VU Research Portal

Language use in intergroup contexts: The linguistic intergroup bias

Maass, A.; Salvi, C.; Arcuri, L.; Semin, G.R.

published in

Journal of Personality and Social Psychology
1989

DOI (link to publisher)

[10.1037/0022-3514.57.6.981](https://doi.org/10.1037/0022-3514.57.6.981)

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Maass, A., Salvi, C., Arcuri, L., & Semin, G. R. (1989). Language use in intergroup contexts: The linguistic intergroup bias. *Journal of Personality and Social Psychology*, 57, 981-993. <https://doi.org/10.1037/0022-3514.57.6.981>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Language Use in Intergroup Contexts: The Linguistic Intergroup Bias

Anne Maass and Daniela Salvi
Università di Padova
Padova, Italy

Luciano Arcuri
Università di Trieste
Trieste, Italy

Gün Semin
University of Sussex
Brighton, Great Britain

Three experiments examine how the type of language used to describe in-group and out-group behaviors contributes to the transmission and persistence of social stereotypes. Two experiments tested the hypothesis that people encode and communicate desirable in-group and undesirable out-group behaviors more abstractly than undesirable in-group and desirable out-group behaviors. Experiment 1 provided strong support for this hypothesis using a fixed-response scale format controlling for the level of abstractness developed from Semin and Fiedler's (1988a) linguistic category model. Experiment 2 yielded the same results with a free-response format. Experiment 3 demonstrated the important role that abstract versus concrete communication plays in the perpetuation of stereotypes. The implications of these findings and the use of the linguistic category model are discussed for the examination of the self-perpetuating cycle of stereotypes in communication processes.

Since the early social-psychological writings on prejudice and intergroup relations (cf. Allport, 1954), theorists have repeatedly commented on the fact that out-group stereotypes and in-group favoritism are apparently highly resistant to change and are likely to persist across generations. In recent years, various explanations have been tendered to account for the persistence of social stereotypes and intergroup biases. Motivational theories (see Tajfel & Turner, 1979; Turner, 1987), on one side, tend to stress the functional role of intergroup bias for self-esteem maintenance. Social-cognitive approaches, on the other side, refer to cognitive principles that contribute to the perpetuation of existing stereotypic beliefs even in the face of disconfirming evidence; among these principles are the preference for expectancy-confirming information (e.g., Snyder & Swann, 1978), superior recall for expectancy-congruent information regarding groups (Hastie & Kumar, 1979; Howard & Rothbart, 1980; Rothbart, Evans, & Fulero, 1979; Srull, Lichtenstein, & Rothbart, 1985), the existence of highly simplistic schemata about the out-group (Linville, Salovey, & Fischer, 1986), the exclusion of atypical group members from the category (Rothbart & Lewis, 1988), and the stereotype-consistent interpretation of ambiguous behavioral episodes (Duncan, 1976; Sagar & Schofield, 1980). For general overviews, see Brewer and Kramer, 1985; Hamilton, 1981; Hamilton & Troler, 1986; and Dovidio and Gaertner, 1986a.

Surprisingly, the role of language in the maintenance and transmission of stereotypes has largely been ignored in this context (for exceptions, see Van Dijk's work; e.g., Van Dijk, 1984,

1987). Although there is ample evidence that language apparently plays an important role in other domains of human cognition, its contribution to intergroup biases in general and to the persistence of social stereotypes in particular has remained unclear.

In this article, we advance the thesis that, in an intergroup context, language is used in a manner that renders disconfirmation of preexisting ideas about in-group and out-group difficult. Applying a recent psycholinguistic model by Semin and Fiedler (1988a, 1988b; Fiedler & Semin, 1988a, 1988b), we argue that the same behavioral episodes are encoded at different levels of abstraction depending on whether such behaviors have positive or negative connotations and whether they are performed by in-group or out-group members.¹ In particular, we argue that socially desirable in-group behaviors and undesirable out-group behaviors are encoded at a high level of abstraction, whereas socially undesirable in-group behaviors and desirable out-group behaviors are encoded at a low level of abstraction. Considering that information encoded at an abstract level is relatively resistant to disconfirmation and implies high stability over time (Semin & Fiedler, 1988a), our approach may contribute to a better understanding of the persistence of stereotypic beliefs. Thus, we propose a model in which existing intergroup biases produce a biased language use, which in turn contributes to the maintenance of existing biases. We briefly outline Semin and Fiedler's model and subsequently discuss its implications for the intergroup context.

¹ Throughout this article, we use the term *encoding* in a somewhat different and more general way than is usually done by cognitive psychologists. In reference to its original meaning as "translating into a code," we define *encoding* as the translation of language-free (visually presented) information into a linguistic code.

Correspondence concerning this article should be addressed to Anne Maass, Dipartimento di Psicologia Generale, Università di Padova, Piazza Capitanio, 3, 35139 Padova, Italy.

Table 1
Semin and Fiedler's (1988a) Linguistic Category Model

Level, category, and characteristic features	Examples
I: Descriptive action verb—Objective description of a specific and observable behavior with clear beginning and end; refers to specific situation and specific object; usually does not have positive or negative connotations	Kiss Look Run Visit Call
II: Interpretive action verb—Describes a general class of behaviors (including various possible behavioral acts), but refers to a specific action with clear beginning and end, to a specific object situation; provides an interpretation beyond the mere description; has positive or negative semantic connotations	Help Offend Inhibit Cheat Threaten
III: State verb—Enduring states (emotional, affective, mental, etc.) beyond specific behaviors or situations; reference to specific object; no clear beginning and end; provides interpretation beyond mere description; does not readily take the progressive form or imperative	Believe Love Admire Desire Envy
IV: Adjectives—Describe highly abstract person dispositions; no object reference or situation reference; highly interpretive, detached from specific behaviors	Honest Impulsive Reliable Helpful Creative

The linguistic category model. Semin and Fiedler's (1988a) linguistic category model distinguishes four linguistic categories that may be used in describing other people (see Table 1). The same behavioral episode may be encoded at four different levels of abstraction. At the most concrete level, it may be encoded in terms of descriptive action verbs (DAVs), such as *to call* or *to touch*, that refer to a single, observable event, defined by at least one physically invariant feature. At the second level of abstraction are interpretive action verbs (IAVs), such as *to help* or *to cheat*, which are no longer bound to physically invariant features although they preserve the reference to a single behavioral episode. In contrast to DAVs, IAVs such as *to help* describe a general class of behaviors without identifying the specific behavior to which they refer in a given context (e.g., *to help* may refer to such different behaviors as lending money, opening the door, or reviving somebody). State verbs (SVs) describe a psychological state such as hate or desire; they have no direct reference to a specific behavioral episode or to a specific situation, but do refer to a specific object. At the highest level of abstraction are adjectives (ADJ), such as *aggressive* or *creative*, that describe highly abstract dispositions or characteristics of a person. ADJs provide generalizations across specific behavioral events, across situations, and across objects.

The following hypothetical example illustrates how the same behavioral episode can be encoded at different levels of abstraction. Imagine that Person A is hitting Person B's arm with his fist. This behavioral event may be encoded at a concrete level as "A is punching B," or at a slightly more abstract level as "A is hurting B." Alternatively, one may interpret the scene as "A hates B," or at the highest level of abstraction, one may conclude

that "A is aggressive." The codification at a higher or lower level of abstraction may have a number of important implications. An abstract statement such as "A is aggressive" implies great stability over time and generality across settings and interaction partners, suggesting that Person A will behave similarly in the future, in different situations, and with other people. He could also be expected to show related aggressive behaviors such as kicking, spitting, or pulling hair. Obviously, none of these conclusions could be drawn from a description at the concrete level.

In fact, Semin and Fiedler (1988a) have demonstrated that abstract statements as opposed to concrete statements are perceived as less verifiable; they also imply greater temporal stability and are perceived as revealing more about the person and less about the situation than are concrete statements; furthermore, they are perceived more readily as giving rise to disagreement.

Implications for intergroup settings. We believe that, as a conceptual framework and methodological tool, the language category model may have a number of interesting implications for intergroup relations and stereotyping. Membership biases are so well documented in the literature on intergroup relations (see Brewer, 1979; Brewer & Kramer, 1985; Hamilton & Trolie, 1986; Tajfel, 1982; Wilder, 1986, for reviews) that a detailed description appears superfluous in this context. This research consistently demonstrates that the mere categorization of people into groups can lead to favoritism toward the in-group and discrimination against the out-group. The tendency to establish a relative superiority of the in-group over the out-group has generally been interpreted in motivational terms. According to social identity theory (Tajfel & Turner, 1979), people strive to enhance the value of their own group relative to the out-group in order to maintain a positive social identity.

Consistent with this general bias, people also tend to hold differential expectancies about the behavior of in-group and out-group members. They expect in-group members to display more desirable and fewer undesirable behaviors than out-group members (Howard & Rothbart, 1980). Furthermore, they are more likely to infer negative dispositions from undesirable out-group behaviors than from undesirable in-group behaviors and are less likely to infer positive dispositions from desirable out-group behaviors than from desirable in-group behaviors (e.g., Taylor & Jaggi, 1974; Pettigrew, 1979; Hewstone & Jaspars, 1984). This suggests that the evidence-to-inference (Rothbart & Park, 1986) or act-to-disposition link is much tighter when a behavior episode confirms preconceived ideas about the actor.

This bias is not as illogical as it may initially appear, considering that stereotypic expectancies make probabilistic rather than deterministic predictions. Assume that members of Group A expect members of an out-group, B, to display socially undesirable behaviors. Obviously they do not expect members of Group B to always behave in socially undesirable ways. Rather, they assume that, compared with the overall average rate of such behaviors, members of Group B have an above-chance probability of engaging in undesirable behaviors and a below-chance probability of showing desirable behaviors. Consequently, when members of the out-group do display desirable behaviors, there is no need to revise the general negative conception of out-group members because a probabilistic expectancy explicitly allows for a certain number of instances in which desirable behaviors will occur. According to Wilder (1986), this

tendency is particularly pronounced for unexpected behaviors of out-group members, for which is held an Aristotelian view of lawfulness that allows for a great number of exceptions before the general rule has to be revised.

There are at least two ways to reconcile unexpected behaviors (e.g., a generous act of a single member of the Jewish community) with the general conception of the category (e.g., Jews are stingy), namely, the dissociation of the single, atypical member from the category as a whole (Allport's, 1954, referencing principle; see also Rothbart & Lewis, 1988; Weber & Crocker, 1983) and the dissociation of the single act from the group member. It is this second mechanism that is of interest here. Thus, an out-group member's socially desirable behavior may simply be interpreted as a specific instance that is situationally and temporally bound and that is largely unrelated to the actor's more enduring properties, such as his or her inclinations or psychological state. Basically, expectancy-incongruent behavior episodes are treated as exceptions to the rule that can be reconciled through situational and temporal constraints. We believe that this mechanism is reflected in a corresponding linguistic intergroup bias in which unexpected behaviors (desirable out-group and undesirable in-group behaviors) are described in concrete terms without generalizing beyond the given information.

Although people may occasionally hold positive expectancies about out-groups (e.g., embedded in the overall negative stereotype of Jews is the expectancy that Jews will display exceptionally high levels of intelligence), the general tendency to favor one's own group and the specific expectancies about in-group and out-group behaviors will, in most cases, coincide. For such cases, we suspect that when people observe members of their own category performing desirable behaviors, they will readily associate those acts with the actors' enduring properties or psychological state, and hence will describe the acts in abstract linguistic terms. A similar tendency should emerge for undesirable out-group behaviors. In contrast, desirable out-group and undesirable in-group behaviors are more likely to be encoded as concrete behavioral instances, without abstraction beyond the given information. Such behavioral episodes that are incongruent with the general perception of in-group and out-group and as such are unexpected are treated as the exception that confirms the rule by simply shifting the level of analysis from the general to the specific.

Hypotheses. We predict, therefore, that people will tend to encode (and to communicate) favorable in-group and unfavorable out-group behaviors at a higher level of abstraction than unfavorable in-group and favorable out-group behaviors. In particular, the same socially desirable (e.g., altruistic) act will be encoded at a higher level of abstraction when performed by an in-group member than when performed by an out-group member. On the contrary, a socially undesirable (e.g., aggressive) act will be encoded at a higher level of abstraction when performed by an out-group member rather than when performed by an in-group member. These patterns of behavior are termed the *linguistic intergroup bias*.

To test these hypotheses, in the first 2 experiments we exposed members of mutually exclusive social categories to a series of episodes in which members of either the in-group or the out-group performed socially desirable or undesirable behaviors. The subject's task was simply to encode the visually pre-

sented information. In Experiment 1, four response alternatives were provided, corresponding to the four levels of abstraction in the linguistic category model, from which the subject had to select one (multiple-choice responses). In Experiment 2, subjects were exposed to the same stimulus material, but were instructed to provide a short description of each scene (free description). Thus, Experiment 1 attempted to maximize experimental control and to minimize the likelihood of uncodable responses, whereas Experiment 2 intended to increase the external validity by rendering the subjects' task similar to naturally occurring encoding processes.

Experiment 1

Method

Choice of in-group versus out-group category. Because the experiments were conducted in Italy, we tried to identify social groups that represented meaningful categories within the cultural context of that country. Various cities in Italy, most notably Siena, have yearly horse racing competitions called *palio*, in which members of the various sections or quarters of the city (so-called *contrada*) compete against each other. The horse races, which take place in the central square of the city, are public festivals where members and supporters of the various *contrada* cheer for their teams. Identification with the *contrada* is remarkably high, partially because of the fact that there is a direct competition between the *contrada* and partially because of the long historical tradition of the *palio*. For example, the *palio* of Ferrara—a small city in northern Italy where the present series of studies was conducted—dates back to 1279. The tradition was interrupted only during short periods in which the city was hit by the Black Plague. Not surprisingly, intergroup hostilities are particularly frequent during the weeks before the *palio* and may take more or less playful forms, ranging from water balloon fights to reciprocal flag thefts to secretly drugging the other team's horses or donkeys.

Of the eight *contrada* of Ferrara, we selected two with approximately equal status and comparable probability of winning for the first experiment.² Fifty-one subjects (15 women and 36 men) from the *contrada* San Giorgio ($n = 31$) and San Giacomo ($n = 20$) participated in the experiment. The mean age of subjects was 24.2.

Procedure. Subjects were recruited in the respective clubhouses of the *contrada*, which, during the weeks before the competition, are regularly frequented by the members of the *contrada*. A female experimenter visited the clubhouses in the evenings and asked those present to participate in a psycholinguistic study. Of the six *contrada* (of which two were used for pilot testing, two for Experiment 1, and two for Experiment 2), four were visited only once, whereas two visits on consecutive evenings were necessary for the data collection in the remaining *contradas*. The overall rate of volunteering (for pilot study, Experiment 1, and Experiment 2 combined) was 93%, with only 11 out of 160 subjects refusing to participate in the study. Depending on the number of subjects present, up to 10 subjects were run at the same time. Communication between subjects regarding the study was discouraged until all questionnaires had been completed. Subjects were informed that the episodes were real and had been collected over the past 2 years. The subjects were further informed that each episode referred to different protagonists,

² Status equality was considered an important prerequisite as the degree of intergroup biases has often been found to vary with group status in previous research (e.g., Deschamps, 1972–73). Because the present experiments attempted to provide the very first empirical test of the linguistic intergroup hypothesis, we decided to keep potentially moderating variables such as status difference constant.

but that the protagonists' identities were disguised in the visual representation in order to protect their anonymity.

Members of two contrada were presented with 16 single-frame cartoons in which a member of either their own contrada (in-group) or a competing contrada (out-group) performed a behavior. Half the cartoons portrayed socially desirable behaviors (e.g., helping), and half portrayed undesirable behaviors (e.g., littering). For both behavior types, half the cartoons showed an isolated actor, whereas the other half showed a protagonist interacting with others. Both types of items were included in order to test whether Semin and Fiedler's (1988a) model—originally formulated for interactional verbs—was also applicable to noninteractional episodes such as littering. The cartoons were further divided into general behaviors such as littering or kicking a dog and palio-specific ones such as stealing each other's flags or drugging the other team's horse. On intuitive grounds, one might suspect that intergroup biases are more likely to surface with palio-specific episodes because they are directly relevant to the categorization. Subjects were provided with four response alternatives, each giving a brief description of the protagonist's behavior in the scene. The four response alternatives corresponded to the different levels of abstraction in the linguistic category model. Subjects were simply asked to select the one that, in their opinion, described the scene best. Thus, the design consisted of a 2 (contrada San Giacomo vs. San Giorgio) $\times$ 2 (in-group vs. out-group protagonist) $\times$ 2 (desirable vs. undesirable behaviors) $\times$ 2 (protagonist alone vs. interacting) $\times$ 2 (general vs. palio-specific behaviors) design in which the last three factors represent repeated measures.

Stimulus materials and pilot testing. The episodes were presented in form of visual episodes (cartoons) for two reasons:

1. Compared with written representations, this procedure is language-free and therefore does not bias subsequent encoding due to the initial level of abstraction (see Fiedler & Semin, 1988a, for the impact of prior encoding abstraction on subsequent processing).
2. Compared with other visual representations such as films or videotapes, cartoons allow systematic variations of the protagonist's group membership (in this case, by changing the color of the actor's shirt according to the color of his contrada) while holding all other features constant.

To select an appropriate set of episodes, 24 episodes were generated and subjected to an extensive pretest involving 44 subjects from two contrada (subsequently not used in either experiment). Twelve episodes described positive behaviors and 12 described negative behaviors. Within each group, the episodes depicted 3 behaviors that were specific and interactional, 3 that were general and interactional, 3 that were specific and noninteractional, and 3 that were general and noninteractional. The cartoons were produced by a semiprofessional cartoonist in such a way that the protagonist's age and sex remained ambiguous. Approximately half the pilot subjects were instructed to select one of four response alternatives, whereas the remaining half provided free descriptions of each episode. Furthermore, half of both subsamples received a version in which the protagonists of all scenes were members of the in-group, whereas for the remaining subjects, the protagonists were out-group members. In addition to the encoding task, pilot subjects were asked to rate the social desirability of each act on a 5-point scale (ranging from 1 = *very negative* to 5 = *very positive*).

On the basis of the pilot test, the initial pool of cartoons was reduced to 16, according to the following criteria. First, the only episodes included were those judged positively (>3.5) or negatively (<2.5), respectively, for both in-group and out-group protagonists. This criterion resulted in the exclusion of various items in which, for example, an undesirable behavior was judged negatively only when performed by an out-group, but was judged neutrally or even slightly positively when performed by an in-group member. Second, episodes were excluded (or, in some cases, modified) whenever the descriptions of the free-response

subjects indicated that the scene had been interpreted differently than was intended by the experimenters.

The final series of cartoons consisted of eight socially desirable and eight undesirable behaviors (either interactional or involving an isolated actor and either general or palio-specific) with two items representing each episode type. The protagonist was always identified by the letter *A*, whereas for interactional items, the interaction partner was labeled *B*. The episodes were presented in the same randomly determined order for all subjects. The cartoons were bound into booklets containing black and white reproductions of the original drawings. The protagonist's group membership was identified by the color of his shirt, which corresponded to the color of his contrada. As in other sports, each contrada is associated with a particular color combination (applied to the competitors' dress, saddles, flags, watch, bracelets, etc.). These colors, which have remained constant over the centuries (as testified by antique flags) and are therefore well-known not only to those directly involved in the palio but also to the general population of Ferrara, allowed an easy identification of the protagonist's contrada. Thus, group membership was unambiguous and highly salient, as it constituted the only color in an otherwise black-and-white drawing. Although a misattribution of the protagonist's group membership was highly unlikely, any possibility of error was further reduced by explicitly reminding the subject of the actor's contrada in the formulation of the dependent variable (e.g., "A, a member of the contrada San Giacomo, . . .").

Level of abstraction: Response alternatives. Four response alternatives were provided for each episode, corresponding to the four levels of abstraction in the linguistic category model and presented in the same standardized order, starting from the lowest level of abstraction (DAV, IAV, SV, ADJ; e.g., A, a member of the contrada San Giacomo, hits B; hurts B; hates B; is violent). To be included among the response alternatives, a description had to meet the following criteria:

1. It had to be considered an appropriate and valid description of a given scene as based on the independent judgments of two raters.
2. It had to be assigned to the same level of abstraction by two independent raters familiar with the linguistic category model.
3. It had to be sufficiently simple to be comprehensible to all subjects independent of their social class or educational level (unusual or sophisticated expressions were avoided).
4. Finally, items were unacceptable whenever the distribution across response alternatives was grossly uneven, such that one of the four alternatives attracted almost all choices in the pilot study.

Subjects were instructed to view each episode in the booklet and to select the response alternative that, in their opinion, best described the scene. There was no time limit. As in previous research (see Semin & Fiedler, 1988b), the subjects' responses were scored according to the level of abstraction, such that higher values indicated greater abstraction (ranging from 1 = *DAV* to 4 = *ADJ*). Each subject's responses were averaged across the two items representing the same episode type (e.g., desirable, palio-specific, interactional episodes).

Positivity rating. After completing the encoding task, subjects were asked to view the episodes again and, this time, to rate the positivity or negativity of each behavioral episode on a 5-point scale (ranging from 1 = *very negative* to 5 = *very positive*).

Results

Level of abstraction. We excluded 3 subjects from the analysis owing to missing data. A preliminary 2 (contrada) $\times$ 2 (in-group vs. out-group protagonist) $\times$ 2 (desirable vs. undesirable behavior) $\times$ 2 (isolated actor vs. interaction) $\times$ 2 (general vs. palio-specific behavior) analysis of variance (ANOVA) with repeated measures on the last three variables was conducted in order to test whether the subject's contrada, the episode's interactional character, or the specific reference to the palio would

Table 2
Mean Level of Abstraction as a Function of Group Membership and Social Desirability: Experiment 1

Membership of protagonist	Desirability of behavior	
	Desirable	Undesirable
In-group		
<i>M</i>	2.69	2.51
<i>n</i>	24	24
Out-group		
<i>M</i>	2.47	2.82
<i>n</i>	24	24

Note. Means are based on a 4-point scale, with higher scores indicating higher levels of encoding abstraction.

moderate the predicted interaction between the protagonist's group membership and the act's desirability. Because none of these three variables produced any main or interaction effects, we report the results of a simplified two-way ANOVA in which the protagonist's group membership constituted the independent variable and social desirability the repeated measure. Obviously, the dependent variable for this analysis was the mean level of abstraction for all eight desirable and all eight undesirable episodes, respectively. The results confirm the predicted interaction between group membership and social desirability, $F(1, 46) = 9.49, p < .01$, which is presented in Table 2. As predicted, the same desirable behavior was encoded at a higher level of abstraction when displayed by an in-group member than when displayed by an out-group member, $t(46) = 1.82, p < .05$, one-tailed. Exactly the opposite occurred for socially undesirable episodes, which were encoded at a higher level of abstraction when performed by an out-group member rather than an in-group member, $t(46) = 2.57, p < .01$, one-tailed. A different way to look at these data is to compare the level of abstraction for desirable and undesirable acts within each group. Here the results show that undesirable out-group behaviors were encoded at a higher level of abstraction than were desirable ones, $t(46) = 2.90, p < .01$, one-tailed, whereas there was a nonsignificant trend in the opposite direction when the same behaviors were attributed to in-group members, $t(46) = 1.49, ns$.

Underlying this analysis and the respective coding scheme is the assumption that the four levels of abstraction form a continuum, approximating an interval scale (see also Semin & Fiedler, 1988b). As this assumption may be questioned, we performed a second set of analyses, using as the dependent variable the overall frequency with which each linguistic category was selected to describe positive and negative episodes of in-group and out-group protagonists. Separate log-linear analyses were performed for desirable and undesirable episodes in order to investigate the effect of the protagonist's group membership on the use of the linguistic categories (DAV, IAV, SV, ADJ). Both analyses supported the predicted interaction between category membership and level of abstraction, $\chi^2(3, N = 384) = 7.66, p = .05$, for desirable episodes and $\chi^2(3, N = 384) = 7.30, p = .06$, for undesirable episodes (see Table 3).

For desirable episodes, concrete descriptions (DAVs and

IAVs) were used more frequently for out-group than for in-group episodes, whereas abstract descriptions (SVs and ADJs) were used more frequently for in-group than for out-group episodes. Follow-up comparisons between expected and observed frequencies, however, indicated that these differences reached conventional levels of significance only for DAVs and SVs (see Table 3). Turning to the undesirable episodes, one sees that an exactly opposite pattern emerged. Here, subjects showed a clear preference for concrete descriptions (DAVs and IAVs) when describing in-group rather than out-group behaviors. At the same time, they tended to describe out-group members more frequently in adjectival or state terms than they did in-group members. Again, actual response frequencies deviated from the expected frequencies only for two of the four categories, namely IAVs and ADJs. Summarizing the results of the log-linear analyses, the frequencies with which each linguistic category was selected as a function of group membership and desirability points, in every case, in the predicted direction. However, comparisons are statistically reliable only for DAVs and SVs when the episodes are socially desirable and for IAVs and ADJs when the episodes are undesirable.

Results are even clearer when the two lower levels (DAV and IAV combined) and the two higher levels (SV and ADJ combined) of abstraction are considered together. In this case, the protagonist's group membership had a strong impact on the level of abstraction both for desirable, $\chi^2(1, N = 384) = 5.53, p < .05$, and undesirable episodes, $\chi^2(1, N = 384) = 6.03, p < .05$. Desirable in-group behaviors were considerably more likely to be encoded at an abstract level ($n = 109$ out of 192 entries) than at a concrete level ($n = 83$), whereas desirable out-group behaviors were more frequently encoded in concrete terms ($n = 106$) than in abstract terms ($n = 86$). This tendency reversed for undesirable episodes, where concrete descriptions ($n = 104$) outnumbered abstract ones ($n = 88$) when the episodes were attributed to an in-group member. For out-group episodes, abstract descriptions ($n = 112$) were favored over concrete ones ($n = 80$).

Positivity rating. The positivity ratings were subjected to a 2 (contrada) $\times$ 2 (in-group vs. out-group protagonist) $\times$ 2 (desirable vs. undesirable behavior) $\times$ 2 (isolated actor vs. interac-

Table 3
Frequency of Linguistic Category as a Function of Group Membership: Experiment 1

Episode type and group membership of protagonist	Linguistic category				Total
	DAV	IAV	SV	ADJ	
Desirable					
In-group	27**	56	59**	50	192
Out-group	42**	64	40**	46	192
Undesirable					
In-group	48	56*	31	57**	192
Out-group	40	40*	32	80**	192

Note. DAV = descriptive action verb, IAV = interpretive action verb, SV = state verb, ADJ = adjective.

* $p < 0.06$, one-tailed. ** $p < .05$, one-tailed.

tion) $\times$ 2 (general vs. palio-specific behavior) ANOVA with repeated measures on the last three variables. Four subjects were excluded owing to missing data. Not surprisingly, in-group behaviors were rated more positively ($M = 2.99$) than out-group behaviors ($M = 2.74$), $F(1, 40) = 10.33$, $p < .01$, and desirable behaviors were rated more positively ($M = 4.06$) than undesirable behaviors ($M = 1.67$), $F(1, 40) = 559.55$, $p < .01$. Furthermore, interactional episodes were rated more positively ($M = 2.94$) than those involving single actors ($M = 2.79$), $F(1, 40) = 7.79$, $p < .01$. Finally, a significant In-group–Out-group $\times$ Interactional–Isolated Actor $\times$ General–Palio-Specific Behaviors interaction indicated that in-group behaviors were rated significantly more positively than identical out-group behaviors only for general and interactional behaviors (3.16 for in-group vs. 2.74 for out-group), $t(40) = 2.02$, $p < .05$, and for palio-specific behaviors involving single actors (3.18 for in-group vs. 2.58 for out-group), $t(40) = 2.90$, $p < .05$. For the remaining two types of episodes (general–single actor and specific–interactional), positivity ratings were virtually identical for in-group and out-group.

Correlational analyses. A somewhat different way to conceptualize our main thesis is in correlational terms. One may predict that, for in-group behaviors, the more socially desirable an action is perceived to be, the higher the average level of abstraction at which it is encoded. The opposite should be expected for out-group behaviors, with desirable actions being encoded at a lower level of abstraction. To test this hypothesis, we obtained correlations between the two dependent variables, using the different episodes (rather than subjects) as the unit of analysis. After excluding subjects with missing responses on either dependent variable, the mean favorability and mean level of encoding abstraction were calculated for each of the 16 episodes. This was done separately for the in-group and out-group conditions. Thus, for each of the 16 episodes, the average positivity rating of the episode was calculated for subjects in the in-group and out-group conditions, as was the respective average abstraction level that each group had selected when encoding the episode. Despite the fact that (as we have reported) out-group episodes had overall been rated less favorably than in-group behaviors, the two groups agreed perfectly on the relative evaluation of the episodes. The correlation between the evaluation of the episodes performed by in-group and out-group members was $r(16) = .96$, $p < .01$. More important, as expected, the more positive an episode ascribed to an out-group member, the lower the average abstraction level at which it was encoded, $r(16) = -.55$, $p < .05$. For episodes ascribed to in-group members, a reverse trend was observed, with positively evaluated episodes encoded at a higher level of abstraction, $r(16) = .28$, *ns*. Not surprisingly, correlations for in-group and out-group members differed significantly from each other ($z = 2.33$, $p < .01$).

Discussion

The results of the first experiment clearly confirm the hypothesis that people encode undesirable out-group and desirable in-group behaviors at a higher level of abstraction than they do desirable out-group and undesirable in-group behaviors. Interestingly, this finding was not modified by any higher-order interaction (see preliminary analysis), suggesting that encoding

differences of desirable and undesirable in-group and out-group behaviors occur for general as well as conflict-setting-specific episodes, and for interactional as well as noninteractional episodes. The correlational analysis provides further support for our contention. The more negatively an out-group action is perceived, the higher the level of abstraction at which it is encoded. Interestingly, people show exactly the opposite tendency when observing behavior episodes of in-group members. In this case, the desirable actions tend to be encoded at a slightly higher level of abstraction.

Yet, one may object that the data pattern could possibly be the result of an experimental artifact. It is conceivable that the experimentally provided response alternatives had more extreme evaluative implications at a higher level of abstraction. That is, the selected state verbs and adjectives may have had a more positive connotation than descriptive and interpretative action verbs when the episode referred to a socially desirable act, but a more negative connotation when the episode referred to a socially undesirable act (e.g., *altruistic* may be considered more positive than *to help*, and *aggressive* may be considered more negative than *to hit somebody*). If this was the case, subjects may simply have selected a low level of abstraction for desirable out-group behaviors and a high level of abstraction for undesirable out-group behaviors in an attempt to describe the out-group as negatively as possible.

To test this possibility, we conducted a small follow-up study in which 35 subjects not involved in the palio were asked to rate the positivity or negativity of each response alternative without having been exposed to the cartoons. To avoid an unreasonable demand on the subjects, we created two subsets by randomly assigning one item of each episode type to Set A and the other to Set B. Each subject rated on a 5-point scale (with higher scores indicating higher social desirability) the positivity of 32 sentences (four desirable and four undesirable episodes with four response alternatives each), which were presented in the same order as in the experiment and in which the names of the contrada had been substituted with *Group X* and *Group Y*. Thus, each subject received a list of 32 sentences (e.g., "A, a member of group X, burns the flag of group Y") and was asked to rate the positivity of each action ("In your opinion, how positive or negative is the behavior or attribute described in this sentence?") A 2×2 within-subjects ANOVA with desirability of the episode and low level (DAV, IAV) versus high level (SV, ADJ) of abstraction as variables revealed a main effect for both— $F(1, 34) = 733.10$, $p < .01$, for desirability and $F(1, 34) = 11.16$, $p < .01$, for high versus low abstraction—but no interaction. This suggests that response alternatives at a higher level of abstraction (SV and ADJ, $M = 2.92$) were considered more positively than those at a lower level of abstraction (DAV and IAV, $M = 2.78$), but this was true for both desirable and undesirable episodes. This clearly precludes the alternative interpretation that the level of abstraction selected by the subjects in Experiment 1 was mediated by the differential positivity of the response alternatives.

This suggests that people do, in fact, encode undesirable out-group and desirable in-group behaviors at a higher level of abstraction than desirable out-group and undesirable in-group behaviors and that this tendency is not simply an experimental artifact. Yet it remains to be demonstrated that group members

show the same tendency when freely encoding an observed action. In the interest of experimental control, response alternatives were provided in the first experiment, thereby preventing subjects from generating their own interpretation of the episodes. Besides imposing a specific interpretation of each episode (which may or may not have coincided with the subject's own interpretation), the forced-choice procedure used in Experiment 1 artificially reduced the almost infinite world of alternative interpretations to only four, which in turn may have led to an overestimation of the linguistic intergroup bias: In this experiment, the only possibility of providing differential descriptions of in-group and out-group behaviors was to shift the level of analysis along the linguistic abstractness dimension. In the real world, subjects are free to use alternative and possibly more immediate strategies of language use to distinguish in-group and out-group behaviors. For example, they may simply use different terms within the same linguistic category (e.g., a behavior that is described as "playful" when displayed by an in-group member may be interpreted as "aggressive" when ascribed to an out-group member; see Duncan, 1976; Sagar & Schofield, 1980). It is possible that the linguistic intergroup bias is displayed only in situations in which people are prevented from using other strategies. Therefore, we conducted a second experiment in which subjects were asked to provide free descriptions of each scene.

Experiment 2

Method

In the second experiment, we used a procedure and stimulus material identical to those used in the previous one, with one modification: Rather than selecting a response in a multiple-choice procedure, subjects were asked to briefly describe each scene in their own words. To ensure that the subject of the sentence would refer to the protagonist, a sentence completion task was used, starting with "A member of the contrada S. Maria in Vado [or S. Spirito]."

Subjects. Forty-four subjects from two contrada (Santa Maria in Vado and Santo Spirito) not involved in the pilot study or in Experiment 1 participated in the study. As in Experiment 1, the sample reflected the general sociodemographic characteristics of the contrada, whose members tend to be young (subjects' mean age = 21.1, ranging from 14 to 33) and predominantly male (30 men, 14 women).

Scoring. Responses were scored by two independent raters familiar with Semin and Fiedler's (1988a, 1988b) scoring criteria (interrater reliability, $r = .81$), but blind as to the protagonist's category membership. The original scoring system was modified in three ways: First, nouns functionally identical to adjectives were scored as ADJ (e.g., "he is an altruist," "he is a spy"). Note that the use of nouns instead of adjectives is quite common in the Italian language. Second, sentences with more than one (nonauxiliary) verb or adjective (e.g., "he is happy because he won") received multiple scores, which were then averaged to an overall score. Third, verbs in connection with *always* were coded as ADJ (e.g., "he always hits dogs"). Furthermore, responses in which the subject of the sentence did not refer to the protagonist were considered uncodable (e.g., "that is okay," "I agree"). Twelve percent of the responses were unclassifiable³. Almost half of the subjects ($n = 21$) gave uncodable responses for at least one item, and 10 subjects provided uncodable responses for both items of a given episode type (e.g., socially desirable, palio-specific episodes involving an interaction partner). To avoid an unreasonable subject attrition, the variables interactional-single actor and generic-palio-specific were collapsed in the analysis. In fact, a sepa-

Table 4

Mean Level of Abstraction as a Function of Group Membership and Social Desirability: Experiment 2

Membership of protagonist	Desirability of behavior	
	Desirable	Undesirable
In-group		
<i>M</i>	2.63	2.75
<i>n</i>	19	19
Out-group		
<i>M</i>	2.38	2.67
<i>n</i>	25	25

Note. Means are based on a 4-point scale, with higher scores indicating higher levels of encoding abstraction.

rate treatment of these variables appears unnecessary in view of the results of Experiment 1, showing that neither the interactional character of the episode nor its specificity to the palio had any effect on the level of encoding abstraction, and that neither variable interacted with the main factors under investigation (group membership and desirability). Thus, for each subject we calculated one mean score for the socially desirable episodes and one for the undesirable episodes by averaging all codable responses.

Results

Level of abstraction. As predicted, desirable in-group behaviors were encoded at a higher level of abstraction than desirable out-group behaviors, $t(40, \text{one-tailed}) = 2.62, p < .01$ (see Table 4). However, undesirable in-group and out-group behaviors were encoded at virtually the same level of abstraction. Furthermore, episodes involving out-group members were encoded at a higher level of abstraction when they were undesirable than when they were desirable, $t(40, \text{one-tailed}) = 3.24, p < .01$, whereas episodes involving in-group members were encoded in a largely unbiased manner.

Positivity rating. Four subjects were excluded owing to missing data. In line with Experiment 1, a 2 (in-group vs. out-group membership of the protagonist) $\times$ 2 (desirable vs. undesirable behaviors) $\times$ 2 (contrada) mixed ANOVA indicated that socially desirable episodes ($M = 3.67$) were rated more positively than

³ Compared with previous research reporting a 6% rate of unclassifiable responses (see Semin & Fiedler, 1988b), uncodable responses are relatively frequent in this experiment (12%). We suspect that three factors may have contributed to the high rate of unclassifiable responses: (a) The average educational level of our subjects was clearly lower than that of the university students used by Semin and Fiedler. The sentence completion, using the protagonist as the grammatical subject, may require a prior training period for such subjects. (b) The social setting (evening meetings at the clubhouses as opposed to the classroom setting in Semin & Fiedler's study) may have contributed to the subjects' tendency not to follow the experimentally imposed rules closely. (c) Contrary to Semin and Fiedler's subjects, our subjects were personally involved in the experimental task as they were participating in an important competition. Their high level of ego involvement may have facilitated unclassifiable responses in which subjects provided spontaneous reactions to (e.g., "these bastards, I agree") rather than descriptions of the action.

undesirable ones ($M = 1.72$), $F(1, 36) = 58.33$, $p < .01$, and that in-group behaviors ($M = 2.93$) were rated more positively than otherwise identical out-group behaviors ($M = 2.53$), $F(1, 36) = 10.61$, $p < .01$.

Discussion

The results of Experiment 2 partially confirm those of the Experiment 1. As in that experiment, subjects tended to encode undesirable out-group behaviors at a higher level of abstraction than desirable ones, even when providing their own interpretation of a behavioral sequence, whereas there was no such bias when encoding in-group behaviors. Apparently, the linguistic intergroup bias is much more pronounced for out-group than for in-group actions. This is in line with Wilder's (1986) contention that people tend to apply an Aristotelian view of lawfulness to out-group members but not to in-group members. Rather than revising the perception of the out-group in the face of expectancy-incongruent evidence, people are likely to regard such inconsistencies as single, exceptional episodes largely unrelated to the general rule. In fact, subjects in Experiment 2 shifted their level of analysis to the concrete pole of the abstractness dimension only when encountering expectancy-incongruent episodes involving out-group members.

The second experiment further suggests that the same desirable episodes were encoded at a higher level of abstraction when performed by an in-group member than when performed by an out-group member. However, contrary to our initial hypothesis and to the findings of the previous experiment, this tendency did not reverse for undesirable behaviors. At this point it remains unclear why the results of the second study deviate in this respect from those obtained in Experiment 1.

With this one exception, the findings of the first two experiments provide consistent support for our contention that language may be used in a biased fashion in intergroup settings. Apparently, the encoding of behavioral episodes varies as a function of the protagonist's category membership and the desirability of his or her action. This finding, we believe, may have interesting implications for the interpersonal communication and transmission of stereotypes. Assume that somebody is being observed as he runs into a burning house and returns a few seconds later carrying a small child in his arms. A news reporter may communicate this story at very different levels of abstractions, by simply describing the behavior sequence, by providing an interpretation ("the protagonist saved the child from the flames; he risked his life"), or even by ascribing abstract dispositions to the protagonist, describing him as courageous or as a hero. Our findings suggest that the news reporter may communicate the story at a lower level of abstraction when the protagonist is an out-group member. If this is the case, the news story should provide less information about the protagonist and imply lesser temporal stability than if the same story was communicated at a higher level of abstraction (see Semin & Fiedler, 1988a).

An intriguing aspect of the linguistic category model in our context is the prediction that communications at higher levels of abstraction are more informative about the actor and, in particular, induce the expectation that the actor will display similar behaviors or traits in the future. In the preceding example, the

description of the protagonist as altruistic should induce a greater expectancy that he or she will engage in similar altruistic acts in the future than would a mere description of the behavior sequence. The same reasoning can be applied to undesirable acts as well. For instance, the sentence "A is a liar" implies greater temporal stability and a greater probability that A will lie in the future than an act-specific description such as "A was lying." Thus, communications at higher levels of abstraction (SV or ADJ) are more likely to produce expectancies about future behaviors, which in turn may bias subsequent information processing in a top-down fashion.

In line with the linguistic category model, we predicted a linear trend such that, moving from the lowest to the highest level of abstraction, more information would be revealed about the protagonist. Furthermore, we predicted an increasing expectancy that the act be repeated. To test these hypotheses, we conducted a third experiment in which subjects were asked to rate how much information sentences at different levels of abstraction revealed about the protagonist and how likely subjects thought it was that the protagonist would display the same behavior or trait in the future.

Experiment 3

Method

Twenty subjects not involved in the palio (10 men, 10 women, mean age = 24.2) were asked to read the response alternatives provided in Experiment 1 without having been exposed to the visual representation (cartoons). To prevent motivational problems, the stimulus material was reduced by randomly selecting one item from each pair representing the same episode type (e.g., one desirable, palio-specific, interactional episode). Thus, each subject rated 32 sentences (four response alternatives referring to each of the eight episodes). The selected response alternatives were presented in the same order as in Experiment 1. The palio-specific group membership (name of contrada) was substituted by the more general term *Group X* or *Group Y*. The experiment was described as a psycholinguistic study. The instructions were virtually identical to those used by Semin and Fiedler (1988a, Experiment 1). Subjects were asked to rate each sentence on a 5-point scale with regard to two questions: (a) In your opinion, how much information does the phrase reveal about the protagonist, and (b) in your opinion, how likely is it that the same action or attribute will be repeated in the future? Each subject's responses were collapsed across generic versus palio-specific and interactional versus noninteractional ratings in order to obtain separate scores for how informative (or how stable) each subject considered DAV, IAV, SV, or ADJ sentences describing desirable and undesirable acts.

Results

Information about protagonist. Two subjects were excluded from this and the subsequent analysis owing to missing data. A 2 (desirable vs. undesirable episodes) $\times$ 4 (level of abstraction) repeated measures ANOVA revealed the predicted main effect for level of abstraction, $F(3, 17) = 5.36$, $p = .01$ (see Table 5). A linear trend analysis (Winer, 1971) indicated that, moving from DAV to ADJ, the amount of information about the protagonist increased steadily, $F(1, 51) = 7.51$, $p < .01$. This effect was modified by an interaction with desirability, $F(3, 17) = 4.32$, $p < .05$. There was a strong linear increase in informativeness for desirable episodes, $F(1, 51) = 7.06$, $p < .05$, but not for un-

Table 5
Mean Ratings of Informativeness and Likelihood of Repetition
as a Function of Level of Abstraction and Desirability
of Episode: Experiment 3

Desirability of episode	Level of abstraction			
	DAV	IAV	SV	ADJ
Information about protagonist				
Desirable				
<i>M</i>	2.51	3.04	3.44	3.31
<i>n</i>	18	18	18	18
Undesirable				
<i>M</i>	3.00	3.13	3.28	3.54
<i>n</i>	18	18	18	18
Total	2.76	3.09	3.36	3.43
Likelihood of repetition				
Desirable				
<i>M</i>	3.79	3.86	3.94	3.89
<i>n</i>	18	18	18	18
Undesirable				
<i>M</i>	3.13	3.11	3.64	3.86
<i>n</i>	18	18	18	18
Total	3.46	3.49	3.79	3.88

Note. DAV = Means are based on 5-point scales, with higher scores indicating greater informativeness and greater likelihood of repetition.

desirable ones, $F(1, 51) = 2.82$, *ns*. However, for both desirable and undesirable episodes, the lower levels of abstraction (DAV and IAV combined = 2.78 for desirable episodes and 3.06 for undesirable episodes) were considered less informative about the protagonist than were the higher levels (SV and ADJ combined = 3.38 for desirable episodes and 3.41 for undesirable episodes), $t(17) = 4.56$, $p < .01$, one-tailed, for desirable episodes, and $t(17) = 2.58$, $p < .01$, one-tailed, for undesirable episodes.

Likelihood of repetition. Similar results emerged for the likelihood of repetition ratings. Again, a 2 (desirable vs. undesirable episodes) $\times$ 4 (level of abstraction) repeated measures ANOVA revealed the predicted main effect for level of abstraction, $F(3, 17) = 5.90$, $p = .01$ (see Table 5). On the average, repetition in the future was considered less likely for the lower levels of abstraction (DAV and IAV combined = 3.48) than for the higher levels of abstraction (SV and ADJ combined = 3.84), $t(17) = 2.07$, $p < .05$, one-tailed. An additional interaction with desirability, $F(3, 17) = 7.77$, $p < .01$, indicated that perceived likelihood of repetition increased reliably only for undesirable episodes, $F(1, 51) = 6.66$, $p < .05$, whereas the linear increase was quite weak for desirable episodes, $F(1, 51) = 3.11$, $p < .09$. In fact, actions encoded at lower levels of abstraction (DAV and IAV combined = 3.12) were perceived as less likely to be displayed again in the future than were those encoded at higher levels (SV and ADJ combined = 3.75) only for undesirable episodes, $t(17) = 5.25$, $p < .01$, but not for desirable ones (DAV and IAV combined = 3.83; SV and ADJ combined = 3.92).

Correlational analyses. To further investigate the effect of linguistic abstraction on perceived informativeness and proba-

bility of repetition, first-order and partial correlations were run between these three variables, using single descriptions rather than subjects as unit of analysis. Thus, informativeness and probability of repetition scores were obtained for each of the 32 sentences of the stimulus material by averaging across subjects. Level of abstraction was reliably and positively correlated with both the amount of information provided about the protagonist, $r(32) = .58$, $p < .01$, and the probability of repetition, $r(32) = .35$, $p < .05$. Furthermore, neither correlation disappeared when the third variable was partialled out (correlation between abstraction and informativeness was $.61$, $p < .01$, when controlling for likelihood of repetition; correlation between abstraction and likelihood of repetition was $.42$, $p < .05$, when controlling for informativeness). This suggests that level of abstraction has entirely independent effects on perceived informativeness and likelihood of repetition. In fact, there is no correlation between these two variables, $r(32) = .02$, *ns*.

Another interesting finding emerges when the two variables are correlated with linguistic abstraction separately for phrases describing desirable and undesirable episodes. Whereas the amount of information about the protagonist increases with increasing level of abstraction for both desirable, $r(16) = .63$, $p < .01$, and undesirable episodes, $r(16) = .53$, $p < .05$, the perceived likelihood of repetition increases with abstraction only for undesirable episodes, $r(16) = .67$, $p < .01$, but not for desirable episodes, $r = .10$, *ns*. This is not completely surprising, considering that our subjects indicated a very high overall probability (3.9 on a 5-point scale) that desirable acts would be repeated in the future, suggesting that the low correlation may at least partially be a function of a restriction of range.

Discussion

The results of the third experiment generally confirm our contention that with increasing level of abstraction the amount of information about the actor and the expectancies of repetition increase in a linear fashion. However, the first trend was more pronounced for desirable episodes and the latter for undesirable episodes. The latter finding is particularly interesting as it suggests that, at least for undesirable actions, abstract information may serve as a schema that induces expectancies about future behaviors, which in turn may guide subsequent information processing. It is also interesting to note that the amount of information a phrase provides about the actor and the probability of the act's being repeated in the future appear to be two completely independent consequences of linguistic abstraction.

General Discussion

Taken together, the present experiments provide the first evidence for biased language use in intergroup contexts. In particular, they provide evidence that the same socially desirable behavior is encoded at a higher level of abstraction when performed by an in-group member than when performed by an out-group member. Apparently, desirable in-group behaviors induce generalizations to the actor's character or psychological state, or, in Heider's (1944) terms, a unit formation of actor and act. In contrast, desirable out-group behaviors are more likely

to be encoded as concrete behavioral instances without abstraction beyond the given information.

Results are somewhat less consistent with respect to socially undesirable episodes. Here only the first experiment found the predicted reversal, such that out-group behaviors were encoded at a higher level of abstraction than in-group behaviors. One may argue that undesirable behaviors are generally more diagnostic and informative about the protagonist (Ajzen & Fishbein, 1975; Jones & Davis, 1965) and as such may facilitate abstract encoding independent of group membership. Although the exact reasons need to be addressed in future research, the present experiments suggest that intergroup biases in language use may be more pronounced and robust for desirable than for undesirable episodes.

A different way to look at the data is to compare encoding of desirable versus undesirable episodes within in-group and out-group. Here our findings suggest that people encode and communicate behavioral information involving out-group protagonists at a higher level of abstraction when it is undesirable than when it is desirable (Experiments 1 and 2). Interestingly, in-group behaviors were encoded in a largely unbiased fashion. In terms of linguistic abstractness, desirable and undesirable in-group behaviors were treated in an undifferentiated manner, whereas out-group behaviors were strongly polarized. The same pattern also emerged from the correlational analyses (Experiment 1), which found a strong correlation between social desirability and abstractness for out-group episodes: The more negative the action, the higher the level of abstraction at which it was encoded. Although this correlation reversed for episodes involving in-group members, the correlation became considerably weaker. This, again, suggests that the language bias is more pronounced for out-group protagonists.

Implications of the Linguistic Intergroup Bias

Differential language use—as demonstrated in the first two experiments—apparently has a number of important implications that may contribute to the persistence of stereotypes. There is evidence that once a negative out-group or positive in-group behavior has been communicated in abstract linguistic terms, it influences subsequent information processing of both source and receiver of the communication in various ways. First, linguistically abstract communications are perceived as providing more information about the actor than do concrete ones (Experiment 3; see also Semin & Fiedler, 1988a).

Second, abstract descriptions are perceived as relatively stable over time (Semin & Fiedler, 1988a) and consequently produce the expectation that the (undesirable) action be repeated in the future (Experiment 3).

Third, abstract encoding may have an interesting—yet to be tested—implication: Abstract descriptions may induce a top-down process in which subsequent information processing is guided (and biased) by the initial description. Recently, some evidence for such a top-down process in social judgment tasks has been reported by Fiedler and Semin (1988a). They found that, after having provided an initial abstract description of a person, people tended to produce additional concrete information congruent with their initial description whenever the validity of the initial trait description was challenged. Interestingly,

concrete initial descriptions apparently had much less impact on subsequent processing. The idea that abstract encoding will induce a top-down process is also congruent with Arcuri's (1983) finding that information about people coded at the level of adjectives (personality traits) had greater diagnostic power than information coded at a more specific level (verbs expressing behaviors). Thus, abstract descriptions tend to trigger schemata-driven processes in impression formation.

Finally, there is evidence that abstract descriptions are considered less verifiable than concrete ones (Semin & Fiedler, 1988a). It is considerably easier to confirm or disconfirm the occurrence of a concrete behavior than the existence of a trait or psychological state. For instance, a single observation should be sufficient to disconfirm a false concrete statement such as "A hits B," whereas many behavioral instances should be required before an abstract statement such as "A is aggressive" can be disconfirmed. The lack of verifiability should then make abstract statements more resistant to disconfirmation and change. This idea is quite similar to Rothbart's (see Rothbart & John, 1985; Rothbart & Park, 1986) contention that traits differ in their susceptibility to evidence and that many stereotypic beliefs pertain to exactly those traits that are not easily verifiable. Whereas Rothbart and Park's (1986) analysis focuses on inter-trait differences, the present model applies the same basic principle to differences in confirmability between linguistic categories.

Considering this wide range of implications, one can easily envisage a self-perpetuating cycle in which biased language use maintains or even aggravates initial intergroup biases in applied settings. For instance, subtle language biases may occur in the legal system, where witnesses may reveal quite different information by describing the observed sequence of events in more or less abstract terms. As the likelihood of repetition of a criminal act partially determines the severity of the sentence, abstract communications can easily contribute to more severe sentences. In a similar vein, teachers or parents may inadvertently bias cross-generational communication in much the same way as reporters may choose different levels of abstraction depending on whether the protagonist of a given news story shares their own category membership. The first indirect evidence for this contention comes from a recent international research program about the image of Africa in the mass media (involving six European and eight African countries and supported by the Food and Agricultural Organization; see Pugliese, 1988). Besides the highly biased news story selection and the generally negative tone of the coverage, the report complained about a "linguistic stereotype . . . in which phrases referring to Blacks are characterized by the predominance of nouns and the almost complete absence of verbs, at least in the active form, whereas the exact opposite occurs in reference to Westerners, denoted by action verbs" (Pugliese, 1988, p. 57).

Beyond its practical contributions, the present approach may also have some interesting methodological implications. First, the linguistic intergroup bias may represent a subtle source of error in those research paradigms in which an ongoing interaction is encoded. In this case, the judge's description of an action may depend, in predictable fashion, on the social category membership of the interactants. Similar biases may occur for certain content analytical techniques that require the recoding

of information. Second, the present approach may provide a useful methodological tool for the empirical study of racism and intergroup discrimination. As traditional reactive measures of racism have become largely ineffective in tapping more subtle racial and ethnic prejudice (e.g., Crosby, Bromley, & Saxe, 1980; Dovidio & Gaertner, 1986a, 1986b; McConahay, 1986), the study of language use may provide a less obtrusive, alternative method, particularly for intergroup contexts in which the overt expression of prejudice is normatively unacceptable, as in the case of Jews or Blacks.

Open Problems

Considering that this series of studies represents a first attempt to test the linguistic intergroup bias, it is not surprising that various issues remain open to future investigation. First, at this point it is not entirely clear which mechanism triggers the linguistic intergroup bias. In the introduction, we proposed a model in which people hold differential expectancies regarding in-group and out-group behaviors, which in turn determine the language representation of behavioral episodes. Expectancy-congruent observations are translated into abstract language that generalizes from the single act to more enduring properties of the actor. In contrast, expectancy-incongruent observations are described in concrete terms as single, situationally and temporally bound instances.

Alternatively, one may argue from an attributional perspective that our results simply reflect the linguistic expression of latent causal attributions. If one assumes that concrete codings reflect implicit situational attributions and abstract codings reflect implicit dispositional attributions, our results could easily be interpreted as the linguistic manifestation of the well-known group-serving bias (e.g., Hewstone & Jaspars, 1984). We are inclined to reject this hypothesis mainly because the taxonomy proposed by the linguistic category model is not equivalent to the continuum of situational-to-personal causation. Moving from DAVs to ADJs, there is no linear increase in implicit personal causation, and low levels of abstraction do not necessarily imply situational causation. DAVs generally do not imply any causation at all, but simply provide a noncausal description of behavioral information. In fact, that DAVs lack interpretation is one of their defining features.

Differences between the linguistic category model and the attributional model become even clearer when considering the two intermediate levels of abstraction. Moving from IAVs to SVs, we observe an increase in abstraction but a decline in personal causation. Since Brown and Fish's (1983) and Garvey and Caramazza's (1974) pioneering work on implicit causality, various studies involving such diverse languages as Chinese, Italian, Afrikaans, and English have investigated the implicit causality of verbs embedded in subject-verb-object sentences. These studies have generally found that the causality is regularly attributed to the subject when the verb is an IAV, but to the object when the verb is an SV (e.g., Au, 1986; Brown & Fish, 1983; Franco, Arcuri, & Cadinu, 1988; Garvey & Caramazza, 1974; Van Kleeck, Hillger, & Brown, 1988; Voster, 1985). Thus, if our subjects mainly had implicit causality on their minds when describing desirable episodes, they should have used IAVs more frequently and SVs less frequently when describing in-group

members than when describing out-group members. The results of Experiment 1 clearly contradict this idea (for additional evidence see Arcuri, Maass, & Semin, 1989). This suggests that the linguistic intergroup bias is not simply a function of an implicit group-serving attributional bias and that people may choose abstract versus concrete language representations for reasons other than their attributional implications (e.g., degree of situational or temporal generalization).

Yet another, primarily motivational explanation of the linguistic intergroup bias may be derived from social identity theory. From this perspective, our findings may be interpreted as a subtle strategy of establishing an intergroup difference in favor of one's own group in order to maintain or enhance a distinct and positive social identity. For instance, describing favorable out-group behaviors as single, concrete instances unrelated to the protagonist's enduring properties may represent just one possible strategy of out-group derogation. Contrary to our own account, this explanation does not rely on the mediating role of differential expectancies.

At this point, it remains unclear whether the linguistic intergroup bias observed in the present experiments reflects an attempt to protect one's social identity, whether it is primarily the consequence of differential expectancies, or whether both processes may have contributed. An exact understanding of the underlying mechanism is not just a matter of academic curiosity. Although the two processes may often coincide, they lead in some cases to exactly opposite predictions. If desire for a positive social identity is at the basis of the linguistic intergroup bias, then such bias should emerge in any intergroup setting in which the in-group-out-group categorization becomes salient. It should be particularly pronounced in situations in which the individual's social identity is at stake, such as after self-esteem threat (Lemyre & Smith, 1985) or when groups are in direct competition (as was the case in our experiments).

Yet if encoding differences are mainly a function of differential expectancies, then predictions vary according to type and specificity of such expectancies. In certain intergroup settings such as the one investigated here, people have negative general views of the out-group, and hence expect out-group members to display more undesirable behaviors and fewer desirable behaviors than in-group members without expecting any specific behavior patterns (see also Howard & Rothbart, 1980). For such cases, predictions derived from a differential expectancy viewpoint will largely coincide with those derived from a social identity perspective. Often, however, expectancies reflect well-defined stereotypes referring to very specific behavior patterns (e.g., Jews are expected to be stingy but not to be unfriendly). For such cases, predictions derived from a differential expectancy perspective deviate from the more general intergroup bias perspective, as, according to the former, biases in language use should occur only for those behaviors that are directly relevant to the stereotype. Furthermore, such stereotypes may occasionally contain positive elements that are in contrast to the overall affective tone of the stereotype, as in the case of the intellectual achievements of Jews. Following a differential expectancy viewpoint, behaviors of Jews that are indicative of intelligence should be encoded at a high level of abstraction because they are expectancy congruent. Yet, non-Jews should prefer concrete language representations if they are mainly concerned

with maintaining a positive social identity. There are even cases in which members of mutually exclusive social categories share the same stereotypes. A striking example of these are sex stereotypes that are largely agreed on by male and female subjects (e.g., Deaux, 1976). According to a differential expectancy perspective, both male and female subjects should encode sex role-incongruent behaviors ("feminine" behaviors of male protagonists and "masculine" behaviors of female protagonists) in concrete terms independent of their valence. This is quite in contrast to a social identity perspective, according to which both sexes should encode undesirable in-group and desirable out-group behaviors in concrete terms independent of their stereotypicality.⁴

We believe that, for a better understanding of the exact underlying mechanism, it would be very useful to study the linguistic intergroup bias in situations in which specific expectancies are either experimentally induced or already contained in a well-defined stereotype.

A second problem to be resolved by future research regards the potential impact of a real or imagined audience or communication partner. Similar to previous studies (Semin & Fiedler, 1988b), subjects in our experiments were instructed to provide written descriptions of events without knowing who would receive their message and what that person's social category membership might be. In real settings, people are generally aware of the group membership of their audience and are therefore able to adjust their message accordingly. It is conceivable that the choice of behavioral descriptions will vary not only according to the group membership of sender and protagonist, but also according to the in-group or out-group status of who is expected to receive the communication. A similar argument can be made regarding the interpretation of communications of varying levels of abstraction (see Experiment 3). Here, the implicit subject informativeness and probability of repetition may vary not only according to the level of abstraction, but also according to the sender's category membership. Thus, the communication process may actually be more complex than suggested by our findings.

Although a number of issues remain unresolved at this point, we hope that the present approach will contribute to the understanding of how language—an often neglected aspect of social cognition—contributes to both the intraindividual maintenance and the interindividual communication and transmission of stereotypes. It is this latter aspect that, we believe, distinguishes the present approach from many previous ones, as it shifts the level of analysis from the individual to the medium by which group relations are maintained and negotiated.

⁴ It is even conceivable that the two processes interact with situational variables such that language use reflects the desire to protect one's social identity whenever groups are in direct competition; however, cognitive expectancy processes may prevail whenever a person judges the behavior of in-group and out-group members from the perspective of an uninvolved observer.

References

Ajzen, I., & Fishbein, M. A. (1975). Bayesian analysis of attribution processes. *Psychological Bulletin*, 82, 261–277.

- Allport, G. W. (1954). *The nature of prejudice*. Reading, MA: Addison-Wesley.
- Arcuri, L. (1983). Schemi di persone e potere diagnostico delle informazioni [Person schemata and diagnostic power of information]. *Giornale Italiano di Psicologia*, 10, 483–500.
- Arcuri, L., Maass, A., & Semin, G. (1989). *Linguistic intergroup bias vs. self-serving attributional bias: The role of implicit causality of verbs*. Unpublished manuscript.
- Au, T. K-F. (1986). A verb is worth a thousand words: The causes and consequences of interpersonal events implicit in language. *Journal of Memory and Language*, 25, 104–122.
- Brewer, M. B. (1979). In-group bias in the minimal intergroup situation: a cognitive-motivational analysis. *Psychological Bulletin*, 86, 307–324.
- Brewer, M. B., & Kramer, R. M. (1985). The psychology of intergroup attitudes and behavior. *Annual Review of Psychology*, 36, 219–243.
- Brown, R., & Fish, D. (1983). The psychological causality implicit in language. *Cognition*, 14, 237–273.
- Crosby, F., Bromley, S., & Saxe, L. (1980). Recent unobtrusive studies of black and white discrimination and prejudice: A literature review. *Psychological Bulletin*, 87, 546–563.
- Deaux, K. (1976). *The behavior of women and men*. Monterey, CA: Brooks/Cole.
- Deschamps, J.-C. (1972–73). Imputation de la responsabilité de l'échec (ou de la réussite) et catégorisation sociale [Attribution of responsibility of failure (or success) and social categorization]. *Bulletin de Psychologie*, 26, 794–806.
- Dovidio, J. F., & Gaertner, S. L. (1986a). *Prejudice, discrimination and racism*. Orlando, FL: Academic Press.
- Dovidio, J. F., & Gaertner, S. L. (1986b). Prejudice, discrimination, and racism: Historical trends and contemporary approaches. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination and racism* (pp. 1–34). Orlando, FL: Academic Press.
- Duncan, B. L. (1976). Differential social perception and attribution of intergroup violence: Testing the lower limits of stereotyping of blacks. *Journal of Personality and Social Psychology*, 34, 590–598.
- Fiedler, K., & Semin, G. R. (1988a). *Language use and reification of social information: Top-down and bottom-up processing in person cognition*. Unpublished manuscript.
- Fiedler, K., & Semin, G. R. (1988b). On the causal information conveyed by different interpersonal verbs: The role of implicit sentence context. *Social Cognition*, 6, 21–39.
- Franco, F., Arcuri, L., & Cadinu, M. (1988). *L'attribuzione di causalità nei verbi interpersonali: Una verifica sulla lingua italiana* [Causal attributions implicit in interpersonal verbs: A test in Italian language]. Unpublished manuscript.
- Garvey, C., & Caramazza, A. (1974). Implicit causality in verbs. *Linguistic Inquiry*, 5, 459–464.
- Hamilton, D. L. (1981). *Cognitive processes in stereotyping and intergroup behavior*. Hillsdale, NJ: Erlbaum.
- Hamilton, D. L., & Trolier, T. K. (1986). Stereotypes and stereotyping: An overview of the cognitive approach. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination and racism* (pp. 127–163). Orlando, FL: Academic Press.
- 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.
- Heider, F. (1944). Social perception and phenomenal causality. *Psychological Review*, 51, 358–384.
- Hewstone, M., & Jaspars, J. M. F. (1984). Social dimensions of attribution. In H. Tajfel (Ed.), *The social dimension: European developments in social psychology* (Vol. 2, pp. 379–404). Cambridge, England: Cambridge University Press.
- Howard, J., & Rothbart, M. (1980). Social categorization and memory

- for in-group and out-group behavior. *Journal of Personality and Social Psychology*, 38, 301-310.
- 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.
- Lemyre, L., & Smith, P. M. (1985). Intergroup discrimination and self-esteem in the minimal group paradigm. *Journal of Personality and Social Psychology*, 49, 660-670.
- Linville, P. W., Salovey, P., & Fischer, G. W. (1986). Stereotyping and perceived distributions of social characteristics: An application to in-group-outgroup perception. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination and racism* (pp. 1-34). Orlando, FL: Academic Press.
- McConahay, J. B. (1986). Modern racism, ambivalence, and the modern racism scale. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination and racism* (pp. 91-125). Orlando, FL: Academic Press.
- Pettigrew, T. F. (1979). The ultimate attribution error: Extending Allport's cognitive analysis of prejudice. *Personality and Social Psychology Bulletin*, 5, 461-476.
- Pugliese, C. (1988). Quale immagine? [Which image?]. *Nigrizia*, 106, 56-57.
- Rothbart, M., Evans, M., & Fulero, S. (1979). Recall for confirming events: Memory processes and the maintenance of social stereotypes. *Journal of Experimental Social Psychology*, 15, 343-355.
- Rothbart, M., & John, O. P. (1985). Social categorization and behavioral episodes: A cognitive analysis of the effects of intergroup contact. *Journal of Social Issues*, 41, 81-104.
- Rothbart, M., & Lewis, S. (1988). Inferring category attributes from exemplar attributes: Geometric shapes and social categories. *Journal of Personality and Social Psychology*, 55, 861-871.
- Rothbart, M., & Park, B. (1986). On the confirmability and disconfirmability of trait concepts. *Journal of Personality and Social Psychology*, 50, 131-142.
- Sagar, H. A., & Schofield, J. W. (1980). Racial and behavioral cues in black and white children's perceptions of ambiguously aggressive acts. *Journal of Personality and Social Psychology*, 39, 590-598.
- Semin, G. R., & Fiedler, K. (1988a). The cognitive functions of linguistic categories in describing persons: Social cognition and language. *Journal of Personality and Social Psychology*, 54, 558-568.
- Semin, G. R., & Fiedler, K. (1988b). *Relocating attributional phenomenon within a language-cognition interface: The case of actors' and observers' perspectives*. Unpublished manuscript.
- Snyder, M., & Swann, W. B. (1978). Hypothesis-testing processes in social interaction. *Journal of Personality and Social Psychology*, 36, 1202-1212.
- Srull, T. D., Lichtenstein, M., & Rothbart, M. (1985). Associative storage and retrieval processes in person memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 11, 316-345.
- Tajfel, H. (1982). Social psychology of intergroup relations. *Annual Review of Psychology*, 33, 1-39.
- Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In W. S. Austin & S. Worchel (Eds.), *The social psychology of intergroup relations* (pp. 33-47). Monterey, CA: Brooks/Cole.
- Taylor, D. M., & Jaggi, V. (1974). Ethnocentrism and causal attribution in a South Indian context. *Journal of Cross-Cultural Psychology*, 5, 162-171.
- Turner, J. C. (1987). *Rediscovering the social group: A self-categorization theory*. Oxford: Blackwell.
- Van Dijk, T. A. (1984). *Prejudice and discourse. An analysis of ethnic prejudice in cognition and conversation*. Amsterdam: Benjamins.
- Van Dijk, T. A. (1987). *Communicating racism: Ethnic prejudice in thought and talk*. Newbury Park, CA: Sage.
- Van Kleeck, M., Hillger, L., & Brown, R. (1988). Pitting verbal schemas against information variables in attribution. *Social Cognition*, 6, 89-106.
- Voster, J. (1985). Implicit causality in language: Evidence from Afrikaans. *South African Journal of Psychology*, 15, 62-67.
- Weber, R., & Crocker, J. (1983). Cognitive processes in the revision of stereotypic beliefs. *Journal of Personality and Social Psychology*, 45, 961-977.
- Wilder, D. A. (1986). Social categorization: Implications for creation and reduction of intergroup bias. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 19, pp. 291-355). New York: Academic Press.
- Winer, B. J. (1971). *Statistical principles in experimental design*. New York: McGraw-Hill.

Received September 19, 1988

Revision received May 26, 1989

Accepted June 27, 1989 ■