

On What It Means to Know Someone: A Matter of Pragmatics

Michael J. Gill
Lehigh University

William B. Swann Jr.
University of Texas

Two studies provide support for W. B. Swann's (1984) argument that perceivers achieve substantial *pragmatic accuracy*—accuracy that facilitates the achievement of relationship-specific interaction goals—in their social relationships. Study 1 assessed the extent to which group members reached consensus regarding the behavior of a member in familiar (as compared with unfamiliar) contexts and found that groups do indeed achieve this form of pragmatic accuracy. Study 2 assessed the degree of insight romantic partners had into the self-views of their partners on relationship-relevant (as compared with less relevant) traits and found that couples do indeed achieve this form of pragmatic accuracy. Furthermore, pragmatic accuracy was uniquely associated with relationship harmony. Implications for a functional approach to person perception are discussed.

Our intuitions tell us that we know a great deal about the people around us. In fact, there are few human convictions held with more tenacity than the belief that we have extensive, often privileged, knowledge about our friends, family members, and lovers. Yet the research literature has not treated this conviction kindly. Time and time again, researchers have shown that people's impressions of others are prone to systematic and potentially consequential errors (e.g., Nisbett & Ross, 1980; but see Krueger & Funder, in press), including selectivity biases in attention, interpretation, and recall; the use of misleading heuristics; and a failure to consider the extent to which others' behaviors are situationally constrained. In addition to being limited in their ability to formulate accurate images of others, people also seem to have little awareness of their limitations. As a result, people can become highly confident of the accuracy of their perceptions even when accuracy is quite modest (e.g., Dunning, Griffin, Milojkovic, & Ross, 1990; Gill, Swann, & Silvera, 1998; Swann & Gill, 1997).

Despite these rather gloomy assessments of the accuracy of the person perception process, we believe that people do indeed acquire highly accurate knowledge about their interaction partners (see also Funder, 1999; Zebrowitz & Collins, 1997). We propose that this knowledge, however, is narrower in scope than people's intuitions would have them believe and too subtle to be detected by some commonly used methods of gauging accuracy. A brief summary of the accuracy literature helps set the stage for our argument.

The Quest for Accuracy

Historical Ups and Downs

The accuracy of personality judgment has been a concern of psychologists since the early part of the last century (see Taft, 1955, for a review). Initial attempts to study accuracy made slow progress, however, because of both an absence of theory and empirical generalities (Gilbert, 1998). More discouraging than either of these shortcomings, though, was the insightful methodological critique of accuracy research crafted by Cronbach and colleagues (1955; Gage & Cronbach, 1955). They proposed that researchers' penchant for taking the mathematical difference between one person's judgment (dubbed the "judge") and a measure of another person's personality (dubbed the "other") resulted in accuracy estimates that conflated several conceptually distinct phenomena, only some of which were relevant to accuracy. Rather than rising to this (admittedly formidable) challenge, accuracy researchers fled the scene. Efforts to study accuracy by comparing perceivers' judgments to some measure of "reality" were largely abandoned.

It was not until recently that researchers returned to these issues. One fruitful line of research has exploited advanced statistical techniques to unravel the conceptually distinct components that go

Michael J. Gill, Department of Psychology, Lehigh University; William B. Swann Jr., Department of Psychology, University of Texas.

This research was supported by a University of Texas Continuing Fellowship awarded to Michael J. Gill, by a Franz and Class of 1968 Junior Faculty Fellowship awarded by Lehigh University to Michael J. Gill, and by Grant SBR-9319570 from the National Science Foundation and Grant MH57455-01A2 from the National Institute of Mental Health awarded to William B. Swann Jr. Sincere thanks are due to Sam Gosling and David Kenny for their insightful comments on earlier versions of this article. We also thank Shari Abrams, Taylor Allen, Michael Apgar, Jessica Cheerman, Lauren Dullum, Christina Febbraro, Leigh Fendrick, Sundeep Patel, Alex Perez, Melissa Plesh, Nuanprang Snitbahn, and Melanie Thiel, whose assistance with data collection, coding, and entry made this article possible.

Correspondence concerning this article should be addressed to Michael J. Gill, Department of Psychology, Lehigh University, Bethlehem, PA 18015. E-mail: m.gill@lehigh.edu

into estimates of accuracy (see Kenny, 1994, for a review). Furthermore, Kenny (1991) has offered a detailed theory regarding moderators of consensus and accuracy that can be tested using these advanced methodological tools. Another fruitful line of research has strived to demonstrate that accuracy is generally higher than would be predicted by some of the gloomier accounts of person perception and to identify moderators that influence when accuracy is high or low (see Funder, 1999, for a review). Despite the great strides that these researchers have taken, we suggest that there is still much to learn. In this article, we champion an alternative approach to accuracy that was inspired by Swann's (1984) discussion of the relation between accuracy and the nature and purpose of interpersonal perception.

Knowing What We Need to Know

Swann (1984) began with the assumption that we cannot truly understand the person perception process without first asking what that process is designed to do, much as we cannot fully understand the structure of birds' feathers without knowing something about flight (for related arguments, see Marr, 1982; McHenry, 1971; Zebrowitz & Collins, 1997). With this assumption in hand, he proposed that because most social relationships are limited in scope and involve relationship-specific identities and goals (see Swann, Bosson, & Pelham, 2002), it is often unnecessary for perceivers to know every aspect of a target or to consider the target from an "objective" point of view. Instead, perceivers often need only achieve *pragmatic accuracy*, or accuracy that facilitates relationship-specific interaction goals. Pragmatic accuracy differs from other conceptions of accuracy (see Funder, 1999; Kenny, 1994, for reviews) in at least two critical ways. First, whereas accuracy is generally equated with the ability to make inferences about targets that are valid across contexts, pragmatic accuracy pertains to beliefs that have validity primarily, sometimes solely, within the confines of the perceiver's relationship with a target. Second, whereas accuracy is typically equated with the ability to discern "objective reality," a pragmatic approach to accuracy suggests (following James, 1907) that accuracy is indexed by utility for the belief holder or by whether a belief brings desired consequences and prevents undesired consequences. Of course, pragmatically accurate beliefs will sometimes satisfy criteria for objective validity (whatever those might be; see Funder [1999] for an excellent discussion), but they need not meet these criteria to qualify as pragmatically accurate. Relationship-relevant utility can exist without objective accuracy.

Adherents to traditional conceptualizations of accuracy could argue that if pragmatic accuracy is defined in terms of utility, it might be more appropriately conceptualized as "instrumentality" rather than as a form of accuracy. Although we agree that it is possible to frame pragmatic accuracy this way, we believe that there is an important advantage associated with considering pragmatic accuracy as a form of accuracy: Whereas traditional forms of accuracy have no necessary relation to relationship harmony and adaptive behavior, pragmatic accuracy matters to people by definition. This is especially important because the psychological meaningfulness of traditional standards of accuracy has recently been called into question (e.g., Krueger & Funder, in press). Thus, the pragmatic accuracy framework represents one means of evad-

ing this conundrum by forging direct links between issues of accuracy and issues of relationship harmony and well-being.

Pragmatic accuracy and recent approaches to personality and social behavior. Recent advances in the understanding of human personality complement some of our assumptions. To begin, Mischel and Shoda (1995) proposed that personality does not consist of highly general traits but rather of a profile of characteristic reactions to situations. This is consistent with Swann's notion that human beings often adopt relationship-specific identities that bear little relation to their identities in other circumstances. According to Mischel and Shoda, the goal of the personality theorist is to determine how people will respond in particular contexts rather than to determine their average level of responding across many different contexts. From this perspective, the work of Mischel and Shoda's personality theorist is similar to Swann's person perceiver in that both are striving to understand context-specific behaviors of targets. One difference, however, is that Swann's person perceiver focuses specifically on those contexts he or she shares with the target, whereas Mischel and Shoda's personality theorist focuses on all possible contexts.

Emerging work on the fundamental forms of sociality (e.g., Bugental, 2000; Fiske, 1992; Kenrick et al., 2002) also complements Swann's proposal. This work suggests that human social life can be conceptualized in terms of domains, with each domain comprising different social goals, different orientations toward others, different implicit schemas that guide action, and so on. According to this work, a person's behavior depends on the domain into which he or she categorizes a given social interaction. For example, if a person categorizes an interaction as concerning hierarchical power relations, then he or she will behave quite differently than if that interaction had been categorized as concerning reciprocal relating. What this means is that human beings may adopt very different patterns of behavior across different domains of social life. As with Mischel and Shoda, then, this is consistent with Swann's (1984) notion that people adopt relationship-specific identities and lends credibility to the idea that person perceivers may be in the business of knowing targets only in the subset of domains that they share with those targets.

Previous studies of the magnitude of pragmatic accuracy. As provocative as Swann's (1984) argument may be, to the best of our knowledge there have been only two attempts to test it systematically. Both attempts have focused on the type of pragmatic accuracy emphasized by Swann: Perceivers should show higher levels of *circumscribed accuracy* (accuracy at predicting a target's behavior in contexts shared with the perceiver) than of *global accuracy* (accuracy at predicting a target's behavior in contexts outside those shared with the perceiver).

Of relevance to this hypothesis, Levesque and Kenny (1993) reported that the average of several perceivers' impressions was predictive of the average level of a target's behavior across several situations, whereas any given perceiver seemed unable to predict how a given target would behave in the presence of that perceiver. Although these data seem to challenge the idea that circumscribed accuracy exceeds global accuracy, they may not. Given that the target was a perfect stranger with whom perceivers did not anticipate any future interaction, there was no motivation for perceivers to work to deduce how the stranger would behave in their own presence. Theoretically, circumscribed accuracy should be high

only when perceivers are motivated to know how targets will behave with them.

In another study that seems germane to the circumscribed global accuracy distinction, Kenny, Kiefer, Smith, Ceplenski, and Kulo (1996) reported that the average of fraternity brothers' impressions of a given brother predicted that brother's behavior across several situations, whereas any given brother was somewhat unable to predict how a given brother would behave with him. Although this finding may seem to suggest that global accuracy is greater than circumscribed accuracy, Kenny et al. conceded that this is not necessarily so. Indeed, given the group-based nature of interactions among fraternity brothers, the form of circumscribed accuracy that would be most pragmatic might be different than it would be, for example, between people who principally interact dyadically (e.g., lovers). Specifically, rather than the ability to predict the behavior of a brother in the presence of only oneself (which might rarely be required), circumscribed accuracy might involve predicting a brother's behavior with other members of the fraternity (which participants were able to do). Thus, Kenny et al.'s study can be construed as providing evidence that circumscribed accuracy was high but as providing no evidence regarding the magnitude of global accuracy (which would have involved participants' abilities to predict a brother's behavior in contexts that they do not typically share with him).

In sum, then, although past studies appear to offer little support for Swann's (1984) argument that circumscribed accuracy is likely to exceed global accuracy, careful inspection of the data reveal that they are inconclusive with respect to the more general pragmatic accuracy argument. Notably, past studies have not considered alternative ways of operationalizing pragmatic accuracy, nor have they tested Swann's prediction that pragmatic accuracy would be more closely associated with people's relationship outcomes than other types of accuracy. The present studies were designed to rectify these shortcomings.

Our Approach

Although Swann's (1984) argument has become virtually synonymous with one instantiation of the pragmatic accuracy issue (i.e., X can predict Y 's behavior with X better than Y 's behavior across a variety of interaction partners), his argument was more general. The cornerstone of his argument was that, "the accuracy of social beliefs is . . . determined by how well they serve the goals of perceivers rather than by the extent to which they are accurate in an ultimate sense" (Swann, 1984, p. 461). This position suggests that a fair assessment of the accuracy of personality judgment requires researchers to take two steps. First, they must identify what types of knowledge are likely to serve the goals of perceivers in a given type of relationship. Second, they must determine whether perceivers actually acquire this knowledge and whether doing so enables them to satisfy their interaction goals. We take both these steps in this article.

Our studies examine pragmatic accuracy in two distinct types of naturally occurring relationships. Study 1 examines group entities (i.e., fraternities and families), and Study 2 examines romantic relationships. We reasoned that in the context of groups that must coordinate joint activities and maintain a sense of cohesion, it is pragmatic to develop consensual understandings of how each group member behaves in the group. Thus, we would expect

consensus to be especially high when the group makes judgments of such behaviors (e.g., how extroverted is Fred when he is with the group?). Such consensus should contribute to group cohesion by, for example, promoting efficient, harmonious decision making regarding who should assume various roles in the group (e.g., Who will lead? Who will foster esprit de corps?). In contrast, consensus regarding behaviors outside the group context is of little pragmatic value and, therefore, is unlikely to be obtained and, if it is, is unlikely to contribute to relationship quality. We recognize, of course, that consensus is not necessarily indicative of objective accuracy (the group can agree but be "wrong" in some ultimate sense; e.g., Kenny, 1994), but, as alluded to previously, pragmatic accuracy is based on utility rather than some standard of objectivity. Furthermore, although consensus and objective accuracy are not necessarily the same thing, there is evidence that consensus grows out of objectively accurate perception (see Funder, Kolar, & Blackman, 1995).

With regard to romantic relationships, we reasoned that it is pragmatic to develop an accurate image of a partner's self-understanding on relationship-relevant traits, or traits that govern behaviors with practical and/or hedonic consequences for the perceiver. Thus, we would expect accuracy to be especially high when romantic partners are describing their partner's self-understanding on relationship-relevant traits as opposed to on traits not relevant to the relationship.¹ Of course, knowing someone's self-views is not the same as knowing their true standing on personality traits. Nevertheless, we believe that knowing a partner's self-views is useful in that it facilitates the negotiation of relationship-relevant identities. For example, insight into a partner's self-views enables one to predict what roles the partner will feel comfortable assuming in the relationship (e.g., Swann, 1990) and to be more persuasive when attempting to change the partner's relationship behaviors (e.g., Swann, 1997).

Study 1: Pragmatic Accuracy in Groups

Targets were members of a fraternity whose personalities were rated by several fraternity brothers and by several family members. Both sets of perceivers (i.e., fraternity brothers and family members) attempted to describe how the target behaved both around the fraternity house and at home with his nuclear family. The key prediction was that fraternity brothers would show greater consensus when describing the target's behavior around the fraternity house as compared with at home with his nuclear family. In contrast, we expected that family members would show greater consensus when describing the target's behavior at home with his nuclear family as compared with around the fraternity house. In addition, we expected that relationship quality would be predicted by consensus regarding the target's behavior in shared contexts (e.g., at the fraternity house for fraternity brothers) but would not

¹ Although it might seem as if we are talking about self-other agreement here (see Kenny, 1994), we are not. Self-other agreement concerns whether one partner's view of the other matches the other's self-view (e.g., does Jack see Jill as independent to the same degree that Jill sees herself as independent?), whereas we are concerned with whether one partner can accurately describe the other partner's self-view (e.g., does Jack know that Jill considers herself independent?).

be predicted by consensus regarding the target's behavior in non-shared contexts (e.g., at the fraternity house for family members).

Method

Participants

Participants were 20 fraternity brothers who agreed to participate after being approached by Michael J. Gill. Nineteen of their family members—10 women, 9 men—participated after being contacted by mail. Of the 19 family members, 16 were parents or step-parents of the target, 2 were grandparents, and 1 was a sister.

Procedure

Fraternity brothers. A male experimenter arrived at Lehigh University's Pi Kappa Alpha house and explained that the brothers were to participate in a study concerned with personality impressions. He further explained that each participant would rate several of his fraternity brothers and would be rated by several of his fraternity brothers and some of his family members. Thus, each fraternity brother was both a target and a perceiver, whereas family members were solely perceivers. After obtaining informed consent, the experimenter distributed three surveys to each participant, enabling participants to rate the behavior of three of their fraternity brothers. Targets were randomly assigned to perceivers with the constraint that three perceivers rated every target. Participants completed the surveys at various private places around the fraternity house while the experimenter waited.

Surveys contained six items tapping six behaviors that are commonly assessed in the literature (i.e., each of these items can be easily linked to the Big Five taxonomy; see Costa and Macrae, 1998): sensitive to others' concerns (e.g., shows empathy and compassion, is a good listener), competitive (e.g., always needs to win, get the last word, is a sore loser), extroverted (e.g., outgoing, talkative, socially energetic), cooperative and willing to help out (e.g., doing chores, pitching in whenever needed), worries about a lot of things (e.g., deadlines, relationships, finances), and easy and pleasant to be with (e.g., not argumentative, not "difficult," not overly negative). Ratings were made on 7-point scales with endpoints labeled *strongly disagree* and *strongly agree*. The number of items was kept to a minimum so as not to burden our participants, all of whom were volunteers.

Crucially, each behavior was assessed with respect to two contexts: with his fraternity brothers (e.g., "In terms of his behavior with his fraternity brothers, he tends to be extroverted [e.g., outgoing, talkative, socially energetic]") and with his nuclear family (e.g., "In terms of his behavior with his nuclear family [e.g., parents, step-parents, siblings], he tends to be extroverted [e.g., outgoing, talkative, socially energetic]"). The order of contexts was counterbalanced across participants; this order variable had no effects in any analysis and thus is not discussed further.

In addition to providing behavior ratings, participants completed items concerning their relationship with targets. First, they indicated how long they had known the target (in years and months) and the extent of their face-to-face contact with the target during that time (i.e., "minimal"; "infrequent"; "moderately frequent"; "extensive"). Then, they completed six items tapping relationship quality: "I like the person whose personality I just rated"; "The person whose personality I just rated gets on my nerves a lot"; "I generally get along very well with the person whose personality I just rated"; "The person whose personality I just rated has a way of annoying me from time to time"; "I have a positive relationship with the person whose personality I just rated"; and "I am not very close to the person whose personality I just rated." These ratings were made on 5-point scales with endpoints labeled *strongly disagree* and *strongly agree*. Finally, participants provided names and addresses of several of their family members, so the experimenter could solicit family members' ratings of each target.

Family members. Consent forms and surveys were mailed to 54 family members. Of these, 27 (50%) returned completed surveys. Eight of these could not be used because they came from the only family member who responded for a particular target, rendering it impossible to compute perceiver consensus. This left 19 family members in the final sample. Family members completed exactly the same measures as did fraternity brothers. As with fraternity brothers' responses, the order in which measures were completed had no effect and thus is not discussed further.

Computation of consensus. Intraclass correlations (Shrout & Fleiss, 1979) assessing perceiver consensus were computed for each target. These assessed the extent to which the pattern of responses across the six behavior ratings (e.g., high on extroversion, low on competitiveness) were similar for the set of perceivers who rated a given target. Separate correlations were computed to gauge the following: (a) fraternity brother agreement regarding the target's behavior around the fraternity house, (b) fraternity brother agreement regarding the target's behavior with his nuclear family, (c) family member agreement regarding the target's behavior around the fraternity house, and (d) family member agreement regarding the target's behavior with his nuclear family. The correlations were based on a one-way random effects model, which is appropriate when each target is rated by a different set of perceivers (Shrout & Fleiss, 1979).

When fraternity brothers served as perceivers, 20 intraclass correlations could be computed (i.e., 1 for each target), with each correlation being based on ratings from three judges (i.e., three perceivers rated each target). When family members served as perceivers, 8 intraclass correlations could be computed (i.e., 1 for each target who was rated by at least 2 perceivers). Six of these were based on data from two perceivers, 1 on data from three perceivers, and 1 on data from four perceivers.

Results

Perceiver Consensus

These intraclass correlations were submitted to a 2 (type of perceiver: fraternity brother, family member) \times 2 (context of ratings: with fraternity brothers, with family members) analysis of variance (ANOVA), with the second factor being within-subjects. As illustrated in Figure 1, this ANOVA revealed a main effect of type of perceiver, $F(1, 26) = 16.8, p < .01$, indicating that, collapsing across contexts, family members showed greater consensus than did fraternity brothers when rating the target's behavior. This finding should be interpreted cautiously because the

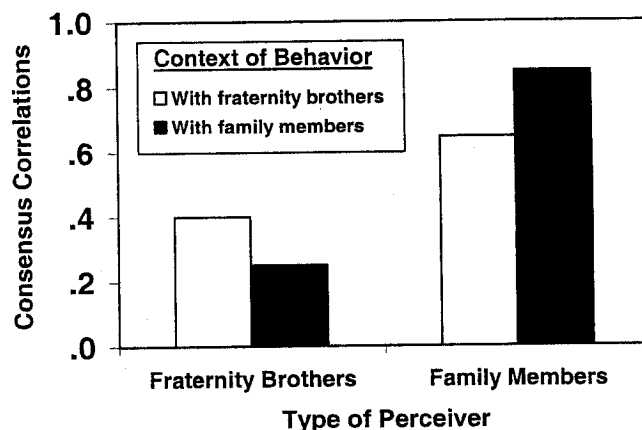


Figure 1. Consensus as a function of type of perceiver and context of behavior.

Table 1
Consensus Correlations Computed Across Targets for Each Trait

Trait	Fraternity brother rater		Family member rater	
	Fraternity	Nuclear family	Fraternity	Nuclear family
Sensitivity	.39	.36	.60	.70
Competitive	.50	.15	-.15	.33
Extroverted	.43	.23	.31	.36
Cooperative	.72	.47	.90	.64
Worry	.50	.20	.62	.81
Pleasant	.43	.17	.05	.51
Average across all traits	.50	.26	.39	.56

response rate for family members was not especially high, and those who responded may have done so because they had the type of relationship that fosters high consensus (e.g., they are especially close-knit).

More important for the pragmatic accuracy formulation, the ANOVA also revealed a significant interaction, $F(1, 26) = 10.5$, $p < .01$ (see Figure 1). Analysis of simple effects revealed that, as predicted, fraternity brothers showed greater consensus when rating the target's behavior around the fraternity house as compared with rating his behavior with his family members, $t(19) = 2.6$, $p < .02$, and family members showed marginally greater consensus when rating the target's behavior with his family members as compared with rating his behavior around the fraternity house, $t(7) = -2.1$, $p < .08$.

Whereas the consensus correlations just described were computed within target and across traits (on the basis of a separate Trait \times Perceiver matrix for each target), it is also possible to compute consensus correlations within traits and across targets (on the basis of a separate Target \times Perceiver matrix for each trait). Computing correlations in the latter manner precludes examining whether consensus is associated with relationship quality—which we examine later—because there is not a consensus score associated with each target. Nevertheless, convergent evidence across different measures of consensus would provide stronger evidence for our hypotheses. Accordingly, we computed perceiver consensus separately for each trait and across targets.² As can be seen in Table 1, these correlations showed a pattern highly consistent with the results described earlier. Thus, virtually identical results were achieved using two different methods of computing consensus, and results using both approaches supported predictions.³

Stereotype Versus Target-Driven Consensus

Ever since Cronbach (1955), researchers have recognized that a measure of accuracy can conflate conceptually distinct phenomena. In the present case, our measure of consensus might include both stereotype consensus and target-driven consensus (see Bernieri, Zuckerman, Koestner, & Rosenthal [1994] for parallel concepts of stereotype accuracy and true accuracy). Stereotype consensus is agreement among perceivers stemming from a shared stereotype of "what targets are generally like" rather than from knowledge of the particular target under consideration. Target-driven consensus, in contrast, is agreement among perceivers stemming from individuating knowledge about a particular target, and it is presumably based on observation of the target's behavior.

There is little precedent in the literature for how to disentangle these components when one examines perceiver consensus. One creative technique, however, was used by Blackman and Funder (1998; see Corsini [1956] for the first apparent use of "pseudo dyads"). They described their approach as establishing a "random baseline" for consensus that captured the extent to which perceivers' ratings were correlated even when they were rating different targets. Blackman and Funder established their random baseline by repeatedly randomly selecting a judge, pairing his or her ratings with ratings from another randomly selected judge who had rated a different target, and computing the correlation. Naturally, if perceivers' ratings correlate highly even when they are rating different targets, this is consistent with the possibility that their ratings across targets are based on a stereotype of what targets are generally like. (Note, however, that this is also consistent with the possibility that targets really are the same and that perceivers accurately detect this.) On the other hand, if perceivers' ratings correlate when they rate the same target but not when they rate different targets, this suggests that perceiver consensus is unlikely to be based on a shared stereotype of what targets are generally like but rather is likely based on knowledge—presumably derived through interaction with the target—that distinguishes the target from targets in general.

On the basis of this logic, we computed estimates of stereotype consensus by randomly pairing different judges' ratings of different targets and computing intraclass correlations. For ratings by fraternity brothers, we created 20 randomly derived sets of three different judges' ratings of three different targets (i.e., so the number of randomly derived trios equaled the number of actual trios). For ratings by family members, we created 16 randomly derived sets of different judges' ratings of different targets (i.e., two times the number of actual sets, because eight sets seemed too few to get a good estimate of stereotype consensus). Twelve of these random sets included two perceivers, 2 included three perceivers, and 2 included four perceivers, which mirrors the proportions in the real sets of family members.

² These correlations require that the same number of perceivers rate each target. Because one target was rated by three family members and one target was rated by four, for those targets, we randomly selected two family members to serve as the perceivers for these correlations.

³ Interestingly, intuitions often suggest that these different methods will produce different results. Kenny and Winquist (2001) have shown how these intuitions can be mistaken.

To get a sense of the magnitude of stereotype consensus under our various conditions, we submitted these intraclass correlations to a 2 (type of perceiver: fraternity brother, family member) \times 2 (context of ratings: with fraternity brothers, with family members) ANOVA, with the second factor being within-perceiver. This ANOVA revealed a main effect of type of perceiver, $F(1, 34) = 13.8, p < .01$, suggesting that, collapsing across contexts, stereotype consensus was higher among family members than among fraternity brothers. This suggests that, overall, family members were more likely to be relying on a shared stereotype than were fraternity brothers. Importantly, however, this main effect was qualified by a significant interaction, $F(1, 34) = 4.3, p < .05$. Analysis of simple effects suggested that this interaction was driven by the fact that fraternity brothers showed less stereotype consensus when rating how the target behaved with his fraternity brothers ($M = 0.00$) than when rating how he behaved with his nuclear family ($M = 0.25$), $t(19) = 2.1, p < .05$, whereas family members did not differ in their level of stereotype consensus when rating how the target behaved with his fraternity brothers ($M = 0.54$) as compared with how he behaved with his nuclear family ($M = 0.47$), $t < 1$.

This pattern of results is interesting in that it suggests that stereotype consensus may contribute more to total consensus (i.e., overall agreement regardless of its source) to a greater degree when perceivers rate targets in unfamiliar as compared with familiar contexts. As a relatively direct test of this possibility, we used t tests to examine the extent to which total consensus—computed using the first method presented above—exceeded stereotype consensus in each condition of our design. Results supported the possibility that ratings of targets in unfamiliar contexts are associated primarily with stereotype consensus, whereas ratings of targets in familiar contexts are associated with consensus in excess of stereotype consensus (i.e., target-driven consensus): Consensus did not exceed stereotype consensus when fraternity brothers rated the target's behavior with his family members ($t < 1$) or when family members rated the target's behavior with his fraternity brothers ($t < 1$); in contrast, consensus did exceed stereotype consensus when fraternity brothers rated the target's behavior with his fraternity brothers, $t(38) = 3.1, p < .01$ ($M_s = 0.40$ and 0.00 for consensus and stereotype consensus, respectively) and when family members rated the target's behavior with his family members, $t(22) = 3.6, p < .01$ ($M_s = 0.85$ and 0.47 for consensus and stereotype consensus, respectively). This suggests that perceivers rely on general stereotypes when rating targets in unfamiliar contexts but distinguish targets from each other when rating them in familiar contexts.

Consensus and Relationship Quality

To test Swann's (1984) prediction that pragmatic accuracy would be associated with greater relationship harmony, we examined the relation between perceivers' responses to our relationship quality measures and their consensus in describing the target within a familiar as opposed to unfamiliar context. These analyses are, of course, based on the measures of consensus computed within target, because we need a separate measure of consensus for each target to correlate consensus with relationship quality (which was measured for each target).

Because these analyses are correlational, they can provide evidence consistent with our prediction, but they cannot support the causal hypothesis that pragmatic accuracy fosters relationship harmony. It could be the case that relationship harmony contributes to pragmatic accuracy.

We first computed an average relationship quality score by reverse-scoring the appropriate items and then averaging responses to the six relationship quality items ($\alpha = .91$) such that a high score indicated a positive relationship. Next, the relationship quality scores for the set of perceivers that had rated a given target were averaged together, creating a measure of how positive relations were between that target and the several perceivers who had rated him. Finally, these relationship quality scores were correlated with our consensus measures of accuracy. Results were fairly consistent with Swann's (1984) argument. First, among fraternity brother perceivers, consensus in ratings of the target with his fraternity brothers was positively associated with relationship quality, $r(18) = .42, p = .06$. This relation was weaker when consensus among fraternity brothers in ratings of the target with his nuclear family were examined, $r(18) = .23, p = .34$. A t test of the difference between these dependent correlations (see Cohen & Cohen, 1983), however, was not significant ($t < 1$), presumably because of the small number of observations involved. Among family member perceivers, the extremely small number of observations ($N = 8$) precludes a definitive test.

Discussion

To the best of our knowledge, these data constitute the first empirical evidence supporting Swann's (1984) pragmatic accuracy formulation. Previously, we argued that pragmatic accuracy in a group-based relationship involves consensual understandings of group members' behaviors in the group context and, thus, that such consensus should be likely to develop. Consistent with this, the present study demonstrated that fraternity brothers showed greater consensus when rating how a fellow brother behaved with his fraternity brothers as compared with his nuclear family, whereas family members showed greater consensus when rating how that fraternity brother behaved with his nuclear family as compared with his fraternity brothers.

Follow-up analyses revealed that, in addition to these differences, consensus when rating a target in an unfamiliar context (e.g., consensus among fraternity brothers who are rating how the target behaves with his nuclear family) appeared to stem from reliance on shared stereotypes, whereas consensus when rating a target in a familiar context (e.g., consensus among fraternity brothers who are rating how the target behaves with his fraternity brothers) significantly exceeded that which could be explained by reliance on shared stereotypes. This suggests that pragmatic accuracy is likely to involve perceivers moving beyond the use of shared stereotypes to recognize distinctions among targets. Finally, correlational analyses were generally consistent with Swann's (1984) suggestion that pragmatic accuracy is associated with relationship quality, although this evidence was weak.

Study 2: Pragmatic Accuracy in Romantic Relationships

In Study 2, we sought to provide additional evidence that perceivers develop pragmatic accuracy in their relationships. Pre-

viously, we argued that pragmatic accuracy in a romantic relationship involves acquiring knowledge of a partner's self-views on relationship-relevant traits. To reiterate, relationship-relevant traits are those that govern important target behaviors in a given relationship—that is, behaviors with practical and hedonic consequences for the target's relationship partner. We assume that different traits are relevant in different types of relationships (cf. the work on social domains cited previously; e.g., Bugental, 2000; Fiske, 1992; Kenrick et al., 2002).⁴ For example, the trait of being an effective public speaker might be highly relevant in one's lawyer but not in one's lover. Conversely, the trait of being emotionally supportive might be highly relevant in one's lover but not in one's lawyer. In Study 2, we test this reasoning in the context of heterosexual dating relationships.

Method

Overview

Perceivers attempted to predict targets' responses to 10 trait items. Some perceivers predicted target self-views on low-relevance traits, some on moderate-relevance traits, and some on high-relevance traits. In addition, on the basis of the assumption that perceivers would prefer to rate a target on relevant as opposed to less relevant traits, within each relevance condition some perceivers freely chose the traits they would rate, whereas the remaining perceivers were assigned the traits they would rate. Perceivers and targets both completed measures of relationship quality. We expected that increments in trait relevance and choice would be associated with higher levels of accuracy, and that accuracy on relevant traits would be more highly associated with relationship quality than would accuracy on less relevant traits.

Participants

Participants were 88 heterosexual dating couples who were recruited by means of an advertisement in the University of Texas newspaper offering \$10 for participation in a dating couples study.

Participants had been dating for between 1 week and 340 weeks, with a median relationship length of 56 weeks. They ranged in age from 18 to 40, with a median age of 20. When asked, 11% described themselves as Asian, 3% as Black, 20% as Hispanic, and 66% as White. Data from 1 couple were not usable because the perceiver had used checkmarks rather than numbers when predicting his partner's responses. This left 87 couples in the final sample.

Procedure

Advance preparation of trait lists. Prior to recruiting couples for Study 2, we created lists of traits that varied in their relevance for dating couples. Low-relevance traits were selected so as to bear no special relevance for dating couples rather than as to be completely irrelevant (indeed, what trait would be completely irrelevant?). To create this list, we selected 20 items that have been used to tap the Big Five personality factors (John, 1990). We reasoned that the breadth of such items (when viewed as a whole) renders them not especially relevant for dating couples. We chose 4 items from each factor and selected items that have shown high factor loadings on the factor they are intended to measure (John, 1990). Moderate-relevance traits were selected so as to be relevant for dating couples in general (although not necessarily for any specific couple). To create this list, we recruited an independent sample of dating couples ($N = 76$) and requested that members provide open-ended descriptions of their partners.

Lists were compiled of all the traits used in these descriptions, and from these lists we selected the 20 most frequently mentioned traits. Different lists of 20 traits were generated for male and female targets because of gender differences in the frequencies of different descriptors. See Table 1 for the traits composing our low-relevance and moderate-relevance trait lists. Finally, high-relevance traits were selected on a couple-by-couple basis as described later.

Laboratory phase. Members of each couple were first randomly assigned to the roles of target and perceiver and were then separated. After providing informed consent, targets learned that we were interested in "the relation between personality and relationship satisfaction/commitment," whereas perceivers learned that we were interested in "how accurately dating partners know one another." We misled targets to prevent them from trying to describe themselves as they thought their partners would describe them rather than as they truly perceived themselves. All participants were assured that their partners would not see their responses.

Immediately thereafter, targets were given a questionnaire packet that would occupy them while the experimenter interacted with perceivers. This questionnaire packet included an eight-item relationship commitment questionnaire (Rusbult, 1980), which asked respondents to indicate how long they wanted their relationship to last, how much they loved their partner, how their relationship compared with their ideal relationship, how likely they were to date a new partner in the near future, and so on.

Perceivers were assigned to the low-, moderate-, or high-relevance condition. Whereas lists of low- and moderate-relevance traits had been prepared in advance (see previous text), high-relevance traits were generated by each perceiver in the high-relevance condition during the laboratory phase of the present study and were thus different for every perceiver in that condition. Specifically, perceivers in the high-relevance condition received a sheet asking them to list 20 personality traits on which they would subsequently be willing to predict their partner's self-rating. We reasoned that perceivers would list traits that were highly relevant to their relationship on the basis of the idea that hedonically relevant stimuli are especially likely to draw attention (e.g., Roskos-Ewoldsen & Fazio, 1992) and to remain highly available in memory (e.g., Rubin & Kozin, 1984).

Next, all perceivers were assigned to either the choice condition or the no choice condition. Perceivers in the choice condition chose from the appropriate list of 20 traits (i.e., low, moderate, or high relevance) the 10 traits on which they thought they could most accurately predict their partners' self-ratings. In contrast, perceivers in the no choice condition were randomly assigned (on the basis of a pregenerated random-numbers table) 10 traits from the list. Thus, perceivers rated targets on 10 traits that were either low, moderate, or high in relevance and that were chosen by perceivers or not. The choice manipulation had no main or interactive effects, and thus is not discussed further.⁵

⁴ Note that our concept of trait relevance is not the same as that proposed by some other authors (see Britt & Sheppard, 1999, for a review). Our concept refers to the notion that certain traits are especially likely to be pertinent in particular kinds of relationships (similar to the social domains concept; e.g., Bugental, 2000; Fiske, 1992; Kenrick et al., 2002). The concept reviewed by Britt and Sheppard (1999) refers to the notion that some people can be meaningfully given a score on a trait—because the trait frequently determines their behavior—whereas other people cannot (see Bem & Allen, 1974).

⁵ We conceptualized our choice manipulation as an indirect manipulation of trait relevance on the basis of the assumption that participants would choose to rate their partners on relevant as opposed to less relevant traits. This manipulation may have failed because participants made choices from lists whose members were roughly equivalent in relevance (i.e., choices were made from any of the low-, moderate-, or high-relevance lists). Thus, the traits might not have varied enough in relevance to enable participants' choices to make a difference in their accuracy levels.

Next, the experimenter created a personality questionnaire based on these 10 traits. The questionnaire was created on a personal computer, using a prefabricated template. While the experimenter did this, perceivers worked on the same questionnaire packet that had been given to targets earlier. After creating the personality trait questionnaire, the experimenter printed out two copies. One copy was brought to targets, who were asked to rate themselves on the 10 traits and to indicate their confidence (0%–100%) that their partners knew where they stood on each of the traits. The other copy was given to perceivers, who were asked to predict their partner's self-ratings and to indicate their confidence (0%–100%) that each prediction was accurate.

Computation of accuracy indices. Accuracy was defined in terms of how well perceivers predicted targets' self-views. Thus, for each couple, we computed a correlation between each target's self-ratings and each perceiver's predictions of those self-ratings across the 10 trait items. When computing correlations within couples, it is important that there be variability in a given respondent's ratings across the set of traits; that is, each respondent must provide some high ratings and some low ratings. If this variability is not present, the resulting range restriction prevents correlations from emerging even if the target and perceiver are largely—or even perfectly—in agreement. For example, imagine that a target and a perceiver responded to 3 items on 7-point scales as follows: "1, 1, 2" for the target, and "2, 1, 1" for the perceiver. The resulting accuracy correlation would be $-.50$. Obviously, this negative correlation is highly misleading, given the facts that the perceiver always correctly identified the target as "quite low" on the three traits (when viewed in the context of a 7-point response scale) and that the target and perceiver were sometimes in perfect agreement and always within just 1 point of each other.

To ensure adequate variability across items, we reverse-scored every other trait item prior to the computation of our accuracy correlations. In the three-item example above, this procedure would involve recoding Item 2 as "7" for both the target and the perceiver, and it would result in the accuracy correlation becoming $.95$. We believe this correlation better captures how well the target and perceiver did at the accuracy task than does a correlation of $-.50$. Also, we must emphasize that, because we are interested in between-condition differences in accuracy rather than the absolute magnitude of accuracy, there would seem to be little need to worry that this procedure makes accuracy appear artificially high.

We did, however, feel that it was important to rule out the possibility that our reverse-scoring procedure might inflate accuracy correlations differentially across our conditions. That is, the variability of participants' responses in one condition might be increased by the reverse-scoring procedure to a greater degree than in another condition, resulting in higher correlations in the condition with enhanced variability (Cohen & Cohen, 1983). Fortunately, this potential bias does not appear to have affected our analyses. When we recomputed our accuracy analyses (described later) and included response variability (i.e., each couple's average standard deviation across items following the reverse-scoring procedure) as a covariate, our effects persisted.

Computation of relationship quality. We computed several measures of relationship quality. First, following appropriate reverse-scoring, we computed measures of target relationship commitment ($\alpha = .91$) and perceiver relationship commitment ($\alpha = .94$) and then averaged these to form a measure of couple commitment (treating them separately produces the same pattern of results). Next, we computed a measure of targets' confidence that their partners could accurately predict their self-ratings by taking the average of the 10 confidence ratings made by targets ($\alpha = .67$). This taps the extent to which targets "feel known" by their partners. The desire to feel known has been shown to be a powerful motivator of behavior (Swann, 1990), and the feeling of being known is an indicator of subjective relationship quality (Swann, Hixon, & De La Ronde, 1992). Finally, we computed a measure of perceiver confidence, as in Study 1, by

averaging perceivers' confidence ratings across their 10 predictions ($\alpha = .70$). Target and perceiver confidence were not averaged together because they are conceptually distinct, and, as is shown later, they do not show the same pattern of correlations with other variables.

Measurement of potential moderators of accuracy. We had independent judges rate all our trait items on several dimensions to both establish the validity of our manipulation of trait relevance and enable us to compare our findings to prior literature. First, to provide evidence for the posited differences among relevance conditions, we had an independent sample of 30 judges rate the low-, moderate-, and high-relevance traits in terms of their relevance to romantic relationships. Second, to examine the role of previously identified moderators of accuracy in contributing to our accuracy findings, we had a separate sample of 34 judges rate the traits on the dimension of evaluativeness (extent to which the trait is either very negative or very positive) and yet another sample of 32 judges rate all traits on the dimension of visibility (extent to which behaviors relevant to the trait are easily visible to a perceiver). Judges rated the items on 5-point scales, with endpoints appropriately labeled (e.g., *not at all visible* and *extremely visible*). Reliabilities of judges' ratings were always very high (all α s $> .90$).

Results

Accuracy of Perceiver Knowledge of Target Self-Views

We predicted that accuracy would increase with increments in trait relevance. To test this prediction, we transformed accuracy correlations using Fisher's r -to- Z transformation, and submitted them to a 3 (trait relevance: low, moderate, high) \times 2 (target gender: female, male) ANOVA. This ANOVA revealed only a main effect of trait relevance, $F(2, 81) = 3.6, p < .04$ (all other F s < 1). As can be seen in Figure 2, as predicted, accuracy was higher in the moderate- and high-relevance conditions than in the low-relevance condition, $t(81) = 2.7, 1.8, ps < .01, .07$, respectively. Contrary to prediction, however, accuracy in the high-relevance condition was no higher than accuracy in the moderate-relevance condition, $t(81) = .91, p = .37$. The failure to find differences between the moderate- and high-relevance conditions might be due to a ceiling effect. Alternatively, it might simply be the case that what we called moderate- and high-relevance traits were both high in relevance, a possibility supported by evidence presented in the following text.

We performed follow-up analyses using the ratings collected from our independent judges to shed light on why these accuracy

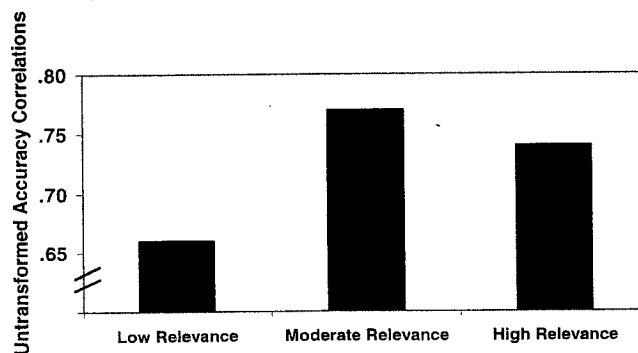


Figure 2. Untransformed accuracy correlations as a function of trait relevance.

differences might have occurred. First, each couple received a score indicating the relevance to romantic relationships, visibility, and evaluativeness of the traits they had rated. We computed this score by averaging together the judges' ratings for each of the 10 traits the couple had rated and then averaging these 10 averages together. Preliminary analysis of these ratings revealed that traits that were relevant to romantic relationships also tended to be evaluative, with the correlation between these two constructs being extremely high, $r(85) = .93, p < .01$. This makes sense in that romantic partners are likely to be concerned with highly evaluative aspects of their partners (e.g., "ability to care for others") as opposed to less evaluative aspects (e.g., "organization"). "Visibility," in contrast, correlated only weakly with evaluativeness, $r(85) = .18, p = .09$, and with relevance to romantic relationships, $r(85) = .12, ns$.

We examined scores on the three potential moderators as a function of trait relevance. These analyses revealed, as expected, that both the moderate-relevance ($M = 4.1$) and the high-relevance ($M = 3.5$) conditions involved traits that were more relevant to romantic relationships than were those in the low-relevance condition ($M = 3.3$), both $t(86) > 3.9, ps < .01$. Unexpectedly, the analysis also revealed that the high-relevance condition involved traits that were *less* relevant to romantic relationships than were those in the moderate-relevance condition, $t(86) = 8.5, p < .01$. On reflection, this result becomes less surprising: The high-relevance traits are "idiosyncratically relevant" in that they were chosen by just one particular perceiver, and there is no reason to expect that independent judges would agree with these idiosyncratic choices. In sum, these analyses support our contention that our relevance manipulation did create sets of traits that differed in terms of their relevance to romantic relationships. Furthermore, it is reasonable to attribute our accuracy findings to this differential relevance given that the effect of manipulated trait relevance on accuracy is reduced to nonsignificance ($F < 1$) when the independent judges' ratings of the relevance to romantic relationships of the traits in each condition are entered as a covariate.

Analysis of trait visibility revealed that the moderate-relevance condition ($M = 3.2$) involved traits that were marginally more visible than those in both the low-relevance ($M = 3.1$) and high-relevance ($M = 3.1$) conditions, both $t(86) > 1.7, ps < .09$. The low-relevance and high-relevance conditions did not differ ($t < 1$). The effect of our trait relevance manipulation on accuracy was still evident even when trait visibility was included as a covariate, $F(2, 83) = 4.1, p < .03$, suggesting that trait visibility cannot account for the effects of that manipulation.

Finally, and not surprisingly given their level of intercorrelation, ratings of trait evaluativeness showed an identical pattern to ratings of relevance to romantic relationships. Specifically, both the moderate- ($M = 3.8$) and the high-relevance ($M = 3.4$) conditions involved traits that were more evaluative than were those in the low-relevance condition ($M = 3.2$), both $t(86) > 4.8, ps < .01$. In addition, the high-relevance condition involved traits that were less evaluative than were those in the moderate-relevance condition, $t(86) = 7.9, p < .01$. As with ratings of relevance to romantic relationships, the effect of manipulated trait relevance on accuracy is reduced to nonsignificance ($F < 1$) when the independent judges' ratings of the evaluativeness of the traits are entered as a covariate. On the surface, this finding seems

to imply that our data can be understood in terms of a previously identified moderator of accuracy. Our results, however, are actually the opposite of what would be expected on the basis of earlier work (John & Robins, 1993). That is, prior work has found a negative association between evaluativeness and accuracy, whereas the present data suggest a positive association. Thus, the present data cannot be interpreted in terms of what is already known about evaluativeness and accuracy.

Disentangling Relationship Relevance and Trait Evaluativeness

The preceding analyses involving independent judges' ratings raise the question of whether it is appropriate to interpret our findings as supporting our pragmatic accuracy formulation. The pragmatic accuracy formulation suggests that relationship relevance per se should foster accuracy. Unfortunately, our relevance manipulation inadvertently manipulated trait evaluativeness as well as trait relevance. Because of this, we wanted to conduct additional analyses to show that relevance per se is sufficient to foster accuracy. Indeed, if trait evaluativeness—rather than relevance—is the critical variable enhancing accuracy in the moderate- and high-relevance conditions, then our pragmatic accuracy effect might simply reflect a tendency for perceivers and targets to paint a rosy picture of the target on highly valenced (i.e., evaluative) items (e.g., "loving," "caring," [not] "selfish," [not] "arrogant"). Obviously, such a tendency is quite distinct from the pragmatic accuracy argument.

Of course, we could shift our focus and argue that such shared rosy pictures constitute a form of pragmatic accuracy: Partners must construct and agree on highly positive relationship identities—that is, they must agree that the target thinks of him- or herself as loving, not selfish—or else the relationship will fail. Yet, however meritorious it might be (see *Discussion* section that follows), this argument makes it difficult to distinguish the pragmatic accuracy argument from the "positive illusions" formulation (e.g., Murray & Holmes, 1997; Taylor & Brown, 1988). The latter formulation argues for the existence and benefits of highly positive self- and partner perceptions on highly evaluative dimensions.

We believe that pragmatic accuracy is more than just a tendency for partners to both endorse positive judgments on evaluative dimensions. To test this idea, we conducted analyses to determine whether relevance would foster accuracy even when the traits are equivalently nonevaluative across relevance conditions. We reasoned that supportive data would provide evidence for the importance of relevance per se and, hence, would support the existence of pragmatic accuracy and the notion that pragmatic accuracy is distinct from shared positive illusions on evaluative traits.

Our approach involved examining whether accuracy on moderate- and high-relevance traits remained higher than accuracy on low-relevance traits even when the traits in all conditions were equivalently nonevaluative. To this end, for each couple we selected a nonevaluative subset of the 10 traits on which they had provided ratings. First, for each couple, we rank ordered in terms of evaluativeness the 10 traits on which they had provided ratings, with "1" referring to the trait lowest in evaluativeness. Next, we selected Traits 1 through 4 for couples within the moderate- and

high-relevance conditions and Traits 2 through 5 for couples in the low-relevance condition.⁶

This selection procedure was only partially successful at creating equivalently nonevaluative traits across conditions. Specifically, following the selection procedure, the nonevaluative subsets in the moderate-relevance condition remained more evaluative ($M = 3.49$) than the nonevaluative subsets in the other conditions ($M_s = 2.92$ and 2.97 for the low- and high-relevance conditions, respectively; $t_s(84) > 8.0$, $ps < .01$). This reflects the fact that virtually all traits in the moderate-relevance condition are highly evaluative (see Table 2). It also means, then, that the moderate-relevance condition is useless for disentangling the effects of relevance and evaluativeness. Critically, however, our selection procedure did yield the desired results with regard to the low- and high-relevance conditions. Specifically, the nonevaluative subsets of traits in the low-relevance condition ($M = 2.92$) were equal in evaluativeness to the nonevaluative subsets of traits in the high-relevance condition ($M = 2.97$; $t < 1$).

The next step involved comparing accuracy across these equivalently nonevaluative relevance conditions. Accordingly, we computed an accuracy correlation for each couple in the low- and high-relevance conditions on the basis of the nonevaluative subset of four traits that had been selected for that couple.⁷ After performing Fisher's r -to- Z transformation on each correlation, we averaged together all the accuracy correlations within the low- and the high-relevance conditions. Then, we compared the average level of accuracy across these two conditions. Results suggested that accuracy remained higher in the high-relevance condition (untransformed $M = .73$) than in the low-relevance condition (untransformed $M = .40$), $t(86) = 2.6$, $p < .02$. Thus, even when we examined equivalently nonevaluative traits across relevance conditions, we saw that relevance continued to foster accuracy. This is, of course, consistent with our pragmatic accuracy formulation and distinguishes it from the positive illusions perspective.

Stereotype Versus Target-Based Accuracy

As discussed in Study 1, researchers are concerned with distinguishing the extent to which agreement in ratings stems from stereotypes about "targets in general" (i.e., stereotype accuracy) as opposed to individuating knowledge of a target (i.e., target-based accuracy). Accordingly, we sought to examine the extent to which the relatively high levels of pragmatic accuracy found in the present study were attributable to each of these phenomena. To disentangle stereotype and target-based accuracy within the context of the present methodology, we needed a measure of the typical target (see Bernieri et al., 1994). To this end, we created questionnaires that included all the traits from the low-, moderate-, and high-relevance conditions, with the traits being divided across different questionnaires so that respondents would not need to rate all of them. These questionnaires were administered to an independent sample ($N = 151$) of college undergraduates. Next, we computed average self-ratings for each trait (there were at least 30 respondents for each trait) separately for male and female respondents. Finally, we computed partial correlations within each of our dating couples in which target self-ratings were correlated with perceiver predictions while controlling for the relation of both of these variables with typical target self-ratings of the appropriate gender.

Table 2
Low- and Moderate-Relevance Traits

Low relevance	Moderate relevance	
	Men	Women
quiet	caring	loving
reserved	loving	affectionate
talkative	nice	warm
assertive	gentle	caring
fault finding	sweet	kind
cold	compassionate	selfish
sympathetic	committed	generous
kind	selfish	outgoing
careless	considerate	shy
disorderly	jealous	charismatic
organized	fun	personable
thorough	outgoing	well liked
tense	friendly	friendly
anxious	likeable	honest
stable	ambitious	committed
calm	motivated	forgiving
commonplace	goal oriented	confident
simple	insecure	proud
imaginative	arrogant	intelligent
intelligent	confident	open-minded

These partial correlations revealed that target-based accuracy followed the same pattern as our total accuracy measure. Average target-based accuracy was .61, .78, and .75 for the low-, moderate-, and high-relevance conditions, respectively. Following an r -to- Z transformation, planned contrasts on these correlations revealed that target-based accuracy in both the moderate-relevance and high-relevance conditions was higher than target-based accuracy in the low-relevance condition, both $t_s(86) > 2.4$, $ps < .02$. In fact, these findings suggest that the effects of our relevance manipula-

⁶ The decision to select four traits was based on the fact that two or three traits seemed too few to compute a meaningful index of accuracy, whereas five or six resulted in the selection of some highly evaluative traits. Furthermore, the selection of traits one through four in the moderate- and high-relevance conditions but of Traits 2 through 5 in the low-relevance condition was necessary to equate evaluativeness across conditions. If we selected traits one through four in the low-relevance condition, we found that evaluativeness in that condition was significantly lower than in the other conditions (which is the confound we are trying to overcome here).

⁷ The nonevaluative high-relevance traits included, for example, "perfectionistic," "influential," "picky," "carefree," "hardworking," "materialistic," "crazy," "introspective," "pessimistic," "secure," "idealistic," "impatient," "creative," "contemplative," "self-conscious," "open-minded," "judgmental," "studious," "fearful," and "spontaneous." Because each couple in the high-relevance condition rated an idiosyncratic list of traits, the full list of nonevaluative high-relevance traits contains over 100 traits. The full list is available from Michael J. Gill. The nonevaluative low-relevance traits involved a subset of the low-relevance traits in Table 2, as follows: (a) the highly evaluative traits "assertive," "cold," "sympathetic," "kind," "careless," and "intelligent" were excluded for all couples, (b) about one half of the couples rated one or more of the moderately evaluative traits "talkative," "fault-finding," "careless," and "calm," and (c) most couples rated some combination of four of the low-relevance traits that remain in Table 2 after the aforementioned traits are deleted.

tion were slightly clearer on our measure of target-based accuracy than on our measure of total accuracy (on which the difference between the high- and low-relevance conditions was marginal). Thus, the relatively high levels of pragmatic accuracy we uncovered seem to reflect target-based accuracy rather than stereotype accuracy.

Relation of Accuracy to Relationship Quality

To test whether pragmatic accuracy was especially predictive of relationship quality, we conducted correlational analyses within each relevance condition. We expected that accuracy within the low-relevance condition would not be predictive of relationship quality, whereas accuracy within the moderate- and high-relevance conditions would be predictive of relationship quality.

We correlated our measures of relationship quality—that is, commitment, target confidence, perceiver confidence—with the accuracy correlations within each relevance condition. See Figure 3 for the correlations. In the low-relevance condition, there were no significant correlations between accuracy and any measure of relationship quality, all p s $> .21$. As predicted by the pragmatic accuracy formulation, however, accuracy within the moderate-relevance condition was associated with relationship commitment ($p < .01$) and target confidence ($p = .05$). The difference between the commitment correlations within the low- and moderate-relevance conditions approached significance ($p < .08$, one-

tailed), but the difference between the target confidence correlations did not ($p > .33$). Also consistent with the pragmatic accuracy formulation, accuracy within the high-relevance condition was significantly related to relationship commitment ($p < .01$) and to perceiver confidence ($p < .01$). The difference between the commitment correlations in the low- and high-relevance conditions was statistically significant ($p < .01$, one-tailed), and the difference between the perceiver confidence correlations approached significance ($p < .08$, one-tailed).

However, did accuracy on relevant traits continue to be especially predictive of relationship quality even when we examined the nonevaluative subsets of traits described earlier (i.e., the subsets designed to create equivalently nonevaluative sets of traits in each condition)? Yes. Consistent with the pragmatic accuracy formulation, accuracy on nonevaluative high-relevance traits was significantly correlated with relationship commitment, $r(27) = .40$, $p < .04$, whereas accuracy on nonevaluative low-relevance traits was not, $r(28) = -.16$, *ns*. These correlations were significantly different from each other ($p < .02$, one-tailed). All other correlations involving nonevaluative traits were nonsignificant (p s $> .14$). Taken together, the results in Figure 3 and the correlations just presented support the pragmatic accuracy prediction that accuracy on relevant traits—even when those traits are nonevaluative—is uniquely associated with psychological benefits such as greater relationship commitment, greater feelings of being known, and greater feelings of knowing one's partner.

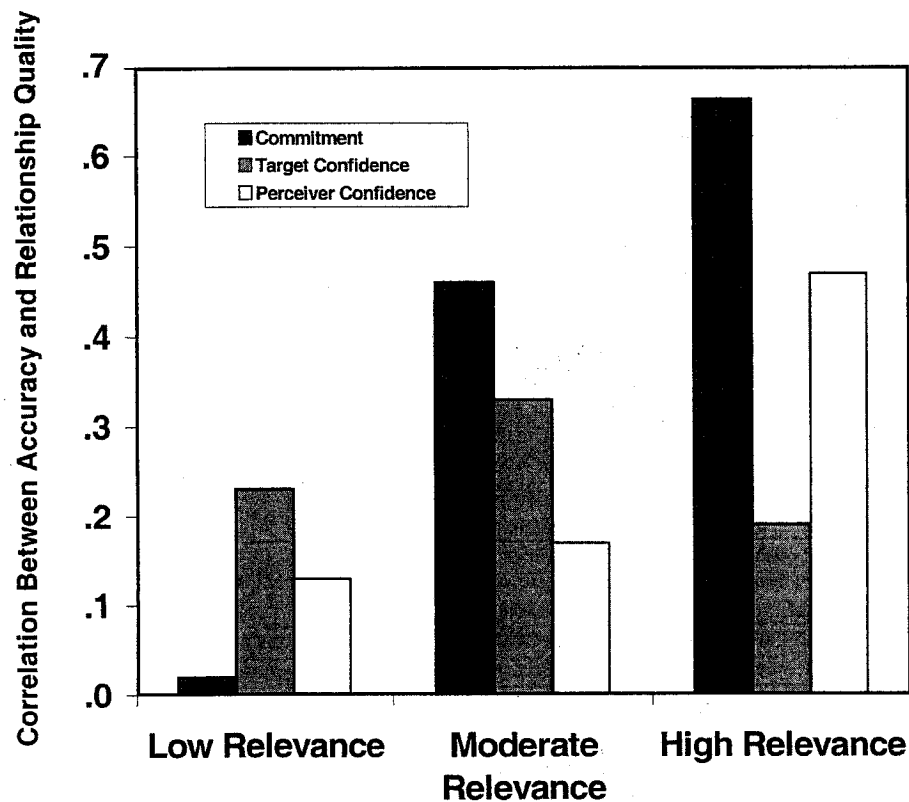


Figure 3. Correlations between accuracy of perceiver knowledge of target self-views and relationship quality in each condition of trait relevance.

Discussion

We argued that in romantic relationships partners should be particularly motivated to develop pragmatic accuracy, that is, accuracy marked by a heightened understanding of a partner's self-views on relationship-relevant traits. Consistent with our predictions, relevance was associated with accuracy such that perceivers demonstrated more accurate knowledge of their partner's self-views on traits that were moderate and high in relevance than on traits that were low in relevance. In addition, *only* pragmatic accuracy correlated with several indices of relationship quality. As in Study 1, of course, the correlational nature of the data precludes causal interpretation.

Ratings collected from independent judges confirmed that moderate- and high-relevance traits were more relevant to romantic relationships than were low-relevance traits (as we expected) and revealed that they were also more evaluative (which we did not expect, but which we do not see as problematic, especially given our follow-up analyses involving the nonevaluative traits). Our finding of high levels of accuracy on evaluative traits might seem inconsistent with work on positive illusions in relationships (e.g., Murray & Holmes, 1997). That literature suggests that relationship quality is enhanced by exceedingly positive—rather than accurate—views of one's partner on evaluative traits. In support of this viewpoint, researchers have shown that relationship quality is higher insofar as relationship partners believe that their partner is better than "the typical partner." We suggest that it is quite possible to both view one's partner as better than the typical partner and accurately describe that partner's relationship-relevant self-views, provided that the partner cultivates a relationship-specific identity that actually is higher in desirability than both his or her general behavior pattern and the behavior patterns of most other potential partners. Thus, positive illusions and pragmatic accuracy can coexist insofar as people negotiate relationship-specific identities.

Support for this possibility comes from a recent investigation of the identities that people negotiated with their relationship partners. Swann et al. (2002) discovered that, within particular relationships, people want exceptionally favorable perceptions of themselves on dimensions that are critical to the well-being of those relationships. Thus, for example, people want members of their sports teams to acclaim their athletic prowess, they want their art instructors to be impressed with their artistic talent, and they want their dating partners to see them as highly attractive. In short, on attributes critical to the survival of relationships, people want evaluations that could reasonably be labeled "positive illusions" if compared with the evaluations that such persons (or other people) typically elicit.

What's more, people enact relationship-specific behaviors that actually elicit evaluations that match their idealized desired evaluations and, more surprisingly, they view these elicited evaluations as highly accurate. For example, even people who thought of themselves as generally rather homely preferred—and actually elicited—highly positive appraisals of their physical attractiveness from their dating partners and subsequently judged these appraisals as accurate. This suggests that people can, on the one hand, view themselves as only moderately attractive ("I am not beautiful in the eyes of many people"), yet, on the other hand, view themselves as quite attractive within the context of a relationship ("I am

beautiful to my partner; indeed, I work hard to make it so"). In short, people construct idiosyncratically skewed, idealized social realities that make their relationships "work." This suggests that a partner's appraisal of a target could simultaneously be a "positive illusion" (i.e., it exceeds the appraisals that the target typically elicits or appraisals elicited by a "typical target") while also showing a high level of pragmatic accuracy (i.e., it faithfully captures the target's relationship-specific self-concept).

Note that although we are arguing here that relationship-specific positive illusions may be pragmatically accurate, our findings suggest that pragmatic accuracy is more than this. In particular, as previously discussed, when we focused our analyses on nonevaluative traits—whose neutral character precludes them from being constructed—we found that accuracy on relationship-relevant traits continued to exceed accuracy on less relevant traits. Therefore, at least some positive illusions can be understood as members of a broader class of pragmatically accurate beliefs that may also include neutral or even negative beliefs.

General Discussion

To the best of our knowledge, ours are the first studies to support Swann's (1984) contention that perceivers acquire pragmatic accuracy—accuracy that facilitates relationship-specific interaction goals—in their social relationships. In Study 1, we tested the notion that a group benefits from consensus about members' behaviors in the group context. We found greater consensus when group members described a target's behavior in the group context as compared with outside the group context. More specifically, we found that fraternity brothers achieved greater consensus when rating a target's behavior around the fraternity house as compared to rating his behavior with his nuclear family, and family members achieved greater consensus when rating a target's behavior with his nuclear family as compared with rating his behavior with his fraternity brothers. Additionally, when we correlated consensus with relationship quality, the only correlation that approached significance suggested that consensus among fraternity brothers regarding targets' behaviors around the fraternity house was positively associated with relationship quality. This is consistent with our expectation that pragmatic accuracy would be especially predictive of relationship quality. Correlations involving consensus among family members, however, were not consistent with our expectations (although this might be attributable to our having a very small sample of family members).

Study 2 was based on the idea that accurate knowledge of a partner's self-views on relationship-relevant traits is pragmatic in the context of a romantic relationship. The study revealed that, in fact, people do acquire more accurate knowledge of their partner's self-views on traits moderate in relevance (those spontaneously generated by many dating couples) and high in relevance (those spontaneously generated by a member of the specific couple whose accuracy was being examined) as compared with traits low in relevance. Furthermore, Study 2 provided clear evidence that accuracy on traits either moderate or high in relevance was more strongly correlated with relationship quality than was accuracy on low-relevance traits. Specifically, there was a tendency for accuracy on both moderate- and high-relevance traits to be more predictive of relationship commitment than was accuracy on low-

relevance traits and a tendency for accuracy on high-relevance traits to be more strongly associated with a perceiver's feelings of knowing his or her partner than was accuracy on low-relevance traits. As in Study 1, the correlational nature of these data precludes a causal argument, and it is plausible that relationship quality fosters pragmatic accuracy rather than vice versa.

Our findings are important because they challenge some widely accepted assumptions regarding the fundamental products of the person perception process and the appropriate criteria for gauging the accuracy of the person perception process. For example, most prominent work in the field assumes that person perception fundamentally involves the formation of inferences about broad personality traits (for reviews, see Funder, 1999; Kenny, 1994). Indeed, research in person perception has focused nearly exclusively on personality trait judgments. We do not dispute that people make inferences about broad personality traits. We do, however, question the idea that such inferences are the most important product of the person perception process. Instead, we propose that inferences about relationship-specific identities (see McHenry, 1971) are critical products of the person perception process, at least as important as are trait inferences. Indeed, if such relatively specific inferences were not critical, it seems unlikely that we would have uncovered evidence that they were associated with high levels of consensus and accuracy and that this consensus and accuracy would be uniquely associated with relationship well-being. In addition to challenging the notion that trait inferences are the essential product of the person perception process, our formulation also questions the assumption that a comprehensive sampling of traits (e.g., covering all of the Big Five factors) provides the most appropriate criterion for gauging accuracy. Our data suggest that people may be especially likely to form accurate judgments about relationship-relevant traits as opposed to traits that capture a broader range of individual differences and that accuracy on relationship-relevant traits is more strongly correlated with relationship functioning than is accuracy on traits that capture a broader range of individual differences. This suggests that a set of traits that captures the entire range of individual differences might not be the best criterion for accuracy, given that accurate knowledge across such traits has little association with relationship quality.

These arguments concerning the most appropriate criteria for gauging accuracy may remind one of the literature on empathic accuracy (see Ickes, 1997, 2003), which has some important features in common with our pragmatic accuracy approach. As does our formulation, the empathic accuracy literature suggests that broad personality trait inferences are not the sole or most fundamental products of the person perception process. Rather, perceivers also want "to accurately infer the thoughts, feelings, motives, and intentions of other people" (Ickes, 1997, p. 1). Notably, these phenomena are more transitory than are enduring personality dispositions, yet a substantial literature now attests to the relationship benefits that can be derived from high levels of empathic accuracy (although low levels are sometimes beneficial, too; see Simpson, Ickes, & Blackstone, 1995). According to the empathic accuracy approach, then, a perceiver's ability to infer the temporary thoughts and feelings of a target is also an appropriate and important criterion for accuracy.

Conclusions and Implications

Our findings provide support for Swann's (1984) suggestion that pragmatic accuracy—accuracy that facilitates relationship-relevant goals—is something that perceivers might achieve with relative frequency. Furthermore, in support of the notion that pragmatically accurate knowledge is psychologically important, we presented evidence that pragmatic accuracy covaries with relationship quality to a greater degree than do other types of accuracy. At the same time, our data suggest that knowledge of a person across a massive array of traits and contexts is relatively unlikely to be achieved by perceivers. We suggest, however, that this lack of knowledge is often unproblematic because it contributes relatively little to successful relationship functioning.

From our vantage point, James's (1907) pragmatic perspective on truth provides a useful perspective on the process of person perception. Apparently, people are not engaged in an indiscriminate attempt to view their relationship partners objectively or to discern every possible aspect of them. Instead, people are content to seek—or construct with their interaction partners (Swann et al., 2002)—smaller "truths" regarding a target in a limited set of contexts and in terms of a limited number of dimensions. Such smaller truths, we believe, are especially likely to be "helpful in life's practical struggles" (James, 1907, p. 96). In the final analysis, this may be what the person perception process is all about.

References

- Bem, D. J., & Allen, A. (1974). On predicting some of the people some of the time: The search for cross-situational consistencies in behavior. *Psychological Review*, *81*, 506–520.
- Bernieri, F. J., Zuckerman, M., Koestner, R., & Rosenthal, R. (1994). Measuring person perception accuracy: Another look at self–other agreement. *Personality and Social Psychology Bulletin*, *20*, 367–378.
- Blackman, M. C., & Funder, D. C. (1998). The effect of information on consensus and accuracy in personality judgment. *Journal of Experimental Social Psychology*, *34*, 161–181.
- Britt, T. W., & Sheppard, J. A. (1999). Trait relevance and trait assessment. *Personality and Social Psychology Review*, *3*, 108–122.
- Bugental, D. B. (2000). Acquisition of the algorithms of social life: A domain-based approach. *Psychological Bulletin*, *126*, 187–219.
- Cohen, J., & Cohen, P. (1983). *Applied multiple regression/correlation analysis for the behavioral sciences*. Hillsdale, NJ: Erlbaum.
- Corsini, R. J. (1956). Understanding and similarity in marriage. *Journal of Abnormal and Social Psychology*, *52*, 327–332.
- Costa, P. T., & Macrae, R. R. (1998). Trait theories of personality. In D. F. Barone, M. Hersen, & V. B. Van Hasselt (Eds.), *Advanced Personality* (pp. 103–121). New York: Plenum.
- Cronbach, L. J. (1955). Processes affecting scores on "understanding of others" and "assumed similarity." *Psychological Bulletin*, *52*, 177–193.
- Dunning, D., Griffin, D. W., Milojkovic, J. D., & Ross, L. (1990). The overconfidence effect in social prediction. *Journal of Personality and Social Psychology*, *58*, 568–581.
- Fiske, A. P. (1992). The four elementary forms of sociality: Framework for a unified theory of social relations. *Psychological Review*, *99*, 689–723.
- Funder, D. C. (1999). *Personality judgment: A realistic approach to person perception*. New York: Academic Press.
- Funder, D. C., Kolar, D. C., & Blackman, M. C. (1995). Agreement among judges of personality: Interpersonal relations, similarity, and acquaintanceship. *Journal of Personality and Social Psychology*, *69*, 656–672.
- Gage, N. L., & Cronbach, L. J. (1955). Conceptual and methodological problems in interpersonal perception. *Psychological Review*, *62*, 411–422.

- Gilbert, D. T. (1998). Ordinary personology. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., pp. 89–150). New York: McGraw-Hill.
- Gill, M. J., Swann, W. B., Jr., & Silveira, D. H. (1998). On the genesis of confidence. *Journal of Personality and Social Psychology*, *75*, 1101–1114.
- Ickes, W. (1997). *Empathic accuracy*. New York: Guilford Press.
- Ickes, W. (2003). *Everyday mind reading: Understanding what other people think and feel*. Amherst, NY: Prometheus Books.
- James, W. (1907). *Pragmatism*. New York: Longman Green.
- John, O. P. (1990). The "Big Five" factor taxonomy: Dimensions of personality in the natural language and questionnaires. In L. A. Pervin (Ed.), *Handbook of personality theory and research* (pp. 66–100). New York: Guilford Press.
- John, O. P., & Robins, R. W. (1993). Determinants of interjudge agreement on personality traits: The Big Five domains, observability, evaluativeness, and the unique perspective of the self. *Journal of Personality*, *61*, 521–551.
- Kenny, D. A. (1991). A general model of consensus and accuracy in interpersonal perception. *Psychological Review*, *98*, 155–163.
- Kenny, D. A. (1994). *Interpersonal perception: A social relations analysis*. New York: Guilford Press.
- Kenny, D. A., Kiefer, S. C., Smith, J. A., Ceglanski, P., & Kulo, J. (1996). Circumscribed accuracy among well-acquainted individuals. *Journal of Experimental Social Psychology*, *32*, 1–12.
- Kenny, D. A., & Winquist, L. (2001). The measurement of interpersonal sensitivity: Consideration of design, components, and unit of analysis. In J. A. Hall & F. J. Bernieri (Eds.), *Interpersonal sensitivity: Theory and measurement*. (pp. 265–302). Mahwah, NJ: Erlbaum.
- Kenrick, D. T., Maner, J. K., Butner, J., Li, N. P., Becker, V., & Schaller, M. (2002). Dynamical evolutionary psychology: Mapping the domains of the new interactionist paradigm. *Personality & Social Psychology Review*, *6*, 347–356.
- Krueger, J., & Funder, D. C. (in press). Towards a balanced social psychology: Causes, consequences, and cures for the problem-seeking approach to social behavior and cognition. *Behavioral and Brain Sciences*.
- Levesque, M. J., & Kenny, D. A. (1993). Accuracy of behavioral predictions at zero acquaintance: A social relations analysis. *Journal of Personality and Social Psychology*, *65*, 1178–1187.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.
- McHenry, R. (1971). New methods of assessing accuracy in person perception. *Journal of the Theory of Social Behavior*, *1*, 109–119.
- Mischel, W., & Shoda, Y. (1995). A cognitive-affective system theory of personality: Reconceptualizing situations, dispositions, dynamics, and invariance in personality structure. *Psychological Review*, *102*, 246–268.
- Murray, S. L., & Holmes, J. G. (1997). A leap of faith? Positive illusions in romantic relationships. *Personality and Social Psychology Bulletin*, *23*, 586–604.
- Nisbett, R. E., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice-Hall.
- Roskos-Ewoldsen, D. R., & Fazio, R. H. (1992). On the orienting value of attitudes: Attitude accessibility as a determinant of an object's attraction of visual attention. *Journal of Personality and Social Psychology*, *63*, 198–211.
- Rubin, D. C., & Kozin, M. (1984). Vivid memories. *Cognition*, *16*, 81–95.
- Rusbult, C. E. (1980). Commitment and satisfaction in romantic associations: A test of the investment model. *Journal of Experimental Social Psychology*, *16*, 172–186.
- Shrout, P. E., & Fleiss, J. L. (1979). Intraclass correlations: Uses in assessing rater reliability. *Psychological Bulletin*, *86*, 420–434.
- Simpson, J. A., Ickes, W., & Blackstone, T. (1995). When the head protects the heart: Empathic accuracy in dating relationships. *Journal of Personality and Social Psychology*, *69*, 629–641.
- Swann, W. B., Jr. (1984). Quest for accuracy in person perception: A matter of pragmatics. *Psychological Review*, *91*, 457–477.
- Swann, W. B., Jr. (1990). To be adored or to be known?: The interplay of self-enhancement and self-verification. In E. T. Higgins & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition* (Vol. 2, pp. 408–448). New York: Guilford Press.
- Swann, W. B., Jr. (1997). The trouble with change: Self-verification and allegiance to the self. *Psychological Science*, *8*, 177–180.
- Swann, W. B., Jr., Bosson, J. K., & Pelham, B. W. (2002). Different partners, different selves: The verification of circumscribed identities. *Personality and Social Psychology Bulletin*, *28*, 1215–1228.
- Swann, W. B., Jr., De La Ronde, C., & Hixon, J. G. (1994). Authenticity and positivity strivings in marriage and courtship. *Journal of Personality and Social Psychology*, *66*, 857–869.
- Swann, W. B., Jr., & Gill, M. J. (1997). Confidence and accuracy in person perception: Do we know what we think we know about our relationship partners? *Journal of Personality and Social Psychology*, *73*, 747–757.
- Taft, R. (1955). The ability to judge other people. *Psychological Bulletin*, *52*, 1–23.
- Taylor, S. E., & Brown, J. D. (1988). Illusion and well-being: A social psychological perspective on mental health. *Psychological Bulletin*, *103*, 193–210.
- Zebrowitz, L. A., & Collins, M. A. (1997). Accurate social perception at zero acquaintance: The affordances of a Gibsonian approach. *Personality and Social Psychology Review*, *3*, 204–223.

Received July 9, 2002

Revision received September 22, 2003

Accepted October 6, 2003 ■

Wanted: Your Old Issues!

As APA continues its efforts to digitize journal issues for the PsycARTICLES database, we are finding that older issues are increasingly unavailable in our inventory. We are turning to our long-time subscribers for assistance. If you would like to donate any back issues toward this effort (preceding 1982), please get in touch with us at journals@apa.org and specify the journal titles, volumes, and issue numbers that you would like us to take off your hands.