Journal of Personality and Social Psychology
1998, Vol. 74, No. 4, 849–864

0022-3514/98/$3.00

The Slime Effect: Suspicion and Dislike of Likeable Behavior Toward Superiors

Roos Vonk
Leiden University

An actor who is likeable toward superiors and dislikeable toward subordinates is judged as extremely dislikeable and slimy (Experiments 1 and 2). Subsequent experiments addressed several theoretical accounts of this *slime effect*. Likeable behaviors toward superiors induce suspicion of ulterior motivation, which is confirmed when dislikeable behaviors toward subordinates are observed (Experiment 3). The operation of a *slime schema* was indicated by the emergence of an illusory correlation between an actor's behavior and the status of the target, such that the actor was erroneously perceived as more likeable toward superiors (Experiment 4). Further, perceivers spontaneously discerned the behavioral pattern of "licking upward–kicking downward," regardless of processing time (Experiment 5). Implications for impression formation and inconsistency resolution, trait inferences and correspondence bias, and lay theories of self-presentational behavior are discussed.

Knowledge of impression management strategies is not confined to social psychologists: Person perceivers in general seem quite aware that people often engage in self-presentational behavior such as ingratiation (cf. Fein, Hilton, & Miller, 1990) and self-promotion (e.g., Godfrey, Jones, & Lord, 1986). As a consequence, many behaviors can be interpreted in at least two ways: by making a correspondent inference, or by inferring some form of self-presentation (i.e., a noncorrespondent inference). For instance, a person who claims to be intelligent may either be intelligent or may be engaging in self-promotion (cf. Reeder & Fulks, 1980), a person who helps somebody and expresses agreement with someone's opinions may either be helpful and hold the same opinions or may be engaging in ingratiation.

In interpreting such behaviors, perceivers appear to be sensitive to the behavioral setting. For instance, self-promoting behavior is more easily identified when perceivers know that the abilities claimed by a person cannot be verified by the audience (Vonk, 1997). Ingratiation is more likely to be recognized when the target of the behavior has power over the actor (Jones, 1964, 1990). Although people can have a variety of reasons to ingratiate themselves with someone, the most prominent one seems to be that the other person has the power to affect their outcomes. Thus, when an actor performs likeable behavior toward a more powerful person, perceivers may correct their inference for the possibility of brownnosing.

Although self-presentational behavior goes on every day and everywhere (e.g., Leary, 1995), social cognition research has paid little attention to the question of how perceivers make inferences from such behavior, and how they integrate an actor's self-presentations with their assumptions about the true inner dispositions and intentions of the actor. The general view in the person cognition field is that behavior is often taken at face value (for a review, see Gilbert & Malone, 1995). This assumption stands in contrast with the common sense intuition that people immediately recognize blatant forms of self-presentation (e.g., a job applicant during an interview who goes on and on about his or her admiration for the interviewer). The present studies examined inferences and impressions of what is perhaps the most ubiquitous form of self-presentation, ingratiation (Jones, 1964). Specifically, it was examined whether perceivers identify ingratiation, and if so, how it affects their evaluations of the actor.

In the experiments reported here, participants read a series of behavior descriptions of an actor. The manipulations involved the covariation between the valence of the behavior (likeable, dislikeable) and the power of the persons toward whom the behavior was performed (e.g., superiors, subordinates). The general hypothesis was that inferences of behavior are corrected for the target persons toward whom it is performed. Specifically, it was expected that the power of these persons is taken into account, such that likeable behavior enacted toward more powerful targets is perceived as a reflection of ingratiation and, thus, is discounted in the impression.

The first goal of Experiment 1 was to provide empirical evidence for an observation from everyday life, namely, that a person whose behavior is described in Dutch as "licking upward–kicking downward" (cf. the authoritarian personality, Adorno, Frenkel-Brunswik, Levinson, & Sanford, 1950) is judged as (a) extremely dislikeable and (b) highly slimy.[1] In

We thank all those who have inspired this research. We also thank Ellen de Bruin and Daan van Knippenberg for useful advice and Walter Uijlenbroek for information on the accountancy business.

Correspondence concerning this article should be addressed to Roos Vonk. Her address is Department of Social Psychology, Leiden University, Wassenaarseweg 52, 2333 AK Leiden, the Netherlands. Electronic mail may be sent to vonk@rulfsw.leidenuniv.nl.

[1] In the Netherlands, where these studies were conducted, there is no general word for ingratiation. The closest resemblance is derived from the noun *slime* (*slijm*), which literally refers to the same slippery substance as the English word, but is more often used to describe persons and behaviors. The verb *to slime* (*slijmen*) refers to the behavior of ingratiating oneself for ulterior motives, regardless of whether these

the remainder of this article, this pattern of judgments is referred to as the Slime effect.

The second goal of Experiment 1 was to acquire more insight into the process of inconsistency resolution. The person memory literature (e.g., Hastie, 1980; Srull, 1981) has provided an abundance of detailed analyses of participants' relatively extensive processing of inconsistent information. Generally, when perceivers observe likeable and dislikeable behaviors by the same person, the usual smoothness of the trait-inference process is disrupted: Perceivers need to think of each behavior in relation to the others, in order to understand the discrepancy. As a result of this elaborative thought, free recall of the behaviors is better than when only consistent behaviors are presented (e.g., Srull, 1981). Although psychologists have a great deal of knowledge about the circumstances that promote or inhibit this elaboration (i.e., the why and when), they know very little about *what* precisely participants are thinking when they are engaging in elaborative thought, that is, when they are trying to resolve an inconsistency. Only a few studies have implicitly addressed this question (e.g., Asch & Zukier, 1984; Vonk, 1993), and their results suggest that participants often attempt to explain inconsistencies by examining whether the actor is likeable in some domain of situations and dislikeable in another (e.g., toward intimates vs. strangers, at home vs. at work; cf. Asch & Zukier's, 1984, notion of situational segregation, and Trafimow's, 1994, evidence of multiple representations of a person in different settings). Inconsistencies within a single individual are more problematic than inconsistencies within a group (e.g., Vonk & Van Knippenberg, 1995; Wyer & Gordon, 1982), because a person is perceived as a psychological unit (Asch, 1946). Thus, it makes sense that participants trying to explain a discrepancy within a person's behaviors attempt to break this unit, by allocating the behaviors to different settings (Vonk, 1994).

Thus, characteristics of the situation may operate as cues in resolving inconsistencies in an actor's behaviors. The hierarchical status of a target person is such a characteristic, because a covariation between behavior valence and target status can explain behavioral incongruencies. This is especially true when the covariation is such that the actor is more likeable toward powerful targets, because this pattern is familiar and intuitively sensible. It follows that when participants read about a person who systematically behaves differently toward superiors and subordinates, they should not have a hard time explaining the inconsistency: They should quickly become aware that it can be accounted for by characteristics of the targets. As soon as they realize this, the inconsistency is resolved, and there is no need to engage in further elaborative thought. As a result, the superior recall effect will be reduced.

In summary, Experiment 1 tested the hypothesis that a person who is likeable toward superiors and dislikeable toward subordinates is (a) judged as highly dislikeable and slimy (the slime effect) and (b) induces less elaborative thought than is usually the case when participants are faced with behavioral inconsistencies, resulting in a reduced superior recall effect.

Experiment 1: "Licking Upward–Kicking Downward"

Participants received information about an actor's behaviors toward both superiors and subordinates. In the central experimental condition, behaviors toward superiors were likeable and behaviors toward subordinates were dislikeable. This condition is referred to as the *slime* condition, that is, the condition in which the Slime effect was expected to emerge, as evidenced by (a) a low likeability rating and (b) a high rating on the trait *slimy*. To establish this effect, the slime condition was compared with the following four others.

In one condition, the actor performed dislikeable behaviors toward superiors and likeable behaviors toward subordinates. This condition is referred to as the *nonslime* condition, because dislikeable behavior toward superiors seems unlikely to result from any ulterior motivation; further, it may be assumed that likeable behavior toward subordinates is perceived as authentic because, even though subordinates usually do have means to affect their superiors' outcomes, their power is less salient. Thus, it was hypothesized that the nonslime actor would be judged as more likeable and as less slimy than the slime actor.

In the mixed condition, participants read the same behaviors, but the behavioral targets were alternated, such that the actor was not systematically more likeable toward superiors or subordinates. It was hypothesized that in this condition, too, the actor would be judged as more likeable and as less slimy than in the slime condition, whereas judgments would not differ from those in the nonslime condition. This result would demonstrate that the hypothesized difference between the slime and the nonslime conditions results from relatively negative judgments in the slime condition, and not from relatively positive judgments in the nonslime condition (e.g., as a result of the nonslime actor being perceived as a "Robin Hood" who stands up for the underdog). An additional demonstration of this can be provided by a comparison of the slime and the nonslime conditions with the following two conditions.

In the negative condition, only dislikeable behaviors were presented; in the positive condition, participants read only likeable behaviors. In both conditions, the targets of the behaviors were superiors as well as subordinates. It was predicted that in the nonslime condition, the actor would be evaluated less favorably than in the positive condition; in the slime condition, likeability judgments were expected to be similar to those in the negative condition, whereas slime judgments were expected to be significantly higher.

In addition to establishing the slime effect, Experiment 1 also tested the assumption that the perceived inconsistency between likeable and dislikeable behaviors is reduced when these behaviors are systematically associated with different targets. In the inconsistent conditions (slime, nonslime, and mixed), the fol-

---

motives are specific and immediate (e.g., *I want this person to consent to my plans*) or unspecified and remote (e.g., *In the long run it may be convenient if this person likes me*). A person who engages in this type of behavior is described by the adjective *slimy* (*slijmerig*) or the noun *slijmbal* (*slimeball*). The word *slime* and its conjugations have a negative connotation and are used frequently to describe flattery, overly friendly behavior, and brownnosing. So, for all practical purposes, these words refer to the same class of behaviors as the term ingratiation, but they are much more informal. In the remainder of this article, the words *slime* and *slimy* are retained because they are closest to the issue under consideration and they are common in the language of the participants in these studies.

lowing process was assumed to occur. First, participants notice the inconsistency between likeable and dislikeable behaviors. Second, in an effort to explain it, they direct their attention to other characteristics of the behavioral field, including characteristics of the targets toward whom the behaviors are performed. In the mixed condition, this does not reduce the inconsistency because the status of the targets is unrelated to behavior valence. Therefore, participants in this condition remain in a state of elaborate processing (perhaps pursuing other strategies of inconsistency resolution; see Vonk, 1994), eventually resulting in superior recall of the items. In the slime and nonslime conditions, on the other hand, participants become aware of the covariation between behavior and target status. At this point, they complete the stage of elaborative thought and enter a third stage, in which they only need to test their hypothesis that the targets' status accounts for the behavior variation. As a result, their free recall rates will be lower than in the mixed condition.

It was also expected that the recall difference between the mixed and slime conditions would be larger than between the mixed and nonslime conditions, because in the slime condition it is easier to quickly identify the behavioral pattern of the actor: Participants may have a highly accessible and well-developed schema about a typical brownnoser, and as soon as this schema is triggered, they can enter the third stage described above. In the nonslime condition, on the other hand, more cognitive effort is required to identify the behavioral pattern: Participants may not have schemas about a person who is more likeable toward powerless than toward powerful people, and if they do, these schemas may be less accessible because they are activated less frequently in everyday life (e.g., the "Robin Hood" type, or a person with an authority problem).

In summary, it was predicted that free recall would be highest in the mixed condition, followed by the nonslime and then the slime condition. In the positive and negative conditions, recall rates should be lowest because these participants do not receive inconsistent information at all.

Method

Participants and procedure. Participants were 56 students (32 women, 24 men) of different levels in the social sciences who were paid for their participation. Upon arrival, they were seated in individual cubicles with a computer that paced them through the entire experiment. By pressing the *Return* key, participants determined the pace of all text on the screen. They were informed that the study was about how people form impressions of other people, and that they were going to read a series of behavior descriptions of a staff member in an organization, named Paul. It was explained that Paul worked at the middle management level and that he, thus, had subordinates as well as superiors in the organization. Subsequently, the behavior descriptions were presented one at a time.

Participants' impressions of Paul were assessed by means of a series of rating scales. The first scale concerned Paul's likeability (1 = *very dislikeable*, 7 = *very likeable*). Further, 12 trait adjectives were presented, including the trait *slimy* (1 = *not at all*, 7 = *highly*). A manipulation check requested participants to indicate whether Paul behaved differently toward different people (1 = *no*, 2 = *yes, toward men vs. women;* 3 = *yes, toward superiors vs. subordinates;* 4 = *yes, toward strangers vs. familiar people*). Finally, participants were given 8 min to recall as many of the behaviors as they could and write them down in their own words. They were encouraged to describe the gist of each behavior and to stay as close as possible to the original description.

Design and stimulus materials. In all conditions, participants read 20 behavior descriptions that had been pretested to be either likeable (M across the likeable items = 5.45 on a 7-point dislikeable–likeable scale; e.g., *Fetching coffee for someone who is working hard*), dislikeable (M = 2.58; e.g., *Refusing to help someone with a problem*), or neutral (M = 4.16; e.g., *Making an appointment with the accountant*). In the slime, nonslime, and mixed conditions, the descriptions consisted of 10 likeable and 10 dislikeable behaviors. In the positive and negative conditions, the same set of likeable or dislikeable descriptions was presented and 10 neutral descriptions were added.

In the pretest, the target of the likeable and dislikeable behaviors was described as *someone* or *a colleague.* In the experiment, these labels were replaced by references to superiors (e.g., *his boss, the managing director, the chairman*) or subordinates (e.g., *his assistant, the cleaner, the secretary*). In the slime condition, the targets of the likeable behaviors were 10 different superiors and the targets of the dislikeable behaviors were 10 different subordinates. In the nonslime condition, this was reversed. In the mixed condition, the superior and subordinate labels were counterbalanced across the likeable and dislikeable behaviors, creating four sets of behaviors; these sets were balanced for evaluative extremity and length. In the negative and positive conditions, the descriptions were the same as the 10 dislikeable (negative condition) and 10 likeable (positive condition) behaviors in the mixed condition. The 10 neutral items in these conditions contained no references to subordinates or superiors.

The 20 descriptions were presented in two different orders. Because the order variation did not produce any systematic effects, this variable was dropped from the analyses below.

Results and Discussion

Manipulation check. All participants in the slime and the nonslime conditions responded correctly that the actor behaved differently toward superiors and subordinates. In the mixed condition, responses were distributed across the different response categories. In the positive and negative conditions, 8 out of 22 participants responded that the actor's behavior varied across superiors versus subordinates or strangers versus familiar people. These responses may be due to the neutral behaviors, which are in fact different from the likeable and dislikeable ones. For instance, after the experiment, 1 participant commented that the actor seemed indifferent and hostile, but did inquire about the account of a client.

Likeability and slime judgments. The first column in Table 1 presents the mean likeability judgment in each condition. The results of Duncan's multiple comparison are entirely in agreement with the hypotheses: The most negative judgment was obtained in the slime condition, where the perceived likeability of the actor was as low as in the negative condition, and significantly lower than in the nonslime and the mixed conditions. These conditions, in turn, differed from the positive condition and not from each other.

Turning to the judgments of how slimy the actor was (Table 1, column 2), it can be seen that in the slime condition, this rating was significantly higher than in any of the other conditions.[2] A relatively high slime rating was also obtained in the

[2] The correlations between likeability and slime judgments differed across conditions: $r = -.24$ in the positive condition, $r = .26$ in the negative condition, $r = .03$ in the mixed condition, $r = -.10$ in the nonslime condition, and $r = -.16$ in the slime condition (all *ns*). Analy-

Table 1
Means and F Tests for Likeability Judgments (1 = Very Dislikeable, 7 = Very Likeable), Ratings on the Trait Slimy (1 = Not at all, 7 = Highly), and Recall Rates in Experiment 1

Condition	Likeable	Slimy	Total no. of items recalled
Negative	2.33_a	1.91_a	5.25_a
Slime	2.08_a	6.50_d	7.40_b
Mixed	3.73_b	4.10_b	10.55_c
Nonslime	4.45_b	2.36_a	$8.91_{b,c}$
Positive	6.00_c	5.80_c	5.30_a
df	4, 51	4, 51	4, 49[a]
F	37.31	25.29	13.97
p	<.001	<.001	<.001

Note. Within columns, means not sharing subscripts are significantly different at $p < .05$.
[a] Two participants did not provide any recall data.

positive condition. This may be explained by the fact that half of the likeable behaviors in this condition were enacted toward superiors. (Note that this result does not imply a slime effect, because likeability ratings were favorable in this condition.) The slime rating in the nonslime condition was as low as in the negative condition, indicating that likeable behavior toward subordinates was not perceived as slimy.

Free recall. The third column in Table 1 presents participants' recall rates. These were computed by counting for each participant the total number of items recalled accurately, according to a gist-of-meaning criterion (cf. Srull, 1981). Items were coded as accurately recalled when the behavior itself was reproduced, even if the participant did not mention the target toward whom the behavior was performed or mentioned the wrong target. This decision was made because students in the positive and negative conditions often did not mention the targets of the behaviors or confused the targets with each other (in agreement with the assumption that they did not pay attention to the targets because they were not engaging in inconsistency resolution). Thus, if only recalled items with the correct target person had been coded as accurate, this would have inflated the difference between consistent and inconsistent conditions.

The data in Table 1 show that recall rates were highest in the mixed condition, in which the items presented were clearly inconsistent with each other. Recall was lower in the nonslime condition, but this difference was nonsignificant. As predicted, recall in the slime condition was significantly lower than in the mixed condition. Recall rates were lowest in the positive and negative conditions. These data suggest that students in the slime condition did not engage in the sort of elaborate processing that usually occurs when behaviors are inconsistent. Apparently, the two distinct categories of persons toward whom the likeable and dislikeable behaviors were enacted allowed students to quickly explain the variation in the actor's behavior.

Altogether, these results demonstrate that perceivers take into account the target of a behavior when making inferences about the actor. When likeable behavior is enacted only toward superiors, it is perceived as utterly uninformative. This seems to imply that a situational correction is made (e.g., Gilbert & Malone, 1995). By considering the relative power of those toward whom the likeable behavior is performed, the perceiver does not fall prey to the correspondence bias. At this point, however, it is not yet clear whether this correction only occurs when perceivers encounter an inconsistency, which draws their attention to other features of the behavioral field, or whether perceivers are generally attentive to the status of the targets of a behavior. This issue will be addressed in Experiment 3.

An important question is which factors precisely are responsible for the slime effect. It was previously mentioned that slime behavior is more likely to be recognized when the target of the behavior can affect the outcomes of the actor, either immediately or in the long run (Jones, 1964, 1990). Although it is likely that the labels used here to refer to superiors are associated with such power (and the subordinate labels with the lack of it), these labels are contaminated with many other connotations such as mere social status, frequency of occurrence (most organizations have a pyramid shape with few employees in high management positions), gender (labels such as *secretary* and *typist* are more likely to trigger an image of a woman than the labels used to denote superiors), and any idiosyncratic features of each label used.

Thus, it remains to be demonstrated that the slime effect is a direct result of perceived differences in the power of the actor and the behavioral targets. Experiment 2 was conducted to provide this demonstration. In this experiment, the dependence relationship between the actor and two categories of target persons was manipulated directly.

Experiment 2: Does the Slime Effect Depend on Dependence?

If we take a closer look at the dependence relationship between the actor in Experiment 1 and his superiors and subordinates, it may be noted that the differences in these relationships consist of two components. First, students presumably inferred that the actor was dependent on his superiors, whereas he was not dependent on his subordinates. Second, students may also have assumed that the actor's superiors were not dependent on him, whereas his subordinates were. In the asymmetrical dependence relationship between leaders and subordinates, these two components are confounded: Subordinates depend on their superiors, who do not depend on them, or at least not to the same extent. But many relations in everyday life are characterized by symmetrical dependence, and it is conceivable that ingratiating behavior in this case is judged more mildly, if only because the target may benefit from it just as much as the actor. In Experiment 2, the dependence of the actor and the dependence of the targets were manipulated orthogonally. These two variables will be referred to as dependence (the actor's dependence upon the

ses of covariance in the latter two conditions (in which the correlations were similar) showed that the difference in likeability judgments persisted when slime judgments were accounted for, $F(1, 20) = 10.12$, $p < .005$, and vice versa, $F(1, 20) = 11.03$, $p < .005$, suggesting that the reported effects on likeability and on sliminess emerged independently of each other.

targets) and power (the targets' dependence on the actor, i.e., his power over them).

Three hypotheses can be formulated. First, if ingratiation is judged more mildly in conditions of symmetrical dependence, a Power × Dependence interaction should emerge.

Second, in Jones's (1964, 1990) analysis, the goal of making others behave in ways that one desires is tantamount to any self-presentational behavior. Jones also assumed that person perceivers are aware that others tend to engage in self-presentation when they depend on someone for desired outcomes. Thus, when an actor is likeable toward people he depends on, perceivers should take into account that the actor may be engaging in ingratiation. On the other hand, when an actor does not depend on a target in any way, no self-presentational motive can be inferred, and a correspond inference will be made. It follows from Jones's analysis that an actor's dependence on a target (i.e., the need to make the target behave in desired ways) is the crucial variable in the slime effect. Thus, it was predicted that dependence would produce a main effect on slime judgments (because likeable behavior toward targets one depends on is seen as slimy) as well as likeability judgments (because behavior toward targets one does not depend on is seen as informative about the actor's true personality).

Third, it is possible that an actor's power over the target contributes to the slime effect as well, and especially to the negative likeability judgments. Dislikeable behavior may be judged more harshly when it is enacted toward those who depend on the actor, because they are not in a position to retaliate or stand up for themselves. Such behavior is likely to be perceived as authoritative, and a person conducting it is probably seen as a bully. This would result in a main effect of power on likeability judgments, such that these judgments are more negative when the actor is dislikeable toward those who depend on him, compared with both a reversed condition and with a control condition in which no information is given on the actor's power.

Method

Participants and design. The experiment was conducted by means of a computer in individual cubicles. Participants were 74 students (52 women, 22 men) from different levels and different departments (mean age 22.8). They were randomly assigned to the conditions of a 3 (dependence of the actor on target group A, on target group B, control condition) × 3 (power of the actor over target group A, over target group B, control) factorial design.

Cover story and manipulations. In this experiment, the actor, Paul, was described as a registered accountant who had recently set up his own accountancy firm. All students were told that an accountant can be hired by an organization for two reasons. First, the government requires that a registered accountant yearly examines the financial records of all large organizations and provides a statement of approval of their annual report. If the accountant has a reason to withhold this approval, this has several consequences for the organization, such as a bad reputation and extra work for many employees in various departments. Second, an accountant can be hired to give advice, for instance, on improving the efficiency of the administrative organization. In this case, the organization is free to decide whether the accountant's advice will be implemented.

Students were told that Paul had recently had an assignment from two different advertizing agencies, called Target and Unique. He had, thus, spent several weeks in both organizations. Students were going to read descriptions of Paul's behavior during his contacts with the people in each of these agencies. Before these descriptions were presented, students were given additional information about the similarities and differences between the two agencies and the two assignments. They were told that both assignments were the same in the amount of work Paul had to do and in his earnings because he received a standard fee. Further, the agencies were similar in the number of people they employed, in profits and turnover, and in type of work.

In one dependence condition, students were told that Paul was planning to establish his own accountancy firm in the same area where Target was located and that it was his goal to focus on the market in this area. Thus, it was important to him to establish good relations with the employees of Target because this would increase his chances of getting more assignments from Target in the future. If Target would not hire him again, his plans for the future would be thwarted. The assignment from Unique, on the other hand, was less important to him: If his plan succeeded, he would have more clients in his own area next year, and he would give up Unique as a client. Thus, Paul did not feel dependent on Unique. In a second condition, the story was the same, but the labels *Unique* and *Target* were switched. In the control condition, no information was given about Paul's dependence on the agencies.

Orthogonal to this manipulation, the actor's power over the two agencies was varied. In one condition, students were told that the employees of Target did not feel dependent on Paul: They had already received an accountants' approval of their annual report for the running year and they had hired Paul only for advisory work; for most of them, Paul's advice would not have any consequences. The Unique agency, on the other hand, had hired Paul to provide the required statement of approval. To emphasize Paul's power over the employees of Unique, it was added that most of them were anxious about his judgment because his disapproval would strongly affect their work and that they, thus, felt dependent on him. Students were also told that Paul was aware of how important his conclusions were for the people in this agency. In the second power condition, the labels *Target* and *Unique* were switched. In the control condition for power, no information was given about the agencies' dependence on Paul.

Before the behavior descriptions were presented, the information was summarized. In the asymmetric-dependence conditions (power over Unique + dependence on Target, and power over Target + dependence on Unique), it was recapitulated that Paul felt dependent on the employees of Target (Unique), but they did not feel dependent on him. Reversely, the employees of Unique (Target) did depend on Paul, but he did not depend on them. In the symmetric conditions (power over Target + dependence on Target, and power over Unique + dependence on Unique), the summary indicated that the actor and the people at one agency were mutually dependent and the actor and the other agency were mutually independent.

Behaviors and dependent variables. The behavior descriptions were nine likeable, nine dislikeable, and two neutral items from Experiment 1 (two likeability-related items, *Taking over a task from someone who is busy* and *Turning over a difficult assignment to someone else,* were dropped because they were less applicable in this setting). The superior and subordinate labels were all replaced by *an employee at Target (Unique)* or *someone in the Target (Unique) office.* In all conditions, the likeable behaviors were enacted toward Target people and the dislikeable behaviors toward Unique people. Because presentation order did not affect the results in Experiment 1, one order was used. The serial positions of the items were rearranged.

After reading the descriptions, students were asked to judge Paul's likeability, followed by a series of adjective trait scales including *slimy.* Two manipulation checks for dependence were administered: Students indicated on a 5-point scale (a) whether Paul felt dependent on the employees at Target, and (b) whether Paul felt dependent on the employees at Unique. Similarly, to check for the power manipulation, students

were asked (a) whether the employees at Target felt dependent upon Paul, and (b) whether the employees at Unique felt dependent on Paul. Finally, students gave their impression of the employees in each of the agencies (1 = *negative*, 7 = *positive*).

Results and Discussion

Manipulation checks. A multivariate analysis of variance (MANOVA), with participants' ratings of the actor's dependence on Target versus Unique (1 = *not at all*, 5 = *highly dependent*) as the within-subjects variable agency, and power and dependence as between-subjects variables, yielded a substantial Dependence × Agency interaction, $F(2, 65) = 55.94$, $p < .001$. The actor's dependence on Target was rated higher ($M = 4.96$) than on Unique ($M = 1.16$) when he had been described as dependent on Target; reversely, his dependence on Target was perceived as lower ($M = 2.52$) than on Unique ($M = 3.40$) in the other experimental condition, although this difference was smaller. Presumably, the manipulation was less successful in the condition where the actor depended on Unique because his behavior toward Target was likeable and his behavior toward Unique was not. Thus, students in this condition may have inferred that the actor must have been somewhat dependent on Target, and could not be all that dependent on Unique. This effect of the actor's behavior is also illustrated by the dependence ratings in the control condition, which were higher for Target ($M = 3.92$) than for Unique ($M = 1.46$). As a result, a main effect of agency, $F(1, 65) = 96.94$, $p < .001$, indicated that, overall, the actor's dependence on Target was perceived as higher than on Unique ($M = 3.80$ vs. 2.01). In part, these results are unfortunate because, ideally, manipulation checks are not affected by anything other than the manipulated variable. At the same time, they are theoretically interesting because they testify to perceivers' beliefs about the effects of dependence on behavior: When an actor behaves likeably toward one group of targets and dislikeably toward another group, perceivers tend to infer higher dependence of the actor on the former group.

A MANOVA with ratings of the actor's power over Target versus Unique produced a Power × Agency interaction, $F(2, 65) = 72.94$, $p < .001$. The dependence of Target on the actor was rated as higher ($M = 3.96$) than that of Unique ($M = 1.80$) when Target had been described as dependent; reversely, the perceived dependence of Unique was higher ($M = 4.44$) than that of Target ($M = 1.80$) in the other experimental condition. In the control condition, the ratings were in between and did not differ significantly ($M = 2.87$ for Target, $M = 2.62$ for Unique). No other effects emerged.

Likeability and slime judgments. A 3 (dependence) × 3 (power) analysis of variance (ANOVA) on the likeability judgments yielded a main effect of dependence, $F(2, 65) = 5.03$, $p < .01$: The actor was judged as significantly less likeable when he depended on Target ($M = 2.32$, $SD = 1.07$) than when he depended on Unique ($M = 3.24$, $SD = 1.16$) or when no dependence information was given ($M = 3.37$, $SD = 1.44$). This effect occurred regardless of power (main effect and interaction, $F < 1$). For the slime judgment, a main effect of dependence, $F(2, 65) = 24.54$, $p < .001$, reflected higher ratings when the actor depended on Target ($M = 6.24$, $SD = 1.13$) than in the other two conditions ($M = 3.44$, $SD = 1.83$ and $M = 3.87$, $SD = 1.65$). The main effect of power ($F = 1.38$) and the interaction ($F < 1$) were nonsignificant.[3]

These results demonstrate that the slime effect obtained in Experiment 1 cannot be explained by particular features of superiors and subordinates but should, instead, be attributed to differences in dependence between the actor and the target persons. Specifically, the data indicate that the crucial variable is the actor's dependence on a target. When an actor performs likeable behaviors toward those he depends on and dislikeable behaviors toward those he does not depend on, he is judged as more dislikeable and more slimy, regardless of his own power over the target persons. Thus, person perceivers appear to share Jones's (1964, 1990) view on the crucial role of dependence in self-presentational behavior.

It might be added that none of the other trait judgments assessed in this study were affected by the power manipulation either, whereas they were affected by dependence: The actor who depended on Target was judged as more manipulative, hostile, and calculating, and as less sincere, impulsive, and stupid. This latter result is interesting because it qualifies participants' overwhelmingly negative judgments of the brownnoser: Whereas they dislike him on the assumption that his likeable behavior is entirely the product of ingratiation, they are aware of the rationality of this behavior.

Experiment 3: From Suspicion to the Slime Effect

We have established that the actor's dependence is a necessary and sufficient condition for the emergence of the slime effect. The next question is whether it is also necessary that perceivers observe both the actor's likeable behaviors toward those he depends on and his dislikeable behaviors toward those he does not depend on. This question is important for both ecological and theoretical reasons. On the ecological side, it is clear that person perceivers in everyday life do not always have all the information that was given to participants in Experiments 1 and 2. They may know how a person treats superiors, but not subordinates, or vice versa.

On the theoretical side, the question can be raised how perceivers' apparent attention to the power of behavioral targets fits with the correspondence bias, according to which perceivers underestimate situational cues such as target characteristics. In part, this bias occurs because perceivers are less attentive to the situation than to the actor's behavior, and because their knowledge of the effects of a situation on behavior is insufficient or incorrect (Gilbert & Malone, 1995). But even if the situation and its behavioral effects are known, correction may fail because it is a relatively effortful stage of the attribution process; it occurs only after completion of the less thoughtful stages of trait identi-

[3] The correlations between likeability and slime judgments were negative in all conditions ($r = -.49$). As in Experiment 1, analyses of covariance indicated that the dependence effect on likeability persisted when slime ratings were accounted for, $F(2, 64) = 3.45$, $p < .04$, and vice versa, $F(2, 64) = 17.34$, $p < .001$. Further, correction for participants' evaluative impression of the two agencies did not alter the effects of dependence on likeability, $F(2, 63) = 3.55$, $p < .04$, or sliminess, $F(2, 63) = 18.46$, $p < .001$.

fication and ascription, and it may not occur at all when cognitive resources are scarce (e.g., Gilbert, Pelham, & Krull, 1988).

As noted earlier, situational correction is more likely when the perceiver encounters an inconsistency in a person's behaviors, as in the present studies: The perceiver examines if the inconsistency can be explained by situational cues, and engages in a thoughtful attributional analysis that defeats the correspondence bias (for a review, see Vonk, 1994). But it is possible that situational correction is generally more likely in cases of socially desirable behavior: Assuming that perceivers are quite aware that such behavior can result from ingratiation motives (Jones & Davis, 1965), they may routinely consider the setting in which it is conducted, because they need to take into account whether self-presentational motives are relevant in the behavioral setting (Vonk, in press). Indeed, the correspondence bias is weaker for socially desirable than for socially undesirable behavior (Reeder & Spores, 1983; Vonk & Van Knippenberg, 1994). Thus, when observing likeable behavior, perceivers may habitually attend to situational cues that point to the possibility of ingratiation. As noted by Fein (1996), person perceivers have a chronic interest in avoiding being duped by others. As a result, situational cues indicating the possibility of ingratiation may disrupt the automaticity of the trait inference process: The perceiver hesitates in taking the behavior at face value and, instead, engages in a more deliberate and thoughtful analysis of the information available (Fein, 1996).

A first implication of this argument was examined in Experiment 3. It is possible that the observation of likeable behavior toward superiors is a sufficient condition for perceivers to become aware that the actor may be engaging in ingratiation. In terms of Fein (1996) and his colleagues (Fein, Hilton, & Miller, 1990; Hilton, Fein, & Miller, 1993), perceivers may become suspicious of ulterior motivation: When an actor voluntarily engages in behavior that is situationally encouraged (such as flattery toward a more powerful person), the behavior is perceived as ambiguous because the perceiver is uncertain whether the behavior is genuine or is driven by ulterior motives. Consequently, judgments of the actor are moderate until the perceiver receives further information that disambiguates the meaning of the behavior.

Self-presentational behavior, such as ingratiation, is precisely the sort of behavior that qualifies as ambiguous, according to Fein et al.'s (1990) analysis: It presents the perceiver with the problem of deciding whether the actor is genuine, in which case a correspondent inference can be made, or is trying to accomplish an ulterior goal, in which case the inference should be corrected. Thus, in Experiment 1, the actor's likeable behaviors toward superiors may have induced a state of suspicion of ulterior motivation. If this is true, the additional dislikeable behaviors toward subordinates in the slime condition should have served as disambiguating information, confirming participants' suspicion. Note that in the positive condition of Experiment 1, participants may have been equally suspicious of the likeable behaviors toward superiors, but in this case, their suspicion was not confirmed because the actor was likeable toward subordinates as well.

The question arises as to how many dislikeable behaviors toward subordinates it takes for the slime effect to emerge. Assuming that likeable behavior toward superiors produces a state of suspicion, perceivers may be highly alert to any information that disambiguates the meaning of such behavior. Thus, the slightest indication that the actor behaves differently toward subordinates may be sufficient for their judgment to shift dramatically toward the negative end of the likeability scale and toward the slimy end of the slime scale. Alternatively, it is possible that the process from suspicion to the slime effect evolves gradually: From the moderate judgments of the suspicious perceiver, judgments may become more negative as more disambiguating information is encountered. These possibilities are examined in Experiment 3, by varying the number of dislikeable behaviors toward subordinates.

The experiment involved five conditions. In the suspicion condition, participants received only likeable behaviors toward superiors. In three other conditions, these behaviors were supplemented with either one, two, or four dislikeable behaviors toward subordinates. Thus, across these four conditions, the number of dislikeable behaviors was progressively increased from zero to four. In a fifth condition, participants were presented with the same set of behaviors as in the suspicion condition, but these were enacted toward persons whose position was not clearly superior or subordinate with respect to the actor. This control condition was included to establish the hypothesized state of suspicion in the first condition: It was predicted that judgments in the control condition would be unambiguously favorable because the behaviors would be taken at face value; in the suspicion condition, on the other hand, judgments would be more moderate as a result of the perceived ambiguity of the likeable behaviors.

Regarding the other three experimental conditions, it was hypothesized that judgments in these conditions would be more negative than in the suspicion condition, as a result of the disambiguating dislikeable behaviors. Specifically, participants in these conditions should be more confident that the actor is dislikeable and slimy, and this should be manifested by more extreme judgments compared with the suspicion condition. As noted, this difference could be reflected either by a discrete shift with one, two, or four dislikeable items, or by a gradual linear pattern across the four experimental conditions.

Method

The procedure, introduction, dependent variables, and manipulation check for behavior were the same as in Experiment 1. Seventy-four students (38 male and 36 female undergraduates from different departments) read 14 behavior descriptions of Paul. In all conditions, 10 of these were the likeable descriptions presented in Experiment 1. In the control condition, the targets of all 10 likeable behaviors were described in the same way as in the pretest (e.g., *a colleague, someone*). In the four experimental conditions, all likeable behaviors were enacted toward superiors, just as in the slime condition of Experiment 1. The additional descriptions consisted of either four neutral items (suspicion and control condition), one dislikeable and three neutral items, two dislikeable and two neutral items, or four dislikeable items. Each of the dislikeable behaviors were enacted toward a subordinate. These descriptions were extracted from Experiment 1, such that the evaluative ratings of the four behaviors were similar to each other. The neutral descriptions did not contain any references to subordinates or superiors.

In all conditions, the series started with three likeable items and ended with one likeable item. The rest of the series consisted of a repeated alternation of one neutral or dislikeable behavior (depending on condi-

tion) and two likeable behaviors. Thus, in the suspicion and the control conditions, the neutral items were presented at the 4th, 7th, 10th, and 13th serial position. In the condition with four dislikeable behaviors, these neutral items were replaced by dislikeable items. In the other two conditions, the serial position of the one or two dislikeable items was varied by three different permutations, with a dislikeable item occurring either at the 4th (and 7th) position, at the 7th (and 10th) position, or at the 10th (and 13th) position. Although this could theoretically make a difference because the assumed state of suspicion is prolonged across these conditions, trend analyses did not reveal any effects of serial position.

Results and Discussion

Manipulation check. In the conditions where two or four dislikeable items were presented, all participants indicated that the actor behaved differently toward subordinates than toward superiors. In the condition with one dislikeable item, only one participant responded incorrectly that the actor's behavior varied across strangers and familiar people. Generally, then, participants required very little dislikeable information to identify the superior–subordinate distinction.

In the control condition, all participants except one, who endorsed the familiar–stranger distinction, responded correctly that the actor behaved in the same way toward everyone. In the suspicion condition, on the other hand, less than half of the participants selected this response: The majority (10 out of 15) responded that the actor's behavior varied across subordinates and superiors. Thus, even though these participants did not receive any information about the actor's behavior toward subordinates, they assumed this behavior would be different.

Likeability and slime judgments. Students' judgments are presented in the first and second columns of Table 2. In the control condition (last row), the actor was perceived as significantly more likeable and less slimy than in the suspicion condition (first row: zero dislikeable behaviors). In the latter condition, the actor's likeability was judged quite moderately, suggesting that participants were uncertain whether the likeable behaviors toward superiors were sincere. The high rating on the slime adjective, on the other hand, indicates that they were more confident about this particular inference. Considered concurrently with the results from Experiment 1, this suggests that the observation of a series of likeable behaviors toward superiors is sufficient to evoke the inference that the actor is slimy, regardless of whether the actor's behaviors toward subordinates are likeable (as in the positive condition of Experiment 1), dislikeable (as in the slime condition of Experiment 1), or unknown (as in the suspicion condition of the present experiment).

Table 2
Means and F *Tests for Likeability (1 = Very Dislikeable, 7 = Very Likeable), Sliminess (1 = Not at all, 7 = Highly), and Reading Times (in Seconds, Across Seven Likeable Items) in Experiment 3*

No. of dislikeable behaviors toward subordinates	Likeable	Slimy	Reading time likeable items
0	4.13_a	6.07_a	3.82_a
1	$3.40_{a,b}$	6.27_a	$3.19_{a,b}$
2	$3.13_{b,c}$	6.33_a	$3.27_{a,b}$
4	2.53_c	6.40_a	2.71_b
$F_{\text{linear effect}}(1, 56)$	10.70**	1.02	7.73**
$F_{\text{quadratic effect}}(1, 56)$	<1	<1	<1
$F_{\text{cubic effect}}(1, 56)$	<1	<1	1.37
Control condition	6.36_d	3.29_b	$3.12_{a,b}$
$F_{\text{overall}}(4, 69)$	20.10***	18.36***	2.51*

Note. Within columns, means not sharing subscripts are significantly different at $p < .05$.
* $p < .05$. ** $p < .01$. *** $p < .001$.

Because the slime ratings in the suspicion condition are quite high to begin with, the predicted increase across the four experimental conditions seems to be restrained by a ceiling effect. The slime ratings do increase across the four conditions (up to the level obtained in the slime condition of Experiment 1), but the differences are quite small and nonsignificant. For the likeability judgment, on the other hand, a significant linear effect was obtained. Starting with the neutral judgment in the suspicion condition, the likeability of the actor gradually decreases as more dislikeable behaviors toward subordinates are encountered; in the condition with 4 dislikeable items, the judgment is severely close to that obtained in the slime condition in Experiment 1 (which contained 10 of these items).

Additional results: Reading times. Although no hypotheses were formulated about reading times, these were also registered for each item presented, as this assessment is entirely unobtrusive. For each student, the mean reading time was computed across the last seven likeable items presented (reading times of the first items were excluded; these tend to be unreliable because participants have to get used to the rhythm of reading and pressing the *Return* key), after excluding nine incidental outliers due to an apparent lapse of attention (for details of this procedure, see Vonk & Van Knippenberg, 1995). Reading times were longest in the suspicion condition (see Table 2, column 3), corroborating the assumed state of suspicion and the concurrent sophisticated processing in this condition. The shortest reading times were found in the condition with four dislikeable behaviors toward subordinates. In the other three conditions, reading times were in between.[4]

For the three experimental conditions in which dislikeable information was presented, the mean reading time across the seven likeable items was compared with the reading time of the one dislikeable item that was presented in all three of these conditions. A main effect of this within-subjects variable, $F(1, 40) = 31.73$, $p < .001$, indicated longer reading times for the dislikeable item than for the likeable items ($M = 4.03$ s vs. 3.02 s), regardless of condition (interaction $F < 1$). Similarly,

[4] Detailed between-subjects and within-subjects trend analyses, in which the likeable items were grouped into different blocks depending on their serial position, did not produce any significant effects. Possibly, this is the result of two competing tendencies: On the one hand, one might expect the likeable items to be studied shorter as more dislikeable items are encountered because these make the likeable items less informative; on the other hand, the alternation of likeable and dislikeable items should encourage participants to remain attentive, whereas attention may decline gradually when no dislikeable items are encountered.

in the condition where four dislikeable items were presented, the mean reading time across these four items was also longer than of the likeable items (M = 3.50 vs. 2.71), $F(1, 14) = 15.00$, $p < .005$.

The present results suggest that the observation of likeable behavior toward superiors evokes a state of suspicion, which is reflected in a relatively long elaboration on the available information and, eventually—when no disambiguating behaviors are encountered—in moderate likeability judgments. In agreement with Fein et al.'s (1990) account of suspicion of ulterior motivation, perceivers in this state appear to be highly attentive to any information that confirms their skepticism: Such information attracts their attention and strongly affects their likeability judgments. Regarding the process from suspicion to dislike, the present data show that this process evolves gradually. As more disambiguating information is received, suspicion gradually turns into confidence that the likeable behaviors are not genuine. The observation of four disambiguating behaviors was sufficient for the slime effect to emerge to its full extent. This progressive linear pattern may be explained in terms of Kelley's (1973) distinctiveness and consistency dimensions: Upon the observation of one dislikeable behavior toward a subordinate, one cannot be entirely certain that the actor is dislikeable, because the behavior could be related to this particular subordinate or this particular occasion. As the number of dislikeable behaviors and the number of different subordinate targets increases, the perceiver learns that the dislikeable behavior is consistently enacted toward multiple subordinates.

The results of this study further demonstrate that perceivers do pay attention to situational cues, even when they are not resolving an inconsistency. Participants in the suspicion condition did not receive one single negative description of the actor, yet their judgments were quite moderate. In agreement with Fein's (1996) analysis, the long reading times of these participants suggested that they were carefully looking for clues that might shed light on the authenticity of the actor's behaviors. In this condition, the power of the target was not considered in order to resolve an inconsistency, but instead, to determine the evaluative meaning of the behaviors, that is, to resolve the attributional ambiguity that is inherent to socially desirable behavior (Jones & Davis, 1965; Reeder & Brewer, 1979; Vonk, in press). Participants who received multiple disambiguating behaviors, on the other hand, had little trouble arriving at their conclusions: Their reading times were shorter than in any of the other conditions. Clearly, the inconsistent dislikeable behaviors were not a problem for these participants; instead, they were the solution to a problem.

Finally, the present data provide a more detailed account of the slime effect. As noted earlier, this effect consists of two components: (a) a high slime rating and (b) a low likeability rating. The results of Experiment 3 demonstrate that the observation of likeable behaviors toward superiors ("licking upward") is sufficient for the first component: These behaviors lead participants to perceive the actor as slimy, regardless of other behaviors. The observation of dislikeable behavior toward subordinates ("kicking downward") is necessary for the second component, however few of these behaviors it takes.

Experiment 4: The Illusory Slime Effect

It has been noted earlier that perceivers may have an accessible social schema for brownnosers and for the behavioral pattern of "licking up–kicking down." Such a schema would be useful in quickly identifying the type of behavior that may be driven by the ulterior motive of ingratiation. The results of Experiment 3 suggest that a target's hierarchical position is an important component of such a schema, and is used by person perceivers to determine whether likeable behavior is suspect. Put differently, it is possible that a *slime schema* is activated when perceivers observe a person who is likeable toward more powerful targets.

This could account for an unexpected finding on the manipulation check of Experiment 3: Two-thirds of the participants who read only likeable behaviors toward superiors (suspicion condition) responded that the actor behaved differently toward subordinates, even though they did not have any information about this. Possibly, the behavior of "licking upward" is sufficient to activate a slime schema—a schema which then goes beyond the information given, by the assumption that the actor may be someone who "kicks downward" as well. Thus, a person who is likeable toward superiors may, at least tentatively, be categorized as the type of person who licks upward *and* kicks downward. As a result, perceivers may erroneously form a link between the person and the "kicking downward" component of this category.

Similarly, it is conceivable that the activation of a slime schema produces the impression that there is a covariation between the actor's behavior and the power of the target, even when in fact there is none. To examine this possibility, the illusory correlation paradigm was used. Traditionally, this paradigm has been applied in research on stereotypes to demonstrate that the cooccurrence of infrequent, distinctive stimuli (such as negative behavior enacted by minority group members) can produce an illusory correlation between behavior and group membership (the paired-distinctiveness effect; e.g., Hamilton & Gifford, 1976). In these studies, participants read a series of behaviors by multiple actors from two different groups. In Experiment 4, this paradigm was used to study the perceived association between behavior (likeable, dislikeable) and group membership of the behavioral targets (subordinates, superiors). Thus, differing from previous illusory correlation studies, there was one single actor, but multiple targets.

In all conditions, 18 behaviors were presented, the majority of which were positive (12 positive, 6 negative). In one condition, 12 behaviors (8 positive, 4 negative) were enacted toward superiors and 6 behaviors (4 positive, 2 negative) toward subordinates. In another condition, this was reversed in terms of subordinates receiving more positive than negative behaviors. Before specifying the central hypothesis, it should be noted that in this particular study, the emergence of an illusory correlation (between behavior and target category) was rather uncertain because there are many differences with the traditional paradigm in which the effect has been established. First, the stimulus consists of one single person rather than multiple group members or individuals. Second, the total number of items presented is much smaller than in the typical illusory correlation study, so that the memory load which contributes to the effect (e.g., Mullen & Johnson, 1990) is lower. And third, participants are instructed to form an impression of the actor, an instruction that attenuates the illusory correlation effect (e.g., McConnell, Sherman, & Hamilton, 1994).

One condition of this experiment, however, contains an ingre-

dient that would seem to provoke an illusory correlation between the actor's behavior and the behavioral targets: When the behaviors toward subordinates are a minority, the actor's negative behaviors share their infrequency with subordinate targets. In this case, the paired-distinctiveness effect would produce the impression that the actor's dislikeable behaviors are more frequently enacted toward subordinates, and that his likeable behaviors are more frequently associated with superiors. This association converges with the assumed slime schema, which could bolster an illusory correlation in the same way that stereotypic expectancies enhance the illusory correlation effect (cf. McArthur & Friedman, 1980). In the second condition, on the other hand, the effect seems less likely because it would produce the impression that the actor's likeable behavior is more often enacted toward subordinates and his dislikeable behavior more often toward superiors. Presumably, this kind of association is unlikely to be perceived erroneously because it does not fit into a readily available cognitive category.

Thus, given that negative behaviors are more infrequent than positive ones, it was predicted that the illusory correlation would be stronger in the condition where subordinate targets occur more infrequently than superior targets because of the combined result of paired-distinctiveness and schema-based expectancies (i.e., distinctiveness-based and schema-based illusory correlation). Specifically, participants in this condition should (a) overestimate the occurrence of negative behaviors toward subordinates and (b) infer more strongly that the actor behaves differently (i.e., more likeably) toward superiors than toward subordinates. As a consequence, the slime effect was expected to emerge in this condition; that is, the actor was expected to be judged as more dislikeable and slimy than in the other condition.

Method

Participants were 38 high school students in the 11th and 12th grades (16 boys, 22 girls; mean age 17.5). They were given the same background information and instructions as in Experiments 1 and 3. Differing from these experiments, the presentation time of each behavior description was fixed at 4 s, with 0.5 s in between each description.

The descriptions consisted of 12 likeable and 6 dislikeable items. The 6 dislikeable items were extracted from the 10 items used in the other studies. The 12 likeable items consisted of the 10 items from the other studies and 2 additional items selected from the behavior pretest. Also, two additional labels were introduced for superiors and subordinates (*supervisor, executive* and *junior employee, clerk*) because one condition required 12 superior labels and the other condition required 12 subordinate labels. In both conditions, the behaviors comprised four sets: (a) eight positive behaviors toward majority targets (superiors or subordinates, depending on condition), (b) four positive behaviors toward minority targets, (c) four negative behaviors toward majority targets, and (d) two negative behaviors toward minority targets. The behaviors were distributed across these sets in such a way that behavioral extremity in each set was the same.

For each participant, the presentation order of the items was determined randomly. After reading the descriptions, participants judged the actor on likeability and on 10 traits, including *slimy*. Subsequently, they were asked whether the actor behaved differently toward superiors and subordinates (1 = *not at all*, 7 = *very much so*) and whether he was more likeable toward one of these groups (1 = *much more likeable toward superiors*, 7 = *much more likeable toward subordinates*, 4 = *no difference*). Finally, participants were reminded that they had read 18 descriptions and were asked to estimate the frequency of likeable behaviors toward superiors, likeable behaviors toward subordinates, dislikeable behaviors toward superiors, and dislikeable behaviors toward subordinates. They were told that their estimates did not need to add up to 18 and that for each estimate, they could choose a number from 0 to 18.

Results and Discussion

Participants' estimates of the frequency of behaviors in each of the four sets are presented in rows 1 through 4 of Table 3, along with the actual frequencies of behaviors in each set. In both conditions, students underestimated the frequency of positive behaviors toward the majority targets, and overestimated the frequency of negative behaviors toward minority targets. The latter effect, which reflects a paired-distinctiveness effect, was larger when subordinate targets were the minority than when superiors were the minority (M = 4.81 vs. 3.00; $p < .001$), whereas estimates for the other three sets did not differ between conditions; Behavior × Target × Condition, $F(1, 36) = 15.88$, $p < .001$. From the frequency estimates, a phi coefficient was computed for each student, reflecting the perceived correlation between behavior and target category (Table 3, row 5). As predicted, this coefficient differed significantly from zero when subordinates were the minority (.17; $t[20] = 4.62$, $p < .001$), whereas it was nonsignificant in the other condition (−.04; $t[16] = 1.50$). Further, the conditions differed significantly from each other, $F(1, 36) = 19.36$, $p < .001$.

Participants' judgments of whether the actor behaved differ-

Table 3
Frequency Estimates and Phi Coefficients of Behaviors Presented (Rows 1–5), Judgments of Differences in the Actor's Behavior Toward Superiors and Subordinates (Rows 6–7), and Judgments on Likeability and Sliminess (Rows 8–9) in Experiment 4

Dependent variable	Superiors majority–subordinates minority	Superiors minority–subordinates majority
Positive items toward majority targets (8)	6.52*	6.06*
Positive items toward minority targets (4)	4.33	5.41*
Negative items toward majority targets (4)	3.33	3.94
Negative items toward minority targets (2)	4.81***	3.00**
Phi-coefficient target-behavior	.173	−.044
Differential behavior (1 = *no*, 7 = *highly*)	4.71	2.41
More likeable toward superiors (1) subordinates (7)	3.09	4.00
Likeable	3.62	5.06
Slimy	5.24	3.06

Note. Pairwise t tests (two-tailed) of participants' frequency estimates with the actual frequency of behaviors presented in parentheses.
* $p < .05$. ** $p < .01$. *** $p < .001$.

ently toward superiors and subordinates are presented in rows 6 and 7 of Table 3. As predicted, when subordinate targets were the minority, students rated the actor's behavior as more variable across subordinates and superiors, $F(1, 36) = 19.40, p < .001$, and as more likeable toward superiors, $F(1, 36) = 12.69, p = .001$. In summary, the behavior series in which the subordinate targets were the minority induced an illusory correlation, such that the perceived association between the actor's behavior and the target category was in agreement with the presumed slime schema, even though there was in fact no such association.

This illusory correlation was also reflected in the likeability and slime judgments (Table 3, rows 8 and 9). When subordinate targets were the minority, the actor was judged as more dislikeable, $F(1, 36) = 13.01, p = .001$, and more slimy, $F(1, 36) = 17.70, p < .001$. In part, these effects were mediated by the perceived association between behavior and target category: When the phi coefficient was entered as a covariate, the predicted means revealed smaller differences between conditions ($M = 3.88$ vs. 4.80 for likeability), $F(1, 35) = 3.67, p < .07$, and ($M = 4.75$ vs. 3.55 for slimy), $F(1, 35) = 4.16, p < .05$. Note that the effect for slimy remained significant, in agreement with the conclusion from Experiment 3, that likeable behaviors toward superiors are sufficient for a relatively high slime rating, regardless of the actor's other behaviors.

Generally, students' ratings of the actor's likeability and sliminess were more moderate than in the slime conditions of Experiments 1 and 3, suggesting some reservation. This is not surprising, considering that the information presented was as inconsistent as in the mixed condition of Experiment 1, implying complete orthogonality of behavior and target category. Nevertheless, in one condition, a moderate slime effect emerged which, in this case, was not justified by the actor's behavior. In fact, the mean evaluative judgment of the actor was below the scale midpoint, which is remarkable considering that two thirds of the behaviors were likeable.

In summary, an illusory correlation only emerged when this correlation was in accordance with expectancies about the typical brownnoser, suggesting that the behaviors activated a slime schema.[5] An additional contribution of this experiment is that it testifies to the argument described previously that perceivers often attempt to allocate likeable and dislikeable behaviors to different domains or settings. As a result, it was argued earlier, the psychological unit of the individual person is broken. In a recent study, McConnell, Leibold, and Sherman (1997, Experiment 2) demonstrated that, when the stimulus is a single person, this segregation is a prerequisite for the emergence of an illusory correlation: When participants were instructed to form an integrative impression of the stimulus person, the effect did not emerge. Presumably, this instruction impeded the segregation process. Note that participants in the present study were in fact instructed to form an impression of the actor. Nevertheless, in one condition, the illusory correlation effect was obtained. There is no other plausible explanation for this finding, other than the assumption that the superior–subordinate segregation required for this effect was encouraged by a preexisting schema that readily lays out the distinction.

Experiment 5: Spontaneous Slime Inferences

Let us now recapitulate the evidence for the assumption that person perceivers have an articulate slime schema. If such a schema exists, it should have several effects. First, like other schemas, it should fill the gaps in information that is missing, so that the perceiver can go beyond the information given. As noted earlier, the manipulation check in Experiment 3 provided some suggestive evidence in this direction. Second, it should make perceivers attentive to cues about the power of the target of likeable behavior, even when these behaviors are encountered in an entirely consistent context, so that there is no need for thoughtful attributional analysis. This was demonstrated in Experiment 3. Third, assuming that a slime schema contains the knowledge that dependence can produce ingratiation, the effects of such a schema should work two ways. On the one hand, perceivers infer sliminess from likeable behavior toward more powerful persons; on the other hand, when a person is likeable toward one group of people and dislikeable toward another group, perceivers may infer that the person is more dependent on the former. The dependence ratings in Experiment 2 suggested that this was the case. Fourth, since existing stereotypes bolster the illusory correlation, this effect should be stronger when its direction converges with a slime schema. This was demonstrated in Experiment 4. Finally, because schemas facilitate information processing, a slime schema should help perceivers to quickly and spontaneously discern the prototypical instance of this schema, that is, the behavioral pattern of licking up–kicking down, even when processing capacity is restrained.

This latter implication was examined in Experiment 5, by manipulating participants' processing time. A slime schema should lead participants to spontaneously recognize the licking up–kicking down pattern, even when they do not have time to elaborate. On the other hand, assuming that participants do not have an accessible schema for the reversed behavioral pattern (i.e., the nonslime condition of Experiment 1), this pattern may be recognized only with longer exposure times because it requires more extensive thought.

A 2 (processing time: short, long) × 2 (pattern: slime, nonslime) design was used to test the hypothesis that participants can spontaneously identify a "slimer" regardless of exposure time, whereas a "nonslimer" is only identified when processing time is sufficient. The critical dependent variable consisted of participants' response times on the question of whether the actor treats superiors and subordinates differently. The reason that this particular question was used is that it should logically evoke an affirmative response in both pattern conditions. Presumably, a short response time on this question implies that participants already drew a conclusion about this issue before the question was asked (cf. Smith & Miller, 1983). Thus, it was hypothesized that response times would reveal a Pattern × Processing Time interaction.

Method

Participants were 165 law and history undergraduates (100 women, 65 men, mean age 20.07). The procedure was the same as that used in

[5] Note that it is not clear which mechanism precisely produced the absence of an illusory correlation in the other condition. It is possible that the pattern of behaviors in this condition did not activate any existing schema, so that the paired-distinctiveness effect was not strengthened by prior associations. Another possibility is that a slime schema was activated and undercut the paired-distinctiveness effect because the two had opposite implications.

the other experiments.[6] The behaviors presented were similar to those in the slime and nonslime conditions of Experiment 1, except that the total number of behaviors was 16, and the covariation between behavior and target status was not perfect: In all conditions, participants received 2 likeable behaviors toward superiors, 2 dislikeable behaviors toward superiors, 2 likeable behaviors toward subordinates, and 2 dislikeable behaviors toward subordinates. The 8 other behaviors depended on pattern: In the slime condition, 4 likeable behaviors toward superiors and 4 dislikeable behaviors toward subordinates were presented. In the nonslime condition, 4 likeable behaviors toward subordinates and 4 dislikeable behaviors toward superiors were presented. Thus, altogether participants in the slime condition read 6 likeable behaviors toward superiors, 6 dislikeable behaviors toward subordinates (in agreement with the slime schema), 2 likeable behaviors toward subordinates, and 2 dislikeable behaviors toward superiors (deviating from the slime schema). In the nonslime condition, the reverse applied. This imperfect covariation was introduced to make things a bit more complicated than the clear-cut pattern presented in Experiment 1, so that students' abilities in identifying the slime pattern could really be put to the test. Presentation order was randomized for each student. Crossed with the pattern variation, the exposure time per item was varied: 3 s in the short condition and 5 s in the long condition.

After reading the behaviors, students were informed that several questions would be asked that should be answered with *yes* or *no*. They were told to respond as accurately and as quickly as possible. For this purpose, two keys on the computer board had a sticker on them: a green sticker with *yes* and a red sticker with *no*. Students were instructed to place their right forefinger on the *yes* key and their left forefinger on the *no* key, and to keep their fingers there throughout the question series. This way, they could immediately respond by simply pressing either their right or their left finger. They were also told that if they accidentally hit the wrong key, they should inform the experimenter of this afterward.

After students were given 3 s to place their fingers on the keys, the questions were presented. The first questions served to help students get used to the procedure, and were about general issues, such as *Are you right-handed?, Were the descriptions clear?,* and *Does Paul work at the middle management level?* Subsequently, the critical question was presented: *Does Paul treat his superiors and subordinates differently?* Succeeding questions were about other traits, for example, *Is Paul a likeable person?* and *Is Paul a slimy person?* For each question, the time was registered from presentation of the question to the student's keypress.

After this series of yes–no questions, the same trait scales were presented as in the other studies, but these are not discussed here because they produced the same results.

Results

Four participants' data were discarded because they said they had hit the wrong key on the critical question or a question shortly before that (2 students), or because their response time on the dependent variable was more than 3 *SD*s above the grand mean (2 students).

A 2 (processing time) × 2 (pattern) ANOVA yielded a main effect of pattern, $F(1, 157) = 14.44$, $p < .001$, qualified by the predicted interaction, $F(1, 157) = 4.30$, $p < .05$. Response times on the question of whether the actor behaved differently toward superiors and subordinates were shorter in the slime than in the nonslime condition ($M = 3.48$ vs. 4.38 s). This difference was caused primarily by the short processing-time condition ($M = 3.35$ vs. 4.76, $p < .001$); in the long condition, the difference was nonsignificant ($M = 3.61$ vs. 4.02).

Examination of students' actual responses indicated that in the slime condition, more students responded affirmatively (63 out of 80) than in the nonslime condition (39 out of 81). Responses did not differ as a function of processing time. Because affirmative responses tend to be quicker than negative ones, it was examined whether this difference was responsible for the effects on reading times, by entering students' responses as a covariate in the reading time analysis. However, the effects remained, $F(1, 156) = 10.55$, $p < .001$, for pattern, and, $F(1, 156) = 4.11$, $p < .05$, for Pattern × Processing time.[7]

These results demonstrate that perceivers can spontaneously discern the behavioral pattern of licking up–kicking down, even when the covariation between behavior and target position is imperfect, and even when processing capacity is restrained. This corroborates the assumption of a slime schema that, once it is activated, facilitates information processing.

General Discussion

Summary and Conclusions

Altogether, the present results convey a coherent picture of how the social perceiver forms inferences and impressions of a brown-noser. First, people appear to have a slime schema, containing knowledge of the prototypical authoritarian personality who licks upward and kicks downward. This schema enables them to quickly and spontaneously identify this type of person, even when the covariation between behavior valence and target status is imperfect (as in Experiment 5) and, in some cases, even when it does not exist at all (as in Experiment 4).

Presumably, the slime schema also contains knowledge of situational variables that elicit ingratiation, such as the actor's dependence (as manipulated in Experiment 2) and the hierarchical position of behavioral targets (as used in the other experiments). As a result, the schema is activated when people observe an actor who is likeable toward superiors, without knowing how this person behaves toward subordinates: The actor is tentatively categorized as a brownnoser. Concurrently, the perceiver infers that the actor is slimy and becomes uncertain about the actor's true likeability. The perceiver enters a state of suspicion of ulterior motivation, in which the available information is scrutinized. As in the studies by Fein (1996), this suspicion state is associated with sophisticated processing (as reflected in the

[6] In this experiment, different instructions were used (form an impression, concentrate on whether the actor behaves differently toward superiors and subordinates, memorize the behaviors). Because this variable did not produce any effects whatsoever, it was dropped from the analyses and is not discussed further.

[7] Response times to the question of whether the actor was slimy and likeable were examined as well, because this might provide insight into the question of whether slime judgments are formed spontaneously. Assuming that likeability judgments are made on-line (cf. Hastie & Park, 1986), it was expected that the response time of slime judgments would not differ from that of likeability judgments. Although these results should be interpreted with caution because both questions were asked after the differential-behavior question and the order was not varied, it is remarkable that slime judgments were significantly faster than likeability judgments in all conditions ($M = 2.13$ vs. 2.76), $F(1, 160) = 23.40$, $p < .001$.

longer reading times in Experiment 3), thus abating the correspondence bias and producing conservative evaluations.

Third, the suspicious perceiver uses information about how the actor behaves toward subordinates in order to test the hypothesis that the likeable behaviors toward superiors are insincere and, hence, are uninformative about the actor's likeability. When the actor turns out to be likeable toward subordinates as well (as in the positive condition of Experiment 1), the suspicion of ulterior motives is disconfirmed: The perceiver decides that the actor is a likeable person after all. This results in favorable likeability ratings, but at the same time, the actor is still perceived as slimy. Apparently, the evidence of true likeability does not relieve the actor of the initially inferred sliminess.

Alternatively, the actor's behaviors toward subordinates may turn out to be dislikeable, in which case the perceiver's suspicion is confirmed. This produces the slime effect: In addition to being slimy (because of the likeable behaviors toward superiors), the actor is perceived as dislikeable, because the behaviors toward subordinates are taken as evidence of the actor's true likeability, and the behaviors toward superiors are discounted in the evaluation.

The Slime Effect in Everyday Life

The present experiments started with an observation from everyday life: At some point, we have all had one or more colleagues whom we disliked or at least distrusted, because their behavior toward superiors seemed unusually flattering, supportive, or interested. Most likely, we too have engaged in brownnosing on multiple occasions, either so blatantly that we have ourselves been the object of suspicion or even dislike, or so cleverly that nobody saw through it, not even us. Brownnosing goes on every day and everywhere simply because it is functional to establish a favorable impression with those that can potentially affect one's outcomes.

Different than in the present experiments, however, brownnosing is likely to go unnoticed on many occasions in everyday life. First, it seems fair to assume that, unlike the actor in these studies, most people do not go about grovelling for everyone they depend on and being rude to everybody else. Impression managers are aware that they will be exposed as brownnosers if their differential behavior toward powerful and powerless persons is obvious. In what might be regarded as a "manual for the ingratiater," Jones (1990, pp. 180–184) recommended several strategies that are unlikely to be detected and are probably used frequently. One of these is to conceal one's dependence or reduce its salience. The results from Experiment 2 substantiate this recommendation, by indicating that an actor's dependence is a crucial cue for inferring that likeable behavior reflects ingratiation. Another strategy is to obscure one's behavioral pattern. Indeed, *mainstream brownnosing,* as it was called by an anonymous reviewer, is not characterized by a flagrant contrast between behaviors toward superiors and subordinates: Proficient slimers hide their motives, for instance, by disagreeing with a superior on trivial issues (Jones, 1990) so that the observer does not notice the preferential treatment they bestow on those they depend on. Nevertheless, the results of Experiment 5 suggest that this strategy should be carried out quite rigorously if it is to succeed: If only a minority of one's behaviors are inconsistent with the licking up–kicking down scheme, it won't help at all.

A second restriction of these studies is that the participants had no personal interest in the sincerity of the actor's behavior. In everyday interaction, on the other hand, it may matter much more how a person behaves toward them than toward other people (cf. Jones & Davis's, 1965, principle of hedonic relevance). Thus, if participants are the target of ingratiation, it may be less relevant how the ingratiater treats other people. Moreover, self-esteem is served if participants assume that the ingratiater's behavior is genuine. A recent meta-analysis by Gordon (1996) demonstrated that ingratiation is more effective when the observer is the object of ingratiation than when the observer is an outsider. Thus, sliminess is less likely to be recognized precisely by those who are intended not to recognize it (see also Vonk, 1998). As noted by Jones, "the target person's vanity . . . often comes to the rescue of the self-presenting actor" (1990, p. 194). The effects of this vanity were illustrated in a study by Kipnis and Vanderveer (1971), who found that ingratiation toward leaders with high self-esteem was more effective than toward those with low self-esteem.

A third reason that brownnosing is less conspicuous outside the laboratory is that perceivers usually lack the entire package of behavioral information that is necessary to confidently infer that likeable behavior results solely from self-presentation. In the slime conditions of the present studies, participants perceived multiple likeable behaviors toward superiors and multiple dislikeable behaviors toward subordinates. In everyday life, on the other hand, perceivers often have little knowledge about a person's behavior toward other people than themselves or those close to them. Superiors especially often lack information about how their subordinates behave toward other subordinates, if only because they do not have the time and the motivation to keep track of all the behaviors of their employees (cf. Fiske, 1993). Seeing that a subordinate is friendly and supportive toward them, they are bound to like this person. It is to be feared that many leaders remain in this sweet state of ignorance, while their subordinates gnash their teeth watching how one of their colleagues "butters up" the boss.

A fourth limitation of these studies is that the actor was always a man. It is possible that this strengthens the Slime effect in comparison with women actors because, according to gender stereotypes, men are (a) less likely to be supportive, interested, and flattering, so that this behavior is perhaps more easily perceived as insincere; (b) more ambitious and, thus, more likely to have a higher stake in getting close to powerful persons; and perhaps (c) more likely to ingratiate themselves in this rather blatant, unsubtle manner. Recently, however, several slime studies have been conducted in our laboratory in which the actor's sex was varied, and we found no indication that slimy behavior by a male is more easily identified than by a female.

Finally, it should be noted that there may be cases in which slimy behavior is judged more mildly than in the present studies. For one thing, it is likely that licking upward is perceived as quite adequate and well-adapted behavior in other cultures. Further, there may be circumstances in which one would find it perfectly sensible and forgivable to ingratiate oneself with a more powerful person, for instance, when the actor has the interest of a larger group at heart, rather than his or her own personal interest; when there is a threat of losing something

(e.g., one's job) rather than a promise of gaining something; or in matters of the heart. After all, all is fair in love and war.

Theoretical Implications

Given the reservations above, the present results demonstrate the "social repercussions of a nasty disposition and a possible lack of character." [8] In addition, they are relevant to several theoretical issues in the field of impression formation and trait inferences.

Inconsistency resolution and social schemas. A first contribution concerns the issue of inconsistency resolution in impression formation. As noted previously, psychologists know when and why inconsistent information induces elaborative thought, but they know very little about what precisely these thoughts are about. The problem facing perceivers who observe inconsistencies within a single actor is that these cannot be explained by distinguishing subgroups, as in the case of stereotype-inconsistent information (cf. Vonk & Van Knippenberg, 1995). However, at the level of the individual person, a parallel process exists that has received very little attention in the literature, namely, impression differentiation: Perceivers may try to identify different circumstances in which the person is likeable and dislikeable. If they succeed in segregating these two types of behaviors by some domain of circumstances, they form multiple, situation-specific representations of the person (cf. Trafimow, 1994; Vonk, 1994).

The present results extend this argument by indicating that some forms of segregation are more readily applied than others because they converge with existing social categories. Specifically, when an actor is likeable toward superiors and dislikeable toward subordinates, the segregation between the two appears to be established without much elaboration, thus reducing the superior recall that is usually found when participants receive inconsistent information (Experiment 1). Further, when people observe a person who licks upward, this may be sufficient for them to tentatively categorize the person as someone who kicks downward as well, even when they have no information about this (Experiment 3). This particular segregation may even occur when the available data indicate that there is in fact no covariation between behavior valence and target status, as was demonstrated by the illusory correlation obtained in Experiment 4. Because this effect could not have emerged without any within-person segregation (McConnell et al., 1997), it must be assumed that the activated slime schema produced the impression segregation (i.e., between interactions with superiors and subordinates) that is a prerequisite for this effect. Finally, the response times in Experiment 5 suggested that this segregation is established spontaneously even when processing capacity is restrained.

It is important to note that segregation requires that people actively categorize the targets of an actor's behavior (or other aspects of the behavioral field), just as they categorize actors when they receive information about the behavior of different group members. In these experiments, the behavioral targets were denoted by different labels such as *secretary, mail courier, chairman,* and *administrator.* The results from Experiments 4 and 5 imply that participants spontaneously categorized the targets as subordinates or superiors. Similarly, it seems plausible that perceivers often categorize other situational cues, for instance, cues denoting that a behavior occurs at work versus at home (Vonk, 1993).

In general, when perceivers are attempting to explain an inconsistency by means of impression differentiation, they may consider some categorizations more readily than others, depending on whether they have social schemas for particular distinctions. The licking up–kicking down schema is an obvious candidate, but there may be others, for instance, the schema of people who are reserved and detached toward strangers, but friendly and spontaneous when they get to know someone better; or the introverted and quiet office clerk who lives it up at wild parties. Some distinctions are implied in particular types in our culture, for instance, the Mafia man who is tough and ruthless in business, but sentimental toward his family members (cf. Asch & Zukier, 1984).

Correspondence bias and ulterior motivation. Generally, the present experiments demonstrate that information about an actor's dependence on a target person is a powerful antidote to the correspondence bias. As noted by Gilbert and Malone (1995), overcoming the correspondence bias requires first of all that the perceiver is aware of the actor's situation and has an adequate understanding of the effects of this situation on behavior. It seems that the slime schema of social perceivers enables them to quickly recognize the power of behavioral targets and grasp the effects of this power on the actor's behavior, without much elaboration. This does not imply that the actual correction required to complete the process described by Gilbert and Malone occurs effortlessly as well. It seems that perceivers, once they become aware of the target's power and its possible effects on the actor, enter a state of suspicion characterized by sophisticated attributional thinking (Fein, 1996). It is this thoughtful analysis that leads them to discount the actor's behaviors toward superiors in their likeability ratings, thus producing either a neutral evaluation (when disambiguating information is absent) or a negative evaluation (based solely on the actor's behaviors toward subordinates). Thus, the Slime effect may be regarded as a combined result of both quick-and-dirty, schematic processing, resulting in a tentative categorization of the actor as a brownnoser, and thoughtful deliberations that are initiated upon the activation of the slime category.

As argued by Fein and his colleagues (Fein et al., 1990; Hilton et al., 1993), these deliberations are aimed at identifying information that could disambiguate the meaning of the suspect behaviors. In their studies, participants received a single piece of highly diagnostic information that instantly confirmed their skepticism. This could suggest that people are either in or out of a suspicion state and that the distinction is discrete. However, the results of Experiment 3 suggest that the distinction may often be gradual: Across one, two, and four dislikeable behaviors toward subordinates, suspicious participants became increasingly confident that the actor's likeable behaviors toward superiors were driven by ulterior motivation.

Impression formation and impression management. As noted by Jones (1964, 1990), ingratiation is probably the most common form of self-presentation; as a result, there is an abun-

[8] This description of the results was given by an anonymous reviewer.

dance of cultural wisdom about ingratiation goals and strategies. Our cultural knowledge of brownnosing is illustrated by the sheer number of words we have at our disposal to refer to the type of person who engages in this behavior (e.g., *bootlicker, toady, flunky, soft-soaper, sycophant,* and several informal descriptions that point to the cause of the brown nose). In addition, it may be assumed that social perceivers have developed all sorts of other social schemas that are relevant to ingratiation. A person's dependence on others can be inferred by a variety of cues, such as hierarchical position labels (as used in the present experiments) and information about people's needs and goals. We have articulate expectancies about the behavior of people who depend on someone. Reversely, when a person is extraordinarily friendly toward us we may wonder "Does this person want something from me?"

Nevertheless, ingratiation is not the only form of impression management, and it seems that person perceivers are quite familiar with many other varieties. For instance, perceivers know that people who "play hard to get" may be hiding quite the opposite; that people present themselves as overly busy when they want to avoid a person; that people play down their effort when they succeed at a task so that it will be attributed to high ability; that people behave ignorantly when they are caught on a train without a fare ticket; and that people express confidence about issues they are really not sure about in order to persuade others. In all of these cases, perceivers may face the dilemma of deciding whether the behavior reflects a correspondent trait or an ulterior motive.

In the impression management literature, many of these varieties of self-presentation have been examined (for reviews, see, e.g., Leary, 1995; Schlenker, 1980). In the large majority of these studies, however, the focus was on variables that affect self-presentational behavior, not on variables that affect how such behavior is perceived. The person perception literature has given even less attention to this issue. And yet, it seems likely that person perceivers consider the possibility of self-presentation on a daily basis. Thus, knowledge of self-presentational behavior may affect many inferences about others. In addition, the study of perceptions of self-presentational behavior can have interesting implications for current theories on impression formation and trait inferences. The impression management literature provides a wealth of material on self-presentational behavior. The person cognition field would gain in both depth and scope by examining what person perceivers know about this aspect of behavior, and how they use this knowledge in making trait inferences and forming impressions.

References

Adorno, T., Frenkel-Brunswik, E., Levinson, D., & Sanford, N. (1950). *The authoritarian personality.* New York: Harper.

Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology, 41,* 258–290.

Asch, S. E., & Zukier, H. (1984). Thinking about persons. *Journal of Personality and Social Psychology, 46,* 1230–1240.

Fein, S. (1996). Effects of suspicion on attributional thinking and the correspondence bias. *Journal of Personality and Social Psychology, 70,* 1164–1184.

Fein, S., Hilton, J. L., & Miller, D. T. (1990). Suspicion of ulterior motivation and the correspondence bias. *Journal of Personality and Social Psychology, 58,* 753–764.

Fiske, S. T. (1993). Controlling other people: The impact of power on stereotyping. *American Psychologist, 48,* 621–628.

Gilbert, D. T., & Malone, P. S. (1995). The correspondence bias. *Psychological Bulletin, 117,* 21–38.

Gilbert, D. T., Pelham, B. W., & Krull, D. S. (1988). On cognitive busyness: When person perceivers meet persons perceived. *Journal of Personality and Social Psychology, 54,* 733–740.

Godfrey, D., Jones, E. E., & Lord, C. (1986). Self-promotion is not ingratiating. *Journal of Personality and Social Psychology, 50,* 106–115.

Gordon, R. A. (1996). Impact of ingratiation on judgments and evaluations: A meta-analytic investigation. *Journal of Personality and Social Psychology, 71,* 54–70.

Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology, 12,* 392–407.

Hastie, R. (1980). Memory for behavioral information that confirms or contradicts a personality impression. In R. Hastie, T. M. Ostrom, E. B. Ebbesen, R. S. Wyer, Jr., D. L. Hamilton, & D. E. Carlston (Eds.), *Person memory: The cognitive basis of social perception* (pp. 155–177). Hillsdale, NJ: Erlbaum.

Hastie, R., & Park, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory-based or on-line. *Psychological Review, 93,* 258–268.

Hilton, J. L., Fein, S., & Miller, D. T. (1993). Suspicion and dispositional inference. *Personality and Social Psychology Bulletin, 19,* 501–512.

Jones, E. E. (1964). *Ingratiation.* New York: Appleton-Century-Crofts.

Jones, E. E. (1990). *Interpersonal perception.* New York: Freeman.

Jones, E. E., & Davis, K. E. (1965). From acts to dispositions: The attribution process in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 220–266). New York: Academic Press.

Kelley, H. H. (1973). The process of causal attribution. *American Psychologist, 28,* 107–128.

Kipnis, D., & Vanderveer, R. (1971). Ingratiation and the use of power. *Journal of Personality and Social Psychology, 17,* 280–286.

Leary, M. R. (1995). *Self-presentation: Impression management and interpersonal behavior.* Madison, WI: Brown & Benchmark.

McArthur, L. Z., & Friedman, S. A. (1980). Illusory correlation in impression formation: Variations in the shared distinctiveness effect as a function of the distinctive person's age, race, and sex. *Journal of Personality and Social Psychology, 39,* 615–624.

McConnell, A. R., Leibold, J. M., & Sherman, S. J. (1997). Within-target illusory correlations and the formation of context-dependent attitudes. *Journal of Personality and Social Psychology, 73,* 675–686.

McConnell, A. R., Sherman, S. J., & Hamilton, D. L. (1994). On-line and memory-based aspects of individual and group target judgments. *Journal of Personality and Social Psychology, 67,* 173–185.

Mullen, B., & Johnson, C. (1990). Distinctiveness-based illusory correlations and stereotyping: A meta-analytic integration. *British Journal of Social Psychology, 29,* 11–28.

Reeder, G. D., & Brewer, M. B. (1979). A schematic model of dispositional attribution in interpersonal perception. *Psychological Review, 86,* 61–79.

Reeder, G. D., & Fulks, J. L. (1980). When actions speak louder than words: Implicational schemata and the attribution of ability. *Journal of Experimental Social Psychology, 16,* 33–46.

Reeder, G. D., & Spores, J. M. (1983). The attribution of morality. *Journal of Personality and Social Psychology, 44,* 736–745.

Schlenker, B. R. (1980). *Impression management: The self-concept, social identity, and interpersonal relations.* Monterey, CA: Brooks/Cole.

Smith, E. R., & Miller, F. D. (1983). Mediation among attributional

inferences and comprehension processes: Initial findings and a general method. *Journal of Personality and Social Psychology, 44,* 492–505.

Srull, T. K. (1981). Person memory: Some tests of associative storage and retrieval models. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 440–463.

Trafimow, D. (1994). *Are people really insensitive to situational specificity? A spreading activation account of correspondence bias.* Unpublished manuscript.

Vonk, R. (1993). The negativity effect in trait ratings and in open-ended descriptions of persons. *Personality and Social Psychology Bulletin, 19,* 269–278.

Vonk, R. (1994). Trait inferences, impression formation, and person memory: Strategies in processing inconsistent information about persons. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 5, pp. 111–149). New York: Wiley.

Vonk, R. (1997). *Impression management and impression formation: Inferences of self-promoting and self-deprecating behavior.* Manuscript submitted for publication.

Vonk, R. (1998). *Self-serving interpretations of flattery: The dilemma of being ingratiated.* Manuscript submitted for publication.

Vonk, R. (in press). Differential evaluations of likeable and dislikeable behaviours enacted towards superiors and subordinates. *European Journal of Social Psychology.*

Vonk, R., & Van Knippenberg, A. (1994). The sovereignty of negative inference: Suspicion of ulterior motives does not reduce the negativity effect. *Social Cognition, 12,* 169–186.

Vonk, R., & Van Knippenberg, A. (1995). Processing attitude statements from ingroup and outgroup members: Effects of within-group and within-person inconsistencies on reading times. *Journal of Personality and Social Psychology, 68,* 215–227.

Wyer, R. S., & Gordon, S. E. (1982). The recall of information about persons and groups. *Journal of Experimental Social Psychology, 18,* 128–164.

Received May 1, 1996
Revision received May 9, 1997
Accepted May 12, 1997 ■