



HAL
open science

Converting Verbs into Adjectives: Asymmetrical Memory Distortions for Stereotypic and Counterstereotypic Information

Anne Maass, Mara Cadinu, Marta Boni, Cristiana Borini

► **To cite this version:**

Anne Maass, Mara Cadinu, Marta Boni, Cristiana Borini. Converting Verbs into Adjectives: Asymmetrical Memory Distortions for Stereotypic and Counterstereotypic Information. *Group Processes and Intergroup Relations*, 2005, 8 (3), pp.271-290. 10.1177/1368430205053943 . hal-00571608

HAL Id: hal-00571608

<https://hal.science/hal-00571608>

Submitted on 1 Mar 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Converting Verbs into Adjectives: Asymmetrical Memory Distortions for Stereotypic and Counterstereotypic Information

Anne Maass, Mara Cadinu, Marta Boni and Cristiana Borini
University of Padua

This paper investigated the hypotheses that (a) inferences from behaviors to traits would occur more frequently than vice versa, (b) this induction–deduction asymmetry would be facilitated by stereotype congruence but inhibited by incongruence, and (c) the tendency to draw trait inferences from stereotype-congruent but not from stereotype-incongruent behaviors would become more pronounced with increasing levels of Need for Cognitive Closure. Participants read information about a female or male job applicant that was in part relevant to gender, in part gender-neutral. The gender-relevant information was either stereotype-congruent or incongruent. Half of the information was presented as trait-adjectives, half as behavior-descriptive verbs. A recognition task was constructed so that some of the items (traits and behaviors) had actually been seen, some were entirely new, and some were new but had been implied by the information given. All three hypotheses were supported. Implications for intra-individual and interpersonal stereotype maintenance are discussed.

KEYWORDS deduction, false memories, induction, inferences, stereotypes

THE ROLE of language in transmitting stereotypes and in preserving them in front of disconfirming evidence has been the object of investigation in many recent studies (for overviews see Bourhis & Maass, in press; Maass, 1999; Wigboldus, Spears, & Semin, 1999). One aspect that is of particular relevance to the present series of experiments is the level of abstraction at which stereotypic or counterstereotypic events are communicated. Events

that correspond to stereotypic expectations tend to be communicated in a relatively abstract language (such as trait adjectives or state verbs,

Author's note

Address correspondence to Anne Maass,
D.P.S.S., Università di Padova, Via Venezia 8,
35131 Padova, Italy
[email: anne.maass@unipd.it]

see Semin & Fiedler's, 1988, Linguistic Category Model), whereas unexpected events tend to be communicated in more concrete language, a phenomenon generally referred to as the Linguistic Expectancy Bias (see Wigboldus, Semin, & Spears, 2000). Since abstract language implies greater temporal stability, greater cross-situational generality, and is overall more telling about the protagonist, the Linguistic Expectancy Bias may contribute in a subtle but significant way to the perpetuation of stereotypes. Indeed, research by Douglas and Sutton (2003) convincingly shows that people use abstract language when they want to convey that an act is typical of the protagonist.

In the present line of research, we want to extend this general idea to memory distortions. Whereas the research reported above deals mainly with language use, we are asking here whether a similar phenomenon may emerge when people memorize events that are consistent or inconsistent with their stereotypic expectancies. In other words, will people inadvertently distort their memories by transforming concrete information into abstract representations? Will they do so mainly when they encounter events that are in line with their stereotypic expectancies?

The Induction–Deduction Asymmetry (IDA)

There is strong evidence that people spontaneously infer traits from concrete behavioral information, or, linguistically speaking, that they transform verbs into adjectives. The most convincing evidence for this idea comes from an extensive research project by Uleman and collaborators, showing that such inferences require no intention or explicit instruction and often occur without awareness (Uleman, 1987; Uleman, Hon, Roman, & Moskowitz, 1996; Uleman & Moskowitz, 1994; Uleman, Newman, & Winter, 1986; Uleman, Winborne, Winter, & Shechter, 1986; Winter & Uleman, 1984; Winter, Uleman, & Cunniff, 1985; see also Carlston & Skowronski, 1994; Carlston, Skowronski, & Sparks, 1995; Van Overwalle, Drenth, & Marsman, 1999). Much less clear,

and rarely investigated, is the opposite inference from traits to behaviors (see Wyer & Lambert, 1994). A closer look at the literature reveals a striking imbalance in social psychological research, with great emphasis on inductive inferences (from behaviors to traits) and an almost complete absence of research on deductive inferences (from traits to behaviors, see Beike & Sherman, 1994, for an overview).

In an attempt to investigate the complementary inference processes within a single paradigm, Maass, Colombo, Colombo, and Sherman (2001) have conducted two studies testing the hypothesis that inductive inferences occur more frequently and more spontaneously than deductive inferences. Participants received descriptions of a target person, half of which presented in trait form (adjectives), half in behavior form (verbs). A subsequent recognition task was constructed in such a way that some of the items (traits and behaviors) had actually been seen, some were entirely new, and some were new but had been implied by the information given (traits implied by behaviors or behaviors implied by traits). Results show that participants frequently misidentified traits as already seen if they had been implied by a behavior, suggesting that they had made spontaneous inductive inferences, whereas the opposite type of error was very rare. Inductive inferences from behaviors to traits were about four times as frequent as inferences from traits to behaviors. Assuming that these systematic memory distortions reflect implicit inferences, this phenomenon was labeled Induction–Deduction Asymmetry (IDA). An additional response time measure (Maass et al., 2001, Experiment 2) suggested that traits were inferred from behaviors during encoding in an online fashion (and, hence, inferred traits were stored in memory along with actually learned traits) whereas inferred behaviors had to be ‘reconstructed’ in a memory-based fashion during retrieval.

Subsequent research confirmed that IDA emerges both on recognition and free recall tasks and that asymmetrical memory confusions occur even when adjectives and verbs share the same word stem (e.g. *aggress-aggressive*,

dominate-dominant, see Maass, Cadinu, Taroni, & Masserini, 2004). Importantly, the asymmetry is linked to person perception but does not occur when the same verb and adjective information is presented as a simple wordlist, that is, without reference to a specific target person (Maass et al., 2004, Experiment 4).

Taken together, these results confirm what has previously been shown by Uleman and others, namely that inductive inferences occur spontaneously and largely without awareness (cf. Uleman, 1987; Carlston & Skowronski, 1994; Van Overwalle et al., 1999). At the same time they extend previous work in important ways: They show that there is an asymmetry in inferential processing such that the inverse inference from trait to behavior is not only comparably rare but also occurs at a later stage in information processing. Presumably, deductive processes fail to occur unless they are triggered by an explicit request or by the need to predict behaviors in specific situations; in contrast, inductive inferences seem to be wired into social information processing and to occur by default during encoding. This is in line with the idea advanced by various authors that there is a general tendency toward abstraction in social information processing and interpersonal communication such that people tend to move from the concrete to the abstract, unless problems arise that force them to return to a more specific level of analysis (Fiedler, Semin, & Bolten, 1989; Hastie & Kumar, 1979; Semin & Fiedler, 1988; Semin & Smith, 1999; for a similar argument in cognitive psychology, see also Posner & Keele, 1968, 1970; Strange, Keeney, Kessel, & Jenkins, 1970).

IDA and stereotypes

The main question addressed in the present research concerns the potentially moderating effect of stereotypic expectancies. Maass et al. (2001) have argued that the relative strength of inductive vs. deductive inferences may vary across situations. The hypothesis investigated here is that the mental transformation of verbs into adjectives may be facilitated whenever the behavior corresponds to stereotypic

expectancies, but inhibited when the behavior contradicts such expectancies. To the contrary, deductive inferences should be largely unaffected by stereotypic expectancies.

Why should this be the case? From a semantic network perspective, there is reason to believe that stereotypes associate social groups more with traits than with behaviors. We assume that, in people's minds, social groups are more strongly associated with traits considered typical of those particular groups than with behaviors. In support of this idea, previous research has shown that stereotypes tend to take the form of abstract trait expectations such as '*Women are sensitive*' or '*Men are dominant*'. For example, Maass, Montalcini, and Biciotti (1998) asked research participants to generate a list of characteristics that they thought were typical of different social groups (Jews, Italians, Germans). Despite the fact that participants had explicitly been told that they could either use phrases (hence verbs) or adjectives, adjectives clearly prevailed over any other linguistic form. Interestingly, stereotype researchers tend to share this view. Many researchers define stereotypes as group-trait associations. For example, Brehm and Kassin (1996) define stereotypes as '*beliefs that associate groups of people with certain traits*' (p. 122). The same predominance of trait adjectives is also reflected in the measurement tools typically employed by stereotype researchers; again, traits tend to outnumber behavioral items regardless of whether the focus is on explicit (Katz & Braly, 1933) or implicit measures (Devine, 1989; Perdue, Dovidio, Gurtman, & Tyler, 1990; Hamilton & Rose, 1980; Kawakami, Dovidio, Moll, Hermsen, & Russin, 2000). Thus, researchers and lay persons alike tend to think of stereotypes as typical traits rather than as typical behaviors.

If social categories (e.g. *men*) are more strongly associated with typical traits (e.g. *dominant*) than with typical behaviors (e.g. *imposing one's opinions*), then any reference to the social category should mainly activate related traits. Thinking about a social category will render relevant trait information particularly accessible (Macrae, Stangor, & Milne, 1994) and, hence,

facilitate inferences from behaviors to such highly accessible traits. Therefore, people who encounter information about a group member (e.g. *a man*), together with a behavior (*imposes his opinions*), are likely to spontaneously infer the corresponding trait (*dominant*) for the simple reason that the trait has become highly accessible as a function of the category activation. When traits are *incongruent* with the social group (e.g. *men—submissive*), however, inhibitory processes are likely to come into play that reduce the accessibility of stereotypic knowledge (see Macrae, Bodenhausen, Milne, & Castelli, 1999); hence, inferences from behaviors to such incongruent traits may be inhibited. Taken together, we suggest that stereotypic expectancies will affect inductive processes by facilitating stereotype-congruent trait inferences and, at the same time, inhibiting incongruent trait inferences. In contrast, stereotypic expectancies should exert little or no influence on deductive inferences from traits to behaviors because social group labels have weaker links to behavioral information. Thus, due to differences in associative strength, trait inferences should be facilitated whenever stereotype-congruent behaviors are encountered (but inhibited when stereotype-incongruent information is provided) while deductive inferences should be largely unaffected by stereotypic expectancies.

To our knowledge, the only published research investigating the effect of stereotypes on inductive and deductive inferences is Wigboldus, Dijksterhuis, and van Knippenberg's (2003) recent work. In five studies, these authors have investigated the hypothesis that spontaneous trait inferences are facilitated when behavioral information is stereotype-consistent but obstructed when it is inconsistent. In their paradigm, participants read sentences about a target person that were either presented in the form of trait-adjectives or in the form of behavior-descriptive verbs and that were either stereotype consistent or inconsistent. Each description was followed by a relevant trait or behavior probe. For example, the participant may have read the phrase '*The skinhead hit the saleswoman*' and was subsequently

asked to indicate whether the word '*aggressive*' had been part of the sentence. The argument of the authors is that, if participants had drawn spontaneous trait inferences while reading the first phrase, it should take them longer to reject the implied trait (*aggressive*) that was not contained in the phrase but implied by the information given. If people make trait inferences from stereotype-consistent but not from inconsistent behavioral information, then reaction times should be much longer in the former than in the latter case. This is exactly what was found: it took participants much longer to reject a trait implied by the previous behavior information when the information was congruent rather than incongruent with stereotypic expectancies. Stereotype congruence did not affect the opposite type of inference (from traits to behaviors—Study 1), suggesting that stereotypes affect inductive but not deductive inferences. Another important finding is that, compared to a stereotype-neutral condition, inferences are inhibited by stereotype incongruence, but they are not reliably facilitated by stereotype congruence (Study 3). Together, these findings suggest that people spontaneously draw inductive inferences, from behaviors to traits, but that this process can be inhibited when the behavioral information is incongruent with stereotypic expectations.

Goal of the present experiment

The predictions of our own studies reported here are much in line with Wigboldus et al.'s (2003) findings, but contrary to their research we are focusing on memory distortions. The recognition probe paradigm used by Wigboldus et al. provides an excellent tool for investigating spontaneous inferences, but it is not concerned with memory errors. Not only are errors very rare in the recognition probe paradigm, but, when they occur, they are excluded from the reaction time analyses. In contrast, memory distortions are the main focus of our studies as they represent particularly relevant long-term effects of asymmetrical inferences. Also, contrary to Wigboldus et al.'s research, we investigated inductive and deductive inferences in a

single paradigm which allows us to show directly the asymmetrical status of inductive and deductive inferences. Third, unlike Wigboldus et al., in the present study we wished to show how individual differences, that is, the level of Need for Cognitive Closure, moderate the pattern of asymmetrical inferences, predicting that individuals with a high level of Need for Closure should be particularly sensitive to the stereotype congruency of the presented information.

In the present research we investigated inferences in a simulated case of personnel selection. Personnel selection is one of many applied situations in which asymmetrical inference processes may play a role. A comparative evaluation of applicants generally requires an integration of a wide variety of sources including objective test results, the candidate's self-presentation during interviewing, and the description of the candidate by others in reference letters and the like. The information provided may in part consist of very concrete, behavioral information, in part of rather abstract statements about the personality, potentials, enduring behavioral tendencies and habits of the applicant. An interesting applied question is therefore how evaluators may integrate various pieces of information. IDA suggests that they are likely to spontaneously and unintentionally translate concrete behavioral information into trait representations while it is unlikely that they will translate abstract trait information into concrete behavioral instances. Also, if the above reasoning is correct, trait inferences should mainly occur for information that corresponds to the stereotypical expectations of the person(s) involved in the selection process.

This hypothesis was tested in a simulated case of personnel selection in which a recommendation letter provided information either in the form of behavioral or trait information. Research participants read the information contained in the recommendation letter allegedly written by a former employer and were subsequently tested for their memory of the content of the letter. The first question was whether participants would erroneously remember behavioral (verb) information as traits (adjectives) but not vice versa.

More importantly, we were interested to see whether stereotype congruence would moderate this effect in ways similar to those found by Wigboldus et al. (2003). Therefore, the gender of the applicant was varied systematically together with the stereotypicality of the information so that the characteristics described in the letter either corresponded or did not correspond to those typically associated with women or men. In other words, both male and female applicants were either described in typically masculine or in typically feminine terms. We predicted a facilitation of inductive inferences when the applicant was described in stereotype-congruent terms, but an inhibition of such inferences when the description was incongruent with the applicant's gender. We expected deductive inferences to be unaffected by stereotypic expectancies.

Cognitive closure

As mentioned above, an additional goal of the present series of studies was to test whether Need for Cognitive Closure (Webster & Kruglanski, 1994) would moderate the tendency to make spontaneous behavior-to-trait inferences, especially when the behavior coincided with stereotypic expectancies. There is evidence that people with high need for closure prefer abstract descriptions to specific ones. For example, Boudreau, Baron, and Oliver (1992) found that an enhanced need for structure was associated with increasing use of global trait labels when describing others. Along the same line, Mikulincer, Yinon, and Kabili (1991) found an increase of global and stable self-attributions after failure experiences as a function of need for structure. These findings suggest that people with a high need for closure may quickly translate behavioral data into trait representations, possibly because traits reduce the complexity of behavioral information and create order and coherence. However, Maass et al. (2001) found only partial support for this prediction in one of their studies and no support in the other. Thus, there is very little evidence that people with greater Need for Cognitive Closure would make more inductive inferences. This failure may be attributable to

the fact that the behavioral information presented in Maass et al.'s study may not have been congruent with participants' expectations. Indeed, a study by Webster, Kruglanski, and Pattison (1997) has shown an increase in linguistic abstraction as a function of increasing need for closure *only* in those situations in which the observed behavior corresponded to expectations. One may therefore hypothesize that individuals with a high need for closure will be more likely to draw spontaneous behavior-to-trait inferences, but only when behaviors are congruent with their expectancies. On the contrary, they may be unlikely to draw inferences when the behavioral information contradicts their expectancies. This possibility was investigated in the present study, hypothesizing that, with increasing Need for Cognitive Closure, people may become more prone to draw trait inferences from stereotype-congruent behaviors, but also more reluctant to draw inferences from incongruent behaviors.

Overview of experimental design

Taken together, the main goals of the present study were to (a) replicate IDA in a hypothetical applied setting, namely personnel selection, (b) to test the moderating function of stereotypic expectancies, and (c) to investigate the joint influence of Need for Cognitive Closure and stereotype congruency. In order to investigate these issues, fictitious job applications were shown to research participants that included a photograph of the applicant together with a recommendation letter from a previous employer. The photograph depicted either a male or a female applicant and the description of the applicant consisted in part of gender-neutral information, in part of information that was either typical of males or typical of females. Thus, half of the participants received a recommendation letter that described the applicant in gender-role congruent terms (masculine male or feminine female), whereas the other half received a gender-role incongruent description (feminine male or masculine female). Across conditions, half of the descriptions in the letter were provided in behavioral terms, half in trait terms. Subsequently, the

participants' task was to recognize the information provided in the letter from a list of items consisting of the actually seen behaviors or traits as well as implied behaviors or traits along with entirely new items.

Hypotheses

First, we predicted that participants would make more *inferences* from behaviors to traits than vice versa, resulting in a greater number of false alarms for traits that had been implied by a behavior than vice versa.

Second, if behaviors are transformed into traits in the mind of our participants, then they should also be less able to correctly recognize actually seen behaviors than actually seen traits, resulting in lower *hit* rates for behaviors than for traits.

Third, concerning *stereotype congruence*, it was predicted that behavior-to-trait inferences would be (a) more frequent for stereotype-congruent than for gender-neutral information but (b) less frequent for stereotype-incongruent than for gender-neutral information. Thus, typically feminine descriptions of female applicants and typically masculine descriptions of male applicants should induce greater inductive inferences than stereotype-incongruent descriptions with gender-neutral descriptions occupying an intermediate position. Deductive inferences were *not* expected to be affected by stereotype congruence.

Fourth, it was predicted that, with increasing *Need for Cognitive Closure*, participants would be more likely to infer traits from stereotype-congruent behaviors but less likely to infer traits from stereotype-incongruent behaviors. No relation between Need for Cognitive Closure and trait inferences were expected for gender-neutral behaviors, nor did we expect a relation between Need for Cognitive Closure and deductive (trait-to-behavior) inferences.

Method

Design

The experiment consisted of a 2 (participant gender) \times 2 (applicant gender) \times 2 (gender-relevance: gender-neutral vs. gender-relevant

information) \times 2 (stereotype congruence of gender-relevant information: congruent vs. incongruent) \times 2 (trait vs. behavior presentation during recognition) design in which the last two factors were within-participant variables.

Participants A total of 280 participants (140 males and 140 females), living in Northern and Central Italy volunteered for the study. The mean age was 29.94 years. The majority were university students while the remaining participants were employed. The educational level was relatively high with 83 holding a university degree, 195 a high school diploma, and only 43 holding diplomas below high school level.

Procedure Participants were told that the research investigated 'how people use information provided in a letter of recommendation in order to evaluate a hypothetical applicant for a high level job'. The first part of the questionnaire consisted of the letter of recommendation supposedly written by a previous employer in support of an applicant applying for a high level job (level 8 in the Italian employment system). Participants were encouraged to read the letter as often as they wished in order to form an impression of the applicant. The reference letter was accompanied by a close-up photograph that allowed us to manipulate gender in addition to other variables (physical attractiveness and clothing that had absolutely no effect and that, for reasons of space, will not be reported here). After reading the letter, participants were asked to respond to the Italian version of the Need for Cognitive Closure Scale (Pierro et al., 1995). Subsequently, participants engaged in a recognition task in which they were asked to tick the information that they had read, in the exact same form, in the letter. Finally, after having viewed the picture of the candidate for a second time, participants were asked to express a judgment about the suitability of the applicant for the job. At the end of the experiment, participants were informed about the purpose of the experiment.

Applicant gender Half of the applicants were young males, half young females. On the basis

of a pretest, photos had been matched for physical attractiveness (2.3 for females and 2.3 for males, on a 5-point scale).

Stimulus material

Each candidate was described in the recommendation letter by 12 pieces of information. Six characteristics described in the letter were gender-neutral, the remaining six were (depending on the condition) either typically feminine or typically masculine. Obviously, in the case of a male candidate, the candidate appeared stereotype-congruent when described as having a typically masculine personality, but incongruent when described in feminine terms. The opposite was true for female candidates. In each case (neutral, feminine, or masculine characteristics), four traits were positive and two negative. Importantly, half of the information contained in each letter was provided in trait form, half in behavior form, so that each participant received six behaviors and six traits. For example, one participant may have received dominance information in trait form (e.g. *is dominant*) and athletic ability information in behavior form (e.g. *practices sports*), whereas another may have received dominance information in behavior form (e.g. *expects to be obeyed*) and athletic ability information in trait form (e.g. *is athletic*).

The experimental material was developed in a step-wise pretesting procedure. Initially, masculine and feminine characteristics were selected from the Italian version of the Bem Gender Role Inventory (De Leo & Villa, 1986) and gender-neutral characteristics were in part taken from the social desirability scale of the same questionnaire, in part generated by the authors and pretested for gender-independence. Since the present study required that characteristics be uncorrelated, all characteristics (masculine, feminine, and neutral) were pretested by asking participants to judge the degree to which one trait would be diagnostic of other traits (e.g. if a person A is *independent*, what is the likelihood that s/he is also *educated, sociable, charismatic*, etc.). Different combinations of 12×12 matrices were judged by a total of 120 participants.¹ Based on this pretest, positively as well

as negatively correlated traits were excluded, leaving only those traits that were considered uninformative of the remaining traits.² Based on this selection criterion, the following traits were included in the final material: *aggressive, dominant, analytical, athletic, courageous, authoritative* as typically masculine traits; *timid, ingenuous, loyal, intuitive, diplomatic* and *cheerful* as typically feminine traits; and *messy, pessimistic, concrete, dynamic, up-to-date*, and *precise* as gender-neutral traits.

Importantly, our design was such that half of the neutral and half of the gender-typed information was presented in trait form, half in behavior form. This required that for each selected trait there would be a matching behavior that was diagnostic of that particular trait *without* being diagnostic of any of the remaining traits.

In order to satisfy these criteria, three behavioral descriptions were generated for each trait that was considered diagnostic for the respective trait by two independent raters. The material was subsequently subjected to a pretest in which 64 participants had to judge the diagnosticity of each behavior for all traits on a 5-point scale (from 1 = *not at all characteristic* to 5 = *very characteristic of a person having trait 'x'*). Behaviors were considered acceptable only if they had been judged highly diagnostic of the target trait (lowest diagnosticity = 4.1, mean diagnosticity = 4.8 on a 5-point scale) and at the same time *not* diagnostic of the remaining traits (mean difference between target trait and closest subsequent trait = 1.8). Following this procedure, one behavior was selected for each trait (for example, *dominant—expects to be obeyed; pessimistic—tends to make negative predictions*; for a complete list of stimulus material, see appendix).

This procedure, however, was asymmetrical since it only tested the diagnosticity of each behavior for the respective traits but not vice versa. In order to assure that traits were as diagnostic of behaviors as behaviors were of traits, an additional pretest was conducted ($n = 38$). Half of the participants rated the diagnosticity of each behavior for all traits (Person A did *behavior*. How likely is it that s/he will have the following traits) whereas the

remaining half rated the diagnosticity of traits for all behavior (Person A is *trait*. How likely is it that s/he will show the following behaviors). The rating was provided on a scale from -2 (indicating that the person will *not* have trait x /show behavior x) to $+2$ (indicating that the person *will* have trait x /show behavior x), with 0 indicating that the trait/behavior provides no information about the target behaviors/trait. Ideally, values should be high and positive (close to $+2$) for the target and 0 for the non-target stimuli. A 2 (trait vs. behavior) \times 2 (diagnosticity for target trait or behavior vs. diagnosticity for non-targets) analysis of variance (ANOVA) revealed a very strong effect for target vs. non target ($F(1, 34) = 953.76, p < .001$). Both traits and behaviors were highly diagnostic of their respective target (behaviors: $M = 1.80$; traits: $M = 1.77$), but they were close to zero and, hence, non-diagnostic for the non-target traits or behaviors (behaviors: $M = .17$; traits: $M = .12$). More importantly, traits were just as diagnostic of behaviors as behaviors were diagnostic of traits and no interaction emerged between the two variables (both F 's < 1). Thus, it is quite evident that the material was balanced with respect to behavior-to-trait and trait-to-behavior diagnosticity.

A final pretest investigated the possibility that traits and behaviors may differ with respect to the ease to which they can be memorized— independently of inferential processes (for example, due to differential vividness). It was therefore important to assure that—in the absence of inferential processes—traits and behaviors included in the study would have a similar likelihood of being remembered. We therefore had research participants ($n = 30$) read either a list of the 18 traits (*analytical, athletic, etc.*) or a list of the 18 behaviors (*to rely on logic when evaluating situations, to practice sports, etc.*), this time not referring to any specific person and under explicit memory instructions. In both cases one additional item (trait or behavior) was added at the beginning and one at the end of the list to reduce chances of primacy and recency effects. After one hour, participants were asked to recognize the previously seen stimuli from a new list containing all

the old traits (or behaviors) and six distracter traits (or behaviors). Correct recognition was quite similar and not significantly different for the trait list ($M = 9.71$) and the behavior list ($M = 10.29$), suggesting that neither stimulus type had an a priori memory advantage.

The above material allowed us to manipulate three independent variables of our factorial design, namely gender-relevance vs. irrelevance, masculinity vs. femininity of gender-relevant information, and trait vs. behavior representation.

Gender-relevance The first variable was gender-relevance, manipulated within participants. Each participant read a recommendation letter in which the applicant was described by six gender-neutral statements (two negative and four positive) and by six gender-relevant statements (either all stereotypically masculine or all stereotypically feminine, again including two negative behaviors and four positive behaviors).

Stereotype congruency of gender-relevant information: congruent vs. incongruent The second factor, manipulated between participants, was masculinity vs. femininity of the gender-relevant information. Half of the participants received a letter in which all gender-relevant information was typically masculine, whereas for the other half it was all feminine. Thus, the gender-relevant information provided in the letter was either stereotype-congruent or incongruent with the gender of applicant.

Trait vs. behavior representation The last factor, manipulated within participants, was the trait vs. behavior representation of the information. Half of the gender-neutral and half of the gender-relevant information were presented in trait form, half in behavior form. For example, in the case of gender-neutral information, each participant received one of the two negative and two of the four positive pieces of information in trait form, the other in behavior form. Trait vs. behavior presentation was counterbalanced across two versions to which participants had been randomly assigned. We

had prepared two versions of each recommendation letter so that the same information (e.g. *dominance*) appeared in trait form (e.g. *dominant*) in one version and in behavior form (*expects to be obeyed*) in the other version.

Also, a standard introduction (*I am writing on behalf of . . .*) and ending (*I would be happy to provide any additional information . . .*) was added to the recommendation letter in order to enhance its credibility. For the same reason, the different pieces of information were integrated into a coherent text by the use of transition adverbs and the like (*at the same time . . ., despite this fact . . .*).

Need for Cognitive Closure Scale After having read the letter at their own pace, participants were asked to complete the Need for Cognitive Closure Scale. This scale, originally developed by Webster and Kruglanski (1994) assesses the tendency to reach cognitive closure. The Italian version by Pierro et al. (1995) used in this research consists of 42 items assessing the desire for predictability, preference for order and structure, discomfort with ambiguity, decisiveness, and close-mindedness (item examples: *'Usually I take important decisions quickly and with certainty'*, *'I don't like to be in situations without knowing what to expect'*). Participants indicated their agreement on a 4-point scale (from *'completely false'* to *'completely true'*). Cronbach's alpha of the entire scale was .98, with alphas for subscales ranging from .85 to .97.

Dependent variables

Recognition The main dependent variable consisted of a recognition measure in which participants were presented with a list of trait adjectives and behaviors and asked to tick those terms that they had read *in the exact same formulation* in the reference letter. The recognition measure included a list of 18 characteristics, including 12 old and 6 new characteristics, which were presented in random order. The old items were all those characteristics that had been presented in the reference letter, including 6 gender-neutral and 6 gender-relevant items (depending on the condition, either all masculine or all feminine). The new items (2 of

negative and 4 of positive valence) were all unrelated to gender stereotyping. Importantly, each characteristic was presented in trait *and* in behavior form. Thus, 18 item pairs were presented as illustrated in this example:

<i>dominant</i>	<i>expects to be obeyed</i>
<i>believes anybody out of</i>	<i>ingenuous</i>
<i>inexperience</i>	

The order of trait and behavior representation within each pair was counterbalanced, as was the order of the 18 item pairs. Participants were explicitly told that they could tick either the first or the second alternative of each pair, or neither or both. In reality, this last option was practically absent among the responses ($n = 4$).³

The recognition measure allowed us to analyze not only the number of correctly recognized traits and behaviors (hits) and false alarms (new items), but, more importantly, the number of inferences from traits to behaviors and vice versa. Thus, if a characteristic had originally been presented in the form of a behavior (e.g. *examines every situation in detail*), but recalled as a trait adjective (*precise*), it was scored as a behavior-to-trait inference. Responses were scored as trait-to-behavior inferences when the participant incorrectly ticked a behavior (e.g. *examines every situation in detail*) that had originally been presented as trait (*precise*).

Suitability The second dependent variable was the subjective suitability of the candidate for the job. Participants were asked to rate on a 4-point scale from 1 = *not at all suitable* to 4 = *very suitable* the degree to which they thought the candidate was suitable for the high level job.

Results

The recognition task provided three measures of interest: (a) the number of hits indicative of the participant's ability to correctly recognize the information s/he had read in the reference letter, (b) the number of inferences, that is false alarms for old items indicative of inferences from traits to behaviors or vice versa and (c) the number of false alarms for entirely new items. (Considering the large sample size, we decided

to correct for inflated α by setting the probability level to $p < .01$.)

Hits

A 2 (participant gender) \times 2 (applicant gender) \times 2 (gender-relevance: gender-neutral vs. gender-relevant information) \times 2 (stereotype congruence of gender-relevant information: stereotype-congruent vs. incongruent) \times 2 (trait vs. behavior presentation during recognition) ANOVA with repeated measures on the third and last variables was conducted for hits. Values could vary from a minimum of 0 to a maximum of 3 in each trait/behavior \times gender-relevance combination. In line with Hypothesis 2, traits ($M = 2.18$) were correctly recognized with greater frequency than behaviors ($M = 1.30$) ($F(1, 272) = 162.58, p < .001, \eta^2 = .37$). Also, gender-relevant information ($M = 1.81$) was recognized more frequently than gender-neutral information ($M = 1.68$) ($F(1, 272) = 8.68, p < .001, \eta^2 = .03$), but this was only true for stereotype-incongruent information. Indeed, a reliable interaction between stereotype congruence and gender-relevance ($F(1, 272) = 14.88, p < .001, \eta^2 = .05$) revealed that stereotype-incongruent information ($M = 2.00$) was better remembered than either stereotype-congruent ($M = 1.62$) or gender-neutral information ($M = 1.66$ if presented together with stereotype-congruent and $M = 1.70$ if presented together with stereotype-incongruent information) (all $t_s > 3.5, p < .01$), whereas no differences emerged between the remaining three conditions. Together, these data indicate a memory advantage for information (both behaviors and traits) that is incongruent with stereotypic expectancies. More importantly, in line with our second hypothesis, they suggest that, compared to traits, recognition of behaviors is poorer presumably because behaviors have been transformed into trait representation.

Inferences

Of greater theoretical interest are the false alarm rates for items that had not been seen but that had been implied by the information given. We had predicted two effects for this measure: In line with IDA, a main effect was

predicted for type of inference such that behavior-to-trait inferences would outnumber trait-to-behavior inferences (Hypothesis 1). We had also predicted a three-way interaction between stereotype congruence, gender-relevance, and type of inference, such that the greatest asymmetry would be observed for items that were relevant to and stereotypically congruent with the applicant's gender.

The 2 (participant gender) × 2 (applicant gender) × 2 (gender-relevance: gender-neutral vs. gender-relevant information) × 2 (stereotype congruence of relevant information: congruent vs. incongruent) × 2 (type of inference: behavior-to-trait vs. trait-to-behavior) ANOVA for inferences confirmed both predictions. In line with Hypothesis 1, a very strong main effect emerged for type of inference, indicating that participants were more likely to falsely recognize a trait adjective that had been implied by a behavior ($M = 1.58$) than vice versa ($M = .75$) ($F(1, 272) = 164.35, p < .001, \eta^2 = .38$).

Moreover, inferences were somewhat more frequent for gender-neutral ($M = 1.23$) than for gender-relevant information ($M = 1.23$) ($F(1, 272) = 8.12, p < .01, \eta^2 = .03$). Also, more inferences were drawn in the condition in which the gender-relevant information was congruent ($M = 1.28$) rather than incongruent ($M = 1.05$) with gender stereotypes ($F(1, 272) = 12.91, p < .001,$

$\eta^2 = .05$). But the above main effects were specified by an interaction between gender-relevance and stereotype congruence ($F(1, 272) = 19.11, p < .001, \eta^2 = .07$). Inferences (including both inductive and deductive inferences) were more likely to occur for stereotype-congruent ($M = 1.32$) than for stereotype-incongruent information ($M = .90$) ($t(272) = 4.94, p < .01$). In contrast, the tendency to draw inferences from gender-neutral information was the same regardless of whether the neutral information occurred together with stereotypical ($M = 1.25$) or together with counterstereotypical information ($M = 1.21$). This is not surprising considering that the gender-neutral information was exactly the same in the two conditions (stereotype-congruent or incongruent).

Most importantly and in line with Hypothesis 3, the above interaction was specified by a significant three-way interaction between gender-relevance, stereotype congruence, and type of inference ($F(1, 272) = 9.01, p < .01, \eta^2 = .03$) (see Table 1). Although behavior-to-trait inferences greatly outnumbered trait-to-behavior inferences in *all* experimental conditions, the two types of inferences were differentially affected by relevance and stereotype congruence. Deductive (trait-to-behavior) inferences did not vary as a function of either relevance or stereotype congruence. In contrast, inductive

Table 1. Mean number of behavior-to-trait and trait-to-behavior inferences as a function of gender-relevance and stereotypicality of information

	Behavior-to-trait inferences		Trait-to-behavior inferences	
	Stereotype-congruent	Stereotype-incongruent	Stereotype-congruent	Stereotype-incongruent
Gender-relevant	1.87 _a <i>n</i> = 140	1.27 _b <i>n</i> = 140	.76 _d <i>n</i> = 140	.53 _d <i>n</i> = 140
Gender-neutral	1.55 _c <i>n</i> = 140	1.65 _c <i>n</i> = 140	.95 _d <i>n</i> = 140	.77 _d <i>n</i> = 140

Notes: Means with different subscripts differ significantly from each other ($p < .01$). Behavior-to-trait inferences refer to false alarms of trait stimuli that had been implied by a behavior, trait-to-behavior inferences to false alarms of behavior stimuli that had been implied by a trait. Gender-relevance represents a within-participant factor, with all participants receiving both gender-relevant and irrelevant information. Stereotype congruence represents a between-participants factor so that the gender-relevant information was either congruent or incongruent with gender stereotypes. The gender-neutral information was kept constant across the congruent and incongruent conditions.

(behavior-to-trait) inferences occurred much more frequently when the behavior matched stereotypic expectancies ($M = 1.87$) than when it was counterstereotypical ($M = 1.27$), with inferences for gender-neutral information occupying an intermediate position. Indeed, inferences from stereotype-congruent behaviors to corresponding traits ($M = 1.84$) were significantly more likely than inferences from gender-neutral behaviors ($M = 1.55$, reported together with stereotype-congruent information), whereas inferences from stereotype-incongruent behaviors ($M = 1.27$) were less likely than inferences from gender-neutral behaviors ($M = 1.65$, reported together with stereotype-incongruent information).

A simplified way to represent these data is to collapse responses to gender-neutral descriptions (which were identical regardless of whether they occurred together with stereotype-congruent or incongruent information) and to compare them to inferences for stereotype-congruent and stereotype-incongruent descriptions (see Figure 1). As can be seen,

inductive (but not deductive) inferences decreased in a linear fashion from stereotype-congruent over gender-neutral to stereotype-incongruent information.

False alarms for old and new items

Theoretically, the greater frequency of behavior-to-trait inferences rather than vice versa could reflect an artifact due to a generalized tendency to give more false-positive responses to traits than to behaviors. If this was true, then false alarm rates should also be greater for new traits that had not been included in the recommendation letter. To control for this possibility, we therefore ran a 2 (participant gender) \times 2 (applicant gender) \times 2 (stereotype congruence) \times 2 (trait vs. behavior presentation during recognition) ANOVA with repeated measures on the last variable, using the false alarms for new items as dependent variable. The analysis revealed absolutely no difference between false alarms for traits ($M = .28$) and false alarms for behaviors ($M = .29$) ($F(1, 272) = .03$, *ns*, $\eta^2 = .00$), nor any interaction with

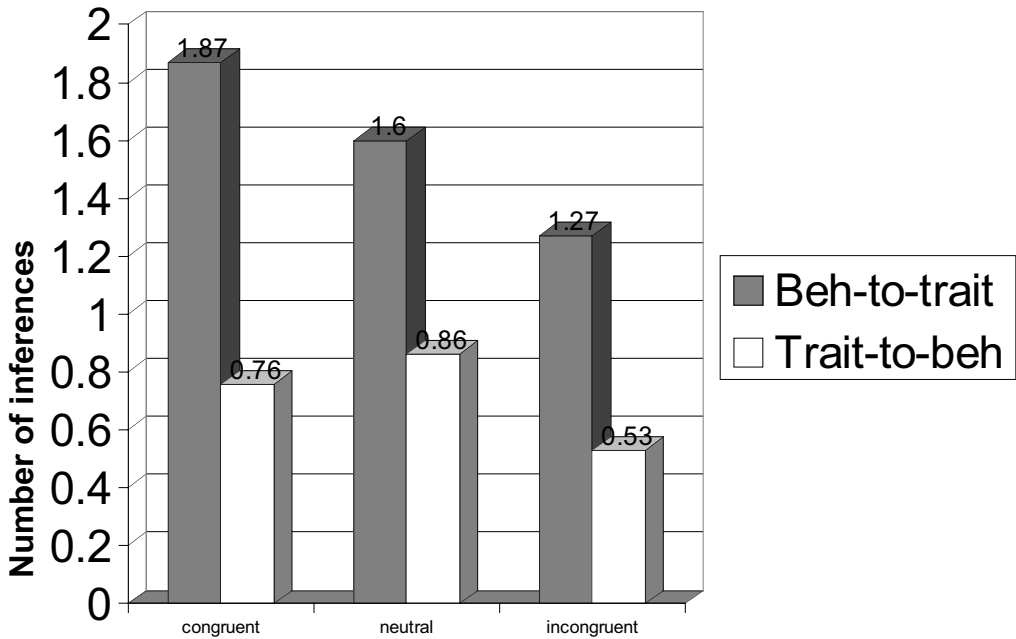


Figure 1. Inferences for stereotype-congruent, gender-neutral, and stereotype-incongruent information.

additional variables (all F s < 1). Hence, the tendency to erroneously identify implied traits as already seen seems to reflect an inductive inference from previously learned behavior information, and not just a tendency toward false alarms for just any trait adjective.

Need for cognitive closure

In the fourth hypothesis, we predicted that the tendency to draw trait inferences from stereotype-congruent behaviors would be particularly pronounced for those participants with a high Need for Cognitive Closure. At the same time, these participants were expected to be reluctant to draw inferences from stereotype-incongruent behaviors. On the other hand, Need for Cognitive Closure was not expected to be predictive of inductive inferences for gender-neutral information nor of deductive inferences.

In order to test this hypothesis, we ran a series of regression analyses using either the inductive (behavior-to-trait) or the deductive (trait-to-behavior) inferences as the dependent variable and using Need for Cognitive Closure as the predictor variable (with all variables z -transformed). This was done separately for gender-relevant and gender-neutral information and

for stereotype-congruent and incongruent information.

Looking first at the inductive (behavior-to-trait) inferences, Need for Cognitive Closure was not predictive of inferences for gender-neutral information ($\beta = -.04$; $F(1,127) = .38$, ns). By contrast, as can be seen in Figure 2, the greater the Need for Cognitive Closure, the more likely participants were to draw trait inferences from stereotype-congruent behaviors ($\beta = .34$; $F(1,138) = 18.35$, $p < .001$), but less likely to draw such inferences from stereotype-incongruent behaviors ($\beta = -.25$; $F(1,138) = 9.15$, $p < .01$). This finding confirms Hypothesis 4, by showing a positive relation between Need for Cognitive Closure and inductive inferences for stereotype-congruent, and a negative relation for stereotype-incongruent behaviors, with inferences from gender-neutral behaviors being unrelated to Need for Closure.

Turning to deductive inferences, Need for Cognitive Closure was unrelated to either stereotype-congruent or stereotype-incongruent inferences. For deductive inferences, Need for Closure emerged only as a negative modest predictor ($\beta = -.14$; $F(1,278) = 5.80$, $p < .05$) in the case of gender-neutral information: the greater the Need for Cognitive Closure, the fewer the

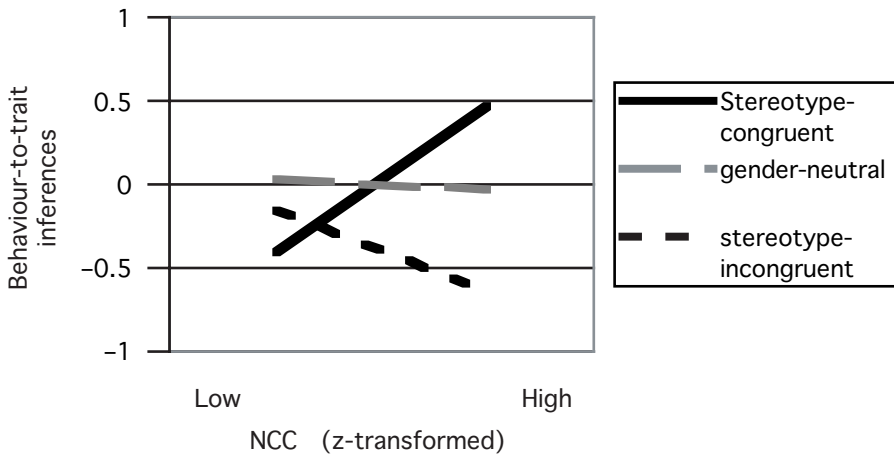


Figure 2. Trait inferences from stereotype-congruent, gender-neutral, and stereotype-incongruent behaviors as a function of Need for Cognitive Closure (NCC).

trait-to-behavior inferences participants tended to draw.

Suitability

An additional measure assessed the degree to which candidates were judged suitable for the job. A 2 (participant gender) \times 2 (applicant gender) \times 2 (stereotype congruence of relevant information: congruent vs. incongruent) ANOVA was performed on the suitability judgments. The only significant effect emerging from this analysis was a main effect for stereotype congruence ($F(1,727) = 8.16, p = .01, \eta^2 = .03$), indicating that candidates that conformed to stereotypic expectancies, that is masculine males and feminine females ($M = 2.96$) were judged as more suitable for the managerial job than those that contradicted such expectancies ($M = 2.66$). Interestingly, this was equally true for male and female candidates, as evidenced by the lack of interaction between applicant gender and stereotype congruence ($F(1,727) = 1.18, ns, \eta^2 = .01$).

At this point, one may ask whether behavior-to-trait inferences, that is, generalizations beyond the specific behavior, are predictive of suitability ratings. One may hypothesize that, the more trait inferences people draw from stereotype-congruent behaviors, the more suitable they may perceive the candidate to be. An inverse relation may be expected for trait-inferences from incongruent behaviors: the fewer inferences that are drawn from such behaviors, the more suitable the candidate should appear. In order to test this interactive function of inferential processes and stereotype congruence, four separate regression analyses were run in which inferences (z -transformed) and stereotype congruence of the description (dummy-coded) were entered in the first step, while the interaction was added in the second step. This was done separately for inductive and deductive inferences and for gender-relevant and gender-neutral information. The predicted interaction between inferences and stereotype congruence emerged, as an almost significant tendency, only for inductive inferences and when the information was relevant to gender ($\beta = .25; t = 1.93, p < .06$). The interaction is

represented in Figure 3. When behavior information contradicted stereotypic expectancies, the greater the number of behavior-to-trait inferences, the lower the suitability of the candidate ($\beta = -.20; F(1,138) = 4.39, p < .05$). For stereotype-congruent information, there was a nonsignificant tendency in the opposite direction such that more trait inferences were associated with higher suitability ratings ($\beta = .05, ns$).

Discussion

The aim of the present experiment was to test asymmetrical inference processes in the context of personnel selection, together with the moderating role of stereotypic expectancies and Need for Cognitive Closure. Results of the present experiment are highly supportive of predictions. First of all, the hypothesized asymmetry was clearly confirmed. Participants were much more likely to erroneously identify an unseen but implied trait (rather than an unseen but implied behavior) as already seen. Note that this occurred only for traits that had been implied by a behavior, but not for entirely new traits, suggesting that we are not simply dealing with a biased guessing strategy of the sort '*since we generally describe people by traits, it is more likely that I have seen a trait rather than a behavior*'. Also, behavior-to-trait inferences occurred despite the explicit instruction to only identify those items that had been presented in the letter in the *exact same form*. Apparently, our participants unintentionally transformed behavioral information into trait adjectives; subsequently, they retrieved this information from memory without being aware that the original information had been provided in a different, less general format. This finding is perfectly in line with previous research suggesting that inductive inferences occur unintentionally (Carlston & Skowronski, 1994; Carlston et al., 1995; Uleman, 1987; Uleman & Moscovitz, 1994; Uleman, Newman, & Winter, 1986; Uleman, Winborne, Winter, & Shechter, 1986; Uleman et al., 1996; Van Overwalle et al., 1999; Winter & Uleman, 1984; Winter et al., 1985). At the same time, our study extends previous



Figure 3. Suitability as a function of trait inferences from stereotype congruent vs. incongruent behaviors.

research by showing, within a unifying experimental paradigm, that this inference process is asymmetrical.

A curious finding of this research is the fact that participants encountered considerable difficulty in recognizing behaviors despite the fact that behaviors are generally considered more vivid and hence easier to retrieve. Two different explanations may be offered for this finding. First, the experimental material may have been biased such that behaviors, for some unknown reason, may have left weaker memory traces than the more abstract trait descriptions. Our pretest argues against this interpretation considering that behaviors and traits led to comparable recognition rates when presented separately and out of context. More plausible is a second explanation according to which the mental transformation of behaviors into traits (which is likely to occur spontaneously during encoding, see Maass et al., 2001, Experiment 2; Widboldus et al., 2003) may have cancelled the original information. It is noteworthy in this context, that our design contained a peculiar feature that should have rendered recognition of the originally seen items relatively easy. The

recognition task was structured so that the original behavior information was presented *together* with the implied trait (side-by-side, either immediately preceding, or immediately following it), and participants were explicitly allowed to tick both, behavior and trait. Despite this instruction, none of the 280 research participants ticked the actually seen behavioral information *together* with the inferred trait, suggesting that the behavioral information was substituted (rather than complemented) by the more abstract trait representation. Rather than storing the episodic behavior information along with the inferred trait, our participants seem to have transformed the information in a more abstract representation without having preserved the concrete behavioral information.

This finding is curious as it contradicts some of the observations on false memories. For example, Roedinger and McDermott (1995), using a very different paradigm, found that implied targets were more likely to be falsely recognized, the better the memory for the actually seen information that triggered the false memory. In Roedinger and McDermott's (1995) case, false memories were stored along

with accurate memory traces whereas in our research false memories seem to have substituted the original traces. This raises the challenging question of when false memories will complement and when they will substitute accurate memory traces. Only future research will be able to give a definite answer to this question.

The second, and probably most important finding of the present study, is that inductive, but not deductive inferences are influenced by stereotypic expectancies. Using gender-neutral information as a baseline, our results show that trait inferences were facilitated when the information matched the applicant's gender, but inhibited when the applicant was described in ways that contradicted socially shared gender-stereotypes. In contrast, deductive inferences did not vary as a function of stereotype congruence. These findings are nicely in line with our third hypothesis, in which we had predicted that both facilitative and inhibitory forces would affect inductive, but not deductive processes.

It is important to point out that the greater inductive potential of stereotype-congruent compared to incongruent behaviors is *not* attributable to any bias in the experimental material. Note that the same items (for example, 'is able to understand a fact before reasoning about it') occurred in the stereotype-congruent and in the incongruent condition, since stereotype congruence depended on the gender of the applicant (what was stereotypical for male candidates was counterstereotypical for females, and vice versa). Since gender of applicant did not interact with any of the remaining variables, one can conclude that the greater inductive potential of stereotypic behaviors is attributable to the fact that they confirm gender-role expectations and not to specific semantic features of the information selected for this experiment.

Taken together, our findings clearly support the idea that IDA will occur when evaluating job applications and that its strength depends on prior category-based expectancies. The mental transformation of verbs into adjectives tends to be facilitated when prior expectancies are confirmed but inhibited when they are discon-

firmed. In contrast, deductive inferences are unrelated to stereotypic expectancies. The fact that events are falsely recognized (or recalled), when they are semantically linked to the information given, is by no means new. False memories have repeatedly been demonstrated by both cognitive and social psychologists. What is new here is the *asymmetry* of these memory illusions, depending on whether the original information is provided in the form of a trait-adjective or a behavior-descriptive verb. Contrary to previous research, our design was such that the same items served either as actually seen stimuli or as associated, implied targets. Such a balanced design allows *comparative* conclusions about the relative likelihood of behaviors and traits to create false memories. What this research shows is a one-sided tendency to transform episodic behavior information into trait information that is not matched by a complementary deductive inference process. In a more general sense, this provides one possible answer to the question of why some stimuli produce false memories while others do not—a question that has been discussed in the literature since the fifties. Our hypothesis is that false memories may be facilitated (a) when moving from the concrete to the abstract, rather than vice versa, and (b) when information is expectancy-congruent rather than incongruent.

Another interesting finding of the present study is that such memory distortions may exert a direct influence on judgment. Job candidates in our study were judged more suitable for a managerial job, the more people tended to infer traits from stereotype-congruent behaviors and the less they inferred traits from incongruent behaviors. In other words, generalizing beyond the specific behavioral episode may have positive effects on the evaluation of the target person if s/he conforms to expectancies, but generalizations from atypical behaviors may lead to negative evaluations. Although additional testing is needed, this suggests that spontaneous memory distortions may have important consequences in relevant social contexts (such as job selection), although the individual may well be unaware of these biases.

An additional aim of our study was to investigate the moderating role of Need for Cognitive Closure. In line with our hypotheses, the present study supported the idea that people with a high Need for Cognitive Closure, and hence a strong need for clarity, predictability, and a strong preference for order and structure, make more trait inferences when the information converges with stereotypic expectancies, but fewer inferences when the information contradicts such expectancies. Taken together, these findings are in line with Webster *et al.*'s (1997) observation that individuals with a high need for closure tend to use relatively abstract language when describing expectancy-congruent events. The language measure used by Webster *et al.* and the memory measure used in the present study nicely converge in producing conceptually equivalent results. In both studies, high need for closure was associated with a shift toward abstraction (from verbs to adjectives or from behaviors to traits, respectively) whenever prior expectancies were confirmed. This suggests that the spontaneous transformation of behavioral into trait information is reliably affected by the desire for definite knowledge or the motivation to avoid ambiguity. It is likely that such an asymmetrical inference process contributes to the maintenance of a relatively stable worldview in people with a high Need for Cognitive Closure, considering that expectancy-congruent episodic information is immediately transformed to a more abstract level, whereas such a transformation is inhibited when expectations are violated. It may be interesting for future research to investigate whether temporary variations in Need for Closure (such as exposure to noise or time pressure) affect inferential processes in much the same way as stable, trait-like differences in Need for Closure seem to impact upon these processes.

Finally, an important question to be investigated in future research is how language bias in memory is related to language bias in communication. Based on Grice's principle of cooperation and, in particular, the maxim of quality, communication partners are not expected to say anything they believe to be false. In the case of the present findings, if a person

has observed, or learned about, a specific behavior of a target person, s/he will be likely to remember this behavior in trait (adjective) form, particularly if the behavior is in line with socially shared stereotypes. Our study also suggests that the observer will be unaware of this transformation and may have lost access to the original (verb) information. Therefore, according to the maxim of quality, it appears likely that our observer will communicate the event to a third person in the same (abstract) way in which the information is stored in memory. Thus, we believe that, in the absence of other motives, interpersonal communication will follow the same (asymmetrical) pattern that was observed in our study. In line with this idea is the great similarity of many results obtained with communication and memory measures. For example, events that correspond to stereotypic expectations tend to be communicated in a relatively abstract language, whereas unexpected events tend to be communicated in more concrete language (see Wigboldus *et al.*, 2000), matching nicely what has been found for memory distortions in the present study. Also, Need for Cognitive Closure appears to affect language use (Webster *et al.*, 1997) and memory in very similar ways. Thus, unless additional motivational processes are operating (e.g. self-presentation, impression management), there is reason to believe that spontaneous memory biases such as those observed here will also affect subsequent interpersonal discourse. However, the exact link between memory distortions and communicative biases, and the conditions under which the two may or may not concur, can only be investigated in research employing both memory and communication measures within the same experimental paradigm.

In conclusion, our research suggests that the unintentional transformation of verb-type information into trait-adjective representations in memory is shaped by stereotypic expectancies, whereas the inverse memory distortion (from adjectives to verbs) is not altered. When behavioral information corresponds to stereotypic expectancies, it is quickly transformed into traits, whereas such memory distortions are

inhibited when the behavioral information contradicts such expectancies. This finding nicely parallels the Linguistic Expectancy Bias (see Wigboldus et al., 2000). People not only use an increasingly abstract language the closer events correspond to their stereotypic expectancies, but they also show the same tendency toward abstraction in their own minds, when memorizing information. In both cases (interpersonal communication and memory), this shift toward abstraction is inhibited when behaviors or events contradict what would be expected on the basis of shared social stereotypes. This suggests that stereotypes and language abstraction are intrinsically linked both in intrapsychic processes and in interpersonal communication.

Notes

1. Responses were provided on a scale from -2 (extremely unlikely) to +2 (extremely likely), with '0' implying that the trait 'x' (e.g. *independent*) provided no information about the likelihood of being also 'y' (e.g. *sociable*).
2. No more than $\pm .5$ from the neutral midpoint.
3. Only in four cases did participants check both behavior and trait, an insignificant amount considering the total of 2880 responses given. Three of the dual responses referred to new trait-and-behavior pairs, and one was a deductive inference case in which the participant had seen the trait and falsely inferred also having seen the behavior.

References

- Beike, D. R., & Sherman, S. J. (1994). Social inference: Inductions, deductions, and analogies. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (2nd ed., Vol. 1, pp. 209–285). Hillsdale, NJ: Erlbaum.
- Boudreau, L. A., Baron, R., & Oliver, P. V. (1992). Effects of expected communication target expertise and timing of set on trait use in person description. *Personality and Social Psychology Bulletin*, 18, 447–452.
- Bourhis, R. Y., & Maass, A. (in press). Linguistic prejudice and stereotypes. In U. Ammon, N. Dittmar, K. J. Mattheier, & P. Trudgill, P. (Eds.), *Sociolinguistics. An international handbook of the science of language and society* (2nd ed). New York: De Gruyter.
- Brehm, S. S., & Kassin, S. M. (1996). *Social psychology*. Boston: Houghton Mifflin.
- Carlston, D. E., & Skowronski, J. J. (1986). Trait memory and behavior memory: The effects of alternative pathways on impression judgment response times. *Journal of Personality and Social Psychology*, 50, 5–13.
- Carlston, D. E., Skowronski, J. J., & Sparks, C. (1995). Savings in relearning: II. On the formation of the behavior-based trait associations and inferences. *Journal of Personality and Social Psychology*, 69, 420–436.
- De Leo, D., & Villa, A. (1986). *Il problema del rilevamento delle tipologie sessuali e il Bem Gender Role Inventory*. [The problem of measuring gender typologies and Bem's gender role inventory]. Firenze: Organizzazioni Speciali.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, 56, 5–18.
- Douglas, K. M., & Sutton, R. M. (2003). Effects of communication goals and expectancies on language abstraction. *Journal of Personality and Social Psychology*, 84, 682–696.
- Fiedler, K., Semin, G., & Bolten, S. (1989). Language use and reification of social information: Top-down and bottom-up processing in person cognition. *European Journal of Social Psychology*, 19, 271–295.
- Hamilton, D. L., & Rose, T. L. (1980). Illusory correlation and the maintenance of stereotypic beliefs. *Journal of Personality and Social Psychology*, 39, 832–845.
- Hastie, R., & Kumar, P. A. (1979). Person memory: Personality traits as organizing principles in memory for behavior. *Journal of Personality and Social Psychology*, 37, 25–38.
- Katz, D., & Braly, K. W. (1933). Racial stereotypes of 100 college students. *Journal of Abnormal and Social Psychology*, 28, 280–290.
- Kawakami, K., Dovidio, J. F., Moll, J., Hermsen, S., & Russin, A. (2000). Just say no (to stereotyping): Effects of training in the negation of stereotypic associations on stereotype activation. *Journal of Personality and Social Psychology*, 78, 871–888.
- Maass, A. (1999). Linguistic intergroup bias: Stereotype-perpetuation through language. In M. Zanna (Ed.), *Advances in Experimental Social Psychology*, 31, 79–121.
- Maass, A., Cadinu, M., Taroni, M., & Masserini, M. (2004). The Induction-Deduction-Asymmetry: Fact

- or artifact? Unpublished manuscript, Università di Padova.
- Maass, A., Colombo, A., Colombo, A., & Sherman, S. J. (2001). Inferring traits from behaviors vs. behaviors from traits: The Induction-Deduction Asymmetry. *Journal of Personality and Social Psychology*, *81*, 391–404.
- Maass, A., Montalcini, F., & Biciotti (1998). On the (dis-) confirmability of stereotypic attributes. *European Journal of Social Psychology*, *28*, 383–402.
- Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Castelli, L. (1999). On disregarding deviants: Exemplar typicality and person perception. *Current Psychology: Developmental, Learning, Personality, Social*, *18*, 47–70.
- Macrae, C. N., Stangor, C., & Milne, A. B. (1994). Activating social stereotypes: A functional analysis. *Journal of Experimental Social Psychology*, *30*, 370–389.
- Mikulincer, M., Yinon, A., & Kabili, D. (1991). Epistemic needs and learned helplessness. *European Journal of Personality*, *5*, 249–258.
- 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.
- Pierro, A., Mannetti, L., Converso, D., Garsia, V., Miglietta, A., Ravenna, M. et al. (1995). Caratteristiche strutturali della versione italiana della scala di bisogno di chiusura cognitiva di Webster e Kruglanski. [Structural characteristics of the Italian version of Webster and Kruglanski's Need for Cognition Scale]. *Testing, Psicometria, Metodologia*, *2*, 125–141.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353–363.
- Posner, M. I., & Keele, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology*, *83*, 304–308.
- Roedinger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology*, *21*, 803–814.
- Semin, G. R., & Fiedler, K. (1988). The cognitive functions of linguistic categories in describing persons: Social cognition and language. *Journal of Personality and Social Psychology*, *19*, 491–508.
- Semin, G. R., & Smith, E. R. (1999). Revisiting the past and back to the future: Memory systems and the linguistic representation of social events. *Journal of Personality and Social Psychology*, *76*, 877–892.
- Strange, W., Keeney, T., Kessel, F. S., & Jenkins, J. J. (1970). Abstraction over time of prototypes from distortions of random dot patterns: A replication. *Journal of Experimental Psychology*, *83*, 508–510.
- Uleman, J. S. (1987). Consciousness and control: The case of spontaneous inferences. *Personality and Social Psychology Bulletin*, *13*, 337–354.
- Uleman, J. S., Hon, A., Roman, R. J., & Moskowitz, G. B. (1996). On-line evidence for spontaneous trait inferences at encoding. *Personality and Social Psychology Bulletin*, *22*, 377–394.
- Uleman, J. S., & Moskowitz, G. B. (1994). Unintended effects of goals on unintended inferences. *Journal of Personality and Social Psychology*, *66*, 490–501.
- Uleman, J. S., Newman, L., & Winter, L. (1986). Can personality traits be inferred automatically? Spontaneous inferences require cognitive capacity at encoding. *Consciousness and Cognition*, *1*, 77–90.
- Uleman, J. S., Winborne, W. C., Winter, L., & Shechter, D. (1986). Personality differences in spontaneous personality inferences at encoding. *Journal of Personality and Social Psychology*, *51*, 396–403.
- Van Overwalle, F., Drenth, T., & Marsman, G. (1999). Spontaneous trait inferences: Are they linked to the actor or to the action? *Personality and Social Psychology Bulletin*, *25*, 450–462.
- Webster, D. M., & Kruglanski, A. W. (1994). Individual differences in the need for cognitive closure. *Journal of Personality and Social Psychology*, *67*, 1049–1062.
- Webster, D. M., Kruglanski, A. W., & Pattison, D. A. (1997). Motivated language use in intergroup contexts: Need-for-closure effects on the linguistic intergroup bias. *Journal of Personality and Social Psychology*, *72*, 1122–1131.
- Wigboldus, D. H. J., Dijksterhuis, A., & van Knippenberg, A. (2003). When stereotypes get in the way: Stereotypes obstruct stereotype-inconsistent trait inferences. *Journal of Personality and Social Psychology*, *84*, 470–484.
- Wigboldus, D. H. J., Semin, G. R., & Spears, R. (2000). How do we communicate stereotypes? Linguistic bases and inferential consequences. *Journal of Personality and Social Psychology*, *78*, 5–18.
- Wigboldus, D. H. J., Spears, R., & Semin, G. R. (1999). Categorization, content and the context of communicative behaviour. In N. Ellemers & R. Spears (Eds.), *Social identity: Context, commitment, content* (pp. 147–163). Oxford, UK: Blackwell Science.
- Winter, L., & Uleman, J. S. (1984). When are social judgments made? Evidence for the spontaneousness

of trait inferences. *Journal of Personality and Social Psychology*, 47, 237–252.

Winter, L., Uleman, J. S., & Cunniff, C. (1985). How automatic are social judgments? *Journal of Personality and Social Psychology*, 49, 904–917.

Wyer, R. S., & Lambert, A. J. (1994). The role of trait constructs in person perception: A historical

perspective. In P. G. Devine, D. L. Hamilton, & T. M. Ostrom (Eds.), *Social cognition: Impact on social psychology* (pp. 109–142). San Diego, CA: Academic Press.

Paper received 19 February 2004; revised version accepted 19 December 2004.

Appendix: Stimulus sentences*

Masculine

aggressive	gets angry easily
analytical	relies on logic when evaluating situations
athletic	practices sports
audacious	does not fear risky tasks
authoritative	exerts power over others
dominant	expects to be obeyed

Feminine

cheerful	is of good humor
diplomatic	knows to treat delicate issues prudently
ingenuous	believes anybody out of inexperience
intuitive	is able to understand a fact before reasoning about it
loyal	does not attempt to deceive others
timid	social relations are embarrassing him/her

Gender-neutral

concrete	considers the concrete aspects of the situation
disorganized	keeps his/her desk in disorder
dynamic	executes his/her tasks with energy
pessimistic	tends to make negative predictions
precise	examines the different facets of a situation in detail
up-to-date	updates his/her knowledge according to the needs of the organization

* The English translation does not always reflect the exact grammatical form (example: the Italian term for the composite 'to get angry' is 'arrabbiarsi', a reflexive verb).