academic Press, Inc.

INTRODUCTION

Assume that you, a psychologist, have to interview several applicants for a job as a used car salesman. To your first question "Do you like to talk and laugh with people?", one candidate replies "Not at all, I am too

This research was supported by Grant FNRS 1.5.323.88F from the Belgian National Science Foundation and facilitated by USIA Grant IA-AEGH-G6192688. The authors thank Wendy Copes, Susan Fiske, Dave Messick, as well as two anonymous reviewers, for their helpful comments on earlier versions of this article. Appreciation is expressed to Pierre Philippot for his assistance in carrying out Experiment 1. Requests for reprints should be sent to V. Yzerbyt or J-Ph. Leyens, University of Louvain, Department of Psychology, Voie du Roman Pays 20, B-1348 Louvain-la-Neuve, Belgium, or to the following BIT-NET/EARN address: YZERBYT@BUCLLN11.

shy." Will you continue the interview? Probably not, but you would have, and extensively so, if the answer had been "yes," and you wanted to be sure you had a good candidate. Suppose now you are a teacher who suspects that several of your students have cheated on an exam. You ask every student if he or she is one of the cheaters. If the answer is "no," there is a good chance that you will probe a little more than if the answer had been "yes."

These two examples show that when it comes to searching for information, the valence of the information received is important, as is its confirmatory status. Indeed, these two variables have already received a lot of attention in the person perception literature (Higgins & Bargh, 1987; Fiske & Taylor, 1984; Schneider, Hastorf, & Ellsworth, 1979; Sherman, Judd, & Park, 1989). Still, in most research on impression formation, the rule has been to provide subjects with fixed sets of items or questions and to observe how subjects' reactions are affected (for a review, see Fiske & Taylor, 1984). Little attention has thus been paid to when and why the social perceiver decides that the evidence provides a satisfactory answer. Such neglect appears surprising if one refers to our daily interactions. Sometimes we make snap judgments; at other times, we are more cautious and rely on a lot of information before making a decision. In its adopted procedure, the present study introduces a new paradigm in person perception research: our subjects are allowed to stop the information search whenever they want. A second innovation of this study is its attempt at extending the phenomenon of negativity effect in impression formation from concrete behaviors to more abstract trait information cues. Finally and most importantly, our paradigm allows the valence and the confirmatory status of the information to be studied in conjunction. Thus revealed is the impact of both variables on social judgment, which, in the context of the present experiments, concerned the suitability for a job.

Negativity Effect in Person Perception

Clearly all traits are not treated equally by the social perceiver (Asch, 1946): Some greatly influence the final impression because of their early position in a list; others have been called central because their presence or absence affects the impression more than that of other traits. The structural relationships between traits seem to have a direct influence on impression formation (Hamilton & Zanna, 1974; Rosenberg, Nelson, & Vivekanathan, 1968; Rosenberg & Sedlak, 1972).

Negative traits seem to be yet another influential category of traits and this phenomenon has been coined the "*negativity effect*" in person perception. Although it is well documented, researchers have not yet reached a consensus about the reason for its existence. Moreover, attention has often been focused on negative behaviors instead of negative traits.

One explanation for the negativity effect comes from correspondent

inference theory (Jones & Davis, 1965): Negative behaviors have a predominant influence because they are counternormative. In other words, being more costly to the actor, they are unexpected (Ajzen & Fishbein, 1975; Jones, Davis, & Gergen, 1961). For example, studies have shown that a single instance of a counternormative behavior has an overriding impact on the overall impression of a target (Birnbaum, 1972, 1973). More recently, Fiske (1980) has argued that since negative (as well as extreme) information is less frequent than positive information (Kanouse & Hanson, 1971), it is more novel and more informative, therefore has greater weight, and thus garners more attention.

As a third approach, Reeder and his colleagues (Reeder, 1985; Reeder & Brewer, 1979; Reeder & Coovert, 1986; Reeder, Messick, & Van Avermaet, 1977; Reeder & Spores, 1983) have defended the idea that negative information generally carries more weight than positive information because of the expected implicit relations between trait categories and behaviors. For some traits (e.g., ability- or morality-related traits), individuals at one extreme of the dispositional continuum are associated with more behaviors than are individuals at the other extreme. For instance, killing someone who passes you on a highway can be performed only by very mean people, whereas donating to the Red Cross can be performed both by nice and mean people. Thus, negativity effects may be expected whenever negative behaviors are seen as more diagnostic than positive behaviors, because positive behaviors are characteristic of people belonging to either the good or the bad categories. More recently, Skowronski and Carlston (1987) formulated a related cue-diagnostics explanation for the negativity, extremity, and also positivity effects. For these authors, diagnosticity means that some attributes are more helpful than others in discriminating between alternative categories.

Clearly the correspondent inference, the novelty, and the schematic (or diagnosticity) models are not totally incompatible with each other (for a detailed discussion, see Skowronski & Carlston, 1989).

Confirmation and Need for Information

A second dimension of the information which appears to be of central concern for the social perceiver is its confirmatory status. Indeed, the question of how people test hypotheses about other people has been extensively investigated over the last decade (Higgins & Bargh, 1987; Klayman & Ha, 1987). The problem has long been whether people look for information biased toward confirmation (e.g., Snyder & Swann, 1978; Snyder & Gangestad, 1981) or for diagnostic information (Trope & Bassok, 1982, 1983; Trope, Bassok, & Allon, 1984). However, Swann and Giuliano (1987) have recently questioned the "utility of contending that people prefer diagnostic information over confirmatory information since they are often one and the same in the eye of perceivers" (p. 511).

Besides generating or choosing questions, subjects are also dealing with the answers they get. As a consequence, the study of hypothesis testing may be enriched by looking at the way subjects react to the information they encounter. When subjects gather evidence in order to test a given hypothesis, at least two types of decision rules might take place (Lewicka, 1988). First, the *sufficiency-oriented* rule corresponds to a subject's privileged search for characteristics which confirm the category. If this rule applies, the hypothesis is deemed correct as soon as a reasonable number of confirming features have been identified. When disconfirming information is encountered, the perceivers keep on searching for additional confirming features. Under the second strategy, the *necessity-oriented* rule, subjects are alert to disconfirmation. More specifically, information that runs counter to one's hypothesis will have a larger influence than confirming information.

Research suggests that these two information-processing strategies do not enjoy equal status (Lewicka, 1988; Peeters & Czapsinski, 1990). Whenever the characteristics of the target object prevail in the decision-making situation, a "disconfirmation effect" would generally be expected; that is, subjects will favor false negative decisions over false positive ones. In other words, they adopt the working hypothesis that (1) a lot of confirmatory evidence is needed for an option to be acceptable and (2) very little disconfirmatory evidence is needed to establish its unacceptability.

Typically, subjects' strategies tend to be anchored in the object when information-gathering behavior is guided by specific motivations, e.g., the "accuracy-goal" (Kruglanski & Ajzen, 1983; Neuberg, 1989). Similarly, we expect the necessity rule to operate when subjects control the amount but not the nature of the target information which is given to them. If the necessity-oriented rule is at work, social perceivers should request more evidence when they encounter information that confirms an a priori hypothesis than when they gather information that contradicts that same hypothesis. This also means that with symmetrically created stimuli, those having an equal number of confirming and disconfirming features, subjects will reject more often than accept any given target. Clearly, this pattern is quite distinct from the traditional confirmation bias and corresponds to the second central idea to be tested in the present study.

Presentation of the Studies and Hypotheses

To test the above ideas, the following original paradigm was employed: subjects were asked to decide whether the candidates for parts in a theater play had the same conception of their future roles as the director, and were thus well suited for these roles. Candidates had supposedly been prompted to indicate their conceptions of how the role should be portrayed by selecting personality traits from a list. Candidates' profiles were positive or negative. Also, whereas some profiles were displaying univalent evi-

dence (the traits were all positive or all negative), others were characterized by mixed evidence (the majority of traits were positive or negative). Finally, half of the subjects were led to believe that the director saw the role as being rather likable, while the remaining subjects learned that the director conceived of the role as being rather unlikable. Subjects knew how much information was available but could make a decision as soon as they felt confident. This paradigm enabled the study of the nature of the decision, the amount of information needed, and the confidence in the decision.

We expected our subjects to accept or to reject the candidates to the extent that they encountered congruent or incongruent information. Still, in line with the necessity-oriented rule, it was predicted that subjects would reject more candidates than they accept. Also, it seemed reasonable to expect the subjects to be more confident about rejection than about acceptance.

As far as the information search is concerned, two processes were predicted to occur. First, in line with the negativity effect, we hypothesized that subjects would request less information when confronted with negative than with positive information. Second, because of the disconfirmation effect, we expected the subjects to request more information when incoming evidence confirmed rather than disconfirmed the role they had in mind.

In our experimental design, all subjects received positive and negative information whether they were informed that the role was seen as likable or unlikable. As a consequence, the two above hypotheses result in an interaction between the role and the type of information provided. Most evidence will be required when subjects receive positive information and have to decide about a likable role. The least information will be needed when they receive negative evidence for a likable role. In between these two extremes, the confirmation and disconfirmation of the unlikable role should induce a request for about an equal amount of information. Moreover, one could also hypothesize a main effect for positivity if this bias is sufficiently strong. Finally, we predicted that negative information following positive items would be more influential for the judgment than the reverse sequence (Birnbau, 1972, 1973). Whereas all subjects of Experiment 1 were presented with the same set of candidates' lists, in Experiment 2 we provided a unique configuration of evidence to each subject.

EXPERIMENT 1

Method

Subjects

Forty-eight female students from the University of Louvain at Louvain-la-Neuve volunteered to participate in this experiment.

Procedure

When the subject arrived at the laboratory, she was seated in front of a microcomputer. There, a female experimenter explained that she wanted to know how lay people select actors for a play. "Professional directors," she went on, "often choose an actor on the basis of the correspondence between the image they have in mind and the way the actor understands the character. Of course, one given character can be understood in several ways, but only one way usually suits the director's idea." The experimenter then asked the subject to imagine that she was a director who saw the role as being a "rather likable" (vs "rather unlikable") person. Fourteen candidates for the role supposedly had been presented with a set of personality traits and had circled the 10 traits that best conveyed the way they believed the role should be portrayed. The subject's task was to select as many actors, out of the 14 candidates, whose conceptions fit her specific reading of the scenario. Also, the subject learned that to the extent she could make a confident decision, she was not expected to use all 10 pieces of information concerning each candidate. The experimenter then answered all questions about the instructions and the use of the computer, started the program and left the room.

For each of the 14 candidates, the subject saw a maximum of 10 traits, one trait at a time. When a trait appeared on the screen, the subject was always reminded of the number of the candidate and had to indicate whether she wanted additional information, i.e., another trait, or whether she considered that the candidate was to be rejected or accepted. Once a decision had been made, the screen cleared and the subject was asked to indicate her level of confidence about the decision. When all candidates had been presented to the subject, a manipulation check appeared on the screen concerning the likability of the role as seen by the director. The subject was then informed that the experiment was over and was thoroughly debriefed.

Materials

To construct the 14 lists of 10 traits, we selected traits out of a set of 121 which had been pretested for positivity (1 = very negative, 9 = very positive) by a sample of 60 subjects taken from the same population.¹ Traits could appear no more than once in any given list and no more than three times in the total set of 14 lists.

First, traits ranging from 4.4 to 5.8 were used to construct two neutral filler lists. Next, three positive lists were made by randomly taking 10 traits which ranged from 5.8 to 8, and three negative lists were made by randomly taking 10 traits which ranged from 2 to 4.4. Last, three lists were constructed by randomly taking six positive and four negative traits (mainly positive lists), and three by randomly taking six negative and four positive traits (mainly negative lists). Within the former set, traits were ordered according to one of three patterns: 1) $++/- -/+ + - -/+ +$; 2) $+++/- - - -/+ + + -$; and 3) $+++ + / - - - - / + +$. The same pattern was used to create the three mainly negative lists. These six mixed lists were introduced mainly to increase realism but also to allow additional tests of the hypotheses. Finally, six orders of presentation were used. The two neutral lists were always put first. One of the three all-positive or one of the three all-negative lists immediately followed.² The eleven remaining lists were then presented in a random order.

¹ The 153 traits used by Rothbart and Park (1986) were translated into French. Only those traits that were clearly understood by 54 out of 56 pilot subjects from the same population were retained for further use, leaving a total of 121 terms.

² We examined the impact of starting with the third list being positive or negative. The Role \times Valence of the List interaction proved to be slightly weaker when the third list was a negative one and was thus qualified by the Order of Presentation, $F(1, 44) = 4.35, p <$

Dependent Measures

For each of the 14 candidates, the subject provided four dependent variables. First, the subject decided either to "reject" or to "accept" the actor. The level of confidence for the decision was also recorded and ranged from "moderate" (=1) to "very high" (=9). The sheer number of traits that subjects requested, ranging from one (=1) to all (=10), constituted the third dependent variable. At the end of the experiment, the subject indicated the likability of the role as seen by the director, on a scale ranging from "very unlikable" (=1) to "very likable" (=9).

Data pertaining to the two filler lists were discarded. The proportions of the four kinds of lists accepted by the subject were calculated for each subject. All three lists of a given kind were thus treated as conceptual replications of each other.³ The basic design was a 2 (Role: likable vs unlikable) \times 2 (Valence of the List: positive vs negative) \times 2 (Nature of the List: univalent vs mixed), the first factor as a between-subjects variable and the last two as within-subject variables. Because the univalent and mixed lists addressed slightly different questions, separate analyses were conducted for both the decisions about acceptance and the proportion of traits dependent variables.

Results

Manipulation Checks: Likability of Candidates

Subjects who thought that the character was "rather likable" rated the role as being much more likable ($M = 7.00$) than subjects who learned that the director saw the part as "rather unlikable" ($M = 2.88$), $F(1, 47) = 110.99$, $p < .001$. Both ratings differed significantly from the scale's midpoint, $t(23) = 8.06$, $p < .001$, and $t(23) = -7.02$, $p < .0001$, for the likable and the unlikable conditions, respectively. Clearly then, subjects remembered how the part was seen by the director.

The Univalent Lists

The univalent lists allowed us to test more precisely the predictions pertaining to the number of positive or negative pieces of information requested by the subjects.

Decisions about acceptance. To check whether subjects accepted those candidates who fit the director's requirements and also rejected those who

.05. However, in light of the huge F value for the interaction of interest and because the pattern was similar in both cases, data pertaining to the two orders of presentation were collapsed. In addition, Order of Presentation did not qualify any of the other effects to be reported for this experiment.

³ To ensure that all 3 representatives of each kind of lists were conceptual replications of each other, we performed a $2 \times 2 \times 3$ nested-design ANOVA on both decision and proportion of traits with Role as a between-subjects variable and Valence and List as within-subject variables, List being nested within Valence. The List within Valence main effect proved significant in both cases, $F(4, 184) = 3.36$, $p < .02$, and $F(4, 184) = 3.55$, $p < .01$, for subjects' decisions and proportion of traits, respectively. This simply indicates that the lists evoke slightly different levels of responses. However, a significant Role \times List within Valence was found, $F(4, 184) = 8.09$, $p < .001$, for the decisions. One problematic list actually worked against our hypotheses but we decided to keep the set of 6 lists.

TABLE 1
PROPORTION OF CANDIDATES ACCEPTED AS A FUNCTION OF ROLE AND VALENCE OF THE LIST:
EXPERIMENT 1

Role	Valence of list			
	All positive	All negative	Mainly positive	Mainly negative
Rather likable	.76 ^a	.10	.38	.08
Rather unlikable	.06	.67	.33	.58

^a Scores range from 0.00 (= all 3 candidates were rejected) to 1.00 (= all 3 candidates were accepted).

did not, subjects' decisions were submitted to a 2×2 mixed-design ANOVA with Role (likable vs unlikable) as a between-subjects variable and Valence of the List (positive vs negative) as a within-subject variable. The ANOVA only revealed a strong Role \times Valence of the List interaction, $F(1, 46) = 243.34, p < .0001$. As expected, when the list corresponded to the role, subjects accepted the candidates, but they rejected the candidates when there was a lack of fit between traits and role (see Table 1).

Proportion of traits used. The proportions of traits used by subjects for each kind of lists were entered in a 2×2 (Role \times Valence of the List) mixed-design ANOVA. Both a main effect of Valence of the List and the interaction between Role and Valence of the List were significant, $F(1, 46) = 11.90, p < .002$, and $F(1, 46) = 13.04, p < .001$, respectively.

The pattern of results strongly supported our hypotheses (see Fig. 1a).

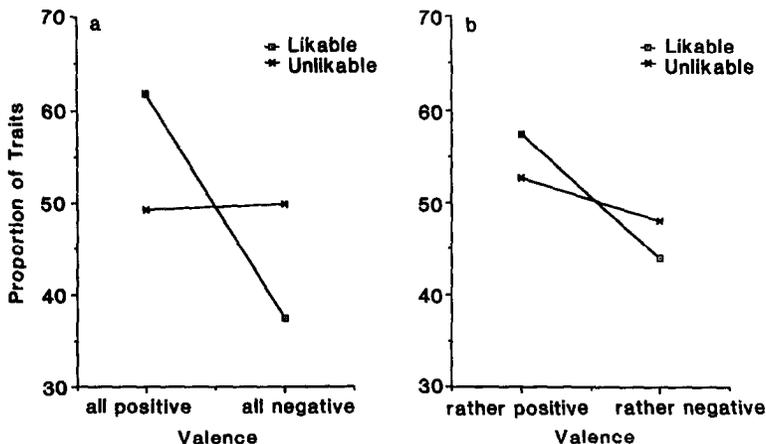


FIG. 1. Proportion of traits requested from the univalent and mixed lists as a function of role and valence of the list: Experiment 1.

On the whole, subjects used more traits when confronted with all-positive candidates ($M = 55.5\%$) than when confronted with all-negative candidates ($M = 43.7\%$). As expected, this phenomenon was qualified by a propensity to request more information when the candidate corroborated the requirements of the director than when evidence clearly contradicted a possible fit between the traits chosen by the candidate and the role. Thus, for the unlikable role, subjects needed as many traits to accept the negative candidates ($M = 49.9\%$) as to reject the positive ones ($M = 49.3\%$), $F(1, 23) < 1$, *ns*. Quite a different picture emerges for the likable role: the subjects asked for much more information when confronted with the all-positive candidates ($M = 61.7\%$) than when confronted with the all-negative ones ($M = 37.5\%$), $F(1, 23) = 22.55$, $p < .001$.

The Mixed Lists

Because confirmation and disconfirmation were both contained in the very same list, the mixed lists offered a means to examine the impact of changing evidence.

Decisions about acceptance. A 2×2 (Role \times Valence of the List) mixed-design ANOVA was run on their decisions. Again, a strong Role \times Valence of the List interaction was found, $F(1, 46) = 28.51$, $p < .0001$, showing that subjects rejected the candidates to the extent that they did not conform to the role (see Table 1). However, a strong Role main effect, $F(1, 46) = 9.31$, $p < .004$, revealed that subjects rejected the candidates more when the role was likable than when the role was unlikable. When they were confronted with a mainly negative candidate, subjects clearly took a different stand depending on the role they had in mind, .08 as opposed to .58, $F(1, 46) = 40.06$, $p < .0001$. No such effect was observed when the candidate was mainly positive, .38 as opposed to .33, $F(1, 46) < 1.00$, *ns*.

Separate analyses were conducted for each of the two least mingled lists, i.e. the ++++/- - - - / ++ and the - - - - / + + + + / - - lists.⁴ Because the valence of information shifted when most subjects made up their mind, the decisions of those subjects who requested no more than four traits and those who needed five or more traits were analyzed using 2×2 ANOVAs with Role (likable vs. unlikable) and Number of Traits (four or fewer vs. five or more) as between-subjects variables (see Table 2). For the mainly positive list, a strong Role main effect emerged, $F(1, 44) = 9.92$, $p < .003$, indicating, as expected, that the subjects

⁴ Only one representative of each pattern of mixed traits had been created within each valence, preventing us from drawing any definite conclusion. Also, subjects were not randomly assigned to conditions and their number in each cell is far from being uniform. Fortunately, the Type I (hierarchical) and Type III (regression) SAS sums of squares led to identical conclusions. We refer to the latter.

TABLE 2
 PROPORTION OF CANDIDATES ACCEPTED FOR THE 4/4/2 LISTS AS A FUNCTION OF ROLE,
 NUMBER OF TRAITS, AND VALENCE OF THE LISTS

Role	Valence of list	
	Mainly positive	Mainly negative
Rather likable ^a		
4 or less	.89 (9) ^b	.00 (19)
5 or more	.20 (15)	.40 (5)
Rather unlikable		
4 or less	.00 (11)	.63 (16)
5 or more	.39 (13)	.63 (8)

^a Scores range from 0.00 (= the candidate was rejected) to 1.00 (= the candidate was accepted).

^b Number in parentheses refer to the number of subjects in the cell.

rejected the positive candidate more in the unlikable role condition. The Role \times Number of Traits interaction was very significant, $F(1, 44) = 23.05$, $p < .0001$, reflecting the impact of the reception of the negative traits on subjects' decisions. When subjects were primed with the unlikable role, the reception of negative information led them to reject the candidate less often. But the most dramatic difference in decisions was observed in the likable role/positive list cell, in which the strong acceptance shifted to strong rejection when the four initial positive traits were followed by negative information. For the negative list, only the Role main effect was significant, $F(1, 44) = 10.58$, $p < .003$. Subjects rejected the candidates more often when the likable rather than the unlikable role was induced. This time, the decisions in the unlikable role/negative list were not influenced by the number of traits requested. The lack of parallelism between these two lists testifies to the larger impact on decisions of the negative inconsistent evidence, as compared to the positive one.

Proportion of traits used. A 2×2 (Role \times Valence of the List) mixed-design ANOVA was run on the proportion of traits requested by subjects (see Fig. 1b). A significant main effect of Valence of the List, $F(1, 46) = 11.10$, $p < .002$, was qualified by a weak Role \times Valence of the List interaction, $F(1, 46) = 2.69$, $p < .11$. In sum, when confronted with mainly positive lists, subjects needed more information to make their final decision ($M = 55.1\%$) than when presented with mainly negative lists ($M = 46.1\%$). However, they also requested more information when the candidate corresponded to the role ($M = 52.9\%$) than when he did not ($M = 48.4\%$). The specific pattern found for the mixed lists closely

parallels the one found for the univalent lists, but it appears to be less pronounced.

Confidence Level

Confidence levels were entered into a $2 \times 2 \times 2$ (Role \times Valence of the List \times Nature of the List) mixed-design ANOVA. None of the effects were significant. Globally, subjects appeared quite confident in their judgments ($M = 7.05$). Surprisingly, the subjects were no less confident when they were confronted with the mixed lists than when the valence of the traits was univocal, $F(1, 46) < 1$, *ns*. Moreover, the lack of second-order interaction indicates that the patterns of data for both the univalent and the mixed lists are not different. Interestingly, and contrary to our hypothesis, the Role \times Valence of the List interaction was not significant, $F(1, 46) = 1.81$, $p < .19$, although subjects tended to report more confidence when the valence of the evidence was inconsistent with that of the role, i.e., when the candidate was rejected ($M = 7.16$), as compared to when the traits confirmed the role, i.e., the target was accepted ($M = 6.94$).

Globally, the results of Experiment 1 strongly supported our main hypotheses. First, for both univalent and mixed lists, the correspondence and the lack of fit between the evaluative tone of the traits and the role resulted in acceptance and rejection of the candidates, respectively. Results from the mixed lists indicated that these candidates were considered less suited for the likable than for the unlikable role. Because for every target with a particular pattern of positive and negative traits there is one for whom the pattern is reversed, the acceptance level was tested against the unbiased rule hypothesis that 50% of the candidates would be accepted. In accordance with the necessity-oriented rule, less than half the candidates were accepted, $M = .40$, $t(47) = -3.99$, $p < .0002$, and $M = .34$, $t(47) = -3.84$, $p < .0004$, for the univalent and mixed lists, respectively. Second, subjects asked for more traits when they encountered positive candidates than when negative information was provided. Also, more information was requested when the evidence conformed to the role than when it did not fit the role.

Still, two weaknesses of Experiment 1 were the lack of replication of the lists for the mixed lists, and the use of lists constructed by the experimenters. These important methodological issues were directly addressed in a second experiment by means of random selection of the traits and replication of the lists. In addition, the confidence data are hardly conclusive, although a trend in the predicted direction was observed. It was important to see whether this effect could be strengthened before abandoning the confidence hypothesis altogether.

EXPERIMENT 2

Method

Subjects

Forty-four female undergraduates from the University of Louvain at Louvain-la-Neuve volunteered to participate.

Procedure

The procedure used in this experiment was almost identical to the one used in Experiment 1. This time a male undergraduate served as the experimenter. Also, 30 instead of 14 candidates were to be judged. The same dependent variables were collected as in Experiment 1.

Materials

The crucial difference from Experiment 1 is that the list construction was repeated for each subject before she actually began the experiment. A computer program was specifically written for this purpose. Traits ranging from 4 to 6 were used to construct six neutral filler lists. Next, 4 lists of 10 positive traits were made by randomly taking traits which ranged from 6 to 8, and 4 lists of 10 negative traits were constructed by randomly taking traits which ranged from 2 to 4. An additional 8 lists were built, 4 by randomly taking 6 positive and 4 negative traits (mainly positive lists), and 4 by randomly taking 6 negative and 4 positive traits (mainly negative lists). Within the former set, the traits were ordered according to the $++++/- - - -/+ +$ pattern, whereas the traits in the latter set were ranged using the $- - - -/+ + + +/ - -$ pattern. To increase realism, 8 lists contained a mixture of positive and negative traits. Specific steps were included in the program to prevent two traits that were the opposite anchors of a bipolar dimension from being selected in the same list. A total of 16 experimental lists pertained to our hypotheses.

The order of presentation of lists was also unique for each subject. It consisted of a blending of systematic and random presentation of the lists. The six neutral lists were always put in positions 1, 2, 11, 12, 21, and 22 to control for both warm-up effects in the beginning of a session and realism during the rest of the experiment. The program then randomly assigned a position to the remaining 24 lists.

Results

Manipulation Checks: Likability of Candidates

Clearly, the role was thought to be much more likable ($M = 7.68$) by subjects who learned that the character was seen as "rather likable," than by subjects who learned that the director saw the part as "rather unlikable" ($M = 1.81$), $F(1, 42) = 109.51$, $p < .0001$. Both ratings were also significantly different from the scale's midpoint, $t(21) = 6.67$, $p < .0001$ and $t(21) = -8.19$, $p < .0001$, for the likable and the unlikable conditions, respectively.

The Univalent Lists

Decisions. A 2×2 (Role \times Valence of the List) mixed design ANOVA was performed on subjects' decisions (see Table 3). A strong Role \times Valence of the List interaction emerged, $F(1, 42) = 241.24$, $p < .0001$.

TABLE 3
 PROPORTION OF CANDIDATES ACCEPTED AS A FUNCTION OF ROLE AND VALENCE OF THE LIST:
 EXPERIMENT 2

Role	Valence of list			
	All positive	All negative	Mainly positive	Mainly negative
Rather likable	.75 ^a	.06	.55	.08
Rather unlikable	.09	.83	.19	.70

^a Scores range from 0.00 (= all 4 candidates were rejected) to 1.00 (= all 4 candidates were accepted).

When the list and the role matched, subjects accepted the candidates. However, subjects rejected the candidates when there was a lack of fit between evidence and role. No other effects were significant.

A check on the overall acceptance levels revealed that, consistent with the necessity-oriented rule, subjects more often rejected than accepted any given target, $M = .43$, $t(43) = -3.09$, $p < .004$.

Proportion of traits used. The proportion of traits requested by subjects for each set of candidates was the dependent variable in a 2×2 (Role \times Valence of the List) mixed ANOVA. Both a main effect of Valence of the List and the interaction between Role and Valence of the List were significant, $F(1, 42) = 7.05$, $p < .02$, and $F(1, 42) = 4.68$, $p < .04$, respectively.

Clearly, the pattern of results found in Experiment 1 was replicated (see Fig. 2a). Subjects requested more traits when confronted with all-

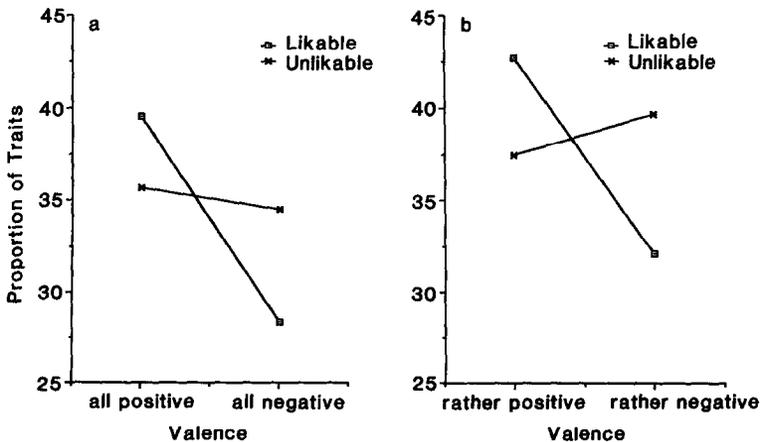


FIG. 2. Proportion of traits requested from the univalent and mixed lists as a function of role and valence of the list: Experiment 2.

positive candidates ($M = 37.6\%$) then when confronted with all-negative candidates ($M = 31.5\%$). As expected, subjects also asked for more information when the candidate corroborated the requirements of the director than when evidence clearly contradicted a possible fit between the traits chosen by the candidate and the role. Thus, for the unlikable role, it took as many traits to accept the negative candidates ($M = 34.5\%$) as to reject the positive ones ($M = 35.7\%$), $F(1, 21) < 1.00$, *ns*. A different picture emerges for the likable role: the subjects requested more traits when presented with all-positive candidates ($M = 39.5\%$) than when confronted with all-negative ones ($M = 28.4\%$), $F(1, 21) = 9.14$, $p < .007$.

The Mixed Lists

Because each subject saw a different set of original lists, it is possible to examine the impact of changing evidence on both the subject's decisions and the number of traits requested. A direct test of the greater impact of negative evidence can also be made through the comparison of the decisions for the two confirmation cells.

Decisions. A 2×2 (Role \times Valence of the List) mixed-design ANOVA was performed and, as expected, a strong Role \times Valence of the List interaction emerged, $F(1, 42) = 73.89$, $p < .0001$. Subjects rejected the candidates when they did not conform to the role (see Table 3). Again, a significant Role main effect, $F(1, 42) = 7.35$, $p < .01$, indicated that subjects rejected the candidates more often when the role was likable than when the role was unlikable.

As was the case for the univalent lists, subjects accepted less than 50% of the candidates they were confronted with, $M = .38$, $t(43) = -4.43$, $p < .0001$. Together with the pattern found for the univalent lists, the acceptance levels for the mixed lists replicate those of Experiment 1 and indicate that subjects may rely on a necessity-oriented rule.

Proportion of traits used. A 2×2 (Role \times Valence of the List) mixed-design ANOVA was conducted on the proportion of traits requested (see Fig. 2b). The Valence of the List main effect was significant, $F(1, 42) = 3.53$, $p < .07$, but was qualified by a significant Role \times Valence of the List interaction, $F(1, 42) = 8.08$, $p < .007$. Again, subjects asked for more information when the candidates conformed to the role ($M = 41.2\%$) than when they did not ($M = 34.8\%$). They also tended to request more traits when confronted with mainly positive candidates ($M = 40.1\%$) than when presented with mainly negative lists ($M = 35.9\%$). As in Experiment 1, these results parallel those found for the univalent lists, albeit in a slightly less pronounced manner.

Confidence Level

Confidence levels were compared using a $2 \times 2 \times 2$ (Role \times Valence of the List \times Nature of the List) mixed-design ANOVA. The subjects were again no less confident when they were confronted with the mixed lists ($M = 7.29$) rather than with the univalent ones ($M = 7.39$), $F(1, 42) < 1$, *ns*. In contrast, both the Valence main effect and the Role \times Valence of the List interaction were significant, $F(1, 42) = 7.48$, $p < .01$, and $F(1, 42) = 17.39$, $p < .0001$. Conforming to the trend observed in Experiment 1, subjects reported more confidence about their decision when they rejected ($M = 7.60$) than when they accepted the candidates ($M = 7.07$). In addition, subjects appeared more confident when the lists displayed were mainly negative ($M = 7.51$) than when mainly positive ($M = 7.16$). The lack of second-order interaction indicates that the data for the univalent and the mixed lists display the same pattern of results.

DISCUSSION

The present results nicely support our hypotheses. First, to make a decision people request more information when confronted with positive rather than negative evidence. In line with previous findings, this phenomenon is apparently due to the fact that unfavorable information carries more weight than does favorable information when one has to decide whether a given person corresponds to a certain profile. Second, subjects request more information when it confirms than when it disconfirms their hypotheses, a finding which agrees with the disconfirmation effect, and runs counter to the "confirmation bias" perspective. The additional finding that subjects rejected the candidates more than they accepted them suggests that a necessity-oriented rule may be at work in the present context.

In previous research on person perception, subjects usually had to make judgments in a fixed-information context. In the present studies, however, they were encouraged to make a decision as soon as they felt confident.

Univalent Evidence

We expected an interaction between the role, i.e., the hypothesis given to the subject (likable vs unlikable), and the information (positive vs negative) for the number of requested traits. In other words, we hypothesized that most traits would be requested when both the role and information provided were positive. This would happen because of the *positivity* and the *confirmatory status* of the traits. We also assumed that fewest traits would be requested when the information was negative for a likable role, as a joint effect of *the negativity* and the *disconfirmatory status* of the traits. According to the same reasoning, the two other cells of the design—for the unlikable role—had to be rather equal in terms of

information requests and fall in between data for the likable role. This is exactly what happened in both experiments.

It may prove helpful to frame the findings in a somewhat different manner in order to clarify the impact of both the valence and the confirmatory status of the evidence. First, Figs. 1 and 2 reveal that the two acceptance conditions (positive list/positive traits and negative list/negative traits) required more information than the two rejection conditions (positive list/negative traits and negative list/positive traits). Thus, unsuitability is more quickly decided than suitability. This highlights the greater influence of the disconfirmatory as compared to the confirmatory evidence. Second, the stronger impact of the negative trait information can be observed within each set of conditions. For the acceptance conditions, subjects need fewer traits to conclude that the target fits for the unlikable than for the likable role. For the rejection cases, however, the subjects request fewer items to judge that the candidate is not suited for the likable compared to the unlikable role.

Another result is important for our hypotheses: in Experiment 2, where the construction of lists was well controlled, subjects were more confident in their judgments when they rejected than when they accepted someone, whether the lists contain univalent evidence or not.

Mixed Evidence

The mixed lists were presented to show that the greater "appetite" for positive traits could be explained by their lesser impact. Although subjects could only know whether the information was mainly positive or negative by looking at all 10 pieces of information, the same interaction pattern was obtained as when the lists were univalent. Thus, even in the case of mixed information, the data provide support for the presence of negativity and disconfirmation effects. Interestingly, data indicate that when subjects received both positive and negative traits, the global impression is rather unlikable.⁵ In other words, negative traits have more weight than positive ones (Briscoe, Woodyard, & Shaw, 1967).

⁵ For both experiments, subjects' decisions were also submitted to a $2 \times 2 \times 2$ (Role \times Valence of the List \times Nature of the List) mixed-design ANOVA. Globally, the Nature of the Lists main effects indicated that subjects tended to reject more candidates when the lists contained mixed evidence, $F(1, 46) = 2.77, p = .10$, and $F(1, 42) = 3.36, p < .08$, for Experiments 1 and 2, respectively. As expected, the Role \times Valence of the List interactions were highly significant, $F(1, 46) = 199.45, p < .0001$, and $F(1, 42) = 184.05, p < .0001$. More interestingly, the three-way interactions, $F(1, 46) = 31.14, p = .0001$, and $F(1, 42) = 18.26, p < .0001$, for Experiments 1 and 2, respectively, revealed that the basic Role \times Valence of the List pattern was significantly less pronounced for the mixed lists than for the univalent ones. The Nature of the List \times Role interaction in Experiment 1, $F(1, 46) = 22.78, p < .0001$, and the combination of the Nature of the List and Role, $F(1, 42) = 6.46, p < .02$, main effects in Experiment 2, point to another consequence of the negativity bias: candidates were rejected most when the role was likable and the information comprised a few negative items.

Reeder and Coovert (1986) tackled the same problem, obtaining an identical conclusion. They used reaction times as a dependent measure to show that the decision was made at the integration rather than at the encoding stage. They hypothesized, and found, that negative information following positive items takes a longer time than the reverse sequence, presumably because unfavorable characteristics lead to more cognitive restructuring work. Fiske (1980, p. 903) was able to show the same pattern when two slides depicting a target follow each other and are highly inconsistent at the evaluative level.

Confidence Levels

In both Experiments 1 and 2, subjects were equally confident in their judgments whether they encountered univalent or mixed evidence. This seems to indicate that our subjects felt quite capable of integrating mixed information in order to make their decision, and certainly no less so than when they received univalent items (for a related argument, see Asch & Zukier, 1984). We predicted that subjects would be more confident in the case of rejection than in the case of acceptance because of the necessity rule. In Experiment 1, the data were in the expected direction but not significantly so. In Experiment 2, however, they were highly significant. This difference in the two experiments may be due to the increased number of lists and/or their better construction. To be sure, confidence levels were requested after the decision had been taken, and therefore cannot be considered as uncontaminated. If indeed there was contamination, we still do not see why it would operate differently for rejection than for acceptance. Obviously, more research is needed in order to uncover the impact of various factors on the confidence level.

Implications for Intergroup Research

On the basis of their correlational study, Rothbart and Park (1986) suggest that the differential confirmatory or disconfirmatory value of positive and negative traits contributes to the maintenance of stereotypes. Our results go exactly in the same direction, but they additionally show the influence of the hypothesis entertained by the subjects. Indeed, we experimentally replicated Rothbart and Park's (1986) results for only the positive hypothesis, i.e., the likable role, not for the negative hypothesis. When the hypothesis is positive, positive information is harder to confirm (i.e., we need more instances to make a decision) than negative information. Such a difference is absent when the hypothesis is negative.

Our results have clear implications if one substitutes in-group members for the positive hypothesis and out-group members for the negative hypothesis. This would correspond to asking subjects whether the described target belongs to their in-group or to an out-group. Although our studies did not allow us to measure the extent to which a target corresponded

to a given profile, but used an all-or-none measure, the data suggest that one easily rejects negative "candidates" for the in-group. Under what they called the "black sheep effect," Marques, Yzerbyt, and Leyens (1988; Marques & Yzerbyt, 1988) have illustrated that such members were judged more extremely than bad out-group members. Although this phenomenon runs counter to the classical expectation for an in-group bias, it clearly indicates that displaying negative features and belonging to the in-group appears incompatible (Leyens & Yzerbyt, 1991; Tajfel, 1969).

CONCLUSION

One clear contribution of our two experiments is that they provide a new research paradigm for the negativity and the disconfirmatory effects. They also extend previous findings concerning the links between behaviors and dispositions to the sheer relation among traits. However, the most significant achievement of the present work is to show that the valence and the confirmatory status of the information clearly have an impact on the path taken by the subjects in their reaching a conclusion about other people. The present task is not unlike real-life situations in which a high number of constraints limits the information search while decisions still need to be made. The procedure apparently more closely replicates actual person perception than does the standard encounter with fixed sets of information, and, as such, may be useful in studying intergroup perceptions. It could also be employed to investigate psychological diagnoses (Rubin & Shontz, 1960; Sines, 1959); such research would complement that of Langer and Abelson (1974) who presented fixed sets of evidence about either a job applicant or a patient. Finally, the paradigm could profitably inspire research to test Kruglanski's (Kruglanski & Ajzen, 1983; Kruglanski & Maysel, 1988) theory of lay epistemology.

REFERENCES

- Ajzen, I., & Fishbein, M. A. (1975). Bayesian analysis of attribution processes. *Psychological Bulletin*, *82*, 261-277.
- Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology*, *41*, 258-290.
- Asch, S. E., & Zukier, H. (1984). Thinking about persons. *Journal of Personality and Social Psychology*, *46*, 1230-1240.
- Birnbaum, M. H. (1972). Morality judgments: Test of an averaging model. *Journal of Experimental Psychology*, *93*, 35-42.
- Birnbaum, M. H. (1973). Morality judgments: Test of an averaging model with differential weight. *Journal of Experimental Psychology*, *99*, 395-397.
- Briscoe, M. E., Woodyard, H. D., & Shaw, M. E. (1967). Personality impression change as a function of the favorableness of first impressions. *Journal of Personality*, *35*, 343-357.
- Fiske, S. T. (1980). Attention and weight in person perception: The impact of negative and extreme behavior. *Journal of Personality and Social Psychology*, *38*, 889-906.
- Fiske, S. T., & Taylor, S. E. (1984). *Social cognition*. Reading, MA: Addison-Wesley.

- Hamilton, D. L., & Zanna, M. P. (1974). Context effects in impression formation: Changes in connotative meaning. *Journal of Personality and Social Psychology*, **29**, 649–654.
- Higgins, E. T., & Bargh, J. A. (1987). Social cognition and social perception. *Annual Review of Psychology*, **38**, 369–425.
- Jones, E. E., & Davis, K. E. (1965). From acts to disposition: The attribution process in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 219–266). New York: Academic Press.
- Jones, E. E., Davis, K. E., & Gergen, K. (1961). Role playing variations and their information value for person perception. *Journal of Abnormal and Social Psychology*, **63**, 302–310.
- Kanouse, D., & Hanson, L. (1971). Negativity in evaluations. In E. Jones et al. (Eds.), *Attribution: Perceiving the causes of behavior* (pp. 47–62). Morristown, NJ: General Learning Press.
- Klayman, J., & Ha, Y. (1987). Confirmation, disconfirmation and information in hypothesis testing. *Psychological Review*, **94**, 211–228.
- Kruglanski, A. W., & Ajzen, I. (1983). Bias and error in human judgment. *European Journal of Social Psychology*, **13**, 1–44.
- Kruglanski, A. W., & Mayseless, O. (1988). Contextual effects in hypothesis testing: The role of competing alternatives and epistemic motivations. *Social Cognition*, **6**, 1–20.
- Langer, E. J., & Abelson, R. P. (1974). A patient by any other name. . . ; clinician group differences in labeling bias. *Journal of Consulting Psychology*, **42**, 4–9.
- Lewicka, M. (1988). On objective and subjective anchoring of cognitive acts: How behavioral valence modifies reasoning schemata. In W. J. Baker, L. P. Mos, H. V. Rappard, & H. J. Stam (Eds.), *Recent trends in theoretical psychology* (pp. 285–301). New York: Springer-Verlag.
- Leyens, J.-Ph., & Yzerbyt, V. Y. (1991). *The ingroup overexclusion effect: Impact of valence and confirmation on stereotypical information search*. Manuscript submitted for publication.
- Marques, J. M., & Yzerbyt, V. Y. (1988). The black sheep effect: Judgmental extremity towards ingroup members in inter- and intra-group situations. *European Journal of Social Psychology*, **18**, 287–292.
- Marques, J. M., Yzerbyt, V. Y., & Leyens, J.-Ph. (1988). Extremity of judgments towards ingroup members as a function of group identification. *European Journal of Social Psychology*, **18**, 1–16.
- 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.
- Peeters, G., & Czapinski, J. (1990). Positive–negative asymmetry in evaluations: The distinction between affective and informational negativity effects. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 1, pp. 33–60). New York: Wiley.
- Reeder, G. D. (1985). Implicit relations between dispositions and behaviors: Effects on dispositional attribution. In J. H. Harvey & G. Weary (Eds.), *Attribution: Basics and applications* (pp. 87–116). New York: Academic Press.
- Reeder, G. D., & Brewer, M. B. (1979). A schematic model of dispositional attribution in interpersonal perception. *Psychological Review*, **86**, 61–79.
- Reeder, G. D., & Coovert, M. D. (1986). Revising an impression of morality. *Social Cognition*, **4**, 1–17.
- Reeder, G. D., Messick, D. M., & Van Avermaet, E. (1977). Dimensional asymmetry in attributional inference. *Journal of Experimental Social Psychology*, **13**, 46–57.
- Reeder, G. D., & Spores, J. M. (1983). The attribution of morality. *Journal of Personality and Social Psychology*, **44**, 736–745.

- Rosenberg, S., Nelson, C., & Vivekanathan, P. S. (1968). A multidimensional approach to the structure of personality impressions. *Journal of Personality and Social Psychology*, **9**, 283–294.
- Rosenberg, S., & Sedlak, A. (1972). Structural representations of implicit personality theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 6, pp. 235–297). New York: Academic Press.
- Rothbart, M., & Park, B. (1986). On the confirmability and disconfirmability of trait concepts. *Journal of Personality and Social Psychology*, **50**, 131–142.
- Rubin, M., & Shontz, F. C. (1960). Diagnostic prototypes and diagnostic processes of clinical psychologists. *Journal of Consulting Psychology*, **24**, 234–239.
- Schneider, D. J., Hastorf, A. H., & Ellsworth, P. (1979). *Person perception*. Reading, MA: Addison–Wesley.
- Sherman, S. J., Judd, C. M., & Park, B. (1989). Social cognition. *Annual Review of Psychology*, **40**, 281–326.
- Sines, L. K. (1959). The relative contribution of four kinds of data to accuracy in personality assessment. *Journal of Consulting Psychology*, **23**, 483–492.
- Skowronski, J. J., & Carlston, D. E. (1987). Social judgment and social memory: The role of cue diagnosticity in negative, positive and extremity biases. *Journal of Personality and Social Psychology*, **52**, 689–699.
- Skowronski, J. J., & Carlston, D. E. (1989). Negativity and extremity biases in impression formation: A review of explanations. *Psychological Bulletin*, **105**, 131–142.
- Snyder, M., & Gangestad, S. (1981). Hypothesis-testing processes. In J. H. Harvey, W. J. Ickes, & R. F. Kidd (Eds.), *New directions in attribution research* (Vol. 3, pp. 171–196). Hillsdale, NJ: Erlbaum.
- Snyder, M., & Swann, W. B. (1978). Hypothesis testing processes in social interaction. *Journal of Personality and Social Psychology*, **36**, 1202–1212.
- Swann, W. B., & Giuliano, T. (1987). Confirmatory search strategies in social interaction: How, when, why and with what consequences. *Journal of Social and Clinical Psychology*, **5**, 511–524.
- Tajfel, H. (1969). Social and cultural factors in perception. In G. Lindzey & E. Aronson (Eds.), *Handbook of Social Psychology* (Vol. 3). Reading, MA: Addison–Wesley.
- Trope, Y., & Bassok, M. (1982). Confirmatory and diagnosing strategies in social information gathering. *Journal of Personality and Social Psychology*, **43**, 22–34.
- Trope, Y., & Bassok, M. (1983). Information gathering strategies in hypothesis testing. *Journal of Experimental Social Psychology*, **19**, 560–576.
- Trope, Y., Bassok, M., & Allon, E. (1984). The questions lay interviewers ask. *Journal of Personality*, **52**, 90–106.