Group Processes & Intergroup Relations
2000 Vol 3(3) 227–256

G
P
I
R

Power Can Bias Impression Processes: Stereotyping Subordinates by Default and by Design

Stephanie A. Goodwin
Yale University

Alexandra Gubin
University of Massachusetts at Amherst

Susan T. Fiske
Princeton University

Vincent Y. Yzerbyt
Université Catholique de Louvain, Louvain-la-Neuve

Powerholders may engage in two stereotyping processes: (a) by default, inattention to stereotype-inconsistent information, due to lack of dependency, and (b) by design, effortful attention to stereotype-consistent information, due to explicit control. Study 1 manipulated control (not dependency) over internship applicants; powerful decision-makers increased attention to stereotypic attributes, consistent with stereotyping by design. Study 2 measured differences in trait dominance as an analog to situational control, replicating Study 1. Study 3 separately manipulated perceiver control and dependency; powerful perceivers increased attention to powerless targets' stereotypic attributes (by design) and also decreased attention to counter-stereotypic attributes (by default). Study 4 compared powerful perceivers' ratings of potential subordinates to their own prior ratings of target categories and target traits. Relative to the powerless, powerful perceivers' impressions were based significantly less on target traits, supporting the attention results.

KEYWORDS impression formation, intergroup relations, power, stereotyping

POWERFUL people are socially sanctioned to judge others, as part of their positions. Jurors, teachers, and managers, for example, must evaluate others to determine guilt, assign grades, or promote employees. Powerful people – people who control other's outcomes – could, like anyone else, rely on stereotypes when they are

Author's note
Address correspondence to either Stephanie A. Goodwin, Department of Psychology, Yale University, New Haven, CT 06520, USA, or Susan T. Fiske, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544–1010, USA.

Copyright © 2000 SAGE Publications
(London, Thousand Oaks, CA and New Delhi)
[1368-4302(200007)3:3; 227–256; 011129]

not especially motivated to be accurate (Fiske, 1998; Hilton & von Hippel, 1996; Neuberg, 1996). However, feeling entitled to judge, powerful people might also make an effort to maintain stereotypes of their subordinates. Stereotyping subordinates could fulfill motivational pressures that accompany powerholders' positions of control and unique authority to judge. At issue here, then, is the relationship between power and motivated stereotyping. By investigating stereotype-based attention, the present research addresses *how* powerful individuals think about their subordinates. In addition, by investigating impression ratings, we address *what* the powerful think about subordinates.

Social power, impression formation, and stereotyping

Power as control

Power can be defined in terms of its inherent social nature (Dépret & Fiske, 1993). More precisely, power reflects the nature of outcome control and outcome dependencies in a social relationship. People who control others' outcomes are relatively *powerful*, and people whose outcomes depend on others are relatively *powerless*. This definition unconfounds common correlates of power (i.e. status, influence, *pace*, Bierstedt, 1950; Cartwright, 1959; French & Raven, 1959) and recognizes gradations in the asymmetry of power relations (Fiske, 1993; Fiske & Dépret, 1996). Asymmetry, one person disproportionately controlling another's outcomes, distinguishes power relations from symmetrical interdependence (Thibaut & Kelley, 1959). In addressing how and what powerful people think about their subordinates, we define power as unilateral outcome control (Thibaut & Kelley's 1959 'fate control'); powerful participants in these studies have disproportionate control over, but are not contingent on powerless subordinates.

Impression formation

In current models of impression formation, perceiver motives predict whether and how stereotypes influence interpretation of available information (Brewer, 1988; Fiske & Neuberg, 1990). People can counteract the apparently automatic activation of stereotypes, but doing so requires considerable motivation and cognitive effort (Fiske, 1998). Assuming the activation of a given group stereotype, the current studies assess whether power alters perceiver attentional processes during impression formation (Studies 1–3) and consequences for these interpretations on judgment (Study 4). The Continuum Model of impression formation is particularly relevant to this question because it specifically addresses interpretation and motivations not to stereotype (Fiske & Neuberg, 1990).

The Continuum Model: Motives to stereotype or not to stereotype People engage in a variety of impression formation strategies that range from effortless to effortful cognitive processes (Fiske & Neuberg, 1990; Fiske, Lin, & Neuberg, 1999). The amount of effort people expend when thinking about others depends on two factors: motivation and cognitive ability (i.e. mental resources).

When uninterested or unable to attend, people engage in effortless category-based impression processes, *stereotyping by default*; this involves relative inattention to potentially disconfirming, stereotype-inconsistent information (Fiske, Neuberg, Beattie, & Milberg, 1987). In contrast, when motivated to form accurate impressions, people engage more effortful impression processes – *individuation* – that can overcome stereotype activation. Information that does not fit stereotypes is relatively more useful than information that does fit (Fiske & Neuberg, 1990); stereotype-consistent information is redundant with what is already 'known'. Consequently, people with accuracy motives attend to stereotype-inconsistent information. Of course, individuating impression processes are necessary but not sufficient for reducing stereotypic impressions; people who attend to stereotype-inconsistent information may instead refute it.

Several factors promote accuracy motives and potential individuation: explicit instruction (i.e. telling people to be accurate; Neuberg, 1989;

Neuberg & Fiske, 1987), accountability (i.e. expecting to justify impressions to a third party; Tetlock, 1992), and dependency on another's performance (Dépret & Fiske, 1999; Erber & Fiske, 1984; Neuberg & Fiske, 1987; Ruscher & Fiske, 1990: Ruscher, Fiske, Miki, & Van Manen, 1991). Outcome-dependent people adopt accuracy motives presumably as a means of prediction and control over their own outcomes. Participants who believe their own chances for winning a prize are contingent on another person's performance, for example, pay significantly more attention to inconsistent information that challenges their stereotypes, as compared to people who are not outcome-dependent.

Power as control in impression formation The definition of power as control applies to impression formation and mutual outcome-dependency, as argued in a tripartite model proposed by Fiske (1993): Powerholders may stereotype subordinates by default – ignoring stereotype-inconsistent information – because (1) they do *not need*, and are therefore unmotivated to attend to subordinates, (2) they may *not want* to attend to subordinates, to the extent that they have dominant personalities, or (3) they perhaps *cannot* attend to subordinates because cognitive demands are high (with many subordinates). This research focuses on the first and (to a lesser extent) second predictions. Powerful people are, by definition, relatively non-contingent on the powerless; outcome-controlling powerful people (and perhaps dominant personalities) are predicted to ignore information that challenges stereotypes, stereotyping subordinates by default. In contrast, as noted, powerless perceivers ought to be motivated to form relatively accurate impressions of powerholders, attending to information that challenges stereotypes.

Might powerholders also actively maintain stereotypes? Social Judgeability Theory Stereotyping subordinates by default is but one strategy whereby the powerful may stereotype subordinates. In addition, the powerful may employ motivated, effortful strategies that also bias their impressions (Goodwin & Fiske, 1996), particularly when they hold an exchange-oriented view of relationships (Chen, Lee-Chai, & Bargh, 2000). Whether or not the powerful stereotype effortfully, in addition to effortlessly, is likely a function of whether they feel entitled to judge, which may be a function either of situation or of person (Goodwin & Fiske, 1996; Leyens, Yzerbyt, & Schadron, 1994).

Simply stated, the fundamental principle of SJT is that people refrain from judging others unless they feel able to do so (Leyens, Yzerbyt, & Schadron, 1992, 1994; Yzerbyt, Dardenne, & Leyens, 1998; Yzerbyt, Schadron, Leyens, & Rocher, 1994). People feel entitled to judge when their judgments: (1) fit the *available information* (e.g. Kunda, 1990), (2) fit *perceivers' theories* about how the social world works (Leyens, Yzerbyt, & Schadron, 1994), (3) meet *cultural expectations* about who can judge and how they ought to do so (Yzerbyt, Leyens, & Schadron, 1997; Yzerbyt et al., 1994), and (4) maintain the *integrity* of important personal beliefs, including values, self-concepts, and group identities (Leyens & Yzerbyt, 1992). When faced with violating any of these four criteria for judging, people feel unable to judge and may hesitate, hedge, or refuse. Applying these judgment criteria to power relations, two factors may motivate powerful people's effort to maintain stereotypes about subordinates: powerholders' potential entitlement to judge and possible desire to maintain power roles (Goodwin & Fiske, 1996).

Power entails entitlement Feeling entitled to judge may increase powerholders' confidence in their own beliefs, including their stereotypes. One explanation for this phenomenon is that entitlement and self-confidence increase with power. Because Western cultures hold that people gain power because they earn it – i.e. because they have requisite skills or expertise – the powerful can feel that their beliefs (in this case, stereotypes) are particularly valid. Unmotivated to go beyond their initial stereotypes, the powerful might then attend only minimally to subordinates, with a net result of stereotyping subordinates by default (i.e. failing

to notice stereotype-inconsistent information). In accord with this idea, people who feel relatively able to judge do expend less effort forming impressions (Leyens et al., 1992).

Stereotyping protects power roles However, cultural expectations about how the powerful *ought* to judge may preclude simple effortless stereotyping by default. Social norms dictate that powerful decision-makers attend to a subordinate enough to justify their decisions. Powerful people who distribute outcomes based on arbitrary or untenable decisions not only risk poor decisions, but also may find themselves unable to defend these decisions; unjustifiable, poor decisions threaten power. Even in short-term power situations, social norms may increase effort; the scripts for decision-making in power contexts may simply demand increased effort, even absent motives to maintain power. Attending effortfully to stereotyped subordinates allows one to justify one's decisions and to satisfy normative pressures to attend to subordinates.

If powerful people are motivated to make justifiable decisions, why not simply predict they would effortfully individuate their subordinates? After all, the most justifiable decisions are accurate decisions. Powerful people sometimes may avoid individuating subordinates because accurate impressions could threaten power-relevant social identities. Individuating subordinates could expose the powerful to information that challenges existing power relations, for example, by showing that a subordinate is especially competent. In contrast, negatively stereotyping subordinates can justify one's power position, the status of one's social groups, and the broader system of power relations between groups (Glick & Fiske, in press; Jost & Banaji, 1994; Yzerbyt, Rocher, & Schadron, 1997). Indeed, as people gain power, their evaluations of others deteriorate, whereas their self-evaluations inflate (Georgesen & Harris, 1998).

Powerful people may protect their individual power identities by strategically protecting relevant beliefs about existing power relations. The powerful may, therefore, attentionally screen out stereotype-inconsistent information, dismissing it because it could threaten beliefs about who should have control. In contrast, stereotype-*consistent* information about subordinates preserves, perhaps even bolsters existing power identities. When powerful people negatively stereotype their subordinates, it justifies their own personal positions of control. Just as individuation can threaten power identities, stereotyping can protect them.

Motivational tug-of-war Powerholders can face a tension between motives to attend to subordinates on the one hand, and motives to maintain stereotypes on the other. One means of satisfying both types of processing motives is to pay effortful attention to stereotype-consistent information, to *stereotype by design*. If the powerful attend effortfully to information that fits their prior expectations, they can maintain the integrity of power-relevant identities while also garnering information necessary to justify these stereotypic judgments. Furthermore, lacking outcome-dependency, the powerful lack incentive to engage in accuracy-based impression strategies when thinking of the powerless. Hence they are likely to stereotype subordinates also *by default*, ignoring non-stereotypic information that could challenge stereotypes, in addition to stereotyping by design. Combining powerholders' potential motives to maintain stereotypes with their lack of outcome dependency, the cards can be stacked against the powerless.

The present studies constitute a first test of the hypothesis that power can lead to motivated stereotype maintenance during impression formation. Four experiments address this possibility in short-term power situations in which a powerholder evaluates information about individuals in the process of distributing resources. Experiments 1–3 demonstrate that power triggers attentional processes biased toward stereotype maintenance. Experiment 4 focuses on the biased impressions that can result from having power in a social situation.

Study 1

Overview

The primary aim of the first study is to establish

whether powerful perceivers increase attention to stereotype-confirming information, stereotyping subordinates by design at the cognitive level. This study's manipulation of power does not manipulate perceiver dependency; in no case were Study 1 perceivers ever contingent on targets for outcomes. Thus, the isolated manipulation of outcome control is hypothesized to affect only stereotyping by design.

The study also examines whether priming internalized egalitarian values could reduce this form of attention bias. Of two possible sources of accuracy motivation (Fiske & Neuberg, 1990) – the self (e.g. personal values), and third parties (e.g., accountability) – personal values are especially likely to be important motivators to the powerful, to the extent that they are indeed autonomous. Prior research indicates that people can intentionally override prejudiced tendencies *if* doing so is central to their self-descriptions (Monteith, Devine, & Zuwerink, 1993). Participants in prior research have adopted accuracy goals when encouraged by situational norms (for high self-monitors) and personality feedback (for low self-monitors; Fiske & Von Hendy, 1992). Perhaps activating personal values to treat others fairly (values shared by many Americans, including prejudiced individuals) would encourage accuracy-oriented impression formation even of subordinates. To ascertain whether internalized values can override the effects of power, we manipulated the value salience (cf. Kiesler, Nisbett, & Zanna, 1969) of treating others in an egalitarian way; we predicted that this manipulation would increase individuated impression strategies by powerful perceivers.

Participants recruited to help select high school students for a local internship learned each applicant's ethnicity and read trait information pre-tested for consistency with ethnic stereotypes. For half of the participants, egalitarian values were made salient prior to evaluating applicants.

Hypotheses Across manipulations, we predicted a three-way interaction between power, responsibility, and type of target information. However, the best test of the stereotyping by design hypothesis can be observed in participants' reactions when they were not primed to feel responsible. The predicted patterns of attention to information are presented below as a function of responsibility condition.

Baseline predictions: Participants not primed to feel responsible Our argument assumes that, on average, feeling responsible for forming fair impressions of others is not a chronic social motive for perceivers in most natural contexts. Perceivers who were not primed to feel responsible, therefore, constitute the baseline condition for testing the hypothesis that power promotes attention processes presumed to reflect stereotyping by design.

We predicted a two-way interaction between power and information consistency in this baseline no-responsibility condition. Non-powerful perceivers are predicted to be minimally invested in attending to targets, and consequently they should pay minimal attention to both types of target information (stereotype-consistent and inconsistent). In contrast, powerful perceivers are predicted to stereotype targets by design. Lacking accuracy motives (like the non-powerful perceivers, being non-contingent), yet more motivated to confirm their stereotype expectations (being in control), powerful participants should increase attention to stereotype-consistent information, relative to the non-powerful.[1] Because neither the powerful nor the non-powerful are outcome dependent on targets in this design, we predict no difference in attention to stereotype-inconsistent information as a function of power (i.e. this design is not diagnostic of stereotyping by default because perceiver contingency is not manipulated, only perceiver control). Thus, powerful perceivers should increase attention to stereotype-consistent information but pay the same attention to stereotype-inconsistent information, relative to the non-powerful.

Responsibility priming effects Priming participants' internal values to feel responsible for outgroup members is predicted to alter participants' motivations, regardless of power, increasing participants' overall attention to targets

(responsibility main effect). Responsibility is further predicted to interact with information consistency (two-way interaction) such that high-responsibility participants should be more likely to use individuating impression strategies, compared to low-responsibility participants.

Finally, as mentioned above, we predict a three-way interaction between responsibility, power, and information consistency. Responsibility should moderate the predicted relationship between power and stereotype-based processing. For powerful participants, feeling responsible should act as an antidote to their motivations to confirm expectations about subordinates. In this case, powerful participants who feel responsible should adopt individuating processes, paying more attention to stereotype-inconsistent information and less attention to stereotype-consistent information.[2]

Method

Design The study employed a 2 (Power) × 2 (Responsibility) between-participants factorial design. Power was manipulated via participants' perceived control over selection of the alleged applicants (30% outcome control, *powerful*, versus no outcome control, *non-powerful*). Responsibility, operationalized as accessibility of shared egalitarian values, was manipulated by priming half of the participants for responsibility and egalitarianism toward outgroup members. Target ethnicity (Anglo, Hispanic) was manipulated within-participants, and trait stereotype consistency (consistent, inconsistent) was manipulated within-target.

The dependent measure of interest was perceiver attention to target trait information. Attention was measured by timing participants' audio-recorded verbal responses to target trait information. Additional dependent measures included the coded content of participants' verbal responses to this information (e.g. types of attributions, elaboration, etc.) and target impression ratings.

Participants Seventy-eight native English-speaking undergraduates were recruited from undergraduate psychology courses in the University of Massachusetts at Amherst to participate in exchange for course credit or US$5. All participants were Anglo-Americans. Four cases were dropped prior to analysis due either to experimenter error in procedures (2), participant error in following procedures (1), or suspicion of the cover story (1). In addition, participant responses to the high-responsibility manipulation were screened for effectiveness. We calculated the hypothetical midpoint of the total responses (31.5) and used this figure as a level of comparison for effectiveness of the manipulation. One participant who scored below this value was assumed to be unaffected by the manipulation and was subsequently eliminated from analyses. The remaining 73 cases were screened for possible outliers on the attention measure. Three participants (from different power × responsibility treatment groups) had attention scores 2.5 standard deviations or more above the cell mean and were eliminated from further analysis. The final sample consisted of 70 participants, 38 men and 32 women.

Procedure

The recruiting cover story A confederate posing as a representative of a fictitious local consulting firm – Gilbert & Swann Consulting – telephoned students who had previously expressed interest in participating. The consulting firm, supposedly under contract with area municipalities, was ostensibly interested in getting college students' opinions about applicants for a high-school internship program slated for reductions. Participants were scheduled for individual one-hour sessions.

The experimental session A conservatively dressed female experimenter posed as a representative of the consulting firm, escorting participants to a lab arranged to resemble an office. Seated at an empty table, participants were handed a letter (on company letterhead) that reminded them of the alleged purpose of their participation: to review and evaluate materials for several applicants.

Pre-testing suggested that participants must truly believe they are evaluating others for the power manipulation to be meaningful, hence the admittedly heavy-handed deception. The

initial consent form included enough information regarding the experimental tasks so as to allow informed consent, but did not explain the true nature of the alleged applicant materials. A second consent (to use their data) was provided at the end of the study. After signing the initial consent form, participants completed a preliminary questionnaire requesting demographic information and assessing self-perceived competence in performing this kind of task.

The power manipulation On a separate page, along with the questionnaire, participants found additional information about the general procedures. The power manipulation was embedded in this information, allowing the experimenter to remain blind to condition. Participants in the *no-power* condition read:

> Your evaluations will not affect our decisions about which students to retain in the program. We are interested in learning your opinions about the student applicants because we believe that your opinions could shed some light on better ways to evaluate applicants for such positions.

In contrast, participants in the *powerful* condition were told:

> Your evaluations will play a major role in determining whether or not each student will be retained in the program. Your overall evaluation of each applicant will be entered into a statistical equation and will account for 30 percent of the final decision to retain the student or not.

As participants read these materials, the experimenter placed a stack of 12 applicant folders in front of the participant. Participants evaluated a total of 6 fictitious applicants. The first 4 filler applications served to habituate participants to the possibly novel situation of being in power, evaluating applications, and speaking into the audio recorder. Each folder contained an application form, six trait information sentences, and a blank impression rating form.

All applicants were female high school seniors applying for generic clerical internships. We adapted the form from a typical job application, including items such as previous work history, honors and awards, job-related skills, and personal references. Target ethnicity was indicated by the applicant's name, and, for the Hispanic target, affiliation with a Hispanic school organization and fluency in Spanish as a primary language. To counterbalance any perceived skills indicated by these manipulations, the Anglo target was presented with similar qualities that were not ethnicity-related (i.e. participating in another school organization, fluent in French as a second language).

Order of the four initial fillers, three Anglos and one Hispanic, remained stable across conditions (Anglo, Anglo, Hispanic, Anglo). Two separate but parallel application forms were developed for the two targets, minimizing differences in the diagnostic skills and work history information on the two forms. Order and ethnicity of the two target forms were counterbalanced across treatments.

Trait information was presented to participants under the guise of handwritten coworker comments, each on separate postcards bearing recent postmarks and addressed to a putative consulting firm at the university. The experimenter allegedly had requested anonymous feedback from people who had previously worked with applicants. The 12 traits used in the target folders were pre-tested to represent uniquely Anglo student stereotypes for Anglos versus Hispanics (all $p < .0005$). For Anglos, stereotypic traits included: *ambitious*, *educated*, *efficient*, *well-mannered*, *industrious*, and *neat*. The Hispanic stereotype was, as noted, negative, including: *emotional*, *feels inferior*, *ignorant*, *loud*, *radical*, and *unreliable*.

Twelve sentences were created to fit a coworker's response (e.g. 'It was helpful to have such an industrious intern working in our office this year', 'We had a constant problem with her tendency to be too loud at work'.). Two randomly split groupings, each containing three Anglo and three Hispanic traits, were counterbalanced across treatments and between targets, presenting all 12 traits to each participant but in different combinations across treatments. Within each grouping, trait consistency consequently depended on the target's ethnicity. Trait order was randomized within target, with the constraints that: (1) no more than two traits of the same consistency appeared consecutively within-target, and (2) order of consistent

and inconsistent information between the two targets was independent.

After reading the applicant materials, participants responded to an 8-item 'Candidate Rating Questionnaire' designed to assess participants' perception of the applicants' likability, competence, and skill. Participants responded to the items by marking an 'X' on a line to indicate their (dis)agreement.

Think-aloud instructions The experimenter maintained that the firm was interested in participants' reactions to applicant materials as well as how participants 'came up with decisions about applicants'. For convenience, the experimenter asked participants to 'think aloud' into an audio-recorder while reviewing the materials.

The experimenter specifically requested that participants read aloud and respond orally to any information on the application form that seemed important to them for any reason. Participants then were to read each trait comment aloud and to say 'whatever comes to mind' about the comment as it pertained to the applicant. The experimenter emphasized the importance of responding to each postcard.

The responsibility priming manipulation Just as the experimenter was about to tell the participant to begin, she glanced through the portfolio she was carrying and, looking surprised, told the participant that she had forgotten about a questionnaire she was to administer. She alleged that the psychology department had allowed the consulting firm access to participants if the firm would include a questionnaire from a 'real' psychology experiment. The experimenter repeated that she had forgotten about the questionnaire, explaining that she had only just received it that day. Glancing anxiously at her watch, she asked participants to complete the form before they began reviewing applications. In reality, the questionnaire was the responsibility manipulation.

The high-responsibility condition primed shared egalitarian values via the Humanitarian-Egalitarian Values Scale (Katz & Haas, 1988). This 10-item scale, specifically designed to prime people to feel responsible toward others, includes items such as 'One should be kind to all people' and 'A good society is one in which people feel responsible for one another'. The present study dropped one item because it might have made participants suspicious. A 6-point scale indicated (dis)agreement with each statement, with higher scores indicating more agreement. Participants in the low-responsibility condition responded similarly to a control questionnaire containing an equal number of filler items pertaining to attitude issues unrelated to the topic of the study (e.g. 'There is not enough emphasis on the arts in our education system'.). To heighten credibility, the questionnaire included a cover-sheet requesting consent to participate in an alleged opinion survey. (No participant expressed suspicion on this point during debriefing.)

While participants completed the responsibility manipulation, the experimenter surreptitiously examined the power manipulation form to discover the experimental condition. After participants completed the questionnaire, the experimenter casually reminded participants of their control over applicants' outcomes (none or 30%) as she instructed them to begin reviewing applicants. The experimenter remained in the room, seated at a second desk with her back to participants, pretending to be involved in other work while participants reviewed applicant folders. If participants failed to respond aloud to the materials for the first two applicants, the experimenter casually leaned back and reminded them to respond aloud to the trait information. Participants were not prompted to respond to the materials after the second applicant folder.

When participants completed the second target folder, the experimenter interrupted, asking participants to stop reviewing folders. Participants next completed a questionnaire to assess cover story credibility and manipulation effectiveness. Participants were then carefully debriefed and provided the second consent form. Although surprised, no participants expressed discomfort. Participants received credit or payment before leaving.

Results and discussion

Manipulation checks Just prior to debriefing,

participants were asked to report how much impact they believed they would have on the hiring decisions on a scale from 1 (no impact) to 6 (a great deal of impact). As expected, the power manipulation significantly increased perceived control over applicant outcomes ($F(1,66) = 8.84$, $p = .004$, powerful $M = 3.79$, no-power $M = 2.95$). Participants also reported how responsible they felt for their decisions (1 = not at all responsible, 6 = very responsible). The responsibility manipulation marginally increased participants' sense of responsibility ($F(1,66) = 3.28$, $p = .07$, high-responsibility $M = 4.89$, low-responsibility $M = 4.39$). There were no other effects or interactions on these manipulation checks.

Attention data Participants' audio-recorded verbal responses to the trait sentences were transcribed verbatim and timed to the hundredth of a second. Timing began when participants read aloud the first word of the trait sentence and ended when they turned the card to proceed to the next sentence. The times were submitted to a mixed-design analysis of variance (ANOVA), with power and responsibility as between-participants factors, and with target ethnicity and trait consistency as within-participant factors. There were no presentation order effects for the two application forms, the trait groupings, or the ethnicity of the applicant.[3] The predicted three-way interaction between responsibility, power, and information consistency failed to reach significance. However, comparable ANOVAs within each of the two responsibility conditions following this omnibus analysis revealed several effects that support our hypotheses.

Baseline effects: No responsibility priming As mentioned, analyses of the effects within the no-responsibility condition constitute the best test of our prediction that powerful perceivers would stereotype subordinates by design. As predicted, for participants who were not primed to feel responsible, a significant interaction emerged between power and information consistency ($F(1,37) = 9.77$, $p = .003$) (see Figure 1).

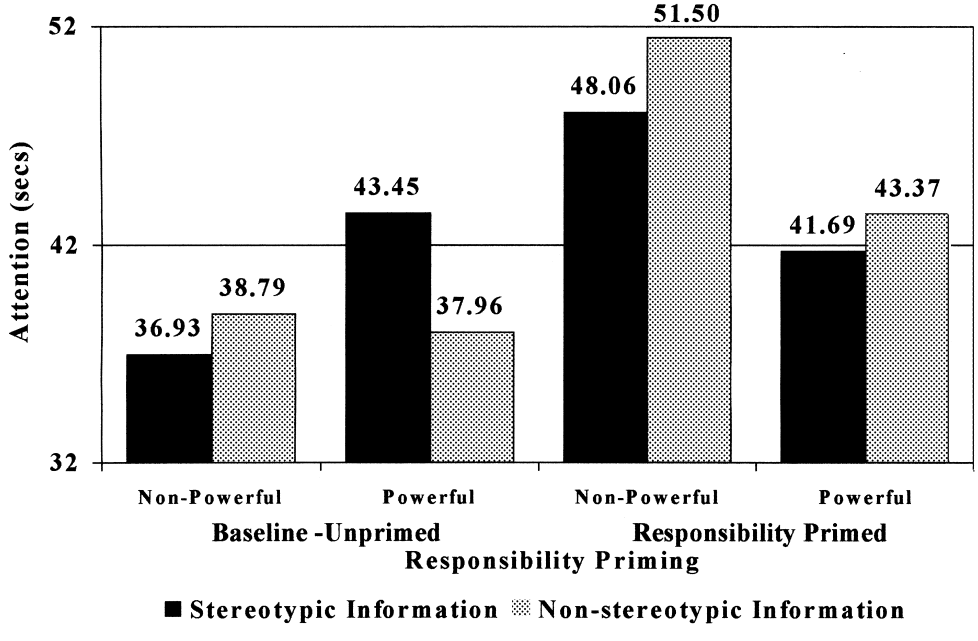

Figure 1. Mean attention to stereotype-consistent and inconsistent information as a function of perceiver's control over target outcomes in the baseline, no-responsibility condition versus the responsibility primed condition: Study 1.

Powerful participants significantly increased attention to stereotype-consistent information (M = 43.45 s), compared to their non-powerful counterparts (M = 36.93 s). Also as predicted, attention to stereotype-inconsistent information remained unchanged. Thus, when not feeling particularly responsible, powerful perceivers stereotyped targets by design, paying much more attention to stereotype-confirming information than did the non-powerful.

In accordance with our social judgeability extension of the power-as-control (PAC) hypotheses, the powerful (who were in control) increased their attention to stereotypic information, compared to the non-powerful participants (who were not in control). Also as predicted by the PAC model, because neither condition was dependent, the powerful did not differ from the non-powerful controls in attention to potentially individuating stereotype-inconsistent information. Thus, powerful participants used attentional processes that were consistent with the idea that they stereotyped subordinates by design.

Responsibility effects Priming responsibility had the anticipated effect of increasing overall attention (F(1,66) = 4.42, p = .04), high-responsibility (M = 45.10 s) low-responsibility (M = 39.06 s). Moreover, the responsibility manipulation had the predicted effect of specifically increasing attention, to individuating information, evidenced by a significant two-way interaction between responsibility and information consistency (F(1,66) = 4.65, p = .03). Participants who felt responsible, as compared to those who did not, significantly increased attention to individuating information (F(1,66) = 6.45, p = .01) while their attention to stereotypic information increased only non-significantly (F(1,66) = 2.03, p = .16). This pattern supports our hypothesis that priming egalitarian values would particularly encourage individuating attention processes. However, as mentioned, the three-way interaction between responsibility, power, and information consistency was not significant. Nevertheless, it is useful to consider the effects within the responsibility priming condition when interpreting the lower-order interactions between responsibility and information consistency.

For high-responsibility participants, analyzed separately, the predicted interaction between power and information consistency was not significant (F(1,31) < 1.00 (Figure 1). Examining the two power conditions separately, the overall two-way responsibility × information interaction occurs primarily in powerful participants, consistent with our predictions. Considering first the data for powerful participants, the interaction between responsibility and information consistency was significant (F(1,35) = 6.18, p = .02). Powerful participants' attention to stereotype-consistent information remained the same under the responsibility manipulation (M = 43.45 s and M = 41.69 s for low- and high-responsibility). In contrast, powerholders' attention to stereotype-*inconsistent* information increased under high responsibility (M = 37.96 and M = 43.37 s for low- and high-responsibility). This influence on powerholders' attention to stereotype-inconsistent information fits our predictions. Feeling responsible did not fully prevent powerful participants from attending to stereotype-consistent information, but it increased attention to potentially individuating, stereotype-inconsistent information about subordinates, thereby disrupting the 'by default' process but not the 'by design' process.

For no-power participants, analyzed separately, a simple main effect appears for responsibility (F(1,31) = 6.02, p = .02), with no-power participants greatly increasing their attention under high versus low responsibility to both stereotype-consistent (M = 36.93 and M = 48.06 s for low- and high-responsibility, respectively) and stereotype-inconsistent (M = 38.79 and M = 51.50 s) information.

These responsibility effects were congruent with most of our hypotheses. Responsibility increased overall attention, and, in particular, increased attention to individuating information, which could potentially alter perceivers' stereotypes. Despite the finding that this manipulation did not significantly *decrease* powerholders' attention to stereotypic information (i.e. the predicted three-way interaction

was not significant), increasing attention to individuating information is arguably more important. Perceivers must at least attend to counterstereotypic information in order to form less stereotypic impressions.

Impression rating data Perplexingly, participants' impression ratings were unaffected by any of our treatment manipulations. Overall, participants seemed to be rating targets at or about the midpoint of the scales, regardless of power, responsibility, or target ethnicity. The lack of impression effects is problematic; although we presume the attention patterns observed under the power manipulation reflect stereotype-based processes, the lack of impression outcomes merited further investigation.

Study 2

Overview, design, and hypotheses

As mentioned, the PAC model hypothesizes that some powerful people may simply be unwilling to attend to their subordinates, regardless of other motivational factors. Personality theorists long ago recognized individual differences in the need for power (Winter, 1973, 1988), authoritarianism (Adorno, Frenkel-Brunswik, Levinson, & Sanford, 1950; Allport, 1954), dogmatism (Robbins, 1974; Taylor & Dunnette, 1974) and, more recently, social dominance (Pratto, Sidanius, Stallworth, & Malle, 1994). Some people are predisposed to seek and acquire power or dominate others. People who are high in trait dominance, as measured by a variety of personality measures (e.g. California Personality Inventory) prefer, perhaps even enjoy, having control over others (Gough, 1987). Need for dominance seems, therefore, particularly germane to the definition of power as outcome control. Possibly, individual differences in trait dominance would lead people to assume certain power roles (powerful or powerless) even when the power relationships are undefined. If so, people high in dominance may be more likely to stereotype others, regardless of their actual power in a given relationship. Study 2 considers whether dominance is a dispositional analog to situational power, leading high-dominance perceivers to seek confirmation of group stereotypes when evaluating individuals.

Although the study of individual differences in person perception has a long history in social psychology (Bruner & Taiguri, 1954; Schneider, Hastorf, & Ellsworth, 1979), surprisingly little research has addressed dominance in impression formation. The few exceptions indicate that high dogmatism leads to quicker impression formation (Taylor & Dunnette, 1974) and requires less information for social judgment (Robbins, 1974). This finding fits the argument that the (dispositionally) powerful feel entitled to judge. With regard to dominance, per se, we know only that dominance interacts with target status (Battistich, Assor, Messé, & Aronoff, 1985). High-dominance perceivers rate high-status targets less favorably than they do low-status targets, whereas low-dominance perceivers do exactly the opposite. This finding fits the argument that the dispositionally powerful are motivated to maintain their position. To date, no research addresses the effects of trait dominance on the cognitive strategies that powerful perceivers employ when forming impressions.

We predict dominance will make a difference even when actual control is ambiguous, that is, when individuals have an opportunity to impose their own expectations about control. To test this possibility, we pre-selected participants for individual differences in need for dominance and asked them to engage in the same job selection task described in Study 1, except for leaving situational power ambiguous. This study also manipulated responsibility as in Study 1. Thus, the study involved a 2 (Need for dominance: high vs. low) $\times$ 2 (Responsibility: high vs. low) between-participants factorial design. Dependent measures were the same as Study 1. We anticipated replicating the effects of Study 1, with need for dominance mimicking power in this study. The high and low dominant are not predicted to differ on perceived contingency, so no differences in attention to stereotype-inconsistent information are predicted.

The design thus does not test stereotyping by default. Compared to their low-dominance counterparts, however, high-dominance perceivers should increase attention to stereotypic information. This pattern should be evidenced by a two-way interaction between dominance and information type. Responsibility was also predicted to interact with information consistency, as in Study 1.

Method

Participants Sixty-four native English-speaking undergraduates were recruited from introductory psychology courses at the University of Massachusetts at Amherst. Students received course credit for participation.

All participants had participated in a mandatory pre-testing session at the beginning of the term during which they completed a battery of questionnaires including the dominance scale of the California Psychological Inventory (CPI; Gough, 1987). The CPI dominance scale is a 36-item measure designed to assess individual differences in need for dominance. Participants respond 'true' or 'false' to a series of self-descriptive sentences (e.g. 'I doubt whether I would make a good leader', 'I would be willing to describe myself as a pretty "strong" personality', and 'People seem naturally to turn to me when decisions have to be made'.). The scale has shown adequate reliability ($\alpha = .79$) in previous research, and has been validated on several diverse samples (Gough, 1987). Students who scored in the upper and lower 30 percent on the scale were eligible to participate in this study.

Of the original 64 participants in the study, data for 8 participants were dropped before analysis. Data for 3 participants were incomplete due to a malfunctioning of the audio-recording equipment. Four participants' data (3 high-dominance, 1 low-dominance) were discarded because the participants expressed significant suspicion about the cover story. Finally, 1 participant's responses to the high-responsibility manipulation suggested that the manipulation was ineffective (see Study 1 criterion). The final sample included 56 participants, 15 men and 41 women, distributed in equal proportions among the experimental groups.

Procedure Eligible participants were contacted by phone, as in the previous study. The cover story and experimental procedures followed those of Study 1, with the exception of the power manipulation. Instead, all participants were informed that the firm was 'interested in their opinions in order to get ideas for better ways to evaluate high school students for intern positions', phrasing similar to that of the no-power condition in Study 1, except that here participants were not given any further indication of how their responses might or might not be used in the selection process. The situation was intended to be fairly ambiguous with regard to amount of outcome control. The remaining materials and procedures were the same as Study 1.

Results and discussion

Manipulation checks To test the success of our ambiguous outcome control directions, participants were asked to estimate how much control they would have over applicant outcomes as in the previous study. As predicted, dominance did not influence participants' perceptions of *actual* outcome control; both high and low dominance perceivers indicated moderately low levels of control over hiring applicants.

The responsibility manipulation check posed to participants prior to debriefing also failed to reach significance. Possibly, this manipulation was simply too weak to override extreme individual differences in dominance. Not surprisingly, then, the responsibility manipulation had no effect on measures of interest and will not be discussed further.

Participants' verbal responses to trait information were transcribed, timed, and entered into a mixed-design ANOVA, as in Study 1.[4]

Dominance effects on attention The predicted two-way interaction between dominance and information consistency was significant and fit the predicted pattern, replicating the effect of situational power in Study 1 ($F(1,52) = 3.92$, $p = .05$). High-dominance participants, compared to low, significantly increased their attention to stereotype-consistent information (high-domi-

nance $M = 41.34$, low-dominance $M = 39.61$), thus replicating the stereotyping by design process observed in Study 1. Also as predicted, and replicating Study 1, there was no difference in attention to stereotype-inconsistent information as a function of dominance (high-dominance $M = 38.22$, low-dominance $M = 38.70$).

As predicted, individual differences in the need for dominance influenced how participants attended to information about others in an ambiguous power situation. High-dominance participants were more likely to adopt stereotype-based attention processes when forming impressions of potential subordinates, attending much more to stereotype-confirming information, relative to low-dominance participants. As with situational power (Study 1), in which only control was varied, but lack of dependency was not manipulated, there were no differences in attention to stereotype-inconsistent information. Thus, the by-design predictions were confirmed.

Alternative interpretations might seem plausible for the increased attention to stereotypic information by high-dominance perceivers (i.e. the stereotyping-by-design process). One could turn our findings on end, arguing that the change in attention to stereotypic information is a decrease in attention on the part of low-dominance perceivers because they are somehow uninvolved and attending only minimally. If this were the case, one would expect low-dominance participants to attend less overall, relative to high-dominance participants, yet there were no differences in overall attention as a function of dominance. Likewise, one could argue that low-dominance participants were less racist and consequently had less developed stereotypes than their high-dominance counterparts, resulting in less attention to stereotypic trait information. However, participants' scores on the dominance scale were surprisingly uncorrelated ($r = .08$, *ns*) with their scores on the Modern Racism Scale (McConahay, 1983), which they had also completed during the pretesting session at the beginning of the semester. Apparently, low-dominance perceivers were simply less willing, compared to high-dominance perceivers, to use their stereotypes to guide attention in a situation where they might control others' outcomes. They acted as if they presumed less control.

Impression data As in Study 1, there were no significant effects of our treatment conditions on participants' expressed applicant ratings. The attention data would better support our contention that powerholders stereotype subordinates by design if the impression ratings corroborated these attention patterns. Perhaps our student participants, unfamiliar with having power, felt too uncomfortable to act on their opinions in the present setting. Study 3 alters the context in which perceivers judge targets, in an effort to create a more familiar context to student participants.

Study 3

Studies 1 and 2 provide initial evidence supporting the hypothesized relationship between power (i.e. outcome control) and cognitive processes presumed to reflect stereotype maintenance: participants who believed they had (Study 1) or merely desired (Study 2) outcome control paid more attention to stereotype-consistent information about job applicants. This pattern of results fits the predicted model of attention stereotyping subordinates by design, by effortfully confirming group stereotypes about subordinates. However, these studies lack an appropriate comparison group for testing the hypothesis that powerful perceivers' attention stereotypes subordinates *by default, by ignoring stereotype-inconsistent information.* Testing attentional stereotyping by default requires comparing the non-dependent powerful to a group of dependent perceivers who are, according to predictions, motivated to individuate by attending to stereotype-inconsistent information. The primary goal of this third study is to include a relevant comparison group for testing both hypothesized stereotype mechanisms – by default and by design – respectively, inattention to stereotype-inconsistency *and* attention to stereotype-consistency. The study also attempted to investigate impressions more carefully.

Overview

We developed a new paradigm to test these hypotheses, asking students to participate in a study of task allocation in working groups. In this within-participant design, groups of participants expected to meet at the end of the experimental session to play three different work roles simultaneously. To unconfound these relationships, participants believed they would play the roles in round-robin fashion, with regard to different targets. Participants would be a (powerful) task *allocator* distributing work tasks to some targets, a (powerless) *receiver* performing tasks for other targets, or an (unrelated) *observer* uninvolved in the distribution or execution of tasks.[5] Thus, participants had separate superiors and subordinates at the same time. Participants also believed that they could win a prize based on their effectiveness at performing the tasks they were assigned by their allocators. The experimenters would judge the powerless receivers, but receivers could increase their chances of winning a prize if the powerful allocator gave them a particular type of task to perform. That is, the powerful controlled the powerless' chances of winning by controlling how the tasks would be allocated (asymmetrical outcome control). Thus, the powerful had control but no contingency, the powerless had no control but were contingent, and observers had neither. (Simultaneous control and contingency constitutes interdependency thoroughly studied already in our previous research.) The effects of control emerge in comparing the powerful and the other two, whereas the effects of contingency emerge in comparing the powerless to the other two.

Hypotheses Participants' impression strategies were predicted to vary as a function of the target's role relative to the perceiver (see Table 1). Overall, the predicted patterns should produce a two-way interaction between power role and information consistency. A priori predictions for each role are presented below.

Power-irrelevant perceivers Perceivers evaluating power-irrelevant observer targets provide the baseline for comparison; these perceivers were neither contingent on targets nor controlled their outcomes. Power-irrelevant perceivers evaluating information about observer targets should pay low overall attention to the information, compared to participants evaluating targets in the other two roles because observer targets were unrelated to the allocation of any tasks (i.e. resources).

Powerless receivers Besides the overall attention main effect relative to power-irrelevant observer targets, attention to stereotype-*consistent* information about powerful allocator targets should not differ much from power-irrelevant targets, because neither are the objects of perceiver

Table 1. Study 3: Power roles and predictions

Perceiver role	Target role	Perceiver contingent on target	Perceiver controls target outcomes	Hypothesized attention
Power-irrelevant	Observer	No	No	• Low overall
Powerless (Receiver)	Allocator	Yes	No	• High to stereotype inconsistency • Low to stereotype consistency
Powerful (Allocator)	Receiver	No	Yes	• Low to stereotype inconsistency • High to stereotype consistency

control, the critical factor in stereotyping by design. However, powerless receivers evaluating powerful targets were predicted to adopt accuracy motives and hence to individuate the powerful: a specific increase in attention to stereotype-inconsistent information.

Powerful allocators The predictions regarding powerful participants' attention to the (powerless) receiver targets are of greatest interest to this study. The by-design hypothesis predicts that the powerful ought to be motivated to process target information stereotypically, and to do so effortfully, increasing their perceived ability to judge. For this reason, powerful perceivers' overall attention to powerless targets should be greater than that of attention to power-irrelevant observer targets, but about equal to that of powerful allocator targets. However, unlike accuracy-motivated perceivers attending to powerful targets, powerful perceivers attending to powerless receivers should be motivated to attend to stereotype-consistent information, relative to the other two targets. Their attention to stereotype-inconsistent information should be much lower than powerless (contingent) perceivers attending to the powerful but equivalent to that paid to power-irrelevant observers, thereby also stereotyping by default.

Method

Design The design of the study was complex, involving the manipulation of independent variables both within- and between- participants, as well as within-targets.

Independent variables

Power role Power role (allocator, receiver, observer) was manipulated within participants, with participants believing they would assume each role in relation to a different participant (see Table 1). When in the (powerful) allocator role, participants believed they controlled powerless receivers' outcomes by means of allocating a task that would improve the receivers' chances of winning a prize. Conversely, when in the (powerless) receiver role, participants believed they were dependent on the allocator for receiving a task that would improve their own chances to win. Finally, vis-à-vis observers, participants were told that they would passively observe a task allocation interaction and therefore neither controlled another nor were dependent in that particular interaction.

Target group membership Gender was selected as the salient target category because it would be easy to manipulate in the context of the experiment without arousing suspicion as to the true nature of the hypotheses. While the proposed model of power and stereotyping would not predict any direct effects of participant or target gender, the literature regarding perceptions of power suggests that men and women view power in different ways (Yoder & Kahn, 1992). Thus, possibly, gender differences might have indirect effects on how the participants responded to being in different power roles. Given possible gender effects, the goal to examine powerful perceivers in a within-participant design, and the unfeasibility of fully crossing gender with power role within-participants (i.e. increasing the number of targets), a compromise design was employed. Target gender was randomized between participants for the powerful and power-irrelevant roles, with the constraint that each participant received one male and one female across these two targets. Gender was randomized *within*-participants for two powerless receiver targets. This design would allow a full within-participant test of any possible interactions between gender effects and the main independent variable of interest (i.e. powerful perceivers' perceptions of powerless receivers).

Dependent variables Attention to trait information and participant ratings of targets served as the primary dependent variables of interest to this study. Both measures were collected and recorded via computer.

Participants Participants were 51 undergraduate psychology students at the Université Catholique de Louvain in Louvain-la-Neuve, Belgium. Participants received optional course credit. Eight of the original participants (4

men, 4 women) were dropped from analyses because their overall attention scores were 3 or more standard deviations above the mean (see results below). In addition, the data for one participant who reported French as a second language were dropped. A total of 42 participants remained in the final sample.

Procedure Participants arrived at the lab in groups of six to twelve and were seated in alternating chairs in front of Macintosh LCii computers. Instructions on the computer screen directed participants to wait quietly until the researcher indicated they could begin. With the exception of a brief introduction by the experimenter, the procedures were carried out on the computer using SuperLab software for the Macintosh computer. All materials and instructions were presented in French. A female experimenter first explained to participants that the majority of the study would take place on the computers, but that participants would break into 'work groups' at the end of the study to distribute and perform some tasks.

After an initial task to familiarize participants with the keyboard, participants learned that in Phase I they would 'tell us about themselves', in Phase II they would learn about other members in their alleged workgroups, and, in Phase III they would break into groups in an adjacent room to decide who would perform the various tasks. Phase III, of course, did not take place.

Phase I In the first phase, participants supplied their gender, age, and language skills. Afterward, the computer instructions explained that participants would next respond to a 140-item test allegedly developed by industrial psychologists to study the characteristics of people in work environments. In its original form, the counterfeit questionnaire was 'too long to administer to each participant', so the computer would supposedly select 20 items at random for each participant to complete. Each item was a self-descriptive statement (e.g. 'Sometimes I am lost in thought and do not realize that others are speaking to me.'). Participants were asked to indicate whether or not each statement described them by indicating yes or no. In fact, the questions were bogus and served only to provide a cover story for the presentation of trait information about other members later in the study.

Phase II Next, participants were told they would have a chance to get to know the other members of their group before the actual interaction. Participants were told they would receive a profile for each member of their group with a codename (presumably to protect the anonymity of participants' responses), and an indication of that person's role (e.g. receiver) in the upcoming interaction with them. The profile also contained the target's gender and year in college. French nouns that were both connotatively and linguistically masculine or feminine were selected as codenames to enhance the salience of the target's gender (e.g. *camion* – truck, for a male target, and *dentelle* – lace, for a female target). In addition, gender-neutral language was avoided and all directions regarding the targets were presented in *both* masculine and feminine form (e.g. 'his/her') to enhance the salience of target gender. This profile screen served, therefore, to identify each target's gender as well as the perceiver's power relationship with regard to the target. Following each target profile screen, participants received a series of eight gender-relevant trait sentences developed from pre-tested traits.

Trait information Trait information was manipulated within-target. For each target, participants received eight trait sentences pre-tested for valence and consistency with gender roles. A total of 18 traits that uniquely described members of each group were selected for the study. Sample female stereotypic traits included *conscientious*, *tender*, *jealous*, and *demanding*, whereas sample male stereotypic traits included *self-confident*, *ambitious*, *authoritarian*, and *too rational*. Valence and consistency were crossed and randomized within target such that for each target half of the information was positive, half was negative, half was consistent with the target's gender stereotype, and half was inconsistent.

Participants believed that each sentence, ran-

domly presented by the computer, had been endorsed as self-descriptive by the target in Phase I. Participants read each statement one at a time on the screen, pushing a designated key to continue when they were ready to advance to the next sentence, at their own pace. The time between presentation of the statement and pressing the key to continue to another statement constituted the attention measure. Immediately after completing the series of eight trait sentences, participants responded on 7-point scales to 13 impression items similar to those used as target trait information.

Participants received information about and evaluated a total of four targets: one powerful, two powerless, and one power-irrelevant observer. The order of target roles was counterbalanced between participants, with the constraint that participants would not evaluate the two powerless receivers in consecutive order. Participants followed this procedure, reading profiles, trait information, and answering impression items, for each of the four targets. Afterward, participants completed a written debriefing questionnaire designed to assess overall suspicion. Having completed the computerized procedures, participants were then fully debriefed as to the purpose of the study, given credit, and dismissed.

Results and discussion

The SuperLab program recorded participants' attention to the trait information for each target to the hundredth of a second. Timing began with the presentation of the stimulus sentence on the screen and ended when participants pressed the designated key (*C*) to continue to the next screen. Timing scores were created for each target by averaging participants' attention to each type of trait information (valence × gender stereotype consistency). Thus, each participant had four attention scores per target: positive/consistent, negative/consistent, positive/inconsistent, and negative/inconsistent. Because target gender had no effect, the times for the two powerless receivers were averaged to create a single powerless receiver score. Participants with scores more than 3 standard deviations above the mean on any of the four (valence × consistency) scores were deleted before further analyses. There were no effects of participant gender for any of the analyses and therefore gender is not considered further.

Attention scores for each perceiver role were submitted first to a within-participant ANOVA using the SPSS-X MANOVA procedure. Perceiver role, information valence, and consistency were entered as within-participant factors. Although this analysis produced effects consistent with our hypotheses, a target-by-target between-participants MANOVA of the attention data suggested that participant fatigue influenced these data. The attention patterns began to fade in the third target and no longer appeared for the fourth target. Because these fatigue effects could be masking the direct effects of the power manipulations, we opted to consider the effects for the first target in a between-participants analysis.

Data for the first target only were submitted to MANOVA with perceiver power role as a between-participants factor and information valence and consistency as within-participant factors. This analysis yielded two two-way interactions (role × consistency; role × valence).

Role × consistency The predicted two-way interaction between role and consistency ($F(2, 39) = 8.67$, $p = .001$) strongly supports our hypotheses that powerful perceivers' attentionally stereotype by design *and* by default. Power-irrelevant perceivers attending to observer targets (Figure 2) paid low and equal attention to stereotype-consistent and inconsistent information, as predicted. The powerful increased their attention to consistent information about the powerless, relative to the power-irrelevant, and as a result, the predicted increase in attention to consistent information for powerful perceivers ($M = 6.29$) relative to power-irrelevant perceivers ($M = 4.84$) was statistically significant ($t(39) = 2.94$, $p = .02$), also as predicted. That is, in the comparison highlighting their control over others, powerful perceivers attentionally stereotyped powerless receivers by design.

The predicted by-default pattern occurred

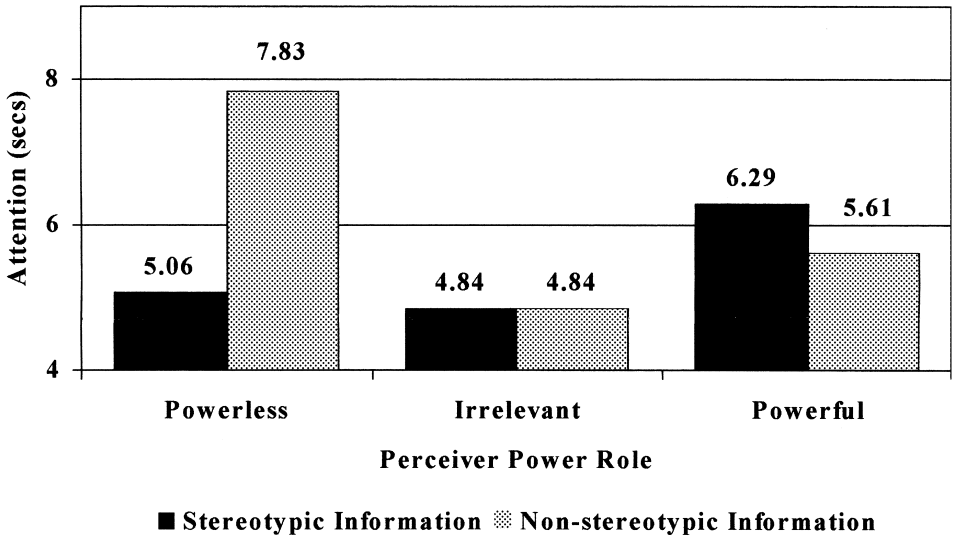

Figure 2. Mean attention to stereotype-consistent and inconsistent information as a function of perceiver's role: Study 3.

for attention to stereotype-inconsistent information. There was no statistical difference in attention to inconsistent information about the targets when participants were not contingent: powerful (M = 5.61) and power-irrelevant perceivers (M = 4.84, $t(39) = 0.90$, *ns*). In comparison, powerless receivers evaluating their powerholders paid significantly more attention to inconsistent information (M = 7.83), compared to perceivers evaluating the other two targets ($t(39) = 2.65$, $p < .01$). As predicted, the powerful stereotyped subordinates by default, as well as by design.

Role × valence The pattern of means with regard to the role x valence interaction ($F(2, 39) = 9.28$, $p = .001$) suggests that the role manipulation had a significant influence on participants' attention to positive information. Attention to negative information was statistically equivalent across the three power roles. Attention to positive information, in contrast, increased in the two power-related roles. The increase in attention to positive information for powerful (M = 5.98) relative to power-irrelevant perceivers (M = 4.71) was only marginally significant ($t(36) = 2.27$, $p < .10$). However, the powerless participants paid significantly more attention to positive information about the powerful (M = 7.81), as compared to the baseline power-irrelevant targets ($t(36) = 5.54$, $p < .001$). This pattern is consistent with previously cited research (Stevens & Fiske, in press) and suggests that when perceivers become outcome dependent, positive information becomes increasingly more important.

Overall attention differences Planned comparisons of the *overall* differences in attention by role revealed that participants were paying equal attention whether they were powerless (M = 6.45) or powerful (M = 5.95), as predicted. Moreover, power-irrelevant perceivers' attention (M = 4.84) was significantly lower than that of the powerless ($t(39) = 2.33$, $p = .02$) or the powerful ($t(39) = 1.61$, $p = .05$). These findings strongly support the hypothesis that participants would be relatively uninterested in targets who were unrelated to the allocation of tasks. More importantly, it suggests that attentional stereotyping by design is effortful; powerful perceivers are actively focusing on stereotype confirming information about subordinates.

As predicted, powerful participants focused attention on information best supporting stereotypes about powerless targets. Powerless participants used equally effortful strategies to attend to information that could individuate the powerful. The results clearly show that perceivers are not only less interested when power is irrelevant to the relationship with the target, but that the powerful and the powerless are equally motivated to attend to one another. What differs for participants in the different power roles is not *how much* they attend, but rather *to what* they attend. For the powerless, attention focused on individuating the powerful, gathering stereotype-inconsistent information presumably to form a more accurate impression. For the powerful, effort concentrated on stereotype-consistent information. In all, these data provide strong evidence that the powerful can adopt attention processes that stereotype subordinates both by default and by design.

Impression data Once again, participants' impression ratings were unaffected by our manipulations of outcome control and also outcome dependency. Much to our concern, participants in the present study did not alter their impressions in accord with their attention strategies. As in the previous two studies, participants appear to be resisting judgment; mean responses across experimental conditions were at or around scale midpoints. Given that prior studies have made a convincing link between attention and judgment, it is important to determine if such a link is observable in power relationships. Study 4 directly addresses this issue.

General discussion: Studies 1–3

Further consideration of participant reactions across the three studies reveals factors in the experimental contexts that may have inhibited participants from expressing any judgments. Mean impression ratings across all three datasets were not significantly different from scale midpoints, regardless of power manipulations. This, coupled with overall low variability in impression ratings indicates that participants were unwilling to express judgments. Reactions during debriefing also suggest that participants were uncomfortable rating targets. Considering the experimental contexts more carefully, one can see why this might be the case. Studies 1 and 2 involved a job selection task wherein participants were asked to choose applicants to retain as interns. More than half of the participants in these studies expressed a desire to meet or interview applicants directly before making decisions. In Study 3, participants were asked to judge targets prior to an alleged meeting. Comments from debriefing indicated that the anticipated meeting with targets led participants to hesitate from making premature ratings. Such a reticence to judge is neither remarkable nor impractical; as described by SJT, social norms dictate that one should refrain from publicly judging those one is about to meet.

An alternative, but equally compelling interpretation of these null impression effects is that the nature of the target information encouraged piecemeal individuation. Recall that all three studies presented participants with equal amounts of stereotype-consistent and inconsistent information. In prior research relevant to our theory, perceivers encountering significant amounts of inconsistent information sometimes engaged in piecemeal trait integration (Fiske et al., 1987). When people are faced with increasing amounts of disparity, they try to resolve the inconsistencies, resulting in less stereotypic impressions. Once more, participant responses during debriefing support this explanation. Participants across all three studies commented that target information was unusually mixed or confusing. Indeed, one participant commented that the targets seemed to have 'personality disorders' because the traits seemed so contradictory. Hence, although powerful participants in our studies were attending more to stereotypic information, the sheer amount of non-stereotypic information may have demanded some resolution of the inconsistencies. If so, one would predict no difference in impression ratings as a function of power role.

This argument would be consistent with prior interdependence studies using similarly mixed stimulus materials (e.g. Ruscher & Fiske, 1990; Ruscher et al., 1991). In those studies, perceivers were presented with balanced presentations of stereotypic and non-stereotypic target information; interdependence manipulations influenced attention and variability in impression ratings, but produced no mean differences in impression ratings. In the present studies, a few variability effects did emerge; however, inconsistencies across studies reduce their interpretability. These types of global impression ratings may simply be insensitive. The next study addresses this issue.

In sum, our design – providing equivalent amounts of dis/confirming trait information – may have subverted our goals to capture effortful stereotype maintenance. A skeptic might argue that this analysis counters our hypothesis that the powerful stereotype by design. We would contend, however, that such a mixture of in/consistent information seemed unusual to our participants because it is rare to encounter such combinations of trait information in naturally occurring social contexts. Moreover, the behavioral confirmation literature indicates that when perceivers interact with targets, they elicit and receive greater amounts of expectancy consistent information (Snyder, Tanke, & Berscheid, 1977; see Snyder & Haugen, 1994 for a review). Indeed, the powerful are especially adept at eliciting behavioral confirmation when interacting with subordinates (Copeland, 1994). These data suggest then that the biased attention patterns observed in our data ordinarily would lead powerholders to form stereotypic impressions in settings where they normally cannot only choose information, but also elicit target responses to fit expectations.

In either case, the absence of any impression data in the present studies leaves open to debate whether powerful perceivers were attending to stereotypic information to refute or to maintain their stereotypes. It is important, therefore, to determine whether power directly alters judgment. Study 4 implements a third paradigm to manipulate power and examine its effects on judgment.

Study 4

Our design draws directly on an idiographic methodology introduced by Pavelchak (1989; see also Fiske et al., 1987) to disentangle category-based vs. piecemeal (attribute-based) processing in social judgment. In this paradigm, individuals first provide ratings of how much they like many specific target categories and specific traits, independent of associations with a particular target person. Later, participants rate how much they like individual targets on the basis of category memberships and traits. Knowing perceivers' own a priori category and trait judgments as well as their subsequent target ratings allows a measure of the degree to which final target ratings correspond (i.e. correlate with) category-based vs. trait-based judgment. Because this technique employs perceivers' own prior judgments of target qualities, it is necessarily a more sensitive measure, as compared to average target ratings.

Applying this paradigm to our hypotheses regarding power, we predict the differential use of category vs. trait information in target judgments to vary as a function of power. Just as the powerful and powerless differentially attend to target information as a function of their roles, their ratings of one another should reflect these motivated attentional preferences. For powerful perceivers, who attend significantly to category-consistent information but relatively ignore inconsistent information, target ratings should be significantly related to perceivers' own prior ratings of target categories but unrelated to their own prior ratings of target attributes. Moreover, powerholders' reliance on category information should be significantly greater than that of powerless perceivers. In comparison, powerless perceivers' target ratings should be unrelated to prior category ratings but significantly related to prior ratings of target attributes. Finally, for powerless perceivers, the strength of the relationship between their own trait attribute ratings and target ratings should be significantly higher, compared to that of powerholders.

We hypothesized an interaction between power role and type of processing. Powerless

assistants were predicted to rate targets on the basis of individuated processing of trait information, not on the basis of targets' category memberships. In contrast, powerful bosses should stereotype by default and by design. Consequently, these perceivers' ratings should reveal significant category-based processing, but little relationship to their own prior ratings of target traits.

Method

Participants One hundred and three undergraduate psychology students were recruited from courses at the University of Massachusetts at Amherst. Participants received course credit in exchange for volunteering. Of the original sample, two participants were not included in analyses because they reported confusion over their power role assignments at the end of the study. A total of 101 participants remained in the final sample, 30 males and 71 females.

Procedure Participants arrived in groups of 8 to 10 and were greeted by a female experimenter. Participants were seated individually at carrels with dividers that extended above eye level. The experimenter explained that the goal of the study was to understand how people from different college majors work together. At this point, participants were told that they would first complete some questionnaires individually and would later be matched with another participant to work on a joint task.

Participants first received a questionnaire labeled 'HMAS' with no further explanation of its significance. The 14 items on the scale focused on leadership qualities and included questions such as 'How good is your ability to supervise others?' and 'How much did you like your co-workers on your last job?'. In reality, the bogus scale was included only as a means of justifying participant assignment to power roles later in the study. Once all participants had completed the questionnaire, the experimenter collected the forms and explained that she would be scoring their responses while they completed the next set of questionnaires.

Measures of category and trait liking We adapted two separate questionnaires from Pavelchak (1989) to assess participants' individual liking for various college majors and personality traits. The Category Liking Questionnaire asked participants to rate how much they liked each of 24 academic majors. Participants responded using a scale from 1 *dislike very much* to 15 *like very much*. The list of majors included the following four majors used as category memberships for targets later in the study: *early childhood education*, *sociology*, *pre-med*, and *psychology*.[6]

The Trait Liking Questionnaire asked participants to use the same 15-point scale to rate how much they would like a person with each of 46 different personality traits. Included were the following traits used as stimulus traits later in the study: *bright*, *closed-minded*, *content*, *efficient*, *ethical*, *helpful*, *impatient*, *nice*, *open-minded*, *polite*, *quiet*, *realistic*, *restless*, *self-disciplined*, *sociable*, *studious*.

Once participants completed the category and trait ratings, respectively, they completed a distracter task that involved estimating the prevalence of 14 personality traits in the general population. The traits used in this task differed from those presented as target information, serving to inhibit participant awareness of the link between their initial ratings and the target information. Finally, participants responded to a series of questions about themselves, providing their own academic majors and four traits that they believed were self-descriptive. In providing self-descriptions, participants were asked to complete the sentence stem 'People say I am . . .' with four separate descriptors. Participants were led to believe that these materials would be distributed to other participants in the next phase of the study. A second female research assistant collected these materials and left the room, presumably to photocopy sets of the materials for the next part of the study.

During this time, the experimenter delivered the power manipulation. She explained that the HMAS was actually the 'Harvard Management Aptitude Scale', allegedly designed to measure a person's potential for being a successful manager, including whether or not a person had potential for assessing others in business situations. The experimenter further purported that the scale was widely used in businesses and

could 'predict quite well people's success at judging and managing other people'. Participants were informed that their own role assignments – boss or assistant – were a function of their performance on the HMAS; participants who allegedly scored high on the scale would play the role of boss, while the remaining participants would play the role of assistant. In reality participants were randomly assigned to role condition. In a third experimental condition, participants were told that a third group of participants would serve as controls, evaluating other control participants while bosses and employees rated one another. Control participants, therefore, had neither control over nor were contingent on the targets they evaluated.[7]

The power role manipulations were first distributed to individual participants in writing. The written form identified the participant's own role assignment and indicated what both boss and employee roles required. The form explained how much control participants in each role would have over partner selection for the performance task (outcome control manipulation) and how lottery tickets for the prizes would be distributed (outcome dependency manipulation). Participants were instructed to read the form silently and to acknowledge the role assignments by signing the bottom of the page.

At this point, the experimenter summarized the role manipulations orally:

> If you are a boss, then your job during the task is to figure out what needs to be done and then to implement it. You will be the leader of the two of you. The assistant's job will be to follow the boss' instructions. Now to make this close to the working world, we will be actually paying you something to take part in this task. This is an added bonus to your credits.
>
> If you are a boss, you will be given 6 lottery tickets toward the cash prize at the end of the semester. If you are an assistant, you will receive 2 to 6 lottery tickets, depending on how well you and your boss have completed the task. We will have two lotteries, each for $50; one for assistants and one for bosses. The winners of the lotteries will receive $50 at the end of the semester.

In the control condition, participants were told about bosses and employees, but were also told:

> If you are in the control condition, you will work on the task independently and receive 3 lottery tickets. We will have three lotteries all for $50, one for assistants, one for bosses, and one for the controls.

When the experimenter concluded the explanation, the research assistant returned with packets of photocopied information, presumably describing the participants in the room. In fact, the packets contained the target information previously prepared by the experimenters.

Target information Participants received category (i.e. college major) and trait information about four different targets. Trait information was ascribed to targets on the basis of consistency with the category major as determined by pre-test. Unlike the previous three studies, where targets were described by equal numbers of stereotype-in/consistent information, each target in the present study was described by three stereotype-consistent and only one stereotype-inconsistent trait (see Table 2). This ratio

Table 2. Study 4: Target category and trait information

	Target 1	Target 2	Target 3	Target 4
Major	Psychology[a]	Early childhood education	Sociology	Pre-med
Traits	closed-minded[b]	impatient[b]	realistic[b]	self-disciplined
	bright	nice	ethical	efficient
	helpful	polite	quiet[b]	content
	restless[b]	studious	open-minded	sociable[b]

[a] Target not included in analyses.
[b] Stereotype-inconsistent traits. All other traits stereotype-consistent.

of traits has the advantage of reducing the possibility that perceivers will engage in individuated processing simply to resolve what are perceived to be significant inconsistencies among the traits. Traits were presented in sentence format, using a fixed sentence stem ('People say I am . . .'). Trait presentation was randomized across targets and between-participants. Order of target category presentation was fixed: psychology as a baseline in-group rating, followed by early childhood education, sociology, and pre-med.

Participants found information regarding one target's academic major and personality traits in each packet. In addition, each of the four packets included a questionnaire labeled 'PLQ' asking participants to respond to the question 'How much do you like this person?' on a scale of 1 *dislike very much* to 15 *like very much*. Participants were instructed to review the packets one at a time. Participants in the control condition were told that control participants would also receive packets so they would have something to do while bosses and assistants rated one another. Control participants believed they were rating other control participants. When participants had completed all four packets, they completed a final questionnaire including manipulation checks before they were probed for suspicion, debriefed, and dismissed with credit for their participation. A random lottery was held at the end of the semester for the $50 prizes.

Results and discussion

Recall that our hypotheses concern the extent to which target ratings are related to target category membership, versus ascribed traits, as a function of perceiver power. Powerless assistants were predicted to individuate potential bosses; hence liking for targets should be significantly related to liking for the traits ascribed to targets, but unrelated to liking for targets' academic majors. In contrast, participants playing the roles of powerful bosses should stereotype potential subordinates by default and by design, making target ratings that related significantly to idiographic ratings of the target's academic major, but that were unrelated to the ratings of the individual target traits. For theoretical reasons, we were interested only in participants' impressions of targets perceived as out-group members. Because the psychology major target would be an in-group member, we did not include these target ratings in our analyses.

To test the relationships, we first created within-participant measures of impression processing for each of the three remaining targets – early childhood education, sociology, and pre-med. For each target, we calculated a *category-based outcome* score for target liking by computing the correlation between the participant's liking for a target and that person's own prior liking for the major ascribed to that target (e.g. how much liking for the pre-med target correlated with liking for pre-med majors in general). Next, we calculated an *individuated outcome* score for target liking by computing the correlation between the participant's liking for each target and that person's own prior liking for each of the four traits ascribed to that target (e.g. how much liking for the pre-med major correlated with general liking for each of the four traits ascribed to the pre-med major). These analyses resulted in six correlations for each participant, each of the three target ratings with corresponding major and corresponding attribute set. Because we were not interested in participants' ratings as a function of academic major per se, we then collapsed across the three target majors by Z scoring and averaging the two types of outcome scores for each participant. This procedure resulted in two average outcome scores (category-based and individuated) for each participant. These correlations reflect the extent to which each participant's target ratings were based on category membership (i.e. stereotyping by default and design) versus personality traits (i.e. individuating outcomes), respectively.

To determine the effects of the power manipulation on participant ratings, we tested the mean outcome processing scores for each power role against zero (Table 3). These analyses revealed a pattern of results consistent with five of six predictions. First, for participants in the control condition, neither category-based

Table 3. Study 4: Mean perceiver processing scores

	Mean score (r)	
Power role	Category-based	Individuated
Control (power-irrelevant)	.17	.24
Assistant (powerless)	.13	.65[a]
Boss (powerful)	.37[a]	.28[b]

[a] $p = .01$; [b] $p = .06$.

($r(36) = .17$; $z = .98$, *ns*) nor individuated ($r(36) = .24$; $z = 1.49$, *ns*) outcome scores significantly differed from zero. In comparison, for powerless assistants, individuated outcome scores ($r(30) = .65$) were significantly greater than zero ($z = 4.00$, $p < .0001$), whereas category-based processing scores ($r(30) = .13$) were not ($z = .67$, *ns*). For the powerless, these two within-participant outcome scores also significantly differed from one another ($t(28) = 3.13$, $p < .005$). As predicted, powerless perceivers' liking for potential bosses was significantly related to how much they liked the individual traits ascribed to those targets, but not at all related to how much they liked the target's academic major. These ratings would be predicted from attentional patterns observed in Study 3, which manipulated outcome dependency.

For participants playing the role of powerful boss, these analyses revealed a pattern of findings that partly confirmed our hypotheses. As predicted, powerholders' category-based processing scores ($r(33) = .37$) were significantly different from zero ($z = 2.13$, $p < .01$), and they were the only group for whom this was true. However, powerholders' individuated processing scores ($r(31) = .28$) were also marginally significant when tested against zero ($z = 1.53$, $p = .06$). Moreover, these two processing scores did not significantly differ from one another ($t(32) = -.45$, *ns*). Thus, for powerful participants, liking for potential subordinates was related to both liking for the target's academic major and liking for the individual traits ascribed to the target. These data only partially support our hypotheses in that we predicted powerholders' ratings to reflect only category-based processing. Comparison of the processing scores across power roles, however, suggests that powerholders were not engaging in the same high level of individuated processing as the powerless, providing further support for our hypotheses. Consistent with predictions, individuated processing was significantly greater for powerless perceivers ($r = .65$) compared to that of the powerful ($r = .28$), ($z = 1.81$, $p < .03$), who did not differ from control participants ($r = .26$), ($z = .12$, *ns*). The overall pattern of processing scores replicates the pattern of attention scores observed in Study 3: powerless assistants had the largest individuation scores, powerful bosses the largest category-based scores; controls were lowest on both.

These data support five of our six hypotheses regarding the relationship between perceiver power and social judgment. As predicted, whether participant ratings reflected stereotype-based impressions versus individuation was a function of perceiver power role. When perceivers were powerless (i.e. lacked control over but were contingent upon targets), their liking for targets was related significantly to their own prior evaluations of target attributes, but unrelated to their evaluations of the categories to which targets belonged. In contrast, when perceivers were powerful (i.e. had control over but were not contingent upon targets), liking for targets was significantly related to liking for target category memberships. This relationship between category ratings and target ratings was significantly higher than that of the powerless. In addition, powerholders' liking for targets was marginally related to how much they liked the traits ascribed to the target. Contrary to our predictions, this suggests that the powerful in this study were not completely ignoring potentially individuating target information. Yet, relative to the powerless, the powerful were engaging in marginally less individuated integration of the trait information. This pattern of results replicates the attention pattern observed in Study 3, strongly supporting the stereotyping by default

hypothesis, and, to a lesser extent, the stereotyping by design hypothesis.

These data go some distance toward describing how powerful perceivers were interpreting target information. Powerholders' ratings were significantly related to target category membership and hence some level of stereotyping is evident in these data. But we did not anticipate that powerholders' ratings would be marginally related to target traits (albeit less than powerless perceivers' trait correlations). The one unanticipated result out of six predictions suggests that this power manipulation may have been somewhat less than fully effective. Recall that Study 4 operationalized power on the basis of perceived control over, as well as outcome contingency on an alleged partner in a performance task. Possibly, our explanation of the roles did not clearly disentangle these two constructs – control and outcome contingency – for our participants. Because participants in both roles were anticipating working together and receiving a joint performance evaluation, powerful participants may have been somewhat concerned about being good bosses in the upcoming task, and this may have led to some individuation of potential partners; thinking more about potential partners' attributes would afford some prediction of how positive the interaction might be, reducing concerns about looking negative in the eyes of the experimenter. Regardless of the prize manipulation, one would expect participants to be motivated to minimize looking 'bad' in the eyes of the experimenters. If so, these presentational concerns could have acted as a mild accountability manipulation, leading the powerful to think more carefully about partner attributes.

General conclusions

Overall, data from the four studies support our theory of power and motivated stereotype maintenance. In Studies 3 and 4, powerholders stereotyped subordinates by default with regard to both the attentional processing of subordinate trait information (Study 3) and judgments (Study 4). Studies 1 through 3 provide evidence that the powerful stereotype subordinates by design with regard to the attentional processing of subordinate trait information, and Study 4 demonstrates that powerholders use category information more than the powerless do. The results of Study 4 are particularly important in that these data provide the first empirical evidence of impression differences as a function of social power. An important next step is to develop a paradigm within a single study directly linking the attention patterns observed in Studies 1–3 with the impression patterns of Study 4.

The pattern of results across the four studies suggests that powerholders may indeed be especially vulnerable to stereotyping their subordinates, unless the powerful are acutely aware of their own responsibility. In addition, these findings open the door to questions about power, impression formation, and discrimination. One question arises with regard to the potential resiliency of powerholders' stereotypes, in light of their effortful processing of stereotypic information and the non-conscious nature of these responses. If powerful people simply stereotyped by default when attending to target information, then overriding their stereotypes might be accomplished by increasing accuracy motives, and getting the powerful to look beyond their stereotypes, as the responsibility data suggest in Study 1. However, to the extent some powerful people work effortfully though perhaps unconsciously to maintain their stereotypes about subordinates, as these data suggest they might, we would expect power to have longitudinal effects on stereotype resistance to inconsistent information.

In our view, power can have deleterious effects on impression formation. We assume that stereotyping subordinates violates normative expectations to judge others on the basis of their individual merits, rather than their social group memberships. Recent theorists have, however, argued that stereotyping others is not necessarily evil, that stereotyping may be pragmatic. For example, stereotyping a child molester is not altogether a bad idea given recidivism rates among this population of criminals and the high cost of giving benefit of the doubt. Stereotyping could pragmatically lead a power-

holder to protect children from possible harm by not hiring the convicted molester as a day-care worker. In other situations, perceivers may apply mutually accepted stereotypic expectations about each other in a way that facilitates social interaction (Leyens, 1983). According to this perspective, stereotypes serve to smooth the way by providing ground rules for social interaction. For example, if two different-sex sexist perceivers adopt and apply the same gender role stereotypes (e.g. that the man should lead the conversation, etc.), the two individuals may be able interact more smoothly than if they entered the interaction sans expectations, or with conflicting expectations. Note, however, that both participants share the respective stereotypes of themselves and each other. Although it is conceivable that stereotyping can be pragmatic, given a major kernel of truth or shared beliefs in the stereotype, we maintain that the stereotyping of subordinates is more often negative and unfounded, facilitating discrimination and in-group favoritism. Moreover, when the powerful stereotype, important economic and life outcomes are often at stake for their powerless subordinates. In our opinion, the potential harm resulting from negatively stereotyping subordinates far outweighs any potential benefits to be gained by kernels of truth or smoothing social interaction via stereotyping.

The present studies constitute a first step in understanding the complex relationship between power and stereotyping, providing several unique contributions to the empirical literature: (1) These studies constitute the first documentation of powerholders' attentional biases – increasing attention to stereotypic information while ignoring non-stereotypic information, and hence, (2) the first evidence that perceivers can be motivated to *confirm* their stereotypes effortfully by engaging in different cognitive strategies. (3) The judgment data (Study 4) are the first data showing that the situationally powerful can be biased in their judgments, relative to the powerless. (4) Finally, Study 2 provides the first link between individual differences in dominance and cognitive processing during impression formation. These dominance effects are also noteworthy because they provide an important anchor for interpreting the attention data in Studies 1 and 3. It is against the responses of the chronically dominant that one gains perspective on the situational effects of power. The finding that the situationally powerful have the same attentional response as the dispositionally dominant, especially in light of the short-term nature of the contexts employed, supports the ecological validity of the present studies. Although the transient nature of the power relationships likely minimized concerns over loss of power, the situationally powerful increased their attention to stereotypic information just as did dominant perceivers, who were likely more concerned with having control over potential subordinates. These data illuminate power as a situational variable. Arguably, motives to maintain power and to protect power-relevant identities would be greater in naturally occurring, long-term power situations. If so, these initial data may be a conservative demonstration of what transpires in naturalistic power relations where the stakes can be much higher.

Notes

1. Because power is manipulated in a transient experimental context in which participants relatively unfamiliar with wielding power (i.e. college students) judge subordinates, desires to maintain power may play little or no role in these studies. Instead, the factors most likely to come into play in the current designs are increased confidence in stereotypes, coupled with the perception that one ought to make decisions effortfully, and desires to maintain beliefs about the self and social identities.
2. We predicted neither a main effect nor interactions for the within-participant ethnicity factor. Power should have the same motivational, and hence the same cognitive consequences, regardless of whether the subordinate is an in-group member. However, due to the nature of our trait information manipulation, we might predict a valence effect manifested as an interaction between ethnicity and information consistency. Our design did not cross valence with consistency because of the nature of the majority students' stereotypes about the chosen

ethnic groups, wholly positive for Anglos and wholly negative for Hispanics. Although this presents a within-target confound, it is not problematic for the factors of interest because of the anticipated lack of interaction between target ethnicity and the two experimental treatments. It is possible though that participants will preferentially attend to negative information, as it is typically perceived to be more informative (Fiske, 1980). If so, we would predict a two-way interaction between ethnicity and information consistency, but no higher-order interactions.

3. The nature of our trait information (positive for Anglos, negative for Hispanics) led to an interaction between ethnicity and information consistency ($F(1,66) = 50.04$, $p < .001$), equivalent to a not surprising valence main effect, with participants attending more to negative information about both Anglo ($M = 46.34$ s) and Hispanic ($M = 45.57$ s) targets than to positive information about either one ($M = 37.68$ s and $M = 37.35$ s, respectively). In addition to this valence effect, ethnicity interacted with the responsibility manipulation ($F(1,66) = 4.31$, $p = .04$), with overall attention increasing more to the Anglo target (mean increase = 7.81 s) than to the Hispanic target (mean increase = 4.29 s) under high responsibility.
4. The Study 1 two-way interaction between ethnicity and information consistency (the valence effect) was replicated in this study. Participants spent more time attending to negative information ($M = 43.17$ s) than to positive information ($M = 36.10$ s) ($F(1,52) = 18.23$, $p < .001$.
5. These role names were chosen to describe the relationships as clearly as possible without using labels that would imply status (e.g. teacher, student). The original French role names were: *allocateur, executeur,* and *observateur.*
6. The target psychology major was included primarily to reduce suspicion about the target materials and to allow participants to practice the task before encountering the targets of interest to our hypotheses. Because most participants were likely to be psychology majors themselves, we were not interested in their evaluations of these in-group targets.
7. Participants in the control condition were run in separate sessions after the initial data collection for bosses and employees. However, participants in these sessions believed that other participants in their sessions could be playing the roles of bosses and employees.

Acknowledgments

This research and the first author were supported by National Institute of Mental Health Grant 41801 and National Science Foundation Grant # 94-21480 to the third author. The first three studies of the paper are based on the masters thesis and dissertation of the first author, who was supported, in part, by a dissertation award from the Association of Women in Science. The first two authors were supported by National Institute of Mental Health Training Grant 18827.

Experiment 2 was presented at the August 1992 annual meeting of the American Psychological Association, Washington, DC. Experiment 3 was presented at the 1995 Joint Meeting of the European Association of Experimental Social Psychology and the Society of Experimental Social Psychology, Washington, DC, and the 1994 Conference on Business Ethics in Organizations, Northwestern University. Experiment 4 was presented at the 1998 annual meeting of the American Psychological Society, Washington, DC.

We thank Icek Aizen, Paula Pietromonaco, Jacques-Philippe Leyens, Caren Jones, Lawrence Zacharias, and Don Operario for their comments on earlier drafts. We also extend our gratitude to Holly Von Hendy for her assistance in developing materials; to Jennifer Canfield and Joelle Karnas for their assistance in executing the studies and coding data; and to Suzanne Klonis for her assistance with coding.

References

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

Allport, G. W. (1954). *The nature of prejudice.* Reading, MA: Addison-Wesley.

Battistich, V., Assor, A., Messé, L. A., & Aronoff, J. (1985). Personality and person perception. In P. Shaver (Ed.), *Review of personality and social psychology* (Vol. 6, pp. 185–208). Beverly Hills, CA: Sage.

Bierstedt, R. (1950). An analysis of social power. *American Sociological Review, 15,* 730–736.

Brewer, M. B. (1988). A dual-process model of impression formation. In T. K. Srull & R. S. Wyer, Jr. (Eds.), *Advances in social cognition* (Vol. 1, pp. 1–36). Hillsdale, NJ: Erlbaum.

Bruner, J. S., & Taiguri, R. (1954). The perception of people. In G. Lindzey (Ed.), *Handbook of social psychology* (Vol. 2, pp. 634–654). Reading, MA: Addison-Wesley.

Cartwright, D. (1959). Power: A neglected variable in social psychology. In D. Cartwright (Ed.), *Studies in social power* (pp. 1–14). Ann Arbor, MI: Institute of Social Research.

Chen, S., Lee-Chai, A. Y., & Bargh, J. A. (2000). Relationship orientation as a moderator of the effects of social power. Unpublished manuscript, University of Michigan, Ann Arbor.

Copeland, J.T. (1994). Prophecies of power: Motivational implications of social power for behavioral confirmation. *Journal of Personality and Social Psychology, 67*, 264–277.

Dépret, E. F., & Fiske, S. T. (1993). Social cognition and power: Some cognitive consequences of social structure as a source of control deprivation. In G. Weary, F. Gleicher, & K. Marsh (Eds.), *Control motivation and social cognition* (pp. 176–202). New York: Springer-Verlag.

Dépret. E., & Fiske, S. T. (1999). Perceiving the powerful: Intriguing individuals versus threatening groups. *Journal of Experimental Social Psychology, 35*, 461–480.

Dovidio, J. F., & Gaertner, S. L. (1993). Stereotypes and evaluative intergroup bias. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping* (pp. 167–193). San Diego, CA: Academic Press.

Erber, R., & Fiske, S.T. (1984). Outcome dependency and attention to inconsistent information. *Journal of Personality and Social Psychology, 47*, 709–726.

Fiske, S. T. (1980). Attention and weight in person perception: The impact of negative and extreme behavior. *Journal of Personality and Social Psychology, 60*, 889–906.

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

Fiske, S. T. (1998). Stereotyping, prejudice, and discrimination. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., Vol. 2, pp. 357–411). New York: McGraw-Hill.

Fiske, S. T., & Dépret, E. (1996). Control, interdependence, and power: Understanding social cognition in its social context. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 7, pp. 31–61). New York: Wiley.

Fiske, S. T., Lin, M., & Neuberg, S. L. (1999). The continuum model: Ten years later. In S. Chaiken, & Y. Trope, (Eds.), *Dual-process theories in social psychology* (pp. 231–254). New York, NY: Guilford 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.P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23, pp.1–74). New York: Academic Press.

Fiske, S. T., Neuberg, S. L., Beattie, A. E., & Milberg, S. J. (1987). Category-based and attribute-based reactions to others: Some informational conditions of stereotyping and individuating processes. *Journal of Experimental Social Psychology, 23*, 399–427.

Fiske, S. T., & Von Hendy, H. M. (1992). Personality feedback and situational norms can control stereotyping processes. *Journal of Personality and Social Psychology, 62*, 577–596.

French, J. R. P., & Raven, B. H. (1959). The bases of social power. In D. Cartwright (Ed.), *Studies in social power* (pp. 118–149). Ann Arbor, MI: Institute for Social Research.

Georgesen, J. C., & Harris, M. J. (1998). Why's my boss always holding me down? A meta-analysis of power effects on performance evaluation.

Glick, P., & Fiske, S. T. (1999). Gender, power dynamics, and social interaction. In M. M. Ferree, & J. Lorber (Eds.), *Revisioning gender. The gender lens* (Vol. 5, pp. 365–398). Thousand Oaks, CA: Sage.

Goodwin, S. A., & Fiske, S. T. (1996). Judge not, lest . . .: The ethics of power holders' decision making and standards for social judgment. In D. M. Messick & A. E. Tenbrunsel (Eds.), *Codes of conduct: Behavioral research into business ethics* (pp. 117–142). New York: Russell Sage Foundation.

Gough, H. G. (1987). *The CPI manual* (rev. ed.). Palo Alto, CA: Consulting Psychologists' Press.

Hilton, J. L., & von Hippel, W. (1996). Stereotypes. In J. T. Spence, J. M. Darley, & D. J. Foss (Eds.), *Annual review of psychology* (Vol. 47, pp. 237–271). Palo Alto, CA: Annual Reviews.

Jost, J. T., & Banaji, M. R. (1994). The role of stereotyping in system-justification and the production of false consciousness. *British Journal of Social Psychology, 33*, 1–27.

Katz, I., & Haas, R. G. (1988). Racial ambivalence and American value conflict: Correlational and priming studies of dual cognitive structures. *Journal of Personality and Social Psychology, 55*, 893–905.

Kiesler, C. A., Nisbett, R. E., & Zanna, M. P. (1969). On inferring one's beliefs from one's behavior. *Journal of Personality and Social Psychology, 11*, 321–327.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin, 108*, 480–498.

Leyens, J.-P. (1983). *Somme-nous tous des psychologues? Approche psychosociale des théories implicites de personnalité* [Are we all psychologists? A social psychological approach to implicit personality theories]. Brussels: Mardaga.

Leyens, J.-P., & Yzerbyt, V. Y. (1992). The ingroup overexclusion effect: Impact of valence and confirmation on stereotypical information search. *European Journal of Social Psychology, 22*, 549–569.

Leyens, J.-P., Yzerbyt, V. Y., & Corneille, O. (1996). The role of applicability in the emergence of the overattribution bias. *Journal of Personality and Social Psychology, 70*, 219–229.

Leyens, J.-P., Yzerbyt, V. Y., & Schadron, G. (1992). Stereotypes and social judgeability. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 3, pp. 91–120). Chichester, UK: Wiley.

Leyens, J.-P., Yzerbyt, V. Y., & Schadron, G. (1994). *Stereotypes and social cognition.* London: Sage.

McConahay, J. B. (1983). Modern racism and modern discrimination: The effects of race, racial attitudes, and context on simulated hiring decisions. *Personality and Social Psychology Bulletin, 9*, 551–558.

Monteith, M. J., Devine, P. G., & Zuwerink, J. R. (1993). Self-directed versus other-directed affect as a consequence of prejudice-related discrepancies. *Journal of Personality and Social Psychology, 64*, 198–210.

Neuberg, S. L. (1989). The goal of forming accurate impressions during social interactions: Attenuating the impact of negative expectancies. *Journal of Personality and Social Psychology, 56*, 374–386.

Neuberg, S. L. (1996). Social motives and expectancy-tinged social interactions. In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition,* (Vol. 3, pp. 225–261). New York: Guilford.

Neuberg, S. L. & Fiske, S. T. (1987). Motivational influences on impression formation: Outcome dependency, accuracy-driven attention, and individuating processes. *Journal of Personality and Social Psychology, 53*, 431–444.

Pavelchak, M. A. (1989). Piecemeal and category-based evaluation: An idiographic analysis. *Journal of Personality and Social Psychology, 56*, 354–363.

Pratto, F., Sidanius, J., Stallworth, L. M., & Malle, B. F. (1994). Social dominance orientation: A personality variable predicting social and political attitudes. *Journal of Personality and Social Psychology, 67*, 741–763.

Robbins, G. E. (1974). Dogmatism and information gathering in personality impression formation. *Journal of Research in Personality, 9*, 74-84.

Ruscher, J. B., & Fiske, S. T. (1990). Interpersonal competition can cause individuating processes. *Journal of Personality and Social Psychology, 58*, 832–843.

Ruscher, J. B., Fiske, S. T., Miki, H., & Van Manen, S. (1991). Individuating processes in competition: Interpersonal versus intergroup. *Personality and Social Psychology Bulletin, 17*, 595–605.

Schneider, D. J., Hastorf, A. H., & Ellsworth, P. C. (1979). *Person perception.* Reading, MA: Addison-Wesley.

Snyder, M., & Haugen, J. A. (1994). Why does behavioral confirmation occur? A functional perspective on the role of the perceiver. *Journal of Experimental Social Psychology, 30*, 218–246.

Snyder, M., Tanke, E. D., & Berscheid, E. (1977). Social perception and interpersonal behavior: On the self-fulfilling nature of social stereotypes. *Journal of Personality and Social Psychology, 35*, 656–666.

Stevens, L. E., & Fiske, S. T. (2000). Motivated impressions of a powerholder: Accuracy under task dependency and misperception under evaluative dependency. *Personality and Social Psychology Bulletin.*

Taylor, R. N., & Dunnette, M. D. (1974). Influence of dogmatism, risk-taking propensity, and intelligence on decision-making strategies for a sample of industrial managers. *Journal of Applied Psychology, 59*, 420–423.

Tetlock, P. E. (1992). The impact of accountability on judgment and choice: Toward a social contingency model. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (pp. 331–376). New York: Academic.

Thibaut, J. W., & Kelley, H. H. (1959). *The social psychology of groups.* New York: Wiley.

Winter, D. G. (1973). *The power motive.* New York: Free Press.

Winter, D. G. (1988). The power motive in women – and men. *Journal of Personality and Social Psychology, 54*, 510–519.

Yoder, J. D., & Kahn, A. S. (1992). Toward a feminist understanding of women and power. *Psychology of Women Quarterly, 16*, 381–388.

Yzerbyt, V. Y., Dardenne, B., & Leyens, J.-P. (1998). The role of naïve theories in social judgment. In V. Y. Yzerbyt, G. Lories, & B. Dardenne (Eds.), *Metacognition: Cognitive and social dimensions.* London: Sage.

Yzerbyt, V. Y., Leyens, J.-P., & Schadron, G. (1997). Social judgeability and the dilution of stereotypes:

The impact of the nature and sequence of information. *Personality and Social Psychology Bulletin*, *23*, 1312–1322.

Yzerbyt, V., Rocher, S., & Schadron, G. (1997). Stereotypes as explanations: A subjective essentialistic view of group perception. In R. Spears, & P. J. Oakes (Eds.), *The social psychology of stereotyping and group life* (pp. 20–50). Oxford, England: Blackwell.

Yzerbyt, V. Y., Schadron, G., Leyens, J.-P., & Rocher, S. (1994). Social judgeability: The impact of meta-informational cues on the use of stereotyping. *Journal of Personality and Social Psychology*, *66*, 48–55.

Paper received 13 August 1988; revised version accepted 19 July 1999

Biographical notes

STEPHANIE A. GOODWIN received her PhD in social and personality psychology from the University of Massachusetts at Amherst. She received awards from the American Psychological Association (APA) and the Association of Women in Science (AWIS) for her dissertation research on power and impression formation. She is currently a Visiting Research Fellow at Yale University where her NSF supported research addresses social power and implicit social cognition.

ALEXANDRA GUBIN is in the PhD program at the University of Massachusetts at Amherst, completing a dissertation on sexual violence. She also conducts research on reconciliation in post-genocide Rwanda.

SUSAN T. FISKE, having served on the faculties of Carnegie Mellon University and the University of Massachusetts at Amherst, is now at Princeton University. An editor of the *Annual Review of Psychology* and the *Handbook of Social Psychology* (4th ed), she was President of the Society of Personality and Social Psychology, won the APA Early Career Award for Contributions to Psychology in the Public Interest, and won the Gordon Allport Intergroup Relations Award. Her research addresses stereotyping, especially issues of interdependence, ambivalence, and content.

VINCENT Y. YZERBYT is Professor of Psychology at the Catholic University of Louvain at Louvain-la-Neuve, Belgium. His research deals with social cognition, stereotyping, and intergroup relations. He currently serves as Associate Editor of the *Personality and Social Psychology Bulletin.*