

ATTITUDES AND SOCIAL COGNITION

Maintaining Stereotypes in the Face of Disconfirmation: Constructing Grounds for Subtyping Deviants

Ziva Kunda
University of Waterloo

Kathryn C. Oleson
Ohio State University

People encountering deviants who violate a stereotype try to maintain the stereotype by subtyping the deviants. They use the deviants' additional attributes to justify subtyping them. Participants read about counterstereotypic targets. Participants who were given no additional information about targets, and so had no grounds for subtyping them, did generalize from them and changed their stereotypes. However, participants who were told that targets had an additional, neutral attribute appeared to use it as grounds for subtyping them; their stereotypes remained unchanged. Participants came to view the neutral attributes as atypical of the stereotype and as associated with deviance, that is, as good reasons for subtyping the deviant. Neutral attributes blocked generalization from truly counterstereotypic targets but not from overly stereotypic ones, suggesting that their effect was due to participants' attempts to explain away individuals who strongly challenge their stereotypes.

This . . . is not a polyhedron at all. It is a monster, a pathological case, not a counterexample.
(Lakatos, 1976, p. 12)

Stereotypes are notoriously difficult to change. Positive contact with members of negatively stereotyped groups can lead to stereotype revision under certain circumstances, but research has shown that negative stereotypes often resist change even in the face of intense manipulations involving cooperation with members of the stereotyped group over extended periods of time (Stephan, 1985). Such positive contact may fail to affect stereotypes because people do not generalize from the positive members whom they have encountered to the group as a whole. Rather, when people encounter group members who violate a

group stereotype—a wealthy African-American or an aggressive housewife—they “fence off” these members by assuming that they constitute a distinct subtype of the group (Allport, 1954; Rothbart & John, 1985; Weber & Crocker, 1983). Consigning deviants to a subtype believed to be atypical and unrepresentative of the group as a whole may enable people to maintain their preexisting global stereotypes even though they are aware that deviants exist. Such subtyping may also explain how one may remain a bigot while proclaiming that “some of my best friends are (Jews, Blacks, etc.)”; the best friends are subtyped as exceptions to the group.

Subtypes constructed to accommodate deviants may become part of one's enduring belief structure. Indeed, it has been shown that stereotypes, like other concepts (Rosch, 1978), are structured hierarchically, so as to include distinct subordinate categories, or subtypes (Brewer, Dull, & Lui, 1981; Cantor & Mischel, 1979). Some subtypes may be quite unrepresentative of the overall stereotype (Devine & Baker, 1991). For example, the group Blacks may be believed to include the subtype of Black businessmen, who may be viewed as quite different from Blacks in general.

There is also evidence that the subtyping of stereotype-disconfirming group members may protect stereotypes from change. In a landmark article on this topic, Weber and Crocker (1983) showed that exposure to group members characterized by stereotype-disconfirming attributes produced less stereotype change when these attributes were concentrated in a small number of group members, who could then be subtyped, than when they were dispersed over a large number of group members, who could not be readily relegated to a subtype. In other words, when people can subtype deviants they do not generalize from

Ziva Kunda, Department of Psychology, University of Waterloo, Waterloo, Ontario, Canada; Kathryn C. Oleson, Department of Psychology, Ohio State University.

This research was supported by Grant 1 RO1 MH42239-01A1 from the National Institute of Mental Health (NIMH), Grant 410-93-0667 from the Social Sciences and Humanities Research Council of Canada, a University of Waterloo Social Sciences and Humanities grant, a National Science Foundation graduate fellowship, and by NIMH postdoctoral fellowship T32 MH 19728). Partial reports of these data were presented at the 1993 conference of the Midwestern Psychological Association and at the 1993 conference of the American Psychological Society. We are grateful to Michael Ross and Mark Zanna for their comments on an earlier version of this article.

Correspondence concerning this article should be addressed to Ziva Kunda, Department of Psychology, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1. Electronic mail may be sent via Bitnet to zkunda@watarts.uwaterloo.ca

them to the overall group and so do not change their global stereotypes. These findings have since been replicated (Johnston & Hewstone, 1992).

The process of fencing off and excluding instances that violate expectancies is not restricted to expectancies that arise from stereotypes. The philosopher of science Lakatos described how mathematicians who encounter a counterexample that challenges an established theorem may engage in a similar process as they struggle to salvage the theorem. In such cases, they set out to determine what is wrong with the counterexample, and they redefine the concept to which the theorem applies by adding ad hoc restrictive clauses that serve to exclude the counterexample. Lakatos gave this process the colorful name of "monster barring"—adding restrictive clauses that permit one to regard the monstrous counterexample as an exception to which the trusty theorem need not apply (Lakatos, 1976).

We propose that subtyping, like monster barring, may result from a constructive attempt to justify dismissing individuals who violate a stereotype so as to salvage the stereotype. People are reluctant to discard established stereotypes simply because they have encountered a deviant, much like Lakatos' (1976) mathematicians who proclaim that "We do not allow wayward counterexamples to destroy respectable proofs at liberty" (p. 29). Rather, people attempt to maintain their stereotypes and the expectancies to which they give rise, just like they attempt to confirm any other expectancy (Snyder, 1984).

When people encounter a person who violates their stereotype of that person's group, they are surprised, much like they are when any other expectancy that they hold is violated. This surprise may trigger an attempt to explain away the violation (Hastie, 1984; Kunda, Miller, & Claire, 1990; Pyszczynski & Greenberg, 1981; Wong & Weiner, 1981). To this end, people may form hypotheses about possible factors that could make the person atypical of his or her group and therefore irrelevant to it: Does the person's age, family background, religion, or any other attribute make him or her an exception that can be dismissed as unrepresentative of the group? The strategies that people use to test such hypotheses may be inadvertently biased toward confirming these hypotheses (Klayman & Ha, 1987; Snyder & Swann, 1978). Therefore, people may come to believe that they have good reason for dismissing the deviant as irrelevant and that they need not change their stereotypes.

The search for ways of maintaining stereotypes in the face of violation may be fueled not only by the surprise triggered by that violation but also by a motivation to hold on to beliefs that serve important needs. People may be driven to preserve their stereotypes because they use the stereotypes to justify their social order, their sense of superiority to others, or their own behavior (Allport, 1954; Sartre, 1946/1948). However, people may attempt to dismiss challenges to stereotypes even when the stereotypes serve no such needs, just like they try to explain away surprising challenges to any other expectancy. Therefore, attempts to subtype deviants are unlikely to be restricted to motivationally loaded stereotypes such as those associated with racial and ethnic groups. Indeed, earlier demonstrations of subtyping have involved stereotypes of lawyers and fraternity members (Rothbart & Lewis, 1988; Weber & Crocker, 1983).

Whatever the source of attempts to subtype deviants—be it predominantly cognitive or motivational factors—to succeed at such attempts, people must be able to justify their reasoning. Much like mathematicians, who need to articulate restrictive clauses to their concepts so as to justify excluding counterexamples, people encountering individuals who violate their group stereotypes feel compelled to find reasons for believing that these deviants are unrepresentative of the group. They search for grounds for subtyping the deviants.

The notion that people feel compelled to justify subtyping deviants has gained support from research on motivated reasoning. It appears that people who want to arrive at a particular conclusion attempt to be rational and to construct justifications of their desired conclusion that would persuade a dispassionate observer. They draw the desired conclusion only if they can justify it (Sanitioso, Kunda, & Fong, 1990). However, the process of justification construction itself is biased by motivation (Kunda, 1987, 1990; Pyszczynski & Greenberg, 1987). The same should be true for people attempting to dismiss deviants who challenge their stereotypes: People should be able to maintain their stereotypes in the face of disconfirming individuals only if they can justify subtyping (and thus dismissing) these individuals. However, people should be capable of creatively justifying the dismissal of such deviants in a biased manner much as they justify the dismissal of evidence that challenges their desired conclusions. People are remarkably adept at generating theories to explain how just about any attribute may be related to just about any outcome. They can, with great ease, explain how each of two opposite attributes may cause the same outcome, and explaining such a causal relation enhances their belief in it (Anderson, Lepper, & Ross, 1980; Anderson & Sechler, 1986; Sanitioso et al., 1990). People should therefore be very good at explaining how just about any attribute that characterizes a deviant constitutes good grounds for subtyping that deviant.

The idea that people use additional attributes that characterize deviants as grounds for subtyping them is consistent with the finding that atypical deviants are more likely to be subtyped and dismissed than are otherwise typical deviants. When members violate their group's stereotype on one dimension, people are more likely to subtype these members and less likely to generalize that stereotype-disconfirming dimension to the group as a whole if these deviant members also violate the stereotype on another dimension (Rothbart & Lewis, 1988; Weber & Crocker, 1983). For example, participants who encountered several poorly dressed lawyers were less likely to modify their stereotype so as to conclude that lawyers are not very well dressed if the poorly dressed lawyers were also atypical in other ways (they were Black and had low incomes) than if they were otherwise typical (Weber & Crocker, 1983).

Earlier theorists have reasoned that people generalize disconfirming attributes (and so modify their stereotype) more from otherwise typical than from atypical group members because the atypical group members are less likely to activate the stereotype, so their disconfirming attributes are less likely to become associated with it (Johnston & Hewstone, 1992; Rothbart & John, 1985; Rothbart & Lewis, 1988). However, subtyping oc-

curs even when the category is already activated and it is clear and salient to the participants that the disconfirming individuals belong to it (as is the case in all previous research on subtyping). Therefore, the failure of atypical members to activate the global stereotype cannot fully account for their reduced impact on the stereotype.

We suggest instead that the additional atypical attributes provide a reason for subtyping the disconfirming members. It is easy to explain why Black lawyers may be different from lawyers in general—perhaps they are pro bono lawyers, perhaps they come from less affluent backgrounds—and it therefore seems reasonable to subtype them. Thus the additional, atypical attribute (Black) serves to define the subtype (Black lawyers). We propose that when people come across a member of a stereotyped category who violates their stereotype, they ask themselves, in effect: “Do I have any grounds for concluding that this person is unrepresentative of the group as a whole?” They then attempt to use any additional information they have about this person to come up with a positive answer. If they succeed, they subtype the individual, and their overall stereotype remains unchanged. If they fail to justify subtyping, however, they do generalize from the individual, thereby changing their stereotype.

If this is the case, the additional attribute said to characterize the deviant need not be atypical of the stereotype to begin with. Even neutral attributes should facilitate subtyping and block generalization from deviants to stereotypes, because people will try to construct accounts of how even such previously neutral attributes provide grounds for subtyping deviants. To explore this possibility we first examined, in Study 1, whether persons who violate a stereotype on a given dimension are more likely to be subtyped if they also are characterized by an additional, previously neutral attribute that may serve as grounds for subtyping them. If the previously neutral attribute is used as grounds for subtyping the deviants characterized by it, then it should come to be viewed as atypical of the stereotyped group and as associated with deviance. We examined this possibility in Studies 2 and 3. If a previously neutral attribute blocks generalization from deviants to stereotypes because of a person’s constructive attempt to justify subtyping the deviants, then it should have this blocking effect only when the deviants are truly counterstereotypic. We addressed this question in Study 4. Whereas in the first three studies we used the relatively mundane stereotype of lawyers as extraverted, in Study 4 we focused on the more affect-laden stereotype of gay men as promiscuous.

Study 1

People apparently believe that they should generalize the behavior of even a single group member to the group as a whole, and they have been shown to do so in several studies (Hamill, Wilson, & Nisbett, 1980; Nisbett, Krantz, Jepson, & Kunda, 1983; Quattrone & Jones, 1980; Rothbart & Lewis, 1988). Therefore, people encountering a member of a stereotyped group who violates their stereotype of that group may feel compelled to generalize from this person to the group and to change their stereotype of that group if they cannot find reasonable grounds for subtyping the disconfirming member. However, people may be able to creatively construct accounts of how just

about any attribute of the deviant renders the deviant atypical of the group. To this end, they may come to view any information about the deviant as grounds for subtyping him or her.

For example, people exposed to an assertive woman who violates their stereotype of women as compliant and unassertive may revise their stereotype of women if they have no additional information about this stereotype-disconfirming woman. However, if these people also know that the woman had (or did not have) brothers, that her parents were supportive (or unsupportive), that she was attractive (or unattractive), or if they have any other kind of background information about her, they may attempt to use this information to explain how women with this attribute, unlike most women, come to be assertive. Thus, they will create a subtype of women, including, for example, women who did (or did not) have brothers, and they will view this subtype as different from and atypical of women in general. They may theorize, for example, that women with brothers are assertive because they learn to emulate their brothers, or that women without brothers are assertive because they are not suppressed at home by dominating males. Having relegated the assertive woman to this subtype, they need not generalize from her to women in general. Their global stereotype may remain unchanged.

To explore these ideas, in Study 1 we examined the extent to which people generalize from a member of a stereotyped group who disconfirms the group stereotype to the global stereotype of that group. Is such generalization, or stereotype change, reduced when participants are also given an additional attribute that they can use as grounds for subtyping the disconfirming individual? We focused on the stereotype of lawyers, shown in pretests to be viewed as quite extraverted. Participants were exposed to an introverted lawyer, who challenged this stereotype. Some participants were given no additional information about this introverted lawyer. We expected them to generalize from him and come to see lawyers in general as more introverted than would control participants. Other participants were given additional information about this lawyer. To ensure that this additional information was not associated *a priori* with lawyers’ introversion, it constituted one of two opposite attributes: The lawyer was said to work for a large firm or for a small firm. Pretests showed both firm sizes to be neutral with regards to lawyers’ introversion or extraversion. To prevent participants from merely dismissing the target as a poor lawyer, the interview also conveyed that he was quite competent. Because pretests indicated that lawyers are stereotyped as competent, portraying the target as typical on this dimension may have increased participants’ belief that they should generalize from him (Rothbart & Lewis, 1988; Weber & Crocker, 1983), and consequently increased the pressure to subtype him.

We expected that participants would be able to use the additional information about firm size to subtype the introverted lawyer. It is easy to explain why lawyers who work for small firms can get by despite being introverted; one could argue, for example, that they need to deal with fewer people. It is just as easy to explain why lawyers who work for large firms can get by despite being introverted; one could argue, for example, that large firms permit division of labor, so they can avoid tasks requiring social skills. In either case, having explained why law-

yers who work for the firm size in question may be particularly likely to be introverted, one may now consign the introverted lawyer to a subtype that is atypical of lawyers in general. One therefore need not generalize from him and change one's stereotypes of lawyers.

Method

Assessing the Stereotype

The stereotype of lawyers was determined from survey data used by Kunda, Miller, and Claire (1990). These authors asked 33 participants drawn from the same participant pool that we used in the present study to rate the likelihood that lawyers had various attributes, on a 7-point scale ranging from 1 (*not at all likely*) to 7 (*very likely*). Lawyers were rated as highly likely to be extraverted ($M = 5.15$). They were also rated as high on traits implying competence and professional success (intelligent, $M = 5.70$, and affluent, $M = 6.03$) and as low on a traits implying incompetence (underachiever, $M = 2.06$). Thus, it appears that lawyers are viewed as extraverted and competent.

Participants

Participants were 70 Princeton University undergraduates of both sexes who participated for pay. Five participants were excluded from the analyses for disbelieving the cover story, and 1 was excluded for failing to recall the manipulation, leaving a total of 64 participants. Participants were randomly assigned to conditions. There were 16 participants in each condition.

Procedure

Experimental participants read a cover sheet explaining that the investigators were examining whether summaries of interviews can convey the same impressions of the interviewees as do full transcripts of the interviews. All participants were told that they had been assigned to read a transcript of an interview. They then read a two-page transcript of an interview with a man named Steve. The interview began with questions about social behavior, and Steve's responses indicated that he was quite introverted. For example, the transcript indicated that he said "I sometimes meet people and don't know what to say beyond hello, so I don't say anything more. Sometimes when people ask me questions I get nervous and don't know what to say, so I just answer with a single word." In response to another question, Steve's reply was "A couple of weeks ago, when I was sitting with this large group of people I didn't know, . . . I was trying to talk to the person next to me. It was difficult. My mind was blank, I couldn't think of anything to say."

The manipulation was introduced in the second-to-last question in the interview. At this point the interviewer asked Steve about his career. In the no-information-about-firm-size condition, Steve's response was only that he is a lawyer and likes his job. In the small-firm and large-firm conditions, he added that he works for a small or large firm. A pretest indicated that lawyers were believed to be about equally distributed among small firms (32%), moderate-sized firms (32%), and large firms (36%).

A different group of 39 participants indicated that small-firm and large-firm lawyers were not viewed as different from lawyers in general or from each other on the introverted and extraverted attributes used as dependent measures in this study. These participants rated how either lawyers working for large firms or lawyers working for small firms compared with lawyers in general on each of the attributes. In each case, the lead-in was: "Compared to other lawyers, lawyers working for large [small] firms are likely to be:" and the scale ranged from -3 (*much less*

[*introverted*] than average) through 0 (*as [introverted] as average*) to +3 (*much more [introverted] than average*). We averaged the attributes into a single measure (reversing the extraverted ones). The mean ratings for small-firm lawyers ($M = -0.20$) and for large-firm lawyers ($M = -0.15$) did not differ from each other, and neither rating differed from 0 (all $ps > .25$). Thus, participants believed, a priori, that small-firm lawyers and large-firm lawyers did not differ from lawyers in general in their introversion.

We were concerned that participants might merely dismiss Steve as a poor lawyer. To avoid this possibility, the interview also conveyed that he was a competent lawyer. The final question asked Steve how he was doing at his job. He responded that he had recently had a very positive annual review and had been given a large raise.

Control participants rated lawyers without having read the interview. They read the same cover story as did experimental participants but were informed that they were not participating in the experiment. Rather, their input was needed because one of the interviewees mentioned being a lawyer. Because people's views about careers may affect their impressions of individuals who have these careers, the experimenters were interested in their beliefs about lawyers in general. They then responded to the same questions about lawyers as did experimental participants.

Dependent measures. Experimental participants first answered four questions about how interesting and informative they found the interview, to bolster their belief in the cover story. Next they read that "People's general views about different careers may affect their impressions of individuals who have these careers. In order to determine if that is the case with your impressions of Steve, we would like to know your beliefs about lawyers in general." Participants then rated the extent to which each of 22 traits characterized lawyers, on average. Among these were embedded 10 extraverted and introverted traits. The ratings were made on 11-point scales ranging from 1 (*not at all*) to 11 (*extremely*). Experimental participants then rated Steve, the interviewee, on the same measures. As a manipulation check, participants were next asked to indicate what sort of firm Steve worked for (by checking "small," "large," or "don't know"). Control participants, who had not read the interview, were given only the questionnaire assessing beliefs about lawyers in general. Finally, participants were probed for suspicion and carefully debriefed about the true purpose of the study.

Results

Identifying Items That Yielded Generalization

Our central hypothesis was that participants who were informed that the introverted lawyer about whom they had read worked for a small or a large law firm would use this information to subtype him. Thus, information about firm size would block generalization from the introverted lawyer to lawyers in general. The question of whether information about firm size could block generalization can be assessed only by examining those items that yielded generalization in the first place, in the absence of this information. It was therefore necessary first to identify those items for which participants who had not been informed about firm size generalized from the target lawyer to lawyers in general and changed their stereotypes of lawyers.

To identify items that yielded generalization, we compared ratings of lawyers' extraversion-introversion made by participants in the no-information-about-firm-size condition with those made by participants in the control condition, who had not been exposed to the introverted lawyer. For 9 of the 10 extraverted and introverted items, no-information participants

rated lawyers in general as more introverted (or less extraverted) than did controls, and the two groups made identical ratings on the remaining item. We focused our analyses on the top half of the items, that is, the 5 items that yielded most generalization, as measured by effect size. Effect sizes for these 5 high-generalization items ranged from 0.39 *SDs* to 0.74 *SDs*, with an average of 0.56 *SDs*. Effect sizes for the remaining 5 low-generalization items ranged from 0 to 0.28 *SDs*, with an average of 0.15 *SDs*.¹

Impressions of Lawyers, in General

We averaged the high-generalization items to form a single measure of introversion, after first reversing the extraverted items (Cronbach's $\alpha = .69$). As can be seen in Figure 1, information about firm size did block generalization from the target lawyer to lawyers in general. Whereas no-information participants rated lawyers as more introverted than did controls, small- and large-firm participants did not—the ratings of both these groups were almost identical to those of controls. A one-way analysis of variance (ANOVA) revealed a significant overall effect for condition, $F(3, 60) = 3.20, p < .05$. A planned contrast in which the no-information condition (contrast weight $[C] = +3$) was pitted against the other three conditions (each $C = -1$) was significant, $F(1, 60) = 9.58, p < .01$. Dunnett's control mean *t* test revealed that the no-information condition also differed significantly from the small-firm condition, $t(60) = 2.50$, and from the large-firm condition, $t(60) = 2.63$, as well as from controls, $t(60) = 2.46$, all *ps* $< .05$.²

Providing participants with information about the size of the law firm for which the introverted lawyer worked appears to have completely eliminated any generalization from this stereotype-disconfirming lawyer to lawyers in general.

Impressions of the Target

We averaged the ratings of the target on the five high-generalization items to form a single measure of introversion, after first reversing the extraverted items (Cronbach's $\alpha = .76$). (These ratings were obtained only from participants in the three experimental conditions, who had read the inter-




Figure 1. Beliefs about lawyers' introversion for participants in each condition. (No info = no information.)

view with the target.) All three groups of participants rated the target as quite introverted (*Ms* = 8.28, 7.84, and 7.80, respectively, for no-information, small-firm, and large-firm participants). A one-way ANOVA revealed, as expected, that the conditions did not differ significantly on these ratings, $F(2, 45) = .64$. A contrast pitting the no-information condition ($C = +2$) against the other two (each $C = -1$) was nonsignificant as well, $F(1, 45) = 1.26, p > .25$.

Study 2

Study 1 showed that when participants had no grounds for subtyping a stereotype-disconfirming group member, they did generalize from him and changed some aspects of their stereotypes. However, when participants believed that the disconfirming member was characterized in addition by one of two previously neutral and opposite attributes, they apparently were able to use either one as grounds for subtyping. In either case they no longer generalized from the deviant to his group. We hypothesized that the additional attribute said to characterize the introverted lawyer—the size of the law firm for which he worked—blocked generalization from him because participants came to believe that lawyers characterized by this additional attribute were atypical of lawyers in general. Such beliefs could permit participants to subtype and dismiss the deviant because people feel less compelled to generalize from atypical than from typical group members to their group (Rothbart & Lewis, 1988; Weber & Crocker, 1983) or to other group members (Rips, 1975).

However, Study 1 provided only indirect support for this process, because we did not assess participants' beliefs about the additional attribute; we inferred that this attribute came to be viewed as atypical and to be used as grounds for dismissing the deviant from the finding that when it was added to the description of the deviant, participants no longer generalized from the deviant to the stereotype. In Study 2 we sought more direct evidence for this process by examining the effects of adding a neutral attribute to the description of a deviant on the extent to which that attribute was subsequently viewed as typical of the stereotype. As in Study 1, participants read an interview with a lawyer who violated the stereotype of lawyers as extraverted by coming across as quite introverted. Half of the participants were told in addition that this lawyer worked for a small law firm, and half received no information about firm size. We expected that participants who were told that the introverted lawyer worked for a small firm would come to view small-firm lawyers as less typical of lawyers in general than would participants who received no information about firm size.

¹ The included items were: awkward, extraverted, introverted, outgoing, shy. The excluded items were: friendly, quiet, reserved, sociable, talkative.

² The average ratings for the five excluded items were: control, $M = 4.31$; no information, $M = 4.52$; small firm, $M = 4.09$; large firm, $M = 4.62$. We subjected these items to the same analyses carried out on the included items, but they yielded no significant results.

Method

Participants

Participants were 27 University of Waterloo undergraduates of both sexes enrolled in Introductory Psychology. They participated to fulfill a course requirement. One participant was excluded from the analyses for disbelieving the cover story, and 3 were excluded for failing to recall their condition, leaving a total of 23 participants. Participants were randomly assigned to conditions.

Procedure

Participants read a cover sheet explaining that the study was intended to examine whether people who watched a videotaped interview with another person would view that person differently than would people who read a transcript of that interview. All participants were informed that they had been randomly assigned to read a transcript of an interview. All participants read the interview with Steve used in Study 1, in which Steve is portrayed as an introverted lawyer. Half of the participants ($n = 11$) read that Steve worked for a small firm, and half ($n = 12$) received no information about firm size.

Dependent Measures

Participants first read that general beliefs about different professions may affect one's impressions of that profession, and that they would therefore be asked for some of their beliefs about lawyers in general first. To ensure that no-information participants would not find the measure to be out of place, the instructions added that "Some of the questions may not be relevant to the particular interview that you have seen, but please answer them anyway as best you can." Participants first responded to the question "How typical are lawyers who work for small firms of lawyers in general?" on a scale ranging from 1 (*not at all*) to 11 (*extremely*). They then rated the personality of the target, Steve, on the five introversion-extraversion items that had yielded high generalization in Study 1. These were embedded among five traits unrelated to extraversion or introversion. Finally, participants completed the same manipulation check used in Study 1, were probed for suspicion, and were debriefed.

Results

As in Study 1, we averaged the five items pertaining to the target's extraversion-introversion, after first reversing the extraverted ones. Both small-firm ($M = 8.02$) and no-information participants ($M = 7.30$) viewed the target as quite introverted, and the two groups did not differ significantly from each other, $F(1, 21) = 1.21, p > .25$.

As expected, participants who had been informed that the introverted lawyer worked for a small firm came to view lawyers working for small firms as considerably less typical of lawyers in general ($M = 4.34$) than did participants who had received no information about firm size ($M = 6.42$) $F(1, 21) = 7.72, p = .01$. Note that the typicality ratings obtained from no-information participants were very close to the midpoint of the scale (6), which suggests that, a priori, lawyers working for small firms were viewed as neither typical nor atypical of lawyers in general.

It appears that when a deviant is characterized by a neutral attribute, that attribute ceases to be neutral. Rather, it comes to be viewed as atypical of the group. Because people do not feel

compelled to generalize from atypical group members to the group (Rothbart & Lewis, 1988; Weber & Crocker, 1983), such reductions in the perceived typicality of the deviant's additional attributes permit one to dismiss the relevance of this deviant to the group. One need not change one's stereotype of a group if the individual who challenged the stereotype is atypical of that group.

Study 3

Study 2 showed that neutral attributes characterizing a deviant come to be viewed as atypical of the deviant's group. We next examined whether such attributes also come to be viewed as particularly likely to be associated with the observed deviance. This may occur because to justify dismissing the relevance of a person who challenges a stereotype to that stereotype, it may be useful to conclude not only that this deviant belongs to an atypical subtype, but also that the subtype in question is atypical precisely along the dimension on which the deviant challenges the stereotype. Thus, in attempting to explain how the stereotype can be correct despite the existence of the deviant, people may theorize that the subtype to which the deviant belongs is particularly likely to have the stereotype-disconfirming attributes that characterize the deviant. Accordingly, they may conclude not only that, like the introverted lawyer they had encountered, lawyers who work for a small (or large) firm are atypical of lawyers in general, but also that they are particularly likely to be introverted. In other words, the deviant's additional, formerly neutral attribute may become associated with deviance and so be used to explain away the deviant.

To explore these ideas, we exposed participants to an introverted lawyer. We told participants that this lawyer worked either for a large firm or for a small firm, or gave them no information about firm size. We expected that participants given information about firm size would come to view lawyers working for that firm size as relatively more likely to be introverted. We also assessed beliefs about the introversion of lawyers in general, so as to replicate our finding in Study 1 that information about firm size blocks generalization from the introverted lawyer to lawyers in general.

Method

Participants

Participants were 72 University of Waterloo undergraduates of both sexes who were enrolled in Introductory Psychology. They participated to fulfill a course requirement. Four participants were excluded from the analyses for disbelieving the cover story, and 2 were excluded for failing to recall their condition, leaving a total of 66 participants. Participants were randomly assigned to conditions.

Procedure

The cover story and procedure were the same as in Study 2. All participants read the interview with Steve, the introverted lawyer. Participants read that Steve worked for a small firm, or for a large firm, or received no information about firm size. They were then asked for their beliefs about how lawyers who worked for small firms and lawyers who worked for large firms might differ in their personalities. They rated

each of 11 attributes on a scale ranging from -5 (*small-firm lawyers are more likely to have this attribute than large-firm lawyers*) through 0 (*small- and large-firm lawyers do not differ on this attribute*) to $+5$ (*large-firm lawyers are more likely to have this attribute than small-firm lawyers*). The list of attributes included the five introverted and extraverted traits shown in Study 1 to yield high generalization from the target to the stereotype. The remaining attributes were unrelated to extraversion or introversion. On a separate page, participants rated the extent to which lawyers, on average, were characterized by each of these attributes, on the same 11-point scales used in Study 1. The order of these two measures was counterbalanced. Finally, all participants rated the target on these measures.

In sum, the design was 3 (firm size: small, large, or no information) $\times 2$ (order: beliefs about firm size first or beliefs about lawyers first). Cell sizes ranged from 10 to 13.

Results

Beliefs About Firm Size

We averaged beliefs about the extent to which the extraverted and introverted items were relatively more (or less) likely for large- than for small-firm lawyers, after reversing the extraverted attributes (Cronbach's $\alpha = .81$). The means of this measure are presented in Figure 2. As expected, in comparison to participants who received no information about firm size, participants who were told that the target worked for a small firm believed that introversion was relatively more likely for small-firm lawyers, whereas participants who were told that the target worked for a large firm believed that introversion was relatively more likely for large-firm lawyers. This pattern was obtained under both order conditions. A 3 (firm size) $\times 2$ (order) ANOVA revealed that the effect of firm size was significant, $F(2, 62) = 6.15, p < .01$. The order effect was not significant, though there was a tendency for participants who provided beliefs about firm size first to rate introversion as more likely for small-firm law-

yers than did participants who rated lawyers first, $F(2, 62) = 2.15, p < .15$. The interaction did not approach significance ($p > .50$). Planned comparisons revealed that the difference between small- and large-firm lawyers was significant, $F(1, 62) = 11.79, p < .01$. Small-firm participants believed that introversion was more likely for small-firm lawyers ($M = -1.45$) than did no-information participants ($M = -0.35$), $F(1, 62) = 5.17, p < .01$. Large-firm participants believed that introversion was more likely for large-firm lawyers ($M = 0.13$) than did no-information participants, but this difference was not significant ($p = .25$). We were surprised that the effect appeared stronger when the lawyer was said to work for a small firm than when he was said to work for a large firm, because both types of information had been shown in Study 1 to be equally effective at blocking generalization from the introverted lawyer. Perhaps, when one is forced to explicitly compare small- and large-firm lawyers with each other, as participants were in this study, it is easier to explain why small-firm lawyers would be more introverted than large firm lawyers than it is to explain the reverse pattern. This difficulty may not have arisen in Study 1 because in that study, participants considered either small or large firms, without having to compare the two with each other.

It appears that information about the size of a firm for which the introverted lawyer worked can affect beliefs about the relation between lawyers' firm size and their introversion, especially when the firm is said to be small. Participants came to view lawyers who worked in a firm of that size as relatively more introverted than other lawyers. Such a belief could serve to justify not generalizing from that introverted lawyer to lawyers in general.

Beliefs About Lawyers in General

As in Study 1, we averaged the five items pertaining to the extraversion-introversion of lawyers, in general, after first re-

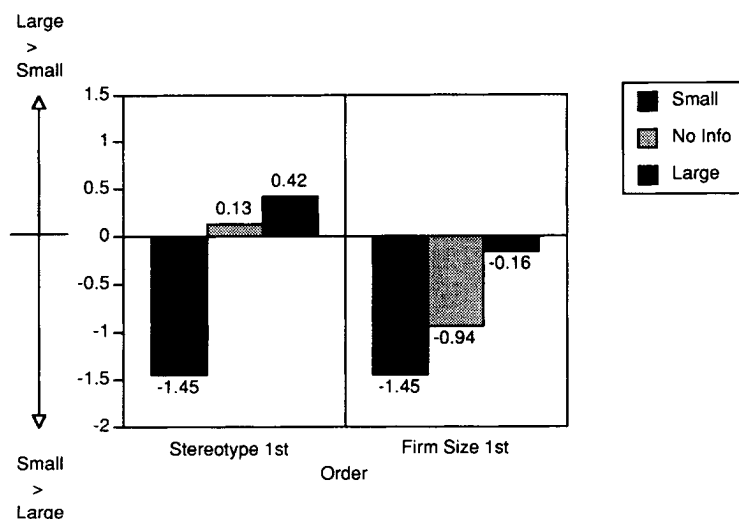


Figure 2. Beliefs about the extent to which introversion is more (or less) likely for large- than for small-firm lawyers as a function of information on firm size and order of questionnaires. (No info = no information.)

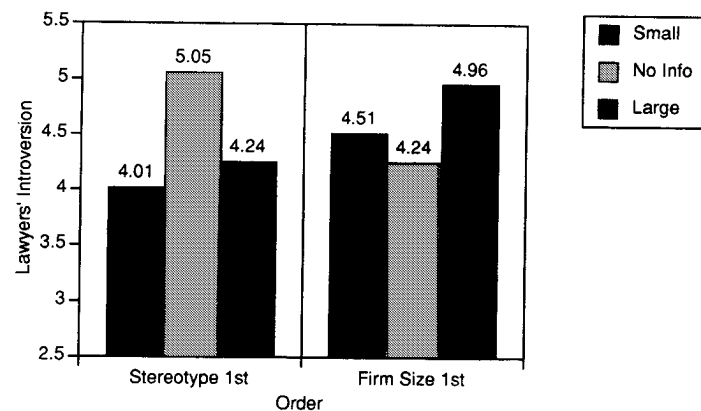


Figure 3. Beliefs about lawyers' introversion as a function of information on firm size and order of questionnaires. (No info = no information.)

versing the extraverted items (Cronbach's $\alpha = .67$). The means of this measure are presented in Figure 3. A 3 (firm size) \times 2 (order) ANOVA revealed a significant interaction, $F(2, 62) = 3.14, p = .05$. The main effects were not significant (both F s < 1). As can be seen in Figure 3, when participants provided beliefs about lawyers in general first, the results replicated those obtained in Study 1: Large- and small-firm participants believed lawyers in general to be less introverted than did no-information participants. In other words, participants who were given information about firm size were less likely to generalize from the introverted lawyer to lawyers in general than were participants who were given no such information. A planned contrast in which the no-information condition ($C = +2$) was pitted against the other two conditions (each $C = -1$) was significant for these participants, $F(1, 62) = 5.67, p < .05$.

However, when we assessed beliefs about lawyers in general after beliefs about firm size, the above effect was not obtained. In this case, ratings made by no-information participants were somewhat lower than those made by participants in the other two conditions, but the contrast testing this difference did not approach significance, $F(1, 62) = 1.41, p = .24$. Apparently, these ratings were contaminated by the ratings of the relations between introversion and firm size that preceded them. We can only speculate on why the contamination occurred. Contrasts testing order effects within each information condition revealed that none were significant alone, but the effects in the large-firm condition ($p = .12$) and in the no-information condition ($p < .09$) were marginal and larger than in the small-firm condition ($p > .25$). As can be seen in Figure 3, large-firm participants tended to generalize more from the target after they had answered the question about firm size than before. Perhaps, as suggested above, when participants were focusing only on large-firm lawyers they were able to explain why this type of lawyer could be particularly introverted, and so they did not generalize from him. When participants were first required explicitly to compare small- and large-firm lawyers to each other, however, they found it difficult to justify viewing large-firm lawyers as more introverted than small-firm lawyers, as indicated by their beliefs about the relation of firm size to introversion. Lacking

justification for subtyping the target, these participants were compelled to generalize from him. One can also see in Figure 3 that no-information participants generalized less from the target to the stereotype after they were asked about firm size than before. Perhaps the questions about different firm sizes sensitized these no-information participants to the fact that lawyers may be a diverse and variable group, and thereby blocked generalization; people generalize less from individuals to groups when the groups are viewed as more variable (Nisbett et al., 1983; Quattrone & Jones, 1980).

We had hoped that this study would allow us to also assess whether beliefs about firm size mediated beliefs about lawyers. Unfortunately, such assessment is not possible because of the order effects.

Beliefs About the Target

As in Studies 1 and 2, we averaged the five items pertaining to the target's extraversion-introversion, after reversing the extraverted items (Cronbach's $\alpha = .69$). A 3 (firm size) \times 2 (order) ANOVA revealed a significant effect for firm size, $F(2, 62) = 5.73, p = .05$. No-information participants ($M = 8.34$) rated the target as more introverted than did small-firm participants ($M = 7.53$) or large-firm participants ($M = 7.37$). The effects for order and the interaction were not significant (both F s < 1). Because beliefs about the target were assessed last, the unexpected firm-size effect may have resulted from contamination by the earlier measures.

Study 4

In Studies 1–3 we showed that when a neutral attribute is added to the description of a deviant, it blocks generalization from the deviant to the stereotype: Participants generalize from the deviant and change their stereotypes in the absence of the neutral attribute but not in its presence. We have suggested that this occurs because the neutral attribute comes to serve as grounds for subtyping the deviant: Participants use it to explain why the deviant is unrepresentative of his or her group. This

interpretation gains support from the findings, in Studies 2 and 3, that, following its association with the deviant, the neutral attribute loses its neutrality: It comes to be viewed as atypical of the stereotype and as associated with deviance.

However, one could argue that the neutral attribute blocks generalization from the deviant to the stereotype for alternative reasons that have little to do with subtyping. One possible alternative account is based on Tversky's (1977) model of similarity. That model assumes that the similarity between two objects (in this case, the deviant and the stereotype) is decreased as the number of attributes that are unique to either of them increases. Therefore, endowing the deviant with an additional, even neutral, attribute that is not included in the stereotype reduces the similarity between the deviant and the stereotype (Tversky, 1977) and so makes the deviant appear less typical of the stereotype. The reduced generalization may result from a straightforward and unbiased reduction in feature-based similarity rather than from a constructive attempt to justify dismissing the deviant.

It is also possible that the neutral attribute blocks generalization from the deviant to the stereotype because it calls participants' attention to the fact that the stereotyped group can vary on an additional dimension. This would reduce generalization from the deviant to the group because people generalize less from group members to their groups when the groups are believed to be more variable (Nisbett et al., 1983; Quattrone & Jones, 1980).

Another possible alternative account is based on the pragmatics of the experimental situation. Participants, who mistakenly believe that the experimenter's communications follow standard conversational norms, may conclude that they should use the additional, neutral information in their judgments because they assume that all information provided by the experimenter is useful (Schwarz, 1994). This account is incomplete in that it does not explain why the neutral information is used in the particular way that it is—that is, to block rather than, say, to enhance generalization. None of these accounts can readily explain why the neutral attributes also come to be seen as atypical of the stereotype and as associated with deviance, and this renders them somewhat less plausible than the constructive subtyping account that we have proposed. Nevertheless, it seemed important to rule out the alternative accounts more thoroughly. We designed Study 4 to do this.

We reasoned that if the neutral attribute blocks generalization because it is used to explain away a group member who challenges the stereotype, then such blocking should occur only when the member is truly counterstereotypic, that is, possesses a trait that is opposite to that suggested by the stereotype. The introverted lawyer described in Studies 1–3 provided such a challenge to the stereotype of lawyers as extraverted. But neutral attributes should be less likely to block generalization when the member's trait differs from the stereotypic one only by a matter of degree, without actually being its opposite. Thus, neutral attributes should not block generalization from a lawyer who comes across as even more extraverted than the stereotypic lawyer—such a person, although different from the stereotype, should not be as surprising to participants and so should not trigger subtyping attempts.

In contrast to these predictions, the similarity-based account, the variability-based account, and the pragmatics-based account all predict that neutral attributes should decrease generalization from an overly stereotypic target just as they do for a truly counterstereotypic one. From the perspective of the similarity-based account, the addition of the neutral attribute should still constitute an increase in the number of features unique to the overly stereotypic member and so should still decrease similarity between this group member and the group stereotype. From the perspective of the variability-based account, the neutral attribute should still call attention to variability, and from the perspective of the pragmatics-based account, participants should still feel just as compelled to make use of the additional, neutral information provided by the experimenter. Therefore, if, as we expect, the addition of a neutral attribute blocks generalization only when the group member is truly counterstereotypic, this will strengthen our interpretation that this effect is due to participants' use of the neutral attribute in their constructive attempts at subtyping the member rather than to the attribute's effect on the similarity between the member and the stereotype, to its role in calling attention to the group's variability, or to the importance that the attribute gains through participants' reliance on conversational norms.

Another aim of Study 4 was to expand the generality of our findings by using a different stereotype and a different kind of neutral attribute. Studies 1–3 all focused on the relatively mundane stereotype of lawyers as extraverted, which is unlikely to be imbued with much affect. One could argue that our findings are restricted to such stereotypes, because more affect- and value-laden stereotypes may be less likely to be modified in the face of challenge in the first place (Johnson & Eagly, 1989); participants may fail to generalize from deviants who challenge such stereotypes even in the absence of neutral attributes. To ensure that our work has implications for how negative stereotypes of stigmatized groups may be maintained and changed, we focused in Study 4 on the highly negative, affect-laden stereotype of gay men (Haddock, Zanna, & Esses, 1993; Herek, 1992). College undergraduates have been shown to rate homosexuals negatively, that is, below 50 on a 100-point evaluation scale—indeed, more negatively than they rate highly stigmatized ethnic groups (Esses, Haddock, & Zanna, 1993). In research conducted on the same participant population used in this study, promiscuity was one of the most-often-mentioned attributes in open-ended descriptions of beliefs about gay men, mentioned spontaneously by over one third of the 63 participants surveyed. Only 1 of these participants described gay men as monogamous.³ We therefore decided to focus on this aspect of the stereotype.

Unlike Studies 1–3, in which participants were exposed to a single target, in Study 4 we used multiple targets. This further expands the generality of our findings. Participants read interviews with eight gay men who came across as either highly unpromiscuous (that is, truly counterstereotypic) or highly promiscuous (that is, stereotypic, but more extreme than the stereotype). Half of the experimental participants were given

³ We are grateful to Meg Rohan for providing us with these data.

additional information about these gay men that could be used to subtype them: They were all said to be accountants. This information was shown in pretests to be neutral in that gay accountants were believed to be no different from other gay men in their promiscuity. The design also included a nonorthogonal control group.

We expected that participants would use the additional information about the targets' shared profession to subtype them only when the targets were truly counterstereotypic. Therefore, we expected that participants would generalize less from the unpromiscuous gay men when they were said to be accountants, and so could be readily subtyped, than when they were not. However, we expected this information to have no impact on generalization from the promiscuous gay men—participants were expected to generalize from them and come to view gay men in general as more promiscuous regardless of whether these targets were or were not said to be accountants.

Method

Participants

Participants were 59 University of Waterloo undergraduates of both sexes who participated for pay. One participant was excluded from the analyses for disbelieving the cover story, leaving a total of 58 participants (28 men and 30 women). Participants were randomly assigned to conditions, with the sexes about equally distributed in each condition.

Procedure

The study was presented as being concerned with how people form impressions of others on the basis of varying degrees of information. Participants read that they would be given excerpts from interviews with men approached while standing together at a conference on "Gay lifestyles in the 90s" who all described themselves as gay. Experimental participants then read excerpts from alleged interviews with eight men. Each was presented on a separate page, under the heading "Gay lifestyles in the 90s." Participants were asked to read the information about each target, take a moment to form an impression of him, and then turn to the next page, on which the next target was described. The information on each target included his first name and initial, sex (male in all cases), and age (target ages ranged from 25 to 47 with a mean of 37).

Profession manipulation. For participants in the accountants condition, the next piece of information about each target was an excerpt from an interview in which the interviewer asked "What do you do for a living?" and the target responded with his profession. In all cases, the profession was some form of accountant (e.g., "I just recently got a job as an accountant," "I'm an accountant for a large firm," "I run my own accounting consulting firm"). For no-information participants this excerpt was omitted.

In a pretest designed to determine whether this information was neutral, 20 participants rated whether gay men employed in each of several professions, including accountants, were likely to be more or less promiscuous than gay men, on average. They did so on a 7-point scale ranging from -3 (*much less promiscuous than average*) through 0 (*as promiscuous as average*) to $+3$ (*much more promiscuous than average*). Accountants were rated, on average, as -0.40 , which was not significantly different from 0 .

Behavior manipulation. All experimental participants then read an excerpt in which the interviewer asked "Can you tell me a bit about the intimate relationships in your life?" In the unpromiscuous condition, six of the eight targets responded by describing themselves as quite un-

promiscuous (e.g., "My partner and I have been together for over two years now. I consider myself married," "I met my partner at university and we've been together ever since," "I have been in an exclusive relationship for about 3 years"). The remaining two targets, included to bolster believability, described themselves in ways that implied higher promiscuity (e.g., "I've never really been in a long term relationship"). In the promiscuous condition, six of the eight targets described themselves as quite promiscuous (e.g., "I usually meet people in gay bars, and we strike up a relationship that can last anything from a few days to a few months," "In a typical month I have two or three different sexual partners," "I travel a lot, and usually end up meeting someone on the road and spending the night with him"). The remaining two targets described themselves as less promiscuous ("I've stopped going to gay bars since I met my current partner. I'm not looking for that kind of adventure any more").

Control participants did not read any interviews. They read the same cover story as did experimental participants, but they were informed that they were not actually participating in the experiment. Rather, their input was needed because the interviewees were known to be gay men. Control participants were told that because people's general views about different groups may affect their impressions of members of that group, the experimenters were interested in their beliefs about gay men in general. Control participants then responded to the same questions about gay men as did experimental participants.

In sum, the design was 2 (promiscuity: unpromiscuous vs. promiscuous) $\times 2$ (profession: accountant vs. no information) with an additional, nonorthogonal control group. Each cell had 11 or 12 participants.

Dependent measures. Participants read next that people's general views about different groups might affect their impressions of group members. Therefore, because the men they had read about were all gay, they were asked for their beliefs about what gay men in general were like. They then rated gay men, on average, on 14 items. These included 5 items related to being promiscuous (promiscuous, unprincipled, likely to engage in one-night stands, likely to sleep around, unethical) and 6 items related to being unpromiscuous (likely to be committed to their intimate partners, likely to have long-term relationships, moral, likely to be faithful to their intimate partners, monogamous, trustworthy). All ratings were made on 7-point scales with endpoints labeled 1 (*not at all*) and 7 (*extremely*). Experimental participants were next asked to recall their impressions of the men they had read about and rate what they were like, on average. These target ratings were on the same measures used to rate gay men in general.

Finally, experimental participants were asked whether they had been given information about the targets' occupations. If they responded yes, they were asked to list as many occupations as they could recall. All participants recalled their conditions, and all participants in the accountant condition reported correctly that all targets were accountants. Control participants, who had not read the interviews, were given only the questionnaire about gay men.

At the very end of the experiment, participants were asked to write down any hypotheses they might have developed about what the study was really about. The experimenter then queried participants about their responses, and prompted them for any suspicions about the nature of the study.

Results

Identifying Items That Yielded Generalization

To provide the fairest test of our hypothesis that neutral attributes block generalization from unpromiscuous targets but not from promiscuous ones, it was necessary to select those items on which, in the absence of neutral attributes, participants did

generalize from both the unpromiscuous and the promiscuous targets: One can only assess whether neutral attributes block generalization if it occurs in their absence. To assess generalization, we compared, for each item, ratings of gay promiscuity made by the control group with those made by each of the two groups (unpromiscuous and promiscuous) that had received no information about the targets' profession. Unpromiscuous, no-information participants rated gay men as less promiscuous than did controls on 10 of the 11 items. The remaining item yielded no difference between the groups. Promiscuous, no-information participants rated gay men as more promiscuous than did controls on 10 of the 11 items. The remaining item yielded the opposite pattern. We thus had two measures of generalization for each item: one from the unpromiscuous targets and one from the promiscuous targets. For each item we looked at the size of the smaller of these two generalization effects (regardless of whether the smaller effect was obtained for the unpromiscuous or for the promiscuous targets), and we ordered the items by the size of their weaker effects. We focused our analyses on the top half of these items, that is, the six items that yielded most generalization on the weaker of the two effects associated with them. This procedure guaranteed that the included items would all yield generalization from both sets of targets, because items that yielded little generalization from either set of targets were excluded. Effect sizes for the six included items ranged from 0.37 *SDs* to 1.17 *SDs*. The average effect size for generalization from unpromiscuous and promiscuous targets were comparable: 0.56 *SDs* and 0.54 *SDs*, respectively. The remaining, excluded items all yielded minimal generalization (less than 0.25 *SDs*) for at least one set of targets. Effect sizes for these excluded items ranged from -0.35 *SDs* to 0.77 *SDs*, with an average of 0.34 *SDs* for unpromiscuous targets and 0.28 *SDs* for promiscuous targets.⁴

Impressions of the Targets

We averaged the ratings of the targets on the six high-generalization items into a single measure of promiscuity, after first reversing the unpromiscuous items (Cronbach's $\alpha = .92$). A 2 (promiscuity) \times 2 (profession) \times 2 (sex of participant) ANOVA revealed only a significant effect for promiscuity, $F(1, 38) = 43.69$, $p < .0001$, indicating that promiscuous targets were viewed as more promiscuous than were unpromiscuous targets. Each set of targets was given practically identical ratings by participants who did and did not receive information about their profession. Accountant and no-information participants rated the unpromiscuous targets, respectively, as 2.56 and 2.60 and rated the promiscuous targets, respectively, as 4.37 and 4.41. Ratings of both the promiscuous and unpromiscuous targets differed significantly from the stereotype, as assessed by controls' ratings of gay men in general ($M = 3.62$, both $ps < .05$). Thus, the manipulation of target promiscuity was highly effective.

Impressions of Gay Men, in General

We averaged the ratings of gay men on the six high-generalization items into a single measure of promiscuity, after first

reversing the unpromiscuous items (Cronbach's $\alpha = .82$). Because the design included a nonorthogonal control group, we first conducted a 5 (condition) \times 2 (sex of participant) ANOVA, followed by planned comparisons. The ANOVA revealed main effects for condition, $F(4, 48) = 5.99$, $p < .001$, and for sex of participant, $F(1, 48) = 10.38$, $p < .01$ (men viewed gay men as more promiscuous, $M = 4.01$, than did women, $M = 3.52$). The Condition \times Sex interaction did not approach significance ($F < 1$).

As seen in Figure 4, the information that the targets were all accountants blocked generalization from the unpromiscuous, truly counterstereotypic targets but did not affect generalization from the promiscuous, overly stereotypic targets. A planned comparison revealed this 2 (promiscuity) \times 2 (profession) interaction to be significant, $F(1, 48) = 5.51$, $p < .05$. Participants who were exposed to unpromiscuous targets but were given no additional information about them generalized from these targets and came to see gay men as less promiscuous than did controls, $F(1, 48) = 4.86$, $p < .05$. However, the addition of the information that all targets were accountants sufficed to block this generalization: unpromiscuous, accountant participants did not differ significantly from controls ($F < 1$). In contrast, participants exposed to the promiscuous targets generalized from them not only in the absence of information about the targets' professions but also in the presence of such information. Both of these groups came to see gay men as more promiscuous than did controls, $F(1, 48) = 4.26$, $p < .05$, for promiscuous, no-information participants, and $F(1, 48) = 3.56$, $p = .06$, for promiscuous, accountants participants. The two groups provided almost identical ratings of gay men.⁵

The findings for unpromiscuous targets provide a conceptual replication of the findings obtained in Studies 1 and 3: The addition of a neutral attribute to the description of counterstereotypic group members blocked generalization from such members to the stereotype of their group. The results for promiscuous targets shed light on the process underlying these findings. If neutral attributes blocked generalization from counterstereotypic targets because they reduced similarity between these targets and the stereotype, because they increased the perceived variability of the group, or because participants assumed that the experimenter intended them to use this information, then neutral attributes should also have blocked generalization from the overly stereotypic targets, for the same reasons. However,

⁴ The included items were: promiscuous, likely to engage in one-night stands, likely to be committed to their intimate partners, monogamous, moral, and trustworthy.

⁵ We conducted the same set of analyses on the full set of items (i.e., including the low-generalization ones). The pattern of results was essentially the same as that obtained for the high-generalization items. The overall ratings of gay promiscuity were: control, $M = 3.48$; unpromiscuous-no information, $M = 2.99$; unpromiscuous-accountant, $M = 3.77$; promiscuous-no information, $M = 3.91$; promiscuous-accountant, $M = 3.96$. The pattern of significance was comparable as well: The Promiscuity \times Profession interaction remained significant at $p < .05$. The only difference was that, for the three groups that did generalize, the significance of the generalization was only marginal when the low-generalization items were included (ps ranged from .07 to .12).




Figure 4. Beliefs about gay promiscuity as a function of targets' promiscuity and information about their professions.

this did not occur. Neutral information had no impact on generalization from promiscuous targets; participants generalized from both sets of targets to the same extent. It appears, then, that neutral attributes block generalization only when participants attempt to subtype group members who challenge their group's stereotype. Evidently, unpromiscuous gay men triggered such attempts, whereas promiscuous ones did not.

General Discussion

In the present studies we challenged participants' well-established stereotypes by exposing them to counterstereotypic individuals. Our findings suggest that people recognize that such deviants question the accuracy of their stereotypes, but they attempt to salvage the stereotypes by subtyping the deviants. Our findings go beyond earlier work on subtyping (Rothbart & Lewis, 1988; Weber & Crocker, 1983) by suggesting that such subtyping results from a constructive and biased process of justification construction. Apparently people feel that if they can justify concluding that the deviants belong to an atypical subtype of the group, one that is particularly likely to foster deviance, they need not generalize from them to the stereotype. They therefore attempt to use any information they have about the deviants as grounds for subtyping them.

Accordingly, we found in Studies 1 and 3 that participants who were exposed to an introverted lawyer who challenged their stereotype of lawyers as extraverted did generalize from him and modified their stereotypes when they were given no additional information about him and so had no grounds for subtyping him. These participants came to view lawyers in general as less extraverted than they had previously believed. However, participants who were informed that the lawyer was also characterized by one of two previously neutral and opposite attributes—that he worked for a small or a large law firm—were apparently able to use this information as grounds for subtyping the lawyer and thus as an excuse for not generalizing from him. Their stereotypes remained unchanged. This finding was replicated in Study 4, this time with the more negative, affect-laden stereotype of gay men. Participants who were exposed to unpromiscuous gay men who challenged their stereotype of gay men as

promiscuous did generalize from them when they had no additional information about these targets; they came to view gay men in general as less promiscuous. However, participants who were also told that the unpromiscuous gay men all shared a profession they had not previously believed to be associated with promiscuity were apparently able to use this information to subtype the unpromiscuous gay men. These participants did not change their stereotypes.

We also found, in Studies 2 and 3, that when a neutral attribute was believed to characterize the deviant, it was no longer viewed as neutral. Rather, it came to be viewed as atypical of the stereotype and as relatively more likely to be associated with deviance. In other words, the attribute came to be viewed as a good reason for subtyping the deviant. This loss of neutrality provides further support for the view of subtyping as a creative process of justification construction. It suggests that the neutral attributes blocked generalization from the deviant to the stereotype because they were used as grounds for subtyping the deviant.

Further support for this view comes from the finding, in Study 4, that although a neutral attribute blocked generalization from truly counterstereotypic targets—unpromiscuous gay men who violated the stereotype of their group as promiscuous—it did not block generalization from overly stereotypic targets—highly promiscuous gay men, whose behavior was more extreme than the stereotypic one but not opposite to it. This pattern helps rule out several alternative accounts of why neutral attributes block generalization from deviants. These include the possibilities that the neutral attribute blocked generalization because it reduced the feature-based similarity between the targets and the stereotype (Tversky, 1977), because it increased the perceived variability of the stereotyped group (Nisbett et al., 1983; Quattrone & Jones, 1980), or because participants felt compelled to use any information provided by the experimenter (Schwarz, 1994). If these accounts were correct, neutral attributes should have blocked generalization from the overly stereotypical targets as well. The finding that they did not therefore strengthens our view that neutral attributes block generalization because they are pressed into service as reasons

for subtyping members who challenge stereotypes. It appears that only truly counterstereotypic members constitute a challenge powerful and surprising enough to trigger attempts at subtyping.

Because people appear to use neutral information about deviants to block generalization from them only when they are driven to attempt to subtype these deviants, it is important to determine what fuels these attempts. It is difficult to tell from these studies exactly what gave rise to the pressure to subtype deviants, but other research points to some likely sources. One probable source is the surprise engendered by the deviants. People are particularly likely to try to explain away information when they find it surprising (Hastie, 1984; Kunda et al., 1990; Pyszczynski & Greenberg, 1981; Wong & Weiner, 1981). Support for the notion that surprise may lead to attempts at subtyping deviants comes from research that has shown that extreme and therefore surprising deviants provoked less stereotype change than did moderate deviants, presumably because the surprisingly extreme deviants were more likely to be subtyped and dismissed (Oleson & Kunda, 1994). Participants' attempts to subtype and dismiss deviants may have also been fueled by their motivation to maintain their stereotypes. When people want to maintain a given belief, they are particularly likely to attempt to justify the dismissal of evidence that challenges it (Kunda, 1987, 1990; Pyszczynski & Greenberg, 1987). Because our participants may have been both surprised by the deviants and motivated to maintain their stereotypes, the independent roles of these two factors are yet to be determined.

There is some evidence that individuals may vary in the extent to which they attempt to subtype deviants. In particular, it appears that people who are highly prejudiced against a given group may be especially likely to try to use neutral information to subtype deviant members of that group and may be particularly sensitive to manipulations that facilitate the use of such information to subtype deviants (Adams & Kunda, 1994). These findings, too, imply that people are more likely to use neutral information to subtype and dismiss deviants if they feel pressure to do so. Once again, however, the source of the pressure felt by highly prejudiced people is difficult to determine. Such prejudiced participants tend to have particularly extreme stereotypes and so may be particularly surprised by the deviant, but they may also be particularly motivated to maintain their stereotypes. In future research, it would be interesting to assess the independent roles of surprise and motivation in promoting subtyping. It may be possible, for example, to increase the motivation to maintain one's stereotype without affecting the stereotype's extremity, perhaps by increasing people's commitment to their stereotypes through manipulations similar to those used by dissonance researchers (Aronson, 1968). If motivation creates pressure to subtype, such manipulations should increase the use of neutral information to subtype deviants.

The notion that people engage in constructive attempts to subtype and exclude individuals who challenge their stereotypes has broad implications; it may shed light on how people deal with challenges to any of their well-established and cherished beliefs. For example, there is evidence that people engage in similar processes of exclusion to deal with behavior that challenges

their expected and desired views of themselves, their peers, or their spouses; they tend to attribute such inconsistent behavior to situational rather than dispositional factors (Bradbury & Fincham, 1990; Kulik, 1983; Kulik, Sledge, & Mahler, 1986). From our perspective, such situational attribution of behavior is analogous to subtyping deviants. It can be viewed as using the situation as grounds for dismissing the relevance of a person's behavior to his or her personality. If, as we propose, such situational attribution results from constructive attempts to dismiss the challenge to one's beliefs, then factors that facilitate justifying such dismissal should increase situational attribution and thereby reduce generalization from the behavior to personality. Our findings suggest that any information about the situational setting of an inconsistent behavior should reduce the impact of that behavior on the perceived personality of the actor and that the same neutral situation may be creatively used to explain away opposite behaviors.

We have suggested that neutral information blocks generalization from deviants because of the role that it plays in justifying their dismissal. The effect of neutral information should therefore be a function of the ease with which people can use it as grounds for subtyping. Some classes of information may be used more readily than others to justify subtyping. It may be easier, for example, to view extraversion as associated with class or profession than with eye color. This ease or difficulty of generating justifications is likely to be determined by people's broader theoretical beliefs about what classes of information are likely to influence personality. It would be interesting to explore whether the blocking effects of neutral information are constrained by participants' ability to use this information to justify subtyping.

Our studies focused on the effects of individuals on stereotypes. An interesting parallel to our findings exists for the effects of stereotypes on the perception of individuals. We found that neutral attributes reduce the impact of an individual on a stereotype. Analogously, neutral attributes have also been shown to reduce the impact of stereotypes on the perception of individuals, a phenomenon termed the *dilution effect* (Nisbett, Zukier, & Lemley, 1981). It is possible that the two phenomena share a common source—the same factors that reduce the relevance of an individual to a stereotype may also reduce the relevance of the stereotype to an individual. Influence in both directions may be blocked once an individual is subtyped and fenced off from the stereotype. Although the dilution effect appears to result at least in part from a reduction in the extent to which the individual is representative of the stereotype (Hilton & Fein, 1989), in some dilution studies subtyping may have further enhanced the effect. As Nisbett et al. (1981) suggested, some of the diluting attributes that have been used may have seemed surprising for the stereotype even though they were non-diagnostic of the particular stereotype-related behavior that participants were asked to predict (e.g., mundane attributes for a member of a group stereotyped as vicious). These attributes may have provided an impetus for subtyping the individual, and any additional neutral information may have then been used as grounds for such subtyping, thereby reducing the impact of the stereotype on judgments about the individual. To date, studies

have examined either the effects of stereotypes on judgments about individuals (Brewer, 1988; Fiske & Neuberg, 1990) or the effects of individuals on stereotypes (Stephan, 1985), but the two processes have not been examined within a single study. Also, there has been no attempt to explore the effects of the same variables on both processes. It would be of great interest to study the two phenomena concurrently so as to determine whether both kinds of influence may be accounted for by the same theoretical framework.

On the practical side, our findings have disturbing implications for the likelihood that people will change their stereotypes in the normal course of their daily lives, as they encounter individuals who disconfirm these stereotypes. Outside the laboratory, one cannot block subtyping by providing people with no information about a deviant other than the deviance; almost invariably, people will have additional information, such as the personal appearance, occupation, or family background of encountered deviants that may be used as grounds for subtyping and dismissing them. It is not surprising, therefore, that stereotypes are so little affected by mere exposure to individuals who disconfirm them (Stephan, 1985). However, our findings do point to promising avenues for provoking stereotype change. It may be possible to block subtyping by making it difficult to group individuals who challenge the stereotype along any single dimension (cf. Wilder, 1986). If the individuals who disconfirm a stereotype all come from different walks of life and represent different occupations, ages, geographic locales, or family backgrounds, it should be difficult to find grounds for consigning them all to an atypical subtype. Therefore, even mere exposure to such diverse deviants may lead to stereotype change. Because the development of successful methods for changing insidious and inaccurate stereotypes would be of great social value, it is important to explore these possibilities.

References

- Adams, B. D., & Kunda, Z. (1994, June). *The role of prejudice in stereotype change and subtyping*. Paper presented at the meeting of the American Psychological Society, Washington, DC.
- Allport, G. W. (1954). *The nature of prejudice*. Reading, MA: Addison-Wesley.
- Anderson, C. A., Lepper, M. R., & Ross, L. (1980). Perseverance of social theories: The role of explanation in the persistence of discredited information. *Journal of Personality and Social Psychology*, 39, 1037-1047.
- Anderson, C. A., & Sechler, E. S. (1986). Effects of explanation and counterexplanation on the development and use of statistical theories. *Journal of Personality and Social Psychology*, 50, 24-34.
- Aronson, E. (1968). Dissonance theory: Progress and problems. In R. P. Abelson, E. Aronson, W. J. McGuire, T. M. Newcomb, M. J. Rosenberg, & P. H. Tannenbaum (Eds.), *Theories of cognitive consistency: A sourcebook* (pp. 5-27). Chicago: Rand McNally.
- Bradbury, T. N., & Fincham, F. D. (1990). Attributions in marriage: Review and critique. *Psychological Bulletin*, 107, 3-33.
- Brewer, M. B. (1988). A dual process model of impression formation. In T. K. Srull & R. S. Wyer (Eds.), *Advances in social cognition* (pp. 1-36). Hillsdale, NJ: Erlbaum.
- Brewer, M. B., Dull, V., & Lui, L. (1981). Perceptions of the elderly: Stereotypes as prototypes. *Journal of Personality and Social Psychology*, 41, 656-670.
- Cantor, N., & Mischel, W. (1979). Prototypes in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 12, pp. 3-52). New York: Academic Press.
- Devine, P. G., & Baker, S. M. (1991). Measurement of racial stereotype subtyping. *Personality and Social Psychology Bulletin*, 17, 44-50.
- Esses, V., Haddock, G., & Zanna, M. P. (1993). Values, stereotypes, and emotions as determinants of intergroup attitudes. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 137-166). San Diego: Academic Press.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23, pp. 1-74). San Diego, CA: Academic Press.
- Haddock, G., Zanna, M. P., & Esses, V. M. (1993). Assessing the structure of prejudicial attitudes: The case of attitudes toward homosexuals. *Journal of Personality and Social Psychology*, 65, 1105-1118.
- Hamill, R., Wilson, T. D., & Nisbett, R. E. (1980). Insensitivity to sample bias: Generalizing from atypical cases. *Journal of Personality and Social Psychology*, 39, 578-589.
- Hastie, R. (1984). Causes and effects of causal attribution. *Journal of Personality and Social Psychology*, 47, 44-56.
- Herek, G. M. (1992). Stigma, prejudice, and violence against lesbians and gay men. In J. C. Gonsiorek & J. D. Weinrich (Eds.), *Homosexuality: Research implications for public policy* (pp. 60-80). Newbury Park, CA: Sage.
- Hilton, J. L., & Fein, S. (1989). The role of typical diagnosticity in stereotype-based judgments. *Journal of Personality and Social Psychology*, 57, 201-211.
- Johnson, B., & Eagly, A. H. (1989). The effects of involvement on persuasion: A meta-analysis. *Psychological Bulletin*, 106, 290-314.
- Johnston, L., & Hewstone, M. (1992). Cognitive models of stereotype change: 3. Subtyping and the perceived typicality of disconfirming group members. *Journal of Experimental Social Psychology*, 28, 360-386.
- Klayman, J., & Ha, Y. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94, 211-228.
- Kulik, J. A. (1983). Confirmatory attribution and the perception of social beliefs. *Journal of Personality and Social Psychology*, 44, 1171-1181.
- Kulik, J. A., Sledge, P., & Mahler, H. I. M. (1986). Self-confirmatory attribution, egocentrism, and the perpetuation of self-beliefs. *Journal of Personality and Social Psychology*, 50, 587-594.
- Kunda, Z. (1987). Motivated inference: Self-serving generation and evaluation of causal theories. *Journal of Personality and Social Psychology*, 53, 636-647.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108, 480-498.
- Kunda, Z., Miller, D. T., & Claire, T. (1990). Combining social concepts: The role of causal reasoning. *Cognitive Science*, 14, 551-577.
- Lakatos, I. (1976). *Proofs and refutations: The logic of mathematical discovery*. Cambridge, England: Cambridge University Press.
- Nisbett, R. E., Krantz, D. H., Jepson, C., & Kunda, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review*, 90, 339-363.
- Nisbett, R. E., Zukier, H., & Lemley, R. E. (1981). The dilution effect: Nondiagnostic information weakens the implications of diagnostic information. *Cognitive Psychology*, 13, 248-277.
- Oleson, K. C., & Kunda, Z. (1994, May). *Maintaining stereotypes in the face of disconfirmation: The role of extremity in inhibiting stereotype change*. Paper presented at the meeting of the Midwestern Psychological Association, Chicago.

- Pyszczynski, T. A., & Greenberg, J. (1981). Role of disconfirmed expectations in the investigation of attributional processing. *Journal of Personality and Social Psychology*, 40, 31–38.
- Pyszczynski, T. A., & Greenberg, J. (1987). Toward an integration of cognitive and motivational perspectives on social inference: A biased hypothesis-testing model. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 20, pp. 297–340). New York: Academic Press.
- Quattrone, G. A., & Jones, E. E. (1980). The perception of variability within ingroups and outgroups: Implications for the law of small numbers. *Journal of Personality and Social Psychology*, 38, 141–152.
- Rips, L. J. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behavior*, 14, 665–681.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Erlbaum.
- Rothbart, M., & John, O. P. (1985). Social categorization and behavioral episodes: A cognitive analysis of the effects of intergroup contact. *Journal of Social Issues*, 41, 81–104.
- Rothbart, M., & Lewis, S. (1988). Inferring category attributes from exemplar attributes: Geometric shapes and social categories. *Journal of Personality and Social Psychology*, 55, 861–872.
- Sanitioso, R., Kunda, Z., & Fong, G. T. (1990). Motivated recruitment of autobiographical memories. *Journal of Personality and Social Psychology*, 59, 229–241.
- Sartre, J. P. (1946/1948). *Anti-Semite and Jew*. New York: Schocken.
- Schwarz, N. (1994). Judgment in social context: Biases, shortcomings, and the logic of conversation. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 26, pp. 123–162). San Diego, CA: Academic Press.
- Snyder, M. (1984). When belief creates reality. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 18, pp. 248–306). New York: Academic Press.
- Snyder, M., & Swann, W. B. (1978). Hypothesis testing in social interaction. *Journal of Personality and Social Psychology*, 36, 1202–1212.
- Stephan, W. G. (1985). Intergroup relations. In G. Lindzey & E. Aronson (Eds.), *Handbook of social psychology* (3rd ed., Vol. 2, pp. 599–658). New York: Random House.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327–352.
- Weber, R., & Crocker, J. (1983). Cognitive processes in the revision of stereotypic beliefs. *Journal of Personality and Social Psychology*, 45, 961–977.
- Wilder, D. A. (1986). Cognitive factors affecting the success of intergroup contact. In S. Worchel & W. G. Austin (Eds.), *Psychology of intergroup relations* (pp. 49–66). Chicago: Nelson-Hall.
- Wong, P., & Weiner, B. (1981). When people ask “why” questions, and the heuristics of attributional search. *Journal of Personality and Social Psychology*, 40, 650–663.

Received June 23, 1994

Revision received October 11, 1994

Accepted October 19, 1994 ■

Correction to Krueger and Clement

In the article “The Truly False Consensus Effect: An Ineradicable and Egocentric Bias in Social Perception,” by Joachim Krueger and Russell W. Clement (*Journal of Personality and Social Psychology*, 1994, Vol. 67, No. 4, 596–610), the equation on p. 605 should have read:

$$P(\text{urn/blue}) = P(\text{urn}) \times \frac{P(\text{blue/urn})}{P(\text{blue})}$$

Correction to Amabile et al.

In the article "The Work Preference Inventory: Assessing Intrinsic and Extrinsic Motivational Orientations," by Teresa M. Amabile, Karl G. Hill, Beth A. Hennessey, and Elizabeth M. Tighe (*Journal of Personality and Social Psychology*, 1994, Vol. 66, No. 5, pp. 950–967), Items 5 and 11 in Table 1 appeared in the incorrect column. The correct version of the table appears below.

Table 1
Work Preference Inventory Items and Scale Placement

Item no.	Item	Primary		Secondary				Primary factor loading rank	
		IM	EM	E	Ch	O	C	Students	Working adults
13	I enjoy tackling problems that are completely new to me.	X			X			1	3
26	I enjoy trying to solve complex problems.	X			X			2	1
3	The more difficult the problem, the more I enjoy trying to solve it.	X			X			3	2
5	I want my work to provide me with opportunities for increasing my knowledge and skills.	X		X				4	12
11	Curiosity is the driving force behind much of what I do.	X		X				5	5
28	I want to find out how good I really can be at my work.	X		X				6	7
7	I prefer to figure things out for myself.	X		X				7	10
30	What matters most to me is enjoying what I do.	X		X				8	13
27	It is important for me to have an outlet for self-expression.	X		X				9	4
14	I prefer work I know I can do well over work that stretches my abilities.	R			R			10	6
8	No matter what the outcome of a project, I am satisfied if I feel I gained a new experience.	X		X				11	15
17	I'm more comfortable when I can set my own goals.	X		X				12	9
23	I enjoy doing work that is so absorbing that I forget about everything else.	X		X				13	11
20	It is important for me to be able to do what I most enjoy.	X		X				14	14
9	I enjoy relatively simple, straightforward tasks.	R			R			15	8
19	I am strongly motivated by the [grades] [money] I can earn.		X				X	1	2
10	I am keenly aware of the [GPA (grade point average)] [promotion] goals I have for myself.		X				X	2	3
24	I am strongly motivated by the recognition I can earn from other people.		X			X		3	5
29	I want other people to find out how good I really can be at my work.		X			X		4	4
16	I seldom think about [grades and awards.] [salary and promotions.]		R				R	5	12
4	I am keenly aware of the [goals I have for getting good grades.] [income goals I have for myself.]		X				X	6	1
6	To me, success means doing better than other people.		X			X		7	7
25	I have to feel that I'm earning something for what I do.		X			X		8	6
22	As long as I can do what I enjoy, I'm not that concerned about exactly [what grades or awards I can earn.] [what I'm paid.]		R				R	9	13
18	I believe that there is no point in doing a good job if nobody else knows about it.		X			X		10	14
15	I'm concerned about how other people are going to react to my ideas.		X			X		11	8
21	I prefer working on projects with clearly specified procedures.		X			X		12	9
12	I'm less concerned with what work I do than what I get for it.		X			X		13	11
1	I am not that concerned about what other people think of my work.		R			R		14	15
2	I prefer having someone set clear goals for me in my work.		X			X		15	10

Note. Items 4, 10, 16, 19, and 22 are worded differently for students and adults. Both are presented here, in brackets. An X indicates that the item falls on that particular scale. An R indicates that it is reverse scored. IM = Intrinsic Motivation Scale; EM = Extrinsic Motivation Scale; E = Enjoyment Scale; Ch = Challenge Scale; O = Outward Scale; C = Compensation Scale.