Running Head: WHEN IT'S BAD TO BE FRIENDLY AND SMART

When It's Bad To Be Friendly and Smart:

The Desirability of Sociability and Competence Depends on Morality

Justin F. Landy,[a] Jared Piazza,[b] and Geoffrey P. Goodwin[c]

Word Count: 12,223

Author Affiliations (Address Correspondence to J.F.L.):

[a] Center for Decision Research
University of Chicago Booth School of Business
5807 S Woodlawn Avenue
Chicago, IL 60637 USA

[b] Department of Psychology
Fylde College
Lancaster University
Lancaster, United Kingdom
LA1 4YF

[c] Department of Psychology
University of Pennsylvania
3720 Walnut Street
Philadelphia, PA 19104 USA

Abstract

Morality, sociability, and competence are distinct dimensions in person perception. We argue that a person's morality informs us about their likely intentions, whereas their competence and sociability inform us about the likelihood that they will fulfill those intentions. Accordingly, we hypothesized that whereas morality would be considered unconditionally positive, sociability and competence would be highly positive only in moral others, and would be less positive in immoral others. Using exploratory factor analyses, Studies 1a and 1b distinguished evaluations of morality and sociability. Studies 2-5 then showed that sociability and competence are positive contingent on morality – Study 2 demonstrated this phenomenon, while the remaining studies explained it (Study 3), generalized it (Studies 3-5), and ruled out an alternative explanation for it (Study 5). Study 6 showed that the positivity of morality traits is independent of other morality traits. These results support a functionalist account of these dimensions of person perception.

Keywords: morality, sociability, competence, person perception, dimensional models

When it's Bad to Be Friendly and Smart:
The Desirability of Sociability and Competence Depends on Morality

Social cognition researchers have posited that there are two "fundamental dimensions" along which we categorize other people (Abele, Cuddy, Judd, & Yzerbyt, 2008; Fiske, 2012; Fiske, Cuddy & Glick, 2007; Judd, James-Hawkins, Yzerbyt, & Kashima, 2005). One, usually referred to as "warmth" (Fiske et al., 2007) or "communion" (Abele & Wojciszke, 2013), captures how a person relates to others. The other, usually referred to as "competence" (Fiske et al., 2007), "agency" (Abele & Wojciszke, 2013), or "ability" (Brycz & Wojciszke, 1992), captures a person's ability to accomplish their goals. Different researchers employ different names for these dimensions, but their theorizing largely overlaps (Abele et al., 2008; for a detailed review of this literature, see Abele & Wojciszke, 2014). One attractive feature of this line of theorizing is that it accounts for why these dimensions are so important to social judgment from a functionalist standpoint (Fiske et al., 2007). Each dimension is said to convey functionally important information: warmth is said to inform us about a person's likely intentions (e.g., is the person benevolent or hostile?), while competence is said to inform us about a person's ability to carry out those intentions successfully.

Recently, it has been noted that the dimension of warmth seems to conflate two distinct aspects of a person: morality – exemplified by traits like honesty, fairness, and sincerity – and sociability – exemplified by traits like friendliness, extroversion, and playfulness (Bauman & Skitka, 2012; Brambilla, Rusconi, Sacchi, & Cherubini, 2011; Brambilla, Sacchi, Rusconi, Cherubini, & Yzerbyt, 2012; Goodwin, Piazza, & Rozin, 2014; Leach, Ellemers, & Barreto, 2007). The fact that two-dimensional theories conflate morality and sociability might not be

problematic if these social judgments were found to function in highly similar ways. However, several lines of evidence indicate that morality and sociability play different roles in impression formation. Moral information tends to trump warmth/sociability information in overall importance (Abele & Brack, 2003; Brambilla et al., 2011; Cottrell, Neuberg, & Li, 2009; Goodwin et al., 2014). Moreover, people seem to consider extroversion (a component of sociability, perhaps the chief component) to be a kind of skill, quite unlike morality (Reeder, Messick, & Van Avermaet, 1977). And, initial judgments of morality based on mixed positively- and negatively-valenced information are negative and resistant to change, whereas such judgments of sociability are positive and change quickly in response to new information (Brycz & Wojciszke, 1992). Morality and sociability information therefore appear to be processed very differently. We argue that they are best thought of as two separate dimensions of person perception, alongside competence, and we provide new empirical evidence for this in Studies 1a and 1b.

We are not the first researchers to draw a distinction between evaluations of morality and sociability. Similar distinctions (using different terminologies) have been made in research on persuasion (trustworthiness versus attractiveness: Kelman, 1958, 1961), personality (honesty versus extroversion and agreeableness: Ashton & Lee, 2001), and face perception (trustworthiness versus likeability: Rule et al., 2010). Yet the field of social cognition has been reluctant to embrace this distinction. Even researchers who differentiate between morality and sociability typically conceptualize them as different components of a single superordinate dimension that is separate from competence, thereby preserving a two-dimensional model of social cognition (e.g., in intergroup relations, Brambilla et al., 2012; Leach et al., 2007; for a review, see Brambilla & Leach, 2014). In the realm of person perception, one of the clearest

articulations of this distinction to date has been made by Brambilla and colleagues (2011). In confirmatory factor analyses, these researchers showed that a three-factor model separating morality, sociability, and competence modeled participants' ratings of the relevance of various traits for forming impressions better than did a two-factor model. Nonetheless, they maintained that "our perspective is not to question the validity and the usefulness of the dual-dimension view of social judgments; warmth and competence are clearly two fundamental dimensions of social perception" (Brambilla et al., 2011, p. 136). Based on the research reviewed above, we disagree; morality and sociability are distinct evaluations. One of the initial goals of the present paper is therefore to provide further substantiation for this claim.

A second goal is to provide evidence for the social function of these separate dimensions of social judgment. Existing accounts stress the separate functions of warmth and competence as being key to their social importance. We extend these accounts by proposing that morality, sociability, and competence each make distinct contributions to person perception. In doing so, we rearticulate the value of morality and competence, while providing a novel account of the information value of sociability. We turn now to this functionalist account.

One critical task of social cognition is to predict others' intentions towards us and towards people we care about (see e.g., Abele & Wojciszke, 2007; Pizarro & Tannenbaum, 2011; Wojciszke, Bazinska, & Jaworski, 1998; Wojciszke, Dowhyluk, & Jaworski, 1998; see also Cottrell et al., 2007; Cuddy, Fiske, & Glick, 2008; Fiske et al., 2007; Leach et al., 2007). In contrast to some prior theorizing, we argue that this information is best conveyed by a person's moral character rather than by their warmth/sociability (see also Goodwin et al., 2014).

A second important task is to predict a person's likelihood of accomplishing what they intend to do. We concur with existing two-dimensional models that a person's competence provides information of this sort (Abele & Wojciszke, 2007; Fiske et al., 2007) – this seems true almost by definition.

However, we propose that sociability conveys a distinct and important sort of information that two-dimensional models overlook, and which fulfills a third task of social cognition. Sociability traits, such as extroversion, convey how effectively a person can build alliances and recruit others to support their moral or immoral intentions (see Ashton, Lee, & Paunonen, 2002), information which is important from an adaptive standpoint (DeScioli & Kurzban, 2009). Someone who is outgoing, friendly, and charismatic will find it easier to recruit allies to support their intentions than will someone who is introverted, cold, and unfriendly, regardless of whether their intentions are good or bad. Indeed, it has been argued that the core element of the personality trait extroversion is not a mere preference for social interaction, but rather a tendency to attract social attention and garner social support, and that, as a consequence, "extraverts tend to win the competition for social attention over introverts and are thereby more likely to attract the most desirable allies, friends, and mates" (Ashton, Lee, & Paunonen, 2002, p. 251). Furthermore, it has been empirically demonstrated that extroverted individuals have larger social networks than do introverted individuals (Pollet, Roberts, & Dunbar, 2011). In turn, individuals with larger support networks are seen as particularly desirable as allies and exchange partners (Curry & Dunbar, 2011). For the same reason, in contexts of group conflict, socially well-connected individuals are seen as more intimidating and formidable foes (Fessler & Holbrook, 2013). Thus, sociability may serve as a reliable cue that an individual or group has the social

support, or the ability to attract the social support, needed to carry out their intentions, whether those intentions are helpful or harmful.

Brambilla et al. (2011) made a somewhat similar point. Following Leach et al. (2007), they proposed that sociability "pertains to cooperation and to forming connections with others" (p. 135). However, our proposal is more specific than previous accounts because it posits that predicting a person's ability to form alliances (their "social strength") is a fundamental social cognitive task, just as predicting others' intentions and capacities are fundamental social cognitive tasks. According to our model, morality, competence, and sociability each convey something unique and functionally important about others in our social worlds, and therefore constitute distinct dimensions in person perception.

It follows that on our model, morality predicts the nature of another person's goals, whereas competence and sociability both predict the likelihood that a person will accomplish their goals, though for different reasons – competence directly predicts the likelihood of goal attainment via one's own actions, whereas sociability indirectly predicts the likelihood of goal attainment via assistance from one's social network (though, of course, there may also be certain sorts of goals that sociability directly aids in accomplishing). Accordingly, whereas morality should always be positive in others, because it is always better for us if others have good rather than bad intentions toward us, competence and sociability should be positive *contingent upon morality*. That is, if a person's intentions toward us are good (e.g., helpful or just), then it is desirable that they be competent and sociable enough to carry out those intentions. However, if a person's intentions towards us are bad (e.g., harmful or unjust), then their being competent or sociable is not desirable, and is unlikely to improve our impressions of them. In some cases, a

person's competence or sociability might even serve to amplify their immorality, leading to more negative overall impressions of them (for related research, see Piazza, Goodwin, Rozin, & Royzman, 2014; Wojciszke, Bazinska et al., 1998).

We therefore made the following three predictions. First, we predicted that morality traits comprise a distinct dimension of person perception, independent from sociability and competence. We refer to this as the *Morality Differentiation Hypothesis*. We tested this hypothesis by factor analyzing trait judgments of real social targets that our participants knew well. Second, we predicted that positive morality traits are always positive in person perception, and negative morality traits are always negative. We refer to this as the *Morality Dominance Hypothesis*. Third, we predicted that positive sociability and competence traits are positive in moral targets, but less positive (and sometimes even negative) in immoral targets. As a consequence, there should be an interaction between a target's morality and their sociability/competence in predicting the positivity of both sociability and competence traits. We refer to this as the *Morality Dependence Hypothesis*. To test the second and third hypotheses, we relied on a variety of dependent measures, including global impressions of others, preferences for various traits in others, and anticipated changes to global impressions following the addition of new trait information.

These predictions build upon earlier research that has partially explored the contingent positivity of competence (though not sociability) and the non-contingent positivity of morality. For instance, Peeters (1992) found that people preferred that their friends possess traits related to competence (e.g., industrious, practical), but preferred that their enemies lack these traits. However, people preferred that both their friends and their enemies possess traits related to morality (e.g., trustworthy, tolerant).

In another relevant study, Wojciszke, Bazinska et al. (1998, Study 4) found that impressions of targets with immoral goals were always negative, but were especially negative when the immoral target exhibited high competence. Similarly, impressions of targets with moral goals were always positive, but were more positive when the targets were also competent. These results are consistent with our theorizing regarding the functional role of competence. However, in this study, the manipulation of competence contained information about whether or not the target actually succeeded in fulfilling his or her goal. For example, one description of a moral and competent target read, "Although himself in a hurry, Andrew stopped on his way seeing a helpless woman; he right away found what was wrong with her car and got it going using an ingenious trick." In contrast, the moral, yet incompetent target also wanted to fix the woman's car, but was unable to do so. The information about the target's competence ("he right away found what was wrong with her car and got it going using an ingenious trick") is therefore confounded with the actual outcome of the scenario (he "got it going"). Consequently, it is unclear whether competent, moral targets were judged more positively more than incompetent, moral targets because of their competence, *per se* (i.e., because of their traits), or simply because they actually caused more positive outcomes in the world. An analogous argument applies to why impressions of competent, immoral targets were more negative than impressions of incompetent, immoral targets in this study.

Therefore, to test for the contingent nature of competence more stringently, we focused only on trait possession, and did not conflate it with goal attainment. We also extended the contingency hypothesis to sociability, which has only recently been theoretically and empirically separated from morality.

Studies 1a and 1b

Our first goal was to test the Morality Differentiation Hypothesis, by examining whether people separately evaluate others' morality and sociability. Existing research has found some degree of separation between these constructs – for instance, Brambilla et al. (2011) found that judgments of the relevance of morality, sociability, and competence traits for forming impressions were described well by a three-factor model. Our aim was to build upon this research by relying on direct judgments of trait possession, which are arguably more natural and less abstract than judgments of trait relevance. We were particularly interested in whether judgments of others' morality and sociability are more closely related to each other than they are to judgments of competence. We had participants rate real people whom they knew on the degree to which they exemplified various traits related to morality, sociability, and competence, then used exploratory factor analyses to uncover the latent constructs underlying these ratings. Two-dimensional models of person perception, which treat morality and sociability as closely related elements of the same prosocial dimension, should predict that morality and sociability traits would factor together, and therefore that two factors would emerge in participants' ratings: a warmth factor and a competence factor. Our model predicts instead that three factors should emerge: separate morality and sociability factors, and a competence factor.

Method

Participants. Participants were recruited through Amazon Mechanical Turk (after exclusions for incomplete studies and failed "Captcha" verifications, Study 1a: $N = 507$ participants, 53% female; Study 1b: $N = 414$, 43% female). We recruited at least 400 participants each in Studies 1a and 1b to ensure that the results of our factor analyses would be reliable (Field, 2005).

Procedure. Participants were told that they would be rating the personalities of several people that they knew on various trait dimensions. In Study 1a, participants thought of six target individuals they knew personally, each one fitting a different criterion: someone that the participant liked, disliked, respected, or did not respect, a parent or parental figure, and a teacher or mentor. These six targets were chosen to cover a wide array of meaningful social relationships. The first four targets were included because liking and respect have both been an important focus of prior research on two-dimensional models of person perception (Wojciszke, Abele, & Baryla, 2009), and the latter two were added to increase coverage across a diverse range of relationships. In Study 1b, participants thought only of liked, disliked, respected, and disrespected individuals. Across both studies, the criteria were presented on separate pages, with the order of presentation randomized for each participant. Participants were instructed not to use the same person for more than one criterion. In order to ensure that participants were thinking of particular people in their life, they were asked to type in the initials of the person they thought of for each criterion.

Participants indicated how much each target possessed eighteen personality traits on 1-9 Likert scales. In Study 1a, six traits each were chosen to instantiate morality (*moral, principled, honest, trustworthy, fair, responsible*), sociability (*sociable, warm, friendly, easy-going, extroverted, playful*), and competence (*competent, capable, intelligent, effective, skillful, talented*) on the basis of prior research (Goodwin et al., 2014, Study 1). In Study 1b, we included traits that instantiate two of these dimensions simultaneously (see Online Appendix for details of a pre-study showing this). Thus, participants rated the targets on traits related to morality (*moral, honest, fair*), sociability (*sociable, friendly, extroverted*), competence

(*competent, effective, talented*), morality and sociability (*humble, respectful, compassionate*), morality and competence (*principled, responsible, disciplined*), and sociability and competence (*cooperative, enthusiastic, dynamic*). These traits were included in order to ensure that if the predicted three-factor solution emerged in Study 1a, this would not be attributable to our selecting trait terms that instantiate only the non-sociable aspects of morality and the non-moral aspects of sociability. After making their ratings, participants completed a brief demographic questionnaire. Aside from basic demographics, no unreported measures were collected in any study reported in this paper.

Results and Discussion

For each target, we factor analyzed participants' trait ratings using Maximum Likelihood Exploratory Factor Analyses (EFAs) with direct quartimin rotation (equivalent to direct oblimin rotation with a delta value of zero; see Fabrigar, Wegener, MacCallum, & Strahan, 1999). We used several approaches to determine how many factors to retain in our models. The Kaiser criterion (i.e., retaining all factors with initial eigenvalues greater than 1.0) retained three factors in all ten analyses. However, this approach has been criticized for being arbitrary (see Fabrigar & Wegener, 2012), so we also used other approaches. First, we conducted a parallel analysis, extracting eigenvalues from 100 randomly simulated data sets with the same specifications as our data and comparing the randomly-generated eigenvalues to those extracted from our data. The idea behind parallel analysis is that any extracted factor that has no more explanatory power than a factor extracted from meaningless, random data should not be retained (O'Connor, 2002). Eigenvalues were extracted from the reduced correlation matrices (i.e., from the common variance among the variables, rather than the total variance, which is appropriate for principal

components analysis, but less so for EFA; see Fabrigar & Wegener, 2012). This method frequently overestimates the number of factors that should be retained (Buja & Eyuboglu, 1992), so we treated the results as establishing an upper limit on the number of retained factors (Fabrigar & Wegener, 2012), and, as a conservative test, we compared our initial eigenvalues to the 95th percentile of randomly generated eigenvalues, rather than the mean (Longman, Cota, Holden, & Fekken, 1989). These analyses indicated that between 3 and 6 factors could not be explained by chance, depending on the target of judgment. We next constructed scree plots of eigenvalues extracted from the reduced correlation matrix for each model. All ten scree plots suggested a three-factor structure, though the plots for the liked target and parent in Study 1a could reasonably be interpreted as suggesting a four- or even five-factor structure as well. Lastly, we compared the fit of two-, three- and four-factor models for each target of judgment using the Root Mean Square Error of Approximation (RMSEA) as our measure of model fit. RMSEAs greater than .10 are generally considered to indicate poor fit, .08-.10, marginal fit, .05-.08, acceptable fit, and .05 or less, good fit. Across all ten targets, a two-factor model fit the data poorly (mean RMSEA: .13, range: .10-.15), while a three-factor model fit the data substantially better (mean RMSEA: .08, range: .07-.09). A four-factor model provided almost no improvement in fit over a three-factor model (mean RMSEA: .07, range: .05-.08). Details of all of these analyses can be found in the Online Appendix.

Importantly, none of these methods supported the two-factor model predicted by two-dimensional theories of person perception. Moreover, when we constrained the analyses to produce only two factors, morality and competence traits, rather than morality and sociability traits, tended to factor together, a result which does not accord with any prior theory of which we

are aware. Conversely, three-factor models produced very clear morality, sociability, and competence factors: in Study 1a, traits loaded most highly on their predicted factors in all but one case (mean factor loading: .71, Range: .42-.90), and cross-loadings were generally low.[1] On average, the retained factors explained 65.59% of the total variance in participants' judgments (Range: 62.21-69.91), with the third factor explaining a substantial amount of variance (*M* = 9.50%, Range: 7.81-11.49) over and above the first two. Moreover, the morality factor was always more highly correlated with the competence factor (mean *r* = .53, range: .44-.63) than with the sociability factor (mean *r* = .30, range: .18-.45), which indicates that there was not an especially close connection between morality and sociability.

In Study 1b, traits that instantiate only one dimension of evaluation always loaded together as predicted. Traits that instantiate more than one dimension showed some variability in their loadings, as would be expected. Nonetheless, the three factors that emerged for all four targets were still clearly interpretable as morality, sociability, and competence in each case. On average, the retained factors explained 63.75% of the total variance in participants' judgments (Range: 60.07-65.40), and the third factor explained a substantial amount of variance (*M* = 8.27%, Range: 7.39-9.61). As above, the morality factor always correlated more highly with the competence factor (mean *r* = .56, range: .50-.60) than with the sociability factor (mean *r* = .26, range: .21-.29). Overall, these results provide novel evidence that judgments of morality and sociability, along with competence, are distinct dimensions in person perception, thereby providing support for the Morality Differentiation Hypothesis. We now turn to testing our other hypotheses: that morality is always positive in others (the Morality Dominance Hypothesis),

[1] "Responsible" loaded slightly higher on the competence factor (.45) than the morality factor (.42) for the disrespected target. This was the only instance in which a term did not load most highly on its hypothesized factor.

whereas sociability and competence both depend on others' morality for their positivity (the Morality Dependence Hypothesis).

Study 2

In Study 2, we provided information about a target person's morality, and either their sociability or competence, and asked participants how positive or negative their overall impression of the target person was. We predicted that impressions of moral and immoral targets would always be positive and negative, respectively (the Morality Dominance Hypothesis), but impressions of social and competent targets would depend on their morality, and that sociability and competence traits would make impressions of moral others more positive, but would do so to a lesser extent for immoral others (the Morality Dependence Hypothesis).

Method

Participants. One hundred undergraduates were recruited through the University of Pennsylvania subject pool to complete a study for partial course credit. One student did not complete the whole study, leaving a final sample of $N = 99$ (62% female). In studies 2-6, we aimed to recruit fairly large samples to provide assurance that our findings were robust and replicable. Indeed, the observed statistical power to detect the critical interaction in each of these studies exceeded .99.

Procedure. The study was conducted online. After consenting to participate, participants were presented with 128 questions asking how positive or negative their overall impression of a hypothetical target person was, on a 1-9 Likert scale. Each target person was described by two trait adjectives, one relating to morality, and one relating either to sociability or competence. Each trait term was either positive or negative. Thus, the 128 items constituted a 2 (Target

Morality: high versus low) by 2 (Level of Non-Morality Trait: high versus low) by 2 (Non-Morality Trait: sociability versus competence) within-subjects design with 16 replications in each cell. These replications were formed by fully crossing four trait terms related to each of the dimensions. The morality terms were *honest/dishonest, trustworthy/untrustworthy, moral/immoral,* and *principled/unprincipled,* the sociability terms were *warm/cold, sociable/unsociable, friendly/unfriendly,* and *extroverted/introverted,* and the competence terms were *capable/incapable, intelligent/unintelligent, competent/incompetent,* and *skillful/unskillful.* These terms were chosen on the basis of prior research demonstrating their relevance to the dimensions of interest (Goodwin et al., 2014; Studies 1a and 1b above).

The order of the 128 questions was randomized for each participant, and we also counterbalanced whether the response scale measuring participants' impressions ranged from "Extremely negative" (on the left) to "Extremely positive" (on the right), or vice versa. The moral (or immoral) trait was always presented first. After responding to all 128 questions, participants completed a brief demographics questionnaire.

Results

Preliminary analyses. Responses were scored such that higher numbers indicate more positive impressions of the target. The replications in each of the eight cells of the design all showed good internal reliability, αs > .91, so we averaged across the sixteen questions in each cell to produce one data point per within-subjects condition per participant. The between-subjects counterbalancing of the response scale had no main effect and it did not interact with the other variables aside from a small, difficult-to-interpret four-way interaction with all three within-subjects variables, $F(1,97) = 4.12$, $p = .045$, $\eta^2_p = .041$. Although this interaction is small,

we report the results of the full model including this between-subjects manipulation; the results do not change meaningfully if this variable is omitted from the analysis.

Within-subjects analyses. We conducted a 2 (Target Morality) x 2 (Level of Non-Morality Trait) x 2 (Non-Morality Trait) repeated measures analysis of variance (ANOVA). We found a main effect of Target Morality, $F(1, 97) = 770.67$, $p < .001$, $\eta^2_p = .89$; as can be seen in Figure 1, in both the sociability and competence conditions, impressions of moral targets were always positive and impressions of immoral targets were always negative. These results support the Morality Dominance Hypothesis. As Figure 1 also shows, impressions of sociable and competent targets were contingent upon morality – positive when the target was also moral, and negative otherwise; similarly, impressions of unsociable and incompetent targets were positive if the target was also moral, and negative otherwise. These results support the Morality Dependence Hypothesis. Moreover, further supporting the Morality Dependence Hypothesis, we observed the predicted interaction between Target Morality and Level of Non-Morality Trait, $F(1, 97) = 123.82$, $p < .001$, $\eta^2_p = .56$. While sociability and competence made large positive contributions to impressions of moral targets (within-subjects *d*s: 1.97 and 1.88, respectively), they made smaller contributions to impressions of immoral targets (*d*s: 1.12 and 1.20). This interaction was also found in separate 2 x 2 ANOVAs for the sociability condition, $F(1, 97) = 158.31$, $p < .001$, $\eta^2_p = .62$, and the competence condition, $F(1, 97) = 56.88$, $p < .001$, $\eta^2_p = .37$. For the sake of brevity, we report all main effects and interactions in Studies 2-6 that are not pertinent to our hypotheses in the Online Appendix.

Insert Figure 1 About Here.

Discussion

Overall impressions of moral targets were always positive, and overall impressions of immoral targets were always negative, thereby supporting the Morality Dominance Hypothesis. In contrast, targets high or low in sociability and competence were evaluated positively only if they were high in morality, but were evaluated negatively if they were low in morality. Furthermore, the positive contributions sociability and competence traits made to overall impressions were smaller for immoral targets than for moral targets. These two results support the Morality Dependence Hypothesis.

Study 3

Study 2 supports our view that morality traits are generally seen as unambiguously positive, whereas the positivity of sociability and competence traits is contingent upon morality. In Study 2, we conveyed information about a target's morality using abstract trait terms, but this method arguably lacks ecological validity. Presumably, in the real world, we typically obtain information about a person's moral character by observing or learning about their actions. Therefore, in Study 3, we sought to replicate the results of Study 2 using fictional scenarios in which a person's morality was indicated by their motivations and behaviors, rather than by abstract personality trait terms (similar to Wojciszke, Bazinska et al., 1998). We also obtained ratings of the likelihood that the target would successfully carry out his or her goal (which was either moral or immoral). Based on our functional model presented above, we predicted that these ratings would mediate the effects of sociability and competence on impressions, which themselves would be moderated by morality.

Method

Participants. Six hundred sixty-three participants were recruited online through Amazon Mechanical Turk. Sixteen failed a "Captcha" verification, suggesting that they were "bot" programs, and seven did not complete the study, leaving a final sample of $N = 640$ (31% female).

Method. After consenting to participate, participants were randomly assigned to one cell of a 2 (Target Morality: high versus low) by 2 (Level of Non-Morality Trait: high versus low) by 2 (Non-Morality Trait: sociability versus competence) between-subjects design. This design is exactly analogous to that of Study 1, except the independent variables were manipulated between-subjects rather than within-subjects. In each condition, participants read five scenarios that each described a different target person attempting to accomplish a goal. Within each condition, the target person's goals were always either moral or immoral, but were otherwise matched in content across conditions (see Methodological Supplement for full scenarios). Each scenario also provided information about the main character's sociability or competence, depending on condition.

For each scenario, participants responded to the main dependent variable, "How negative or positive is your overall impression of [character's name]?", the hypothesized mediator, "How likely do you think it is that [character's name] succeeded in [character's goal]?", and a manipulation check, "How immoral or moral is [character's name]?" on 1-9 Likert scales. The order of the dependent variable and the mediator was counterbalanced between-subjects, and the manipulation check was always presented last. The order of the five scenarios was randomized for each participant. After responding to all five scenarios, participants completed a brief demographics questionnaire.

Results

Preliminary Analyses. Across the five different scenarios, responses to the dependent variable and mediator showed good internal reliability (αs .91 and .83, respectively), so we averaged them together to create one composite dependent variable and one composite mediator. The morality manipulation was successful – across the five scenarios, the target person was seen as more moral in the moral condition (M = 5.87, SD = 1.47) than in the immoral condition (M = 2.36, SD = 1.17), ts(638) > 16.02, ps < .001, ds > 1.26. The order of question presentation showed no main effect and no significant interactions. We therefore collapsed across this variable in all subsequent analyses.

Main Analyses. We conducted a 2 (Target Morality) x 2 (Level of Non-Morality Trait) x 2 (Non-Morality Trait) between-subjects ANOVA, the results of which replicated the findings of Study 1. We again found a main effect of Target Morality, $F(1, 632) = 1397.25$, $p < .001$, $\eta^2_p = .69$; as illustrated in Figure 2, impressions of moral individuals were always neutral-to-positive, while impressions of immoral individuals were always very negative. However, the impressions of sociable, competent, unsociable, and incompetent targets were mixed – impressions of sociable and competent targets were positive only when the target was also moral, but negative otherwise, while impressions of unsociable and incompetent targets were neutral if the target was moral, but negative otherwise. These results support the Morality Dependence Hypothesis.

Insert Figure 2 About Here.

Moreover, as predicted, the critical interaction between Target Morality and Level of Non-Morality Trait was significant, $F(1, 632) = 67.70$, $p < .001$, $\eta^2_p = .070$. This interaction reflects the fact that high sociability or competence contributed positively to impressions of moral individuals (between-subjects ds: 1.03 and 1.63, respectively), but contributed much less to impressions of immoral individuals (ds: .40 and .14). This interaction held in both the

sociability condition, $F(1, 319) = 9.53$, $p = .002$, $\eta^2_p = .029$, and the competence condition, $F(1, 313) = 44.91$, $p < .001$, $\eta^2_p = .13$. This result provides further support for the Morality Dependence Hypothesis. Each of the five scenarios also showed this basic pattern of results when analyzed separately.

Moderated Mediation Analysis. From a functionalist standpoint, morality indicates a person's good or bad intentions, while competence indicates a person's ability to carry out those intentions. Thus, a person's competence should positively predict the perceived likelihood that they will achieve their goals, which in turn, should predict overall impressions of that person. That is, perceived likelihood of success should mediate overall impressions. However, the direction of this mediation – or at least, the size of the indirect effect – should depend on the person's morality. When a person is moral, the perceived likelihood that they will achieve their (praiseworthy) goals should positively predict overall impressions, but when a person is immoral, the perceived likelihood that they will achieve their (blameworthy) goals should less positively predict overall impressions. The strongest version of the Morality Dependence Hypothesis is that for an immoral person, competence would *negatively* predict overall impressions through the likelihood of goal attainment. However, that should only happen in cases where a person's competence serves only to amplify their immorality, and makes no other redeeming contributions to their personhood. Since we were not confident that competence would be welded exclusively to morality in this way, we made a more conservative prediction that competence would *less positively* predict overall impressions for immoral individuals. In sum, a person's morality should moderate the mediated relationship between competence and overall impressions.

In a similar fashion, as we have argued above, sociability provides information about whether a person is likely to be able to recruit allies to help them pursue their goals. The more effectively one can recruit allies, the more likely one is to achieve one's goals in the end. In this sense, sociability functions as a form of social competence, so the same moderated mediation would be expected for sociability as well. Figure 3 models these relationships conceptually. We tested these moderated mediation models using the PROCESS Macro for SPSS (Hayes, 2013), Model 14, with 10,000 bootstrap resamples.

Insert Figure 3 About Here.

Table 1 presents the coefficients for each term in this analysis. The most important result to note is the significant interaction between morality and perceived likelihood of success (in both the sociability and competence conditions), indicating that morality moderates the effect of perceived likelihood of success on overall impressions. Consistent with our theorizing, the indirect effects of sociability and competence on overall impressions through perceived likelihood of success were larger for moral targets than for immoral targets (Sociability Condition: b_{Moral} = .68, $b_{Immoral}$ = .28; Competence Condition: b_{Moral} = 1.79, $b_{Immoral}$ = .77), and these differences across the levels of the moderator (morality) were statistically significant (Sociability Condition: Index of Moderated Mediation: .40, 95% Confidence Interval: [.04, .76]; Competence Condition: 1.02, [.63, 1.40]). In other words, the presence of sociability and competence always had a positive effect on overall impressions, mediated through the perceived likelihood that the target would accomplish their goals, but this relationship was substantially weaker for immoral targets than for moral targets.

Insert Table 1 About Here.

Discussion

Study 3 replicated the results of Study 2 using scenarios describing moral or immoral behaviors rather than abstract trait terms to convey targets' morality. First, moral targets were uniformly viewed neutrally or positively, and immoral targets were uniformly viewed very negatively (supporting the Morality Dominance Hypothesis). Second, impressions of sociable, unsociable, competent, and incompetent targets depended on their morality (supporting the Morality Dependence Hypothesis). Third, how positively sociability and competence traits contributed to overall impressions depended on a target's morality (further supporting the Morality Dependence Hypothesis). A conditional process analysis (Hayes, 2013) of this third result showed that the effects of sociability and competence on overall impressions were mediated through the perceived likelihood that a target would achieve his or her goals, and that this mediated relationship was moderated by the target's morality, such that the indirect effect was substantially and significantly smaller for immoral targets. In other words, sociability and competence increase the perceived likelihood that a person will achieve his or her goals, whether those goals are moral or immoral, but the effect this has on overall impressions depends on the person's moral character.

Study 4

Studies 2 and 3 demonstrated that moral and immoral individuals unconditionally produced positive and negative overall impressions, respectively, whereas impressions of individuals high or low in sociability or competence depended on their morality. These studies also showed that sociability and competence traits contributed more positively to impressions of moral individuals than immoral individuals – the effect of these traits on impressions was moderated by morality. We interpret these patterns of moderation as supporting our Morality Dependence Hypothesis. However, one could also interpret them as supporting a "Competence

Dependence Hypothesis" and a "Sociability Dependence Hypothesis", in that the effects of morality are stronger for competent and sociable targets than for incompetent and unsociable targets (see Figures 1 and 2). Strictly speaking, this interpretation is equally consistent with the interaction reported in these studies. Indeed, consistent with this reading of our results, we did observe significant effects of the manipulations of Level Non-Morality Trait in both Studies 2 and 3, *p*s < .001 (see Online Appendix for details). However, the greater conditionality of sociability and competence is revealed not just by the presence of the main effects and interactions reported above, but also by the overall pattern of observed means. The Morality Dependence hypothesis captures the idea that positive sociability and competence traits are positive in moral targets, but less positive (and sometimes negative) in immoral targets (see Introduction). This idea is clearly supported by the data, as Figures 1 and 2 show – impressions of sociable, unsociable, competent, and incompetent targets were positive or negative largely as a function of the respective targets' morality. In contrast, impressions of moral or immoral targets were not contingent in this way – impressions of moral targets were nearly always positive, and occasionally neutral, and impressions of immoral targets were always very negative – notwithstanding differences in sociability and competence information. As a consequence, we think the data as a whole support the greater conditionality of sociability and competence, as compared with morality.

To bolster this interpretation, we implemented more direct tests of our key hypotheses in the ensuing studies. In Study 4, we elicited participants' preferences for various traits in others, and in Studies 5-6, we elicited judgments of the effect that various traits would have on overall impressions of others. In all three studies, we focused participants' attention more directly on how various traits would shape their impressions of others.

In Study 4, participants were given information either about a person's sociability or competence, or alternatively, about their morality, and then reported whether they would prefer the person to be high or low on traits indexing the trait dimension about which they received no information. We predicted that participants would always prefer others to possess high morality, regardless of those others' sociability or competence (the Morality Dominance Hypothesis), but would only prefer others to possess high sociability or competence when these traits were coupled with high morality (the Morality Dependence Hypothesis). We also predicted that preferences for sociability and competence would be stronger for moral than for immoral others (the Morality Dependence Hypothesis). Finally, we also examined whether sociability and competence might actually be non-preferred in immoral others, rather than simply less preferred – a particularly strong version of the Morality Dependence hypothesis.

Method

Participants. One hundred-thirteen undergraduates (62% female) were recruited through the University of Pennsylvania subject pool to complete an online study for partial course credit.

Procedure. The study was conducted online. After consenting to participate, participants were presented with 128 questions. For each question, a morality trait was paired with either a sociability trait or a competence trait. Thus, one variable in the design was the dimension of the non-morality trait (sociability or competence). A second variable was whether the morality trait was presented first as given information, followed by a question about participants' preference for the non-morality trait; or alternatively, whether the non-morality trait was given, followed by a question about participants' preference for the morality trait. Thus, if the morality trait was given, the question might be (e.g., in the sociability condition): knowing that the person is moral, would you prefer that they be sociable or unsociable? Whereas, if the non-moral trait was given,

the question might be: knowing that the person is sociable, would you prefer that they be moral or immoral? Finally, the third variable in the design was the valence of the given trait, i.e., whether the given trait was high or low on the dimension of interest (e.g., moral or immoral; sociable or unsociable; competent or incompetent). Thus, the study had a 2 (Non-Morality Trait: sociability versus competence) by 2 (Given Trait: morality versus non-morality) by 2 (Level of Given Trait: high versus low) within-subjects design. There were 16 replications in each cell of the design, formed by pairing the same trait terms used in Study 2. For a schematic depiction of this study's design, see Figure A.1 in the Online Appendix.

Responses were made on 1-9 Likert scales, with the endpoints indicating a strong preference that the target possess one trait over the other (e.g., sociable over unsociable), and the midpoint indicating indifference. The order in which the 128 questions were presented was randomized for each participant. We also counterbalanced, between-subjects, whether the high (e.g., sociable) or low (e.g., unsociable) trait terms appeared first or second in the question and, in parallel, on the response scale. After completing the 128 questions, participants responded to a brief demographics questionnaire.

Results

Preliminary Analyses. The replications in each of the eight cells of the design all showed good internal reliability, $\alpha s > .93$, so we averaged across the sixteen questions in each cell to produce one data point per within-subjects condition per participant. Moreover, the between-subjects manipulation of the order in which the high and low trait terms appeared showed no significant main effects or interactions, so we collapsed across this manipulation for all subsequent analyses.

Within-Subjects Analyses. We conducted a 2 (Non-Morality Trait) x 2 (Given Trait) x 2 (Level of Given Trait) repeated-measures ANOVA. As predicted, the critical Given Trait x Level of Given Trait interaction was significant, $F(1, 112) = 138.96$, $p < .001$, $\eta^2_p = .55$. This interaction was also observed in separate 2 x 2 ANOVAs that examined the sociability and competence conditions separately (respectively: $F(1,112) = 102.12$, $p < .001$, $\eta^2_p = .48$; $F(1, 112) = 172.99$, $p < .001$, $\eta^2_p = .54$). As the top panel of Figure 4 shows, in both the sociability and competence conditions, participants preferred that an acquaintance be moral, and there was little difference in preferences whether that person was sociable/competent or unsociable/incompetent, supporting the Morality Dominance Hypothesis. In contrast, as shown in the bottom panel of Figure 4, in the morality conditions there was a much larger difference in people's preferences for the other person to be sociable or competent – sociability and competence were preferred given that the other person was moral, but not when the person was immoral, supporting the Morality Dependence Hypothesis. In fact, if an acquaintance was immoral, participants showed a reversal in their judgments, actually preferring this person to be unsociable and incompetent – thus supporting the strong version of the Morality Dependence Hypothesis articulated above. All condition means depicted in Figure 4 differed significantly from the scale midpoint, $ts(112) > 5.27$, $ps < .001$, one-sample $ds > .49$, which indicates that people were not merely indifferent in the conditions with immoral targets, but actually preferred immoral others to be at least somewhat unsociable and incompetent. This result corroborated by a participant-level analysis – a large proportion of participants showed precisely the pattern of responses predicted by the strong version of the Morality Dependence Hypothesis (see Online Appendix for details).

Insert Figure 4 About Here.

Discussion

The results of this study strongly supported our hypotheses. High-morality traits were always preferred in another person, regardless of their sociability or competence. However, high-sociability traits and high-competence traits were only preferred when the other was known to be moral. When the other was known to be immoral, our participants actually preferred that they lack these traits, at least to a degree.

Study 5

Taken together, Studies 2-4 support our assertion that morality is always seen as positive in others, whereas sociability and competence are positive only in moral others, and not in immoral others. In Study 4, we found that unsociability and incompetence were in fact preferred in immoral others. By contrast, in Studies 2 and 3, we found that sociability and competence still contributed positively to impressions of immoral others, though less positively than for moral others.

One possible reason for this apparent discrepancy is that, in Study 4, the dependent measure asked explicitly about participants' preferences for various traits. In contrast to the overall impression measures used in Studies 2 and 3, this preference measure may have focused participants' attention more concretely on the likely interactive effects of various traits. Accordingly, it may have enabled them to realize that high competence or sociability should amplify the effects of a person's prevailing morality, thereby making a moral person better than they otherwise would be, and an immoral person worse than they otherwise would be. A separate, deflationary possibility is that participants might have desired immoral targets to be unsociable and incompetent in Study 4 because this is what they thought such immoral individuals deserved, rather than because these trait dimensions increase a person's likelihood of

goal attainment (our postulated mechanism). For instance, people may prefer that an immoral target be incompetent, not because this decreases the likelihood that the target will successfully harm them, but because it would be unjust for such a person to reap the benefits of competence.

In Study 5, we aimed to rule out this alternative, justice-based explanation. We provided a general, characterological description of a target person, and asked participants how their impression of the target would change if they knew that he was sociable or unsociable, competent or incompetent. In this way, we kept participants' focus on their impressions, rather than their direct preferences (which might reasonably incorporate these sorts of justice concerns). We also varied the degree of immorality of the target individual, reasoning that support for our predictions may be especially evident when the target is thoroughly immoral rather than more mildly immoral. A thoroughly immoral target will have more immoral goals, in terms of both number and extremity, than a mildly immoral target. Accordingly, because both competence and sociability increase a person's likelihood of attaining their goals, these traits should be especially negative in thoroughly immoral targets.

Method

Participants. Two hundred-thirty participants were recruited through Amazon Mechanical Turk. Two failed a "Captcha" verification, and three more did not complete the survey, leaving a final sample of $N = 225$ (37% female).

Procedure. Participants were randomly assigned to one of three conditions: Very Immoral, Slightly Immoral, or Moral. Each condition provided a description of a target person ("Mike") consisting of three loosely related pieces of information about his moral character, which varied by condition (see Methodological Supplement for full scenarios).

Participants first indicated how positive or negative their overall impression of Mike was, on a 0-100 sliding scale, as a manipulation check. They next indicated how much more positive or negative their impression would be if they knew that Mike possessed each of 18 different traits. These traits were organized in a 2 (Trait Dimension: sociability versus competence) by 2 (Trait Valence: positive versus negative traits) design with three replications in each cell. The traits were: *sociable, extroverted, friendly* (positive sociability), *unsociable, introverted, unfriendly* (negative sociability), *competent, skillful, intelligent* (positive competence), and *incompetent, unskillful, unintelligent* (negative competence). We also included six filler traits, *emotional, adventurous, artistic, unemotional, unadventurous,* and *non-artistic,* to obscure the aspects of personality that we were most interested in, for a total of 18 traits. The filler traits were not included in our analyses. The order of presentation of the traits was randomized for each participant, and responses were made on 1-9 Likert scales ranging from "much more negative" to "much more positive." After responding to all 18 traits, participants answered a brief demographic questionnaire.

We predicted that there would be an interaction between Target Morality and Trait Dimension, such that the positive effect of sociability and competence traits on overall impressions (relative to unsociability and incompetence traits) would increase as the targets became more moral (supporting the Morality Dependence Hypothesis).

Results

Preliminary analyses. The manipulation of morality was successful – initial impressions of the very immoral target were extremely negative (M = 10.05, on a 0-100 scale, SD = 17.55), impressions of the slightly immoral target were somewhat negative (M = 39.05, SD = 21.12), and

impressions of the moral target were quite positive (M = 86.74, SD = 18.73). All pairwise comparisons were significant, ts > 9.12, ps < .001, ds > 1.49.

Responses to the three traits indexing positive sociability, negative sociability, positive competence, and negative competence were averaged together to form composite measures of predicted impression change (αs .80, .58, .88, and .81, respectively). The pattern of means is essentially identical when responses to individual traits, rather than the composite scales, are compared.

Main analyses. Participants' responses to the questions about how their impression of Mike would change if he exhibited a specified trait were analyzed using a 3 (Target Morality: Very Immoral versus Slightly Immoral versus Moral) x 2 (Trait Dimension: Sociability versus Competence) x 2 (Trait Valence: Positive versus Negative) mixed-measures ANOVA with repeated measures on the last two factors. The critical interaction between Target Morality and Trait Valence was observed, thus supporting the Morality Dependence Hypothesis, $F(2, 222) = 20.56$, $p < .001$, $\eta^2_p = .16$. Moreover, this critical interaction was observed for both sociability traits, $F(2, 222) = 19.05$, $p < .001$, $\eta^2_p = .15$, and competence traits, $F(2, 222) = 15.30$, $p < .001$, $\eta^2_p = .12$. Results are graphed in Figure 5. As predicted, the effects of sociability and competence depended on the target's morality. For the very immoral target, positive sociability and competence traits were anticipated to make impressions more negative, similar to the preference ratings in Study 3; in contrast, for the slightly immoral target, positive sociability and competence traits were anticipated to make impressions slightly more positive, similar to the impression ratings in Studies 1 and 2; and, for the moral target, positive sociability and competence traits were anticipated to make impressions moderately more positive. Unsociability and incompetence traits were consistently anticipated to make impressions of a target more

negative. Thus, the difference between sociable and unsociable traits, and between competent and incompetent traits grew larger as the target become more moral, as we had predicted. All means differed significantly from zero (i.e., no change in impression), ts > 2.90, ps < .006, ds > .33, except positive competence in the Very Immoral condition, the negative effect of which was marginally significant, $t(75) = 1.82$, $p = .073$, $d = .21$.

Insert Figure 5 About Here.

Discussion

Sociability and competence are not only more positive in moral people than immoral people, but they are also anticipated to have a negative effect on impressions of thoroughly immoral individuals. In the case of slightly immoral people, sociability and competence are still anticipated to exert positive effects on impressions, though not as strongly as they are for moral people. These results further underscore the contingency of these two dimensions of person perception – specifically, they support the Morality Dependence Hypothesis.

The effects of positive sociability and competence did not mirror those of unsociability or incompetence. Instead, learning that a target was unsociable or incompetent was consistently anticipated to have a negative effect on impressions. This means that, paradoxically, impressions of very immoral targets were anticipated to become less positive, regardless of whether newly-learned trait information was positive or negative. Nonetheless, our focus here is on the ways in which positive sociability and competence interact with morality in impression formation, so we leave in-depth exploration of this result to future research.

Study 6

Studies 2-5 have shown that the positivity of sociability and competence depends on a target's morality. However, it remains plausible that this is true of *any* positively valenced trait,

such that impressions of immoral people are depressed by the addition of any new positive trait information. We think this is unlikely for the reasons outlined above, namely that there are functional reasons why sociability and competence are contingent upon morality, and why morality is always seen as positive. Yet, we have not yet ruled out the possibility that adding any new positive trait information to an immoral person would contribute negatively to impressions of that person. An analysis of the filler traits from Study 5 can shed some light on this: if *emotional, adventurous,* and *artistic* are treated as positive traits, and *unemotional, unadventurous, and unartistic* are treated as negative, an interaction between Target Morality and Trait Valence analogous to the interaction reported in Study 5 does emerge, $F(2, 222) = 7.26$, $p = .001$, $\eta^2_p = .06$, although the effect size is noticeably smaller than for the sociability and competence traits ($\eta^2_p = .15$ and .12, respectively). This suggests that sociability and competence may be *more* contingent on morality than other kinds of traits. Yet, it is not clear that traits like *emotional* and *adventurous* are as positively valenced as traits like *friendly* and *competent*, making this result impossible to interpret definitively.

However, there is one class of traits that is highly valenced, and that our theoretical model predicts should not contribute negatively to impressions under typical circumstances - morality traits. Accordingly, in this final study, we set out to demonstrate that, unlike sociability and competence traits, morality traits always contribute positively to impressions of others, even those who are very immoral. To do this, we examined two fundamental – yet separable – aspects of a person's moral character: honesty and compassion (see Landy & Uhlmann, 2016). Specifically, we manipulated which aspect of a target's morality participants received information about, and then provided additional information, either about the other aspect of his morality, or about his sociability or competence. We also manipulated whether the target was

moral (i.e., honest or compassionate, depending on condition), or immoral (dishonest or uncompassionate). We expected sociability and competence to have a positive impact on impressions of moral targets, but not immoral targets (and possibly a negative impact). In contrast, we expected morality to have a positive impact on impressions regardless of the morality of the target, given that morality should always be positive in others. This is an especially stringent test of the Morality Dominance Hypothesis – if any positively valenced trait becomes neutral or even negative in the presence of negative morality, then honesty and compassion should show this effect as well. If, however, morality is (nearly) always positive, then positive information about honesty or compassion, even in a thoroughly immoral person, should still improve impressions of them.

Method

Participants. Two hundred fifty-six participants were recruited online through Amazon Mechanical Turk. Three failed a "Captcha" verification, and three failed to complete the study, leaving a final sample of $N = 250$ (38% female).

Procedure. After consenting to participate, participants were randomly assigned to receive one of four characterological profiles of a target person, similar to those used in Study 5. Between-subjects, the profiles varied in terms of the Moral Domain being described (Honesty versus Compassion) and the initial morality of the target person (Target Morality: Moral versus Immoral). In the honesty condition, the target person was described either as extremely honest (Moral) or extremely dishonest (Immoral), whereas in the compassion condition, the person was described either as extremely compassionate or extremely uncompassionate. As in Study 5, after reading the initial profile, participants rated their overall impression of the target person, and then indicated whether their impression of the person would become more positive or negative if

the person possessed each of 18 traits. Within this set of 18 traits, two additional variables, Trait Dimension (Sociability versus Competence versus Morality), and Trait Valence (Positive versus Negative), were varied within-subjects. The positive and negative sociability and competence traits from Study 5 were used again here, but instead of filler traits, we included three positive and three negative morality traits. Thus, in the honesty condition, the three rated traits related to the dimension of morality not initially described, namely compassion (Positive: *compassionate, kind, caring;* Negative: *uncompassionate, unkind, uncaring*), and in the compassion condition, the three traits related to honesty (Positive: *honest, trustworthy, sincere;* Negative: *dishonest, untrustworthy, insincere*). The order in which the 18 trait terms were presented was randomized for each participant. After responding to all 18 traits, participants responded to a brief demographics questionnaire.

The key predictions were as follows. For moral targets, the addition of positive trait information (especially morality information) should enhance anticipated impressions across all three trait dimensions (sociability, competence, and morality), and the addition of negative trait information (again, especially morality information) should depress anticipated impressions across all three trait dimensions. In contrast, for immoral targets, only the addition of positive morality information (and not sociability or competence information) should enhance anticipated impressions, whereas (based on previous results) the addition of negative trait information (especially morality trait information) should depress anticipated impressions across all three dimensions. This therefore amounts to the prediction of a three-way interaction – the two-way interaction between Trait Dimension and Trait Valence should be qualitatively different for moral targets as compared with immoral targets.

Results

Preliminary analyses. As in Study 5, we averaged responses to the three traits indexing positive sociability, negative sociability, positive competence, negative competence, positive morality, and negative morality (αs .76, .55, .89, .84, .89, and .82, respectively). The reported pattern of means is essentially identical when responses to individual traits, rather than the composite scales, are compared.

The manipulation of target morality was successful; the moral target elicited much more positive initial impressions (M = 90.86 on a 0-100 scale, SD = 10.42) than the immoral target (M = 11.29, SD = 17.19).

The between-subjects manipulation of Moral Domain (honesty versus compassion) showed no main effect or interactions, so we collapsed across this variable in all subsequent analyses, so as simply to compare Moral and Immoral Targets.

Main analyses. Participants' responses were analyzed using a 2 (Target Morality: Moral versus Immoral) by 3 (Trait Dimension: Sociability versus Competence versus Morality) by 2 (Trait Valence: Positive versus Negative) mixed-measures ANOVA with repeated measures on the last two factors. The predicted three-way interaction was found, $F(2, 496) = 15.88$, $p < .001$, $\eta^2_p = .06$, indicating that the two-way interaction between Trait Dimension and Trait Valence differed by Target Morality. This interaction was in fact larger for moral targets ($F(2, 254) = 139.02$, $p < .001$, $\eta^2_p = .53$) than for immoral targets ($F(2, 242) = 44.06$, $p < .001$, $\eta^2_p = .27$). Figure 6 presents the shape of these interactions. For moral targets, the overall difference between positive and negative traits was quite substantial for all three trait dimensions. This difference was especially pronounced for moral traits, as predicted, and it was smallest for sociability traits, and middling for competence traits. In contrast, for immoral targets, the difference between positive and negative traits was substantial for moral traits, but it was

considerably smaller for both competence and sociability traits. Indeed, consistent with our predictions, participants anticipated that only positive morality information would have a positive effect on their impressions of the immoral target, $t(122) = 4.67$, $p < .001$, $d = .44$, but they did not anticipate that positive sociability or competence information would have any effect on their (very negative) impressions of the immoral target, ts$(122) < 1.48$, ps $> .14$, ds $< .14$.

Insert Figure 6 About Here.

Discussion

Study 6 showed that positive information about one aspect of morality is consistently anticipated to improve impressions, even of people who are thoroughly lacking in other aspects of morality. Specifically, participants anticipated that positive information about an uncompassionate person's honesty, or about a dishonest person's compassion, would improve their impressions of those individuals. This result further supports the Morality Dominance Hypothesis, as it highlights the power that moral traits exert on the impressions we form of others, even when an initial impression is quite negative. However, positive information about an immoral target's sociability or competence had no such effect, supporting the Morality Dependence Hypothesis. This indicates that the contingent effects of sociability and competence found throughout this research do not generalize to *any* sort of trait – morality, at least, is typically unconditionally positive, but sociability and competence are not.

General Discussion

In this research, we provide evidence for a functionalist account of the role of morality, sociability, and competence in person perception. Studies 1a and 1b used exploratory factor analyses of trait ratings of real people to provide new empirical evidence that morality and sociability are separate dimensions of person perception (supporting the Morality Differentiation

Hypothesis). Five subsequent studies showed that whereas the positivity of morality in others is unconditional (the Morality Dominance Hypothesis), the positivity of sociability and competence in others depends on their morality (the Morality Dependence Hypothesis). In Study 2, we assessed overall impressions of others with various qualities and found that moral people were always evaluated positively regardless of their other traits, whereas sociable and competent people were only evaluated positively when they were also moral. Moreover, the positive effects of sociability and competence traits on impressions were considerably greater for moral targets than for immoral targets. Study 3 replicated these findings using descriptions of behavior rather than trait terms to convey a target's morality. Study 3 also showed that the effect of sociability and competence on overall impressions is mediated through the perceived likelihood that the target will achieve their goals, and that this relationship is moderated by the target's morality. In Study 4, we found that participants always explicitly preferred that another person be moral, regardless of their other characteristics. In contrast, they preferred only moral people to be sociable and competent, whereas they preferred immoral people to be at least somewhat unsociable and incompetent. In Study 5, we found that people anticipate that positive sociability and competence can have a negative effect (as opposed to an attenuated positive effect) on their impressions of thoroughly immoral individuals, and ruled out an alternative explanation for the results of Study 4. In Study 6, we found more direct support for the claim that moral traits are unconditionally positive by showing that positive morality information is anticipated to improve impressions of even thoroughly immoral people, whereas positive sociability and competence information is not.

These results therefore offer further support for the claim that morality and sociability are separate dimensions of person perception (e.g., Brambilla et al., 2011; Brambilla et al., 2012;

Goodwin et al., 2014; Leach et al., 2007), rather than highly related subcomponents of one superordinate prosocial dimension (e.g., Fiske et al., 2007). If both morality and sociability were components of the same dimension of judgment, one would expect them to be processed in similar ways when forming impressions of others. Yet our results showed that judgments of others' morality and sociability do not cohere, and, furthermore, that they interact, such that morality is always seen as positive, but sociability is only positive in the presence of morality. This shows a striking divergence in how people use information about these qualities in others when forming impressions, and strongly suggests that they are not part of the same dimension of social cognitive judgments.

Our findings are consistent with the functionalist account outlined in the Introduction. Insofar as morality informs us about other people's likely intentions toward us, we should generally prefer that other people be moral, regardless of their other qualities: if another person has good intentions toward us, rather than bad ones, this will generally produce better outcomes for us. However, both competence and sociability inform us about the likelihood that a person will fulfill those intentions, albeit in different ways – competence informs us about a person's ability to achieve their goals by their own actions, while sociability informs us about their likelihood of recruiting others to help them. Given that we prefer that people with immoral intentions not be able to fulfill them, we should consider sociability and competence to be less positive in such people, and perhaps sometimes truly negative. The present results support this overall picture.

Our findings align with the findings of Wojciszke, Bazinska et al. (1998, Study 4) and with prior theorizing regarding competence. They extend upon this work by showing that sociability, which has often been treated as part of the same superordinate dimension of

judgment as morality, is actually thought about quite differently from morality. Sociability is, in fact, only contingently positive – its positivity depends on a person's morality.

While the present research suggests that moral traits are unconditionally positive, an interesting task for future research would be to explore whether there are some situations in which moral traits are not desired in others. Existing research shows that "moral rebels" are sometimes disliked when their behavior has negative implications for observers' own morality (Minson & Monin, 2012; Monin, Sawyer & Marquez, 2008), but this research has not yet been extended to the broad trait level – whether people might ever actually dislike moral traits in others has not yet been established. It might be the *claim* to morality that people take offense to in moral rebels, not their actual possession of moral traits. Another possibility is that the perceivers' own self-assessed morality may moderate the value of morality in others. Perceivers who consider themselves moral may especially value morality in others, since this would mean that the two parties' fundamental goals align. But, perceivers who consider themselves immoral may not always prefer morality in others, particularly if it means that their goals do not align with those of the target. Thus, the unconditional positivity of morality may itself be conditional on the perceiver viewing themselves as moral (or at least not immoral). Exploring this possibility would be an interesting task for future research.

More broadly, these results speak to the complexity of human social judgment. It is clear that additive linear models cannot fully capture the subtle contingencies in how we think about others in our social worlds. Instead, trait dimensions interact in predictable ways to produce overall impressions. Our results also provide further evidence that one trait dimension in particular is primary: when it comes to person perception, and perhaps most of social cognition,

morality information is dominant, and plays a large role in coloring how we interpret everything else.

References

Abele, A. E. & Brack, S. (2013). Preference for other persons' traits is dependent on the kind of social relationship. *Social Psychology, 44,* 84-94.

Abele, A. E., Cuddy, A. J. C., Judd, C. M., & Yzerbyt, V. Y. (2008). Fundamental dimensions of social judgment. *European Journal of Social Psychology, 38,* 1063-1065.

Abele, A. E., & Wojciszke, B. (2007). Agency and communion from the perspective of self versus others. *Journal of Personality and Social Psychology, 93,* 751-763.

Abele, A.E., & Wojciszke, B. (2013). The Big Two in social judgment and behavior. *Social Psychology, 44,* 61-62.

Abele, A. E. & Wojciszke, B. (2014). Communal and agentic content in social cognition: A dual perspective model. *Advances in Experimental Social Psychology, 50,* 195-255.

Ashton, M. C. & Lee, K. (2001). A theoretical basis for the major dimensions of personality. *European Journal of Personality, 15,* 327-353.

Ashton, M. C., Lee, K., & Paunonen, S. V. (2002). What is the central feature of extraversion? Social attention versus reward sensitivity. *Journal of Personality and Social Psychology, 83,* 245-252.

Bauman, C. W., & Skitka, L. J. (2012). Corporate social responsibility as a source of employee satisfaction. *Research in Organizational Behavior, 32,* 63-86.

Brambilla, M. & Leach, C. W. (2014). On the importance of being moral: The distinctive role of morality in social judgment. *Social Cognition, 32,* 397-408.

Brambilla, M., Rusconi, P., Sacchi, S., & Cherubini, P. (2011). Looking for honesty: The primary role of morality (vs. sociability and competence) in information gathering. *European Journal of Social Psychology, 41,* 135-143.

Brambilla, M., Sacchi, S., Rusconi, P., Cherubini, P., & Yzerbyt, V. Y. (2012). You want to give a good impression? Be honest! Moral traits dominate group impression formation. *British Journal of Social Psychology, 51,* 149-166.

Brycz, H. & Wojciszke, B. (1992). Personality impressions on ability and morality dimensions. *Polish Psychological Bulletin, 23,* 223-236.

Buja, A. & Eyuboglu, N. (1992). Remarks on parallel analysis. *Multivariate behavioral research, 27,* 509-540.

Cottrell, C. A., Neuberg, S. L., & Li, N. P. (2007). What do people desire in others? A sociofunctional perspective on the importance of different valued characteristics. *Journal of Personality and Social Psychology, 92,* 208-231.

Cuddy, A. J., Fiske, S. T., & Glick, P. (2008). Warmth and competence as universal dimensions of social perception: The stereotype content model and the BIAS map. *Advances in Experimental Social Psychology, 40,* 61-149.

Curry, O., & Dunbar, R. I. M. (2011). Altruism in networks: The effects of connections. *Biology Letters, 7,* 651-653.

DeScioli, P., & Kurzban, R. (2009). The alliance hypothesis for human friendship. *PLoS ONE, 4,* e5802.

Fabrigar, L. R. & Wegener, D. T. (2012). *Exploratory Factor Analysis.* Oxford: Oxford University Press.

Fabrigar, L. R., Wegener, D. T., MacCallum, R. C., & Strahan, E. J. (1999). Evaluating the use of exploratory factor analysis in psychological research. *Psychological Methods, 4,* 272-299.

Fessler, D. M. T., & Holbrook, C. (2013). Friends shrink foes: The presence of comrades decreases the envisioned physical formidability of an opponent. *Psychological Science, 24,* 797-802.

Field, A. (2005). *Discovering statistics using SPSS.* London: Sage.

Fiske, S. T. (2012). Warmth and competence: Stereotype content issues for clinicians and researchers. *Canadian Psychology, 53,* 14-20.

Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences, 11,* 77-83.

Goodwin, G. P., Piazza, J., & Rozin, P. (2014). Moral character information predominates in person perception and evaluation. *Journal of Personality and Social Psychology, 106,* 148-168.

Hayes, A. F. (2013). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach.* New York, NY: The Guilford Press.

Judd, C., James-Hawkins, L., Yzerbyt, V., & Kashima, Y. (2005). Fundamental dimensions of social judgment: Understanding the relations between judgments of competence and warmth. *Journal of Personality and Social Psychology, 89,* 899-913.

Kelman, H. C. (1958). Compliance, identification, and internalization: Three processes of attitude change. *Conflict Resolution, 2,* 51-60.

Kelman, H. C. (1961). Processes of opinion change. *Public Opinion Quarterly, 25,* 57-78.

Landy, J. F. & Uhlmann, E. L. (2016). *Morality is personal*. Invited submission to K. Gray & J. Graham (Eds.), *The atlas of moral psychology: Mapping good and evil in the mind.* New York: Guilford.

Leach, C. W., Ellemers, N., & Barreto, M. (2007). Group virtue: The importance of morality (vs. competence and sociability) in the positive evaluation of in-groups. *Journal of Personality and Social Psychology, 93,* 234-249.

Longman, R. S., Cota, A. A., Holden, R. R., & Fekken, G. C. (1989). A regression equation for the parallel analysis criterion in principal components analysis: Mean and 95th percentile eigenvalues. *Multivariate Behavioral Research, 24,* 59-69.

Minson, J. A., & Monin, B. (2012). Do-gooder derogation: Disparaging morally motivated minorities to defuse anticipated reproach. *Social Psychological and Personality Science, 3*, 200-207.

Monin, B., Sawyer, P. J., & Marquez, M. J. (2008). The rejection of moral rebels: Resenting those who do the right thing. *Journal of Personality and Social Psychology*, *95*, 76-93.

O'Connor, B. P. (2000). SPSS and SAS programs for determining the number of components using parallel analysis and Velicer's MAP test. *Behavior Research Methods, Instruments, & Computers, 32,* 396-402.

Peeters, G. (1992). Evaluative meanings of adjectives in vitro and in context: Some theoretical implications and practical consequences of positive negative asymmetry and behavioural-adaptive concepts of evaluation. *Psychologia Belgica, 32,* 211-231.

Piazza, J., Goodwin, G. P., Rozin, P., & Royzman, E. B. (2014). When a virtue is not a virtue: Conditional virtues in moral evaluation. *Social Cognition, 32,* 528-558.

Pizarro, D. A. & Tannenbaum, D. (2011). Bringing character back: How the motivation to evaluate character influences judgments of moral blame. In M. Mikulincer & P. R. Shaver (Eds.), *The social psychology of morality: Exploring the causes of good and evil* (pp. 91-108). Washington, DC: American Psychological Assosciation.

Pollet, T. V., Roberts, S. G. B., & Dunbar, R. I. M. (2011). Extraverts have larger social network layers but do not feel emotionally closer to individuals at any layer. *Journal of Individual Differences, 32,* 161-169.

Reeder, G. D., Messick, D. M., & Van Avermaet, E. (1977). Dimensional asymmetry in attributional inference. *Journal of Experimental Social Psychology, 13,* 46-57.

Rule, N. O., Ambady, N., Adams, Jr., R. B., Ozono, H., Nakashima, S., Yoshikawa, S., & Watabe, M. (2010). Polling the face: Prediction and consensus across cultures. *Journal of Personality and Social Psychology, 98,* 1-15.

Wojciszke, B., Abele, A. E., & Baryla, W. (2009). Two dimensions of interpersonal attitudes: Liking depends on communion, respect depends on agency. *European Journal of Social Psychology, 39,* 973-990.

Wojciszke, B., Bazinska, R., & Jaworski, M. (1998). On the dominance of moral categories in impression formation. *Personality and Social Psychology Bulletin, 24,* 1251-1263.

Wojciszke, B., Dowhyluk, M., & Jaworski, M. (1998). Moral competence-related traits: How do they differ? *Polish Psychological Bulletin, 29,* 283-294.

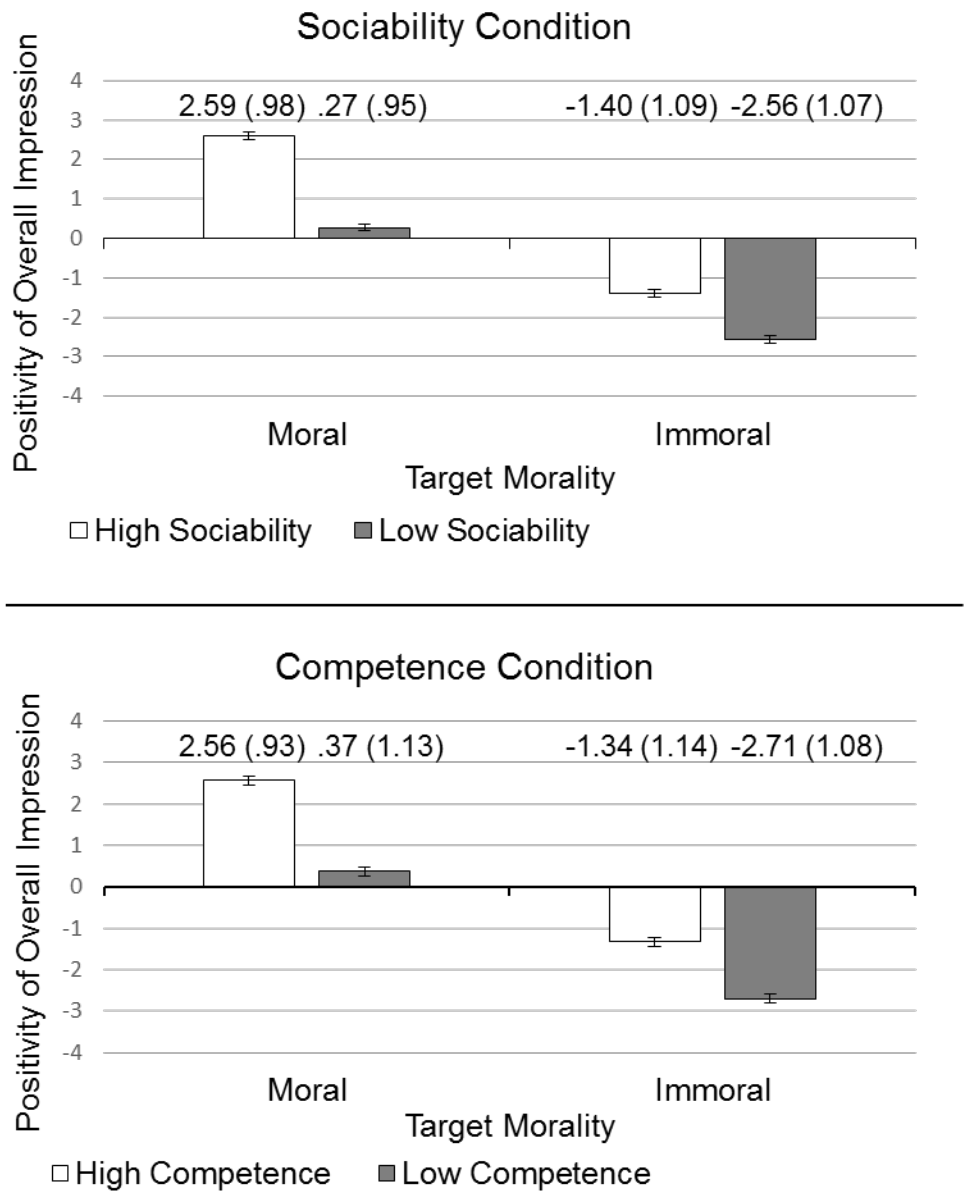

Figure 1. Positivity of overall impressions in Study 2. Error bars ± 1 S.E. Condition means and standard deviations are presented above each bar.

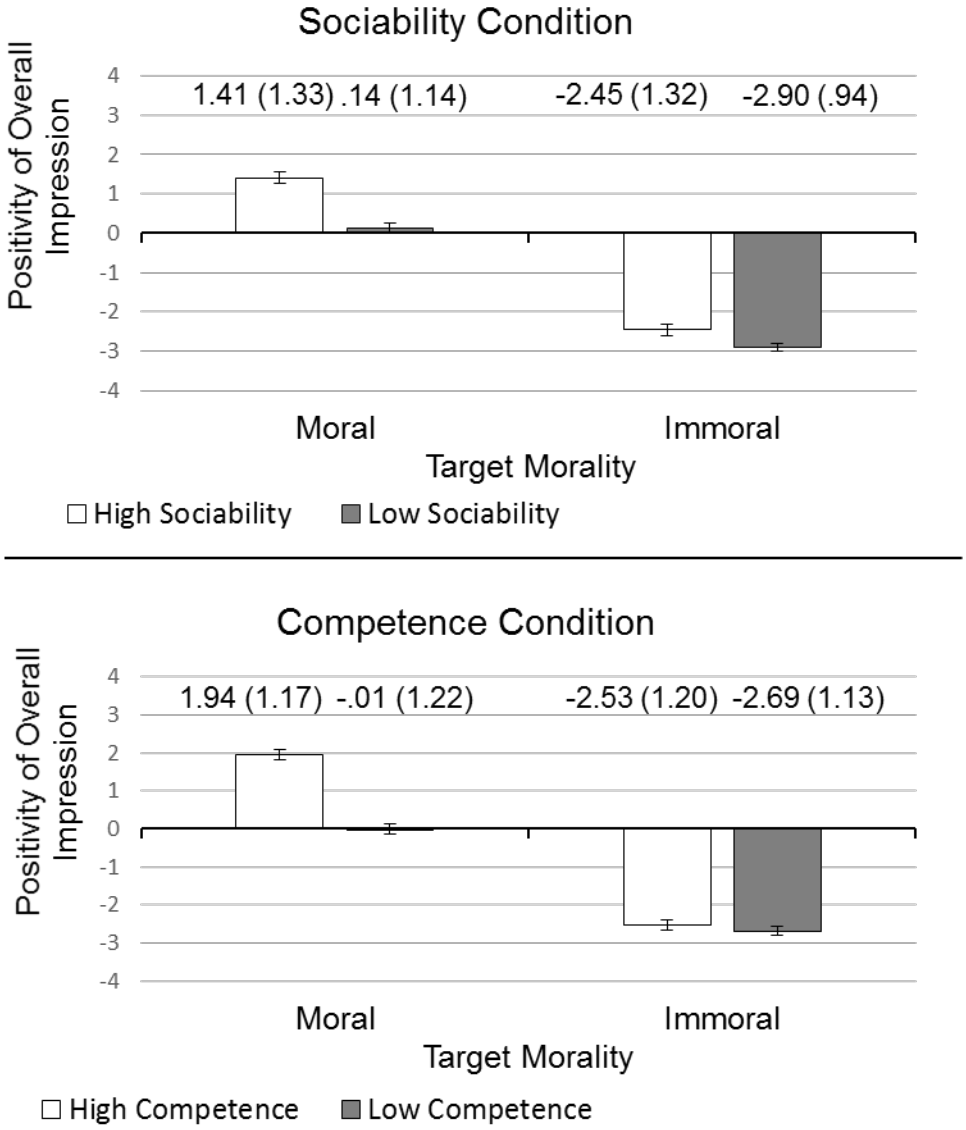

Figure 2. Positivity of overall impressions in Study 3. Error bars ± 1 S.E. Condition means and standard deviations are presented above each bar.

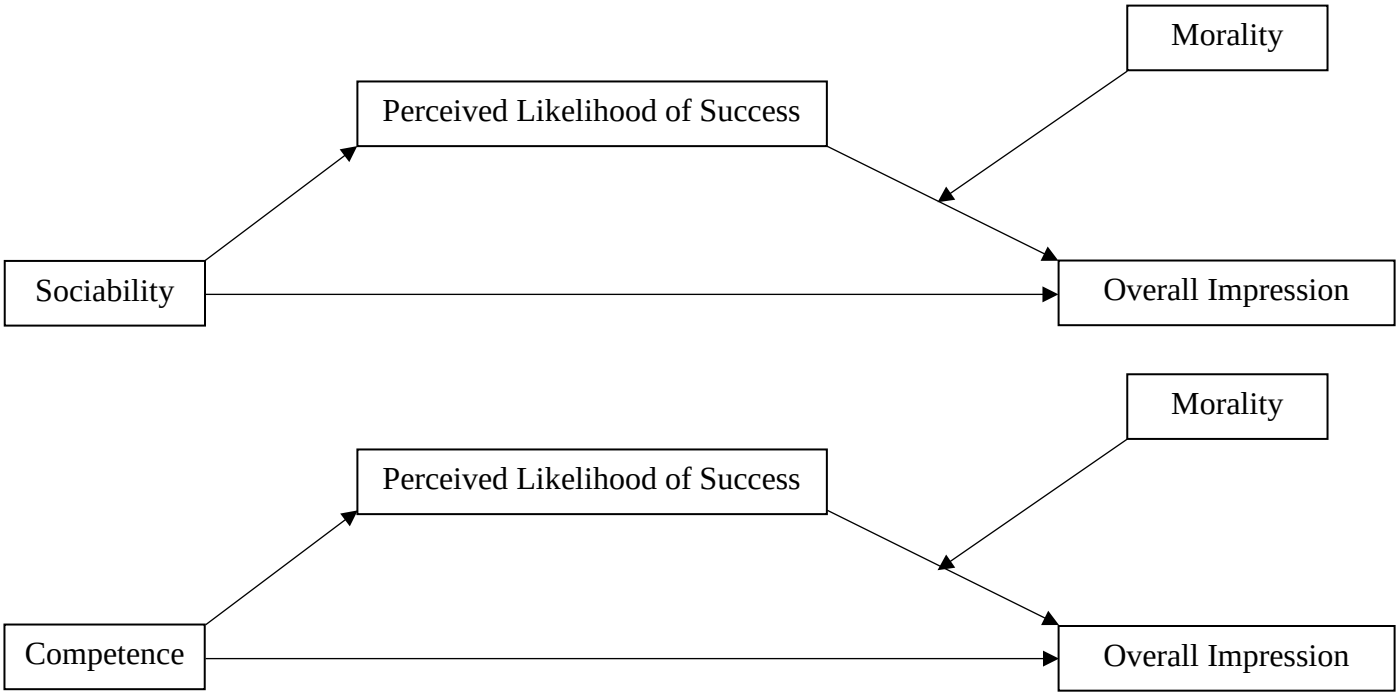

Figure 3. Moderated mediation models for the sociability and competence conditions in Study 3. The target's sociability or competence predicts the perceived likelihood that they will achieve their goals, which in turn predicts one's overall impression of the target. This latter relationship is moderated by the target's morality.

Table 1. Coefficients of moderated mediation models in Study 3, with upper and lower limits of bootstrapped 95% confidence intervals.

Sociability Condition

Predictor Variable	Outcome Variable	Coefficient	Lower Limit	Upper Limit
Level of sociability	Likelihood of success	1.99*	1.72	2.27
Level of sociability	Overall impression	.37*	.03	.72
Likelihood of success	Overall impression	.14	-.004	.28
Target morality	Overall impression	3.26*	2.96	3.56
Morality x Likelihood	Overall impression	.20*	.03	.37

Competence Condition

Predictor Variable	Outcome Variable	Coefficient	Lower Limit	Upper Limit
Level of competence	Likelihood of success	2.92*	2.66	3.18
Level of competence	Overall impression	-.26	-.65	.14
Likelihood of success	Overall impression	.26*	.13	.39
Target morality	Overall impression	3.47*	3.22	3.72
Morality x Likelihood	Overall impression	.35*	.22	.48

Note: Exact p-values were not computed in the bootstrap analysis; asterisks indicate coefficients for which the 95% confidence interval does not contain zero.

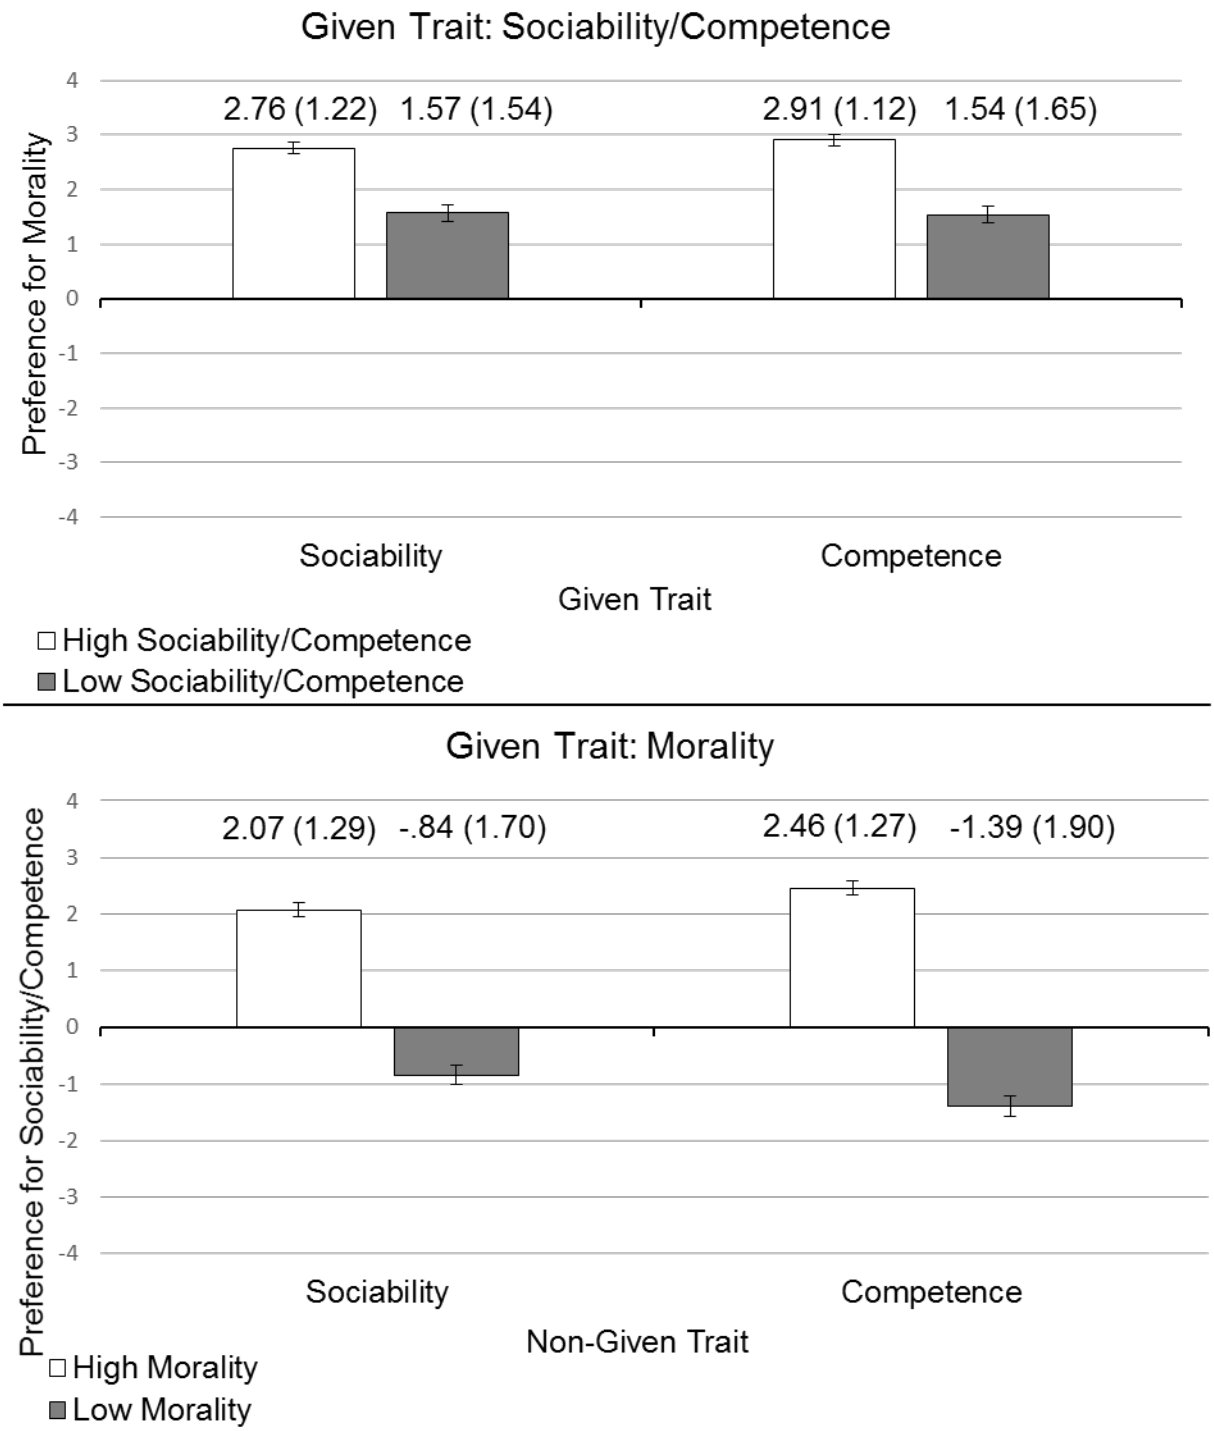

Figure 4. Mean preference ratings for non-provided traits in Study 4. Error bars ± 1 S.E. Condition means and standard deviations are presented above each bar.

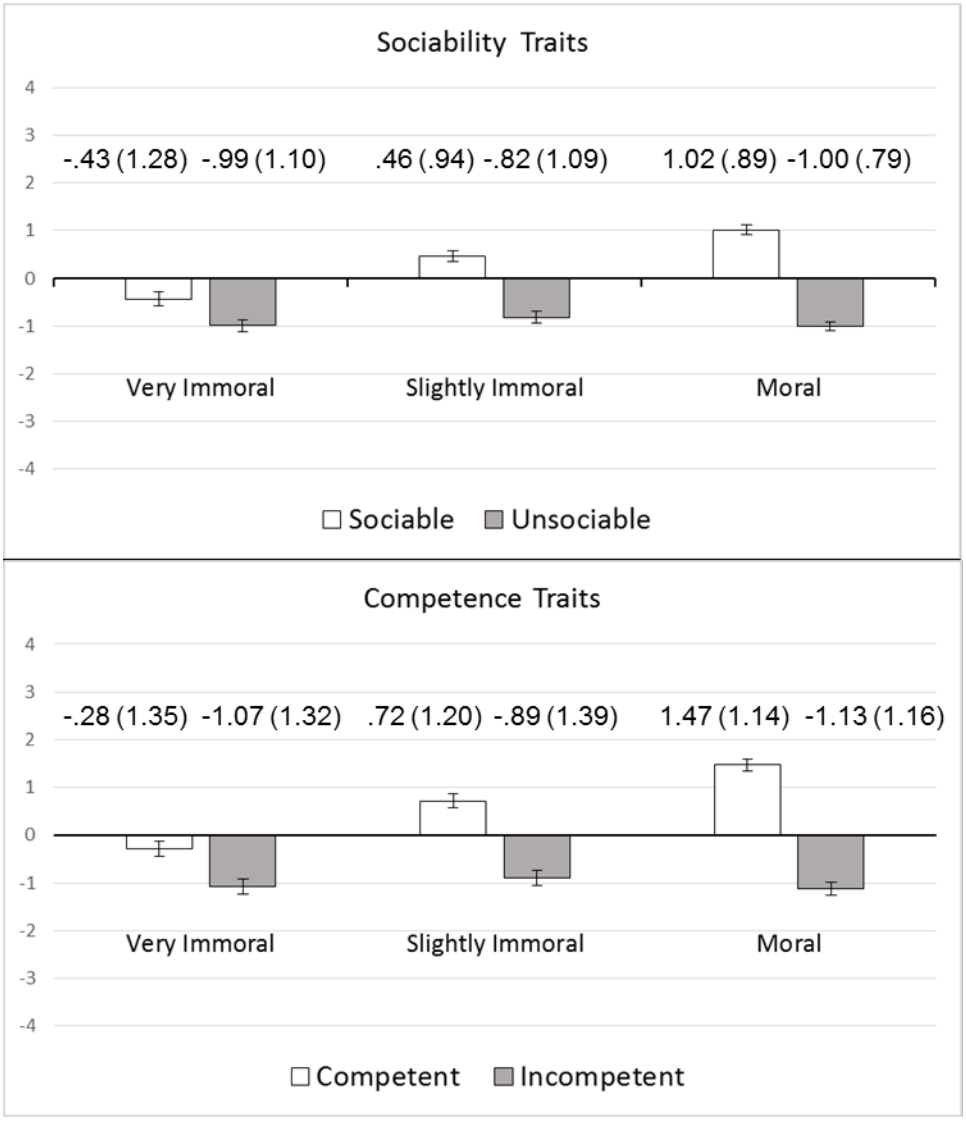

Figure 5. Change in overall impressions in Study 5, by target morality and type of trait information. Condition means and standard deviations are presented above each bar.

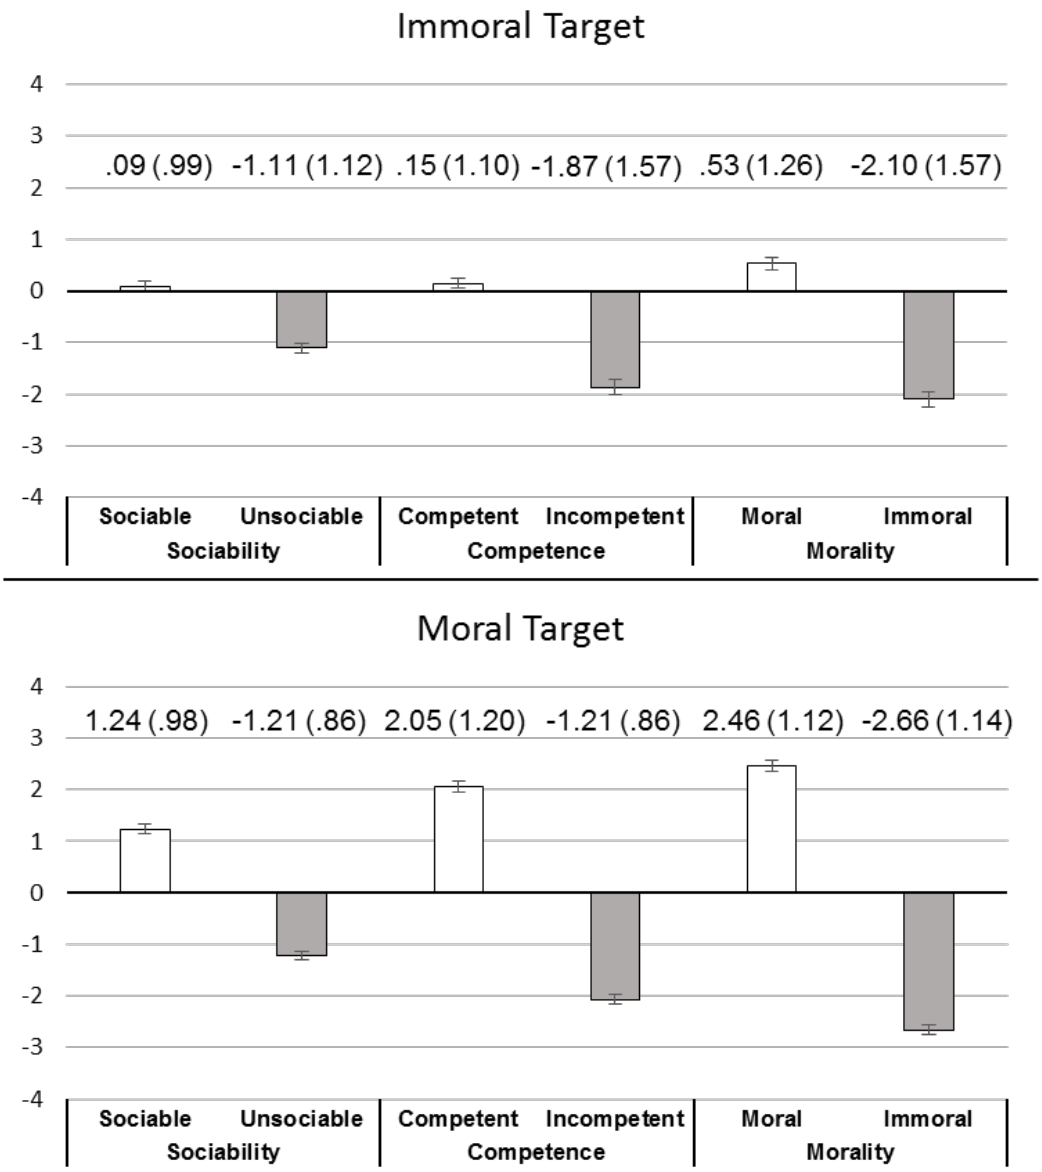

Figure 6. Change in overall impressions in Study 6, by target morality and type of trait information. Condition means and standard deviations are presented above each bar.