

Agreement on the perception of moral character

Erik G. Helzer^a

R. Michael Furr^a

Ashley Hawkins^b

Maxwell Barranti^a

Laura E. R. Blackie^a

William Fleeson^a

^a Psychology Department, Wake Forest University

^b Psychology Department, Duke University

Accepted for publication in *Personality and Social Psychology Bulletin*. Published by SAGE Publications. This article may not exactly replicate the final version published in the SAGE Publications journal. A link to the published article: <http://dx.doi.org/10.1177/014616721455495>

This publication was made possible through the support of grants from the John Templeton Foundation and the Templeton World Charity Organization. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of JTF or TWCO.

Correspondence concerning this article should be addressed to: Erik Helzer, Department of Psychology, Wake Forest University, PO Box 7778, Winston-Salem, NC 27109, email:

helzereg@wfu.edu

WORD COUNT: 9,900

Abstract

This study tested for inter-judge agreement on moral character. A sample of students and community members rated their own moral character using a measure that tapped six moral character traits. Friends, family members, and/or acquaintances rated these targets on the same traits. Self/other and inter-informant agreement was found at the trait level for both a general character factor and for residual variance explained by individual moral character traits, as well as the individual level (judges agreed on targets' "moral character profiles"). Observed inter-judge agreement constitutes evidence for the existence of moral character, and raises questions about the nature of moral character traits.

KEYWORDS: Morality, agreement, personality, character

Agreement on the perception of moral character

People care deeply about moral character. Assessments of character, such as whether a person is *compassionate* or *fair*, are central to the impressions people hold of one another (Fiske, Cuddy, & Glick, 2007; Goodwin, Piazza, & Rozin, 2014; Pizarro & Tannenbaum, 2011; Strohminger & Nichols, 2014; Wojciszke, Bazinska, & Jaworski, 1998). These assessments guide decisions about whom to affiliate with (Goodwin et al., 2014), and whom to trust (van't Wout & Sanfey, 2008). It is, therefore, surprising that so little is currently known about the accuracy of character judgments. Are evaluations of moral character grounded in the real, enduring dispositions of the person, or do they simply exist “inside the head” of judges? The answer to this question is essential to answering fundamental questions about the nature of morality, but to date research in moral psychology provides little instruction.

In this paper, we use *moral character* to refer to a global disposition to engage in moral behavior. We conceptualize moral character as a summation of a person's standing on a number of *moral character traits*, such as honesty, compassion, fairness, etc. We then focus the paper on assessing the extent to which different judges agree in their perceptions of moral character. Does a person's view of her own moral character strengths and weaknesses correspond with how she's seen by others? And do her acquaintances, friends, and family members agree with one another on the structure of her moral character? In answering these questions, this research sheds light on important conceptual and philosophical issues, from the very existence of moral character to the degree of subjectivity present in a person's moral perceptions.

The existence of moral character

Although people regularly describe one another in terms of moral character, there is skepticism stemming from both psychology and philosophy as to whether such a construct actually exists.

Character traits such as honesty or compassion, for example, might simply be folk psychological concepts that people apply in their understanding of others, residing only in the minds of perceivers. This skepticism rests on the assumption that perceivers exaggerate the actual degree of consistency present in people's moral behavior or underestimate the degree to which moral behavior is determined by factors outside of the individual (such as situational forces).

Doubts about character date back to a famous and rigorous study conducted by Hartshorne and May (1928), which tested for consistency in children's honesty behaviors. Famously, this attempt seemed to provide meager evidence for dispositional honesty: Although children's behavior in one honesty context (cheating on a test) was predictive of their behavior in another honesty context (cheating at a game), the relationship was seen as modest at best. This led researchers in psychology to abandon the study of moral character, in favor of the study of situational factors that increased or decreased moral behavior in the average person. This approach has indeed been fruitful, providing many memorable demonstrations of the "power of the situation" to inhibit or promote moral acts (Darley & Batson, 1973; Haney, Banks, & Zimbardo, 1973; Milgram, 1974). To present day, the study of situational effects on moral judgment and action continues to dominate the field of moral psychology (for a review, see Fleenor, Furr, Jayawickreme, Meindl, & Helzer, 2014).

Within moral philosophy, some scholars have looked to these psychological findings to argue strongly against the existence of moral character (Doris, 2002; Harman, 2000; 2003; 2009). If moral behavior is inconsistent and heavily affected by fleeting features of a person's situation, the argument goes, then it makes no sense to conceive of people as possessing stable dispositional tendencies toward moral behavior. As Doris (2002) puts it:

The experimental record suggests that situational factors are often better predictors of behavior than personal factors... In very many situations, it looks as though personality is less than robustly determinative of behaviors. To put things crudely, people typically lack character. (p. 2)

Of course, the fact that situations exert powerful effects on behavior is not, in and of itself, evidence

against dispositions, moral or otherwise. Evidence for moral dispositions comes not from the main effects of situations, but from the existence of stable individual differences in moral behavior (Jayawickreme, Meindl, Helzer, Furr, & Fleeson, in press). To date, however, empirical evidence of robust moral character traits is scarce – a problematic omission in light of the importance that people place on the concept of moral character.

Using agreement to test for moral character

To the extent that individual differences on a given trait exist and are observable, ratings of that trait furnished by different judges should agree with one another (Allport, 1937; Kendrick & Funder, 1988; Funder, 1980, 1991). For this reason, one key piece of evidence for the existence of traits concerns whether independent observers agree both on the structure of a person's personality and on the distribution of traits within a given sample (Kendrick & Funder, 1988). In a typical agreement study, a variety of judges, often including the self, provide reports on a “target person's” personality. These reports are compared across all targets to determine whether different judges agree on a target's standing on a particular trait (for recent reviews, see Kenny & West, 2010; Vazire & Carlson, 2010).

Why does agreement constitute evidence for the existence of traits? There are at least two reasons. Allport (1937) summarizes one reason as follows:

What is most noteworthy in research on personality is that different observers should agree as well as they do in judging any one person. This fact alone proves that there must be something really there, something objective in the nature of the individual himself that compels observers, in spite of their own prejudices, to view him in essentially the same way.

In other words, to the extent that distinct raters, in spite of potential errors and biases in their perception of a target, nonetheless arrive at similar conclusions about a target's personality, there is most likely something in the target driving these shared perceptions.

The second reason agreement is a criterion for trait existence hinges on the fact that trait ratings are summaries of a target's behavior, aggregated across multiple observations by a given judge.

Agreement among different judges, therefore, represents behavioral stability that the target him- or herself manifests in different contexts with different people. Robust self-other agreement similarly reveals that the way the target acts with a particular informant (as indicated by the informant's rating) corresponds with how that target typically acts with others, or when alone (as indicated by his or her self-rating). Thus, agreement reveals the existence of stable and observable traits that are manifested to a similar degree across the distinct contexts embodied by each rater.

According to this logic, the present study provides a critical test for the very existence of moral character. To the extent that different judges show significant agreement in judgments of moral character, there is likely something “really there” in the target that drives these shared perceptions. To date, such evidence is scarce.

The challenge of perceiving moral character

Agreement has been documented for a broad array of personality traits across dozens of studies (Kenny & West, 2010; Vazire & Carlson, 2010), which indirectly implies that different raters may also agree on judgments of moral character. But, moral character traits do differ from many non-moral traits (i.e., Big Five traits) in fundamental ways that are known to affect agreement. Thus, we cannot assume that findings from previous research will generalize to examinations of moral character, although it may well be that moral character traits (if they exist) are embedded within, or built from, more basic traits such as agreeableness or conscientiousness. We must examine moral character traits directly and specifically for evidence of agreement.

Specifically, two features of moral character traits might drive down levels of agreement, compared to what is typical for non-moral traits. We consider each of these in turn.

Moral character traits are evaluative. If moral character judgments are intuitive assessments of a person's underlying goodness, they are, by definition, evaluative (Paulhus & John, 1998).

Inter-judge correlations for ratings of highly evaluative traits (such as conscientiousness or intellect) tend to be lower than correlations on less evaluative traits (extraversion and emotional stability; John & Robbins, 1993; Vazire, 2010). This lower agreement may be due to the fact that ratings of evaluative traits are biased by social desirability or other self-presentational concerns. Ratings of a person's compassion, for example, will likely arise from some true reflection of the person's underlying compassion, as well as error due to a desire to see the target in a particular light. Such errors in self- or other-assessment may degrade the quality of those reports, making it difficult to detect agreement among raters.

It is important to note that evaluativeness is unlikely to have the opposite effect— i.e., to artificially increase agreement. One might argue, for example, that if targets and informants are both influenced by social desirability, agreement among targets and informants can be accounted for simply by this shared bias. Note, however, that trait desirability should affect all raters in our sample, either to the same degree (which would leave the agreement correlations untouched) or to different degrees (which would make it more difficult to detect actual differences between people, and thus more difficult to find agreement). Thus, at the extreme, social desirability would not produce greater agreement, it would actually make agreement impossible to calculate.

There is one way in which evaluative processes could artificially produce agreement: if paired raters share a common top-down bias (such as general positivity, or a “halo”) that differs from pair to pair. For example, raters in one pair may share a general positive evaluation of their target that is different from the general positive evaluation that raters in a different pair share with one another about their target. Provided this general positivity was unrelated to targets’ morality (an unlikely possibility, in our estimation), this pattern of biases could produce spurious agreement on moral character. Thus, in the analyses that follow, we were careful to measure general positivity toward targets to examine its impact on moral character agreement.

Moral character traits are somewhat “internal.” Moral character traits like honesty or compassion operate through many channels, some of which are outwardly observable (e.g., behavior), while others are private and internal (e.g., affect and motivation). Manifestations of compassion, e.g., might take the form of both outward displays of kindness and internal distress felt while observing another individual's suffering. Internal states may play a unique role in evaluations of moral character. For example, research has shown that people use a target's intentions more when rating him or her on attributes like helpful, kind-hearted, sympathetic, considerate, and generous than when rating him or her on attributes like confident, articulate, skillful, talented, and wise (Kruger & Gilovich, 2004).

Traits based upon internal qualities tend to be more difficult to assess, at least for outward observers (Vazire, 2010). Traits such as intellect, which is manifested in intellectual curiosity and openness to ideas, are less visible from an outsider's perspective, and thus, observers' ratings may provide impoverished insight into a target's standing on these traits (Funder & Colvin, 1988; John & Robbins, 1993; Vazire 2010). Similarly, in the moral domain, observers may be unaware of a person's true concern for others (compassion) or their distress at seeing injustice (fairness), and may thus miss a crucial piece of the person's underlying moral character.

Evidence for agreement on moral character

Despite these potential challenges, recent work provides evidence that different raters do agree in their assessments of morally-relevant traits. Self-assessments of Honesty-Humility, a domain of personality relevant to a person's sincerity, fairness, greed-avoidance, and modesty, as well as Guilt Proneness, an individual difference variable related to various kinds of moral behavior, are reliably associated with peer assessments (Ashton & Lee, 2010; Cohen, Panter, Turan, Morse, & Kim, 2013; Lee, Ashton, Pozebon, Visser, Bourdage, & Ogunfowora, 2009). Past work also suggests that different facets of morality may be more observable than others. For example, agreement on the Sincerity facet

of Honesty-Humility tends to be much lower than agreement on the fairness, greed-avoidance, and modesty (Lee et al., 2009).

In terms of inter-informant agreement, only indirect evidence is available with regard to moral character traits. Frimer and colleagues have shown that on moral character traits such virtuousness, principledness, and fairness, different judges agree in their assessments of influential cultural figures (Frimer, Walker, Lee, Riches, & Dunlop, 2012; Frimer, Biesanz, Walker & MacKinlay, 2013). It is unclear from these findings, however, whether informants would agree in their perceptions of a familiar, but not famous, person, with whom they have a great deal of first-hand experience, and about whom no cultural stereotypes exist.

Thus, the present research contributes to this growing literature by examining agreement on a broader assortment of moral character traits than has previously been studied, including both specific (fairness, compassion, honesty, and temperance), and broad (what we are terming *moral concern* and *general morality*) traits, by examining both self/other and inter-informant agreement alongside one another in a single study, and by investigating two distinct conceptualizations of agreement – agreement at the trait level, as well as agreement at the level of moral profiles.

Method

Participants

Our sample consisted of students and community members. “Target” participants were recruited from introductory psychology classes as well as the local community. “Informants” were nominated by target participants, and represented a diverse sample of family members, friends, and coworkers.

Student subsample. Undergraduate students were recruited for a study on “perceptions of personality,” and were compensated with credit in their psychology courses. These target participants were asked to nominate up to eight informants to rate the targets' personalities. To encourage diversity in these nominations, we asked targets to consider family members, college friends, and hometown

friends as potential informants. All nominated informants were contacted by email and asked to complete the study. Informants were compensated with either a \$10 Amazon.com gift card or entry into a raffle for the chance to win an iPad.

Community subsample. The remainder of the sample were community members in the Winston-Salem area. The “community target” participants were recruited through public advertisements and advertisements for ongoing studies at Wake Forest University. They were recruited for a study on “perceptions of personality,” and completed the measures reported here as part of a larger study of morality. Community target participants nominated up to 9 informants from different domains of life (significant others, coworkers, friends).¹ All nominated informants were contacted by email or mail to complete the study. In exchange for participation, targets and informants were entered into a raffle offering a 1-in-3 chance of winning a \$200 gas card.

Initially, 263 target participants completed the study. To be eligible for inclusion in this data set, we required useable data from at least two respondents (e.g., a target and one informant, or two informants). This resulted in a final sample of 173 targets and 493 informants. The usable number of informants ranged between 1 and 7, with a median of 2 informants per target (see Table 4 for sample sizes by subsample).

Materials and Procedure

Target participants completed all measures either online or on paper in the lab. The Moral Character Questionnaire that served as the basis of this study was administered along with other questionnaires not relevant to the current study. At the end of the session, targets supplied names and contact information for potential informants. Informants either completed measures online or on paper (in the latter case, informants returned measures through the mail). In addition to completing the Moral Character Questionnaire, informants indicated their global positivity (from 1 to 5, anchored at *Not at all* and *Very much*) toward Targets using two items, “I like this person,” and “I enjoy spending time

with this person.”

Moral Character Questionnaire. Targets’ moral character was assessed using a measure designed by the research team. In crafting the measure, we sought to capture targets’ standings on a broad array of moral character traits. More established scales, such as Honesty/Humility in the HEXACO (Lee & Ashton, 2004) or the Moral Foundations Questionnaire (MFQ; Graham, Nosek, Haidt, Iyer, Koleva, & Ditto, 2011) include content related to some important facets of morality; however, our aim was to assess characteristics that were uniquely and perhaps uncontroversially moral. For example, although Honesty-Humility captures facets of personality that are clearly relevant to moral character (e.g., behavioral tendencies toward stealing and bribery), it also includes tendencies that are not directly moral (e.g., possessing a sense of entitlement, using ingratiation instrumentally). Similarly, although the MFQ captures individual differences in foundational moral concerns, people (conservatives and liberals alike) do not associate all foundations with moral character (Frimer et al., 2013).

For this reason, we developed a scale that incorporated items from established scales and newly-generated items by the research team. Items from the HEXACO (Lee & Ashton, 2004), Temperament and Character Inventory (Cloninger, Przybeck, Svrakic, & Wetzel, 1994) and the Self-Control Scale (Tangney, Baumeister, & Boone, 2004) were carefully combined by the research team to represent each targeted moral character trait (or “facet”). Items were chosen to maximize face- and construct-validity with regard to common-sense notions of “moral character.” New items were generated primarily for the *moral concern* and *general morality* facets, as measures for these broad moral character traits were not established in the previous literature. These items were written after careful analysis of the moral character trait, strengths, and virtues literature and thorough discussion within the research team.

The resulting 41 items tapped 6 facets of moral character, with six to eight items per facet (See

Appendix for all items and psychometric properties). Targets and informants rated how well each of the 41 statements describes him/herself (targets) or “your friend or family member” (informants), using a five-point scale, from 1 (*Very inaccurate*) to 5 (*Very accurate*). Items assessed Fairness (“Treats everyone in a similar way”), Honesty (“Tells the truth”), Compassion (“Is indifferent to the needs of others” [reverse]), and Temperance (“Rarely overindulges”), Moral Concern (“Makes decisions with ‘doing the right thing’ in mind”), and General Morality (“Is a moral person”).

Assessing evaluativeness and internality. Five independent raters (all psychology graduate students who were blind to hypotheses) were recruited to code each of the 41 items on evaluativeness and observability/internality.

To assess evaluativeness, we asked raters about the extent to which it was “socially desirable to be a person who...” followed by the wording of the item. We described social-desirability as “the degree to which a trait is seen as ‘good’ in the eyes of other people.” Ratings were made on a 5-point scale from *Very undesirable* to *Very desirable*, with a neutral midpoint.

To assess observability/internality, we asked raters about the extent to which they would use information about thoughts and intentions versus observable behavior to assess each item. Ratings were made on a five-point scale from *Just thoughts and intentions* to *Just observable behavior*, with *Both thoughts/intentions and observable behavior* as a midpoint.

Results

Table 1 presents descriptive statistics and reliabilities for the six facets of character. Facet scales generally showed acceptable reliability, with an average alpha of .75 (lowest α = .52 for self-rated Honesty, highest α = .87 for informant-rated Temperance).² In general, ratings by targets and informants were quite high for all six facets (note that means for each scale are above the midpoint of 3).³

As can be seen in Table 2, the six facets were positively correlated with one another (correlations between traits ranged from .17 to .68, mean r = .45). When they were factor analyzed

together, a unidimensional structure (*Global Moral Character*) accounting for 56.95% of the variance emerged. As can be seen in the last column of Table 2, extraction communalities for the six facets ranged from .24 (Temperance) to .71 (General Morality), *Med.* = .52. Thus, each of the six facets includes substantial variance unique from the general factor. For this reason, and because there is value in assessment of individual differences at different levels of breadth (Wood, Nye, & Saucier, 2010), we provide agreement analyses both on Global Moral Character and the six facets.

Prior to performing agreement analyses, ratings for all six facets were examined for outliers. Facet scores falling 2.5 SD away from the group average for the facet were excluded from analyses. Additional analyses confirmed that these exclusions did not meaningfully change the magnitude of the correlations reported.

Two ways of conceptualizing agreement

There are at least two ways of conceptualizing agreement (Funder & Colvin, 1997; Furr, Dougherty, Marsh, & Mathias, 2007), each offering different insights into the meaning and nature of agreement (Furr, 2009). Perhaps the most common conceptualization, trait-level agreement, measures the extent to which raters agree on a target's standing, relative to other people, on a particular personality trait. A second conceptualization considers agreement on a person's pattern of traits, or their "profile" (Furr, 2010). We examine both conceptualizations within our data to better understand the nature of moral agreement.

Agreement at the trait level

We first examined agreement at the trait level for Global Moral Character, as well as its six facets. Resulting correlations reflect the degree to which targets and informants (self/other agreement) or different informants (inter-informant agreement) agreed about which targets had the highest and lowest scores on each dimension. A significant positive self-other agreement correlation indicates, for example, that targets who see themselves as relatively compassionate (compared to the way that other

targets see themselves) also tend to be seen as relatively compassionate by their informants (as compared to the way that other informants see their targets).

Because the number and type of informants varied unsystematically between targets, ratings from informants were treated as exchangeable or “non-distinguishable” (Kenny, Kashy, & Cook, 2006). We thus adapted path analytic procedures outlined by Furr and Wood (2013).⁴ Analyzing each dimension in a separate model, this procedure capitalizes on all available informant ratings to estimate, in essence, the correlation between targets’ self-ratings and any randomly-selected informant (self/other) or between any two randomly-selected informants (inter-informant).⁵ Note that because of this, we were not able to provide the degrees of freedom for estimates reported below. Instead, we report sample sizes and confidence intervals for each estimate in Tables 3-5.

As shown in the top row of Table 3, there was significant positive self/other agreement for Global Moral Character, $r = .36$, and for each of the six facets. The magnitude of facet-level agreement varied by facet. Agreement was lowest for Compassion, $r = .15$, and highest for Temperance, $r = .42$.

The fourth row of Table 3 displays significant positive inter-informant agreement for Global Moral Character, $r = .24$, as well as for five of the six facets. Inter-informant agreement was lowest for Honesty, $r = .09$, and highest for Temperance, $r = .35$. Notably, the pattern of inter-informant agreement for the different facets was highly similar to the pattern of self-other agreement. Indeed, the correlation between Table 3’s self-other and inter-informant rows (excluding Global Moral Character) is robust at $r(4) = .71$, indicating that the facets with highest (or lowest) self-other agreement also tended to have the highest (or lowest) inter-informant agreement.

Agreement on facet residuals. To better understand the relationship between Global Moral Character and its constituent six facets – and in particular, whether these facets capture unique variance in moral behavior over the global dimension – we computed, for each facet, a score that represented the unique information conveyed by ratings of the focal facet over and above ratings of the other five

facets. For example, to examine unique agreement on compassion, we residualized targets' and informants' compassion ratings, respectively, on targets' and informants' ratings of the other five facets.

As can be seen in Table 3, significant positive correlations remained for self/other agreement on Fairness, Honesty, Temperance, and General Morality and for inter-informant agreement on Fairness, Compassion, and Temperance, indicating that these facets carried unique, agreed-upon information about moral behavior over what was conveyed by Global Moral Character. Agreement on the other facets (Compassion and Moral Concern for self-other agreement, and Honesty, Moral Concern, and General Morality for inter-informant agreement) dropped to non-significance, suggesting that shared perceptions of these dimensions resulted from agreement on the other facets of Global Moral Character. More broadly, average agreement on the six facets was reduced from $\bar{r} = .27$ and $.20$ for self/other and inter-informant agreement, respectively, to $.15$ and $.09$ in the residualized models, suggesting, in line with communality estimates provided by the factor analysis above, that a little more than half of the variance in agreement on specific facets was due to agreement on Global Moral Character.

Was agreement driven halo effects? To ensure that agreement was not driven spuriously by informants' global positivity toward targets, we repeated agreement analyses for Global Moral Character and the six facets, controlling for informants' rated liking of the targets. As can be seen in the third and sixth row of Table 3, controlling for liking did not have a meaningful impact on the magnitude of self/other or inter-informant agreement. Thus, observed agreement did not seem to be due to a positive "halo" in informants' ratings of targets.

Effects of evaluativeness and internality on self/other agreement. We next examined whether evaluativeness and internality/observability affected agreement. We conducted these analyses at the item level so as to maximize the number of observations per analysis. Evaluativeness and Observability ratings for each of the 41 items were averaged across raters. On average, raters saw the items as highly desirable (scores ranged from 3.25 to 5, $M = 4.43$, $SD = .46$), and as embodied in a mix of thoughts and

intentions as well as outward behavior (scores ranged from 2 to 4.29, $M = 3.54$, $SD = .58$).

We then computed self-other and inter-informant agreement for each of the 41 items or “moral characteristics” using the same procedure as above, and correlated item-level agreement with ratings of Evaluativeness and Observability. As predicted, Evaluativeness was correlated $r_s(39) = -.37$ and $-.24$ and Observability $r_s(39) = .29$ and $.28$ with self-other and inter-informant agreement, respectively. Thus, echoing past research, weaker agreement was observed for socially desirable and less observable moral characteristics.

Agreement by subsample

Although the sizes of each subsample are relatively low (and thus results are less precise), we examined agreement for the Community and Student samples separately (see Table 4). These subsample analyses are of interest primarily as an admittedly-rough internal gauge of replicability, and three noteworthy findings emerge. First, self-other agreement on Global Moral Character was significant in both subsamples. Self-other agreement was also significantly positive for all facets in the Community sample and for four of the six facets in the Student sample. This provides replication of general significant positive self-other agreement on moral character and its component facets. Second, the pattern of self-other agreement across facets replicated across the two subsamples. For example, in both subsamples, Temperance elicited the highest levels of agreement, while Honesty and Compassion consistently elicited the lowest levels. More formally, the correlation between the two subsample rows of self-other agreement (from Table 3) was robust at $r(4) = .78$, indicating that the facets with highest (or lowest) self-other agreement in the Community sample also had the highest (or lowest) self-other agreement in the Student sample. Third, the overall level of self-other agreement was descriptively higher in the Community vs. Student sample. Thus, although positive and generally significant in both subsamples, self-other agreement was somewhat stronger among Community members.

Turning to inter-informant agreement, those in the Community sample agreed with one another

($r = .36$ for Global Moral Character) to a much greater extent than informants in the Student sample ($r = .08$ for Global Moral Character). Thus, robust evidence for inter-observer agreement was obtained in only one subsample. In addition, the pattern of differences among the six facets did not cleanly replicate across subsamples. Most strikingly, while Temperance elicited the highest agreement in the Community sample ($r = .51$), it elicited the lowest agreement in the Student sample (tied with Compassion at $r = .04$).

Agreement on moral character profiles

In a profile approach, agreement is defined at the level of each target, indexed by the average profile correlation for each pair of raters (e.g., self with informant or informant with informant). Suppose, for example, that Tom sees himself as being very compassionate, somewhat fair and honest, and quite poor at controlling his impulses (temperance). This pattern (high compassion, low temperance, and mid-range honesty and fairness) is Tom's moral character profile. In a profile approach, agreement is conceptualized as the extent to which different judges have a similar view of Tom's pattern of moral character traits (Furr, 2010). Profile agreement is typically quantified via a simple Pearson correlation between profiles as rated by two people – e.g., target and informant. A positive correlation indicates that target and informant agree about the target's moral character strengths and weaknesses (i.e., which facets are relatively high in the target's moral profile and which are relatively low). In contrast, a zero correlation indicates no systematic agreement, and a negative correlation would indicate that target and informant have inverse views of the target's moral strengths and weaknesses.

For all profile analyses, we omitted the two broader facets of Moral Concern and General Morality. Because Moral Concern and General Morality are non-specific and cut across several

domains of morality, they are unsuitable to include along with the more specific facets in targets' moral profiles (Furr, 2010).

To examine self-other agreement, two moral character profile correlations were calculated for each target-informant pair (Furr, 2008). The first was an overall profile correlation: A target's self-ratings for fairness, honesty, compassion, and temperance were correlated with the ratings of each of his or her informants. These pairwise correlations were then subjected to a Fisher's *Z* transformation, and averaged across all informants for each target, to obtain an average profile agreement score for each target. These were then averaged across all targets for a single index of overall agreement. The second profile correlation indexed distinctive agreement, which controls for the moral profile of the average or "normative" person. This provides a very conservative index of inter-rater agreement that estimates agreement on how a particular person differs from the norm (Biesanz, 2010; Furr, 2008). To calculate distinctive agreement, we created a "distinctive" self-rated moral profile for each target by subtracting the average self-rating for each facet from the target's self-rating on the corresponding facet. We similarly created a distinctive moral profile for each informant by subtracting the average informant's rating for each facet from each informant's ratings. We then correlated a target's distinctive self-rated moral profile with each of his or her informants' distinctive moral profiles, indicating the degree to which each informant sees the target's distinctive moral qualities similarly to the target's self-view. As with overall agreement, these pairwise correlations were then *Z*-transformed and averaged into a single index of self-other distinctive agreement.

The first row in Table 5 provides robust evidence of self-other agreement for moral character profiles. The average level of overall agreement was very strong, $\bar{r} = .73, p < .001$, but perhaps even more impressive is the level of distinctive self-other agreement. Although it is a much more conservative estimate of agreement, distinctive self-other agreement was still well above zero, at

$\bar{r} = .32, p < .001$.⁶ This provides compelling evidence of self-other agreement about targets' moral character strengths and weaknesses. Moreover, these significant findings replicate across the two subsamples, though agreement was again somewhat stronger among Community participants than Students (see Table 4).

Similar procedures were used to examine inter-informant agreement. Specifically, for each target, informants' profiles were correlated with each other, and the correlations were then Z-transformed and averaged for each target. Then, those averages were themselves averaged across all targets to obtain single indices of overall and distinctive inter-informant profile agreement

The second row in Table 4 provides evidence of inter-informant agreement for moral profiles, with significant average levels of both overall ($\bar{r} = .66, p < .001$) and distinctive agreement ($\bar{r} = .30, p < .001$). Thus, ratings of moral character converged not only on targets' relative character strengths and weaknesses, but also on how those strengths and weaknesses differentiated targets from the average person. These findings generally replicate across the two subsamples, with the sole exception being non-significant distinctive inter-informant agreement in the student sample (see Table 5).⁷

General Discussion

This study provides evidence for both self/other and inter-informant agreement on perceptions of moral character. Across a variety of targets, informants, moral character traits, and methods of conceptualizing agreement, participants generally showed significant – and sometimes substantial – convergence in their ratings of moral character. These results provide new insights into questions that are of interest to moral, social, and personality psychologists; they address important issues concerning the very nature of moral character, and broaden current understanding of moral perception and judgment.

Implications regarding the existence of moral character traits

These findings provide a strong basis for the study of moral character, by revealing meaningful individual differences in morality. The fact that moral character ratings from different judges agreed with one another suggests, first, that individuals differ from one another in their levels of moral behavior, and second, that those differences are stable and manifested across different situations, with different raters, as well as alone. Agreement also suggests that moral character is sufficiently visible to a range of observers, who, despite their idiosyncratic biases, nonetheless converge in their impressions of who is more or less moral (Allport, 1937; Funder, 1980; 1991). Thus, contrary to the strong claims of situationists, the study of character appears to be based upon solid empirical ground.

These arguments notwithstanding, a critical reader might be inclined to offer three counter-explanations to our findings that, if true, would raise doubts about whether character judgments are tracking actual dispositions. We address these in turn.

Objection 1: Shared, inaccurate heuristics. One alternative explanation for these findings is that agreement was driven by reliance upon shared, inaccurate heuristics. If, for example, targets and informants based their ratings on morally-irrelevant characteristics (such as using expressed positive affect to assess honesty) rather than actual correlates of moral character, then “agreement” could be driven by cues that are unrelated to actual moral character. Importantly, this objection requires that the heuristics be inaccurate, since any accurate heuristics would, by definition, be based on genuine manifestations of moral character.

While we think it is likely that moral character judgments are based upon heuristics, it seems quite unlikely that those heuristics are wholly inaccurate. In fact, research has begun to confirm that people’s moral heuristics are accurate, such that they are predictive of morally-relevant behaviors, in domains spanning compassion, honesty, and fairness (Carre, McCormick, & Mondloch, 2009; Haselhuhn & Wong, 2012; Stirrat & Perrett, 2010). A second reason to doubt this alternative explanation is that trait-level agreement among informants in our study was consistently weaker than

self/other agreement. Because stereotype reliance decreases with target familiarity (Biesanz, West, & Millevoi, 2007), and because the self is a familiar target, a heuristics explanation would likely predict exactly the opposite patterns of agreement.

Objection 2: Agreement based upon shared, positive “halos.” Another alternative explanation for our findings is that agreement was a function of shared top-down positivity toward a target. That is, if informants think positively about targets to the same degree that targets think positively about themselves (and if target-informant pairs differ from one another in the degree of this shared positivity), then “agreement” could simply be the function of these top-down biases. As Table 3 shows, however, partialling out informants’ liking of targets had virtually no effect on observed agreement, suggesting that shared positivity cannot explain agreement.

Objection 3: Self-presentation. A final alternative explanation is that agreement resulted from targets self-presenting to informants. Perhaps some targets talk a great deal about their moral successes, while others talk a lot about their moral failures, and informants use this information to judge character. If so, “agreement” would actually be driven by informants simply parroting back what targets have told them about themselves. Like Objection 1, this would only be problematic if targets’ self-presentation is unrelated to their actual behavior (otherwise self-presentation is simply a mechanism by which people learn about others’ character). It seems unlikely that this is the case. Furthermore, if a self-presentation were an explanation of our findings agreement should have been strongest for the most desirable traits (the traits on which a person is most likely to self-present). However, our results suggest that the desirability of moral character traits decreased agreement, casting doubt on a self-presentation explanation.

Thus, while we do not doubt that heuristics, shared positivity, and self-presentation affect self- and informant-ratings of character, we doubt strongly, and on the basis of patterns in our data, that these were confounding or illegitimate factors that artificially produced agreement. Instead, we believe our

data represent true agreement among various judges, and constitute evidence for the existence of character.

Implications regarding the nature of moral character

Several patterns in our data raise important questions about the nature of moral character. First, these data carry implications for understanding the structure of moral character. Second, although agreement on moral character was comparable to, if not somewhat higher than, levels of agreement John and Robins (1993) found for other highly evaluative traits, it was lower in magnitude than typical agreement seen for Big Five traits. Third, there were noticeable differences in levels of agreement seen in our two subsamples: While community agreement was substantial, agreement in the student sample was weaker—and informants in the student sample simply did not agree to a significant degree with one another. It is possible that any one of these patterns could be accounted for by methodological or measurement issues. For example, our study employed a novel measure of morality, making it difficult to compare the magnitude of our results to past work. However, this measure was psychometrically sound (with one exception: the low reliability for self-reported Honesty); as a result, we will focus our discussion away from measurement or methodology and toward the nature of moral character traits.

Structure of moral character. Looking across all indicators (factor analysis, extraction communalities, intercorrelations among moral character traits, residual agreement) our data suggest that moral character, like other broad traits, has a strong core dimension (represented by Global Moral Character), built from some number facets or narrower traits that are themselves meaningful and observable. Relevant analyses suggest that about half of the variance in observable moral character is accounted for by a global factor, leaving substantial variance to be explained by specific moral character traits (in particular, fairness and temperance seemed to consistently elicit interrater agreement over and above this general factor). Future research will be needed to elaborate the exact nature of

these more specific traits. In particular, one interesting question for future research is whether the moral character traits derived from self- and peer-ratings actually conform to the traditional categories used in virtue ethics (a literature that inspired our design), or represent altogether different dimensions on which people differ from one another on moral behavior (see, for example, Miller, 2013).

Lower agreement for moral character traits. We found evidence that differing levels of agreement for the 41 items in the Moral Character Questionnaire were accounted for by differences in the internality and evaluativeness of those traits. Although we could not directly compare moral character traits to other, non-moral traits on these dimensions, we suspect that part of the reason agreement in this study was lower than agreement for Big Five traits was that moral character traits are particularly internal and evaluative (Vazire, 2010).

The internality of moral character traits is what makes them both interesting and difficult to study from the perspective of observers. Relative to assessments of non-moral traits, moral character traits may require additional inferences about the contents of targets' minds, such as their motivational states and intentions. For example, volunteering to help a coworker may be seen as an act of compassion only if the intention is to alleviate the coworker's suffering, rather than to curry favor with one's boss or to springboard into an after-work date. By comparison, judgments of whether a person is talkative or emotionally stable may not require deeper inferences about mental states: A person is talkative regardless of whether he is talkative for the sake of being talkative or talkative for the right reasons (cf. Kammrath, Mendoza-Denton, & Mischel, 2005). Thus, motives may be more central to judgments of moral character traits, requiring additional inferences on the part of perceivers, and degrading agreement.

Evaluativeness also affected agreement, with more socially-desirable moral characteristics eliciting lower levels of agreement than less socially-desirable moral characteristics. This replicates previous research (John & Robbins, 1993), and raises interesting questions about how to best elicit

valid ratings of moral character traits – ratings that will provide the truest estimate of actual differences between people on these moral dimensions.

Interestingly, ratings of internality and evaluativeness were themselves correlated with one another. Indeed, the most observable moral characteristics were rated as the least socially desirable, and so it is difficult to determine whether evaluativeness or internality *per se* is actually driving differences in agreement. This non-independence also provides an intriguing avenue for future research: Perhaps the reason a particular moral characteristic is seen as highly desirable is because it is reflected in a person's internal processes, such as their intentions and desires. Thus, moral characteristics that are seen as more internal may be seen as better indicators of a person's "true" character.

Lower agreement in the student subsample. Another striking result is that agreement varied noticeably by subsample. In particular, within the student sample, inter-informant agreement was nonsignificant using both trait-level and profile analyses, and self/other agreement, though stronger than agreement among informants, was still weaker than agreement seen in the community sample. This difference does not appear to be due to statistical power – if anything, the student sample included a greater number of informants per target than did the community sample. It also did not seem to be due to restriction of range among targets' self-ratings in the student sample. Interestingly, informants in the student sample did show descriptively more restricted range than informants in the community sample (on five of the six facets the range of student informants' ratings was smaller than that of community informants), and the observed range for each of the six facets was correlated at $r(4) = .45$ with the magnitude of self/other agreement.

Thus, it is possible that informants in the student sample contributed to lower agreement by offering restricted ratings of moral character. The direction of the restriction (ratings were skewed toward the top of the scale) suggests that informants had difficulty identifying the less moral targets, but why? One possibility is that breaches in moral behavior for the average undergraduate were not

chalked up to moral character, but to developmental or social forces. An act of callousness, for example, may be attributed to immaturity or peer pressure, leaving favorable impressions of the target's trait compassion untouched. Given that a sizeable proportion of informants in the student sample were parents of the targets, it may be that some informants were motivated to maintain charitable impressions.

A different possibility is that informants in the student sample were from broader life domains than informants in the community sample. In addition to parents, the student informant sample consisted of targets' hometown friends and college friends. Thus, informant reports in the student sample represented not only distinct social contexts, but different periods across the lifespan (with parents' ratings summarizing moral character throughout life, college friends summarizing moral character since arriving at the university, and hometown friends somewhere in between). If moral character follows a predictable developmental trajectory, such that moral character traits are more consistent in mid- to late-adulthood than in late-adolescence and emerging-adulthood, it would make sense that reports based upon observation of behavior during this earlier period would be less consistent.

What is revealed by a profile approach to character?

Although agreement at the trait-level was lower in this study than what might be expected given past research, levels of profile agreement were comparable to those seen elsewhere using similar analytic approaches (Furr et al., 2007). To our knowledge, this is the first study to employ a profile approach to studying agreement on moral character; so, the robustness of both overall and distinctive agreement observed here offers promising directions for future research. Indeed, a profile approach to the study of moral character may offer certain conceptual advantages over the traditional trait-level approach, perhaps by better capturing the psychology of moral character perception. It is possible (and

maybe even likely) that when thinking about moral character, people categorize one another and themselves more in terms of intra-individual trade-offs between various moral character traits, and less in terms of the target's absolute standing on a broad moral character dimension. If so, one reason agreement on moral profiles was higher than trait-level agreement in this study might be that this analytic tool better captures the underlying process with which people perceive morality.

Using profiles to understand moderators of agreement. Results from the moral profile analyses reveal another important insight about perceptions of moral character: Namely, that there exists a high degree of variability in the magnitude of agreement. At the level of dyads (i.e., any two raters), profile correlations ranged from -1 to +1, suggesting that different raters agreed with one another to varying degrees. This variability amidst overall agreement raises important questions about the correlates of agreement. We suspect that agreement on moral character is related to a number of relational factors already identified in past agreement research, such as acquaintanceship (Funder & Colvin, 1988), the quality of the judge, and the information available about the target (Funder, 1995).

Perhaps because morality is so central to people's understanding of themselves and others, agreement on a person's moral character is likely to have far-reaching interpersonal consequences (Tenney, Vazire, & Mehl, 2013). Self/other agreement, for example, may strengthen interpersonal closeness by allowing people to verify their self-views (Swann, 1983), or simply by minimizing interpersonal disagreement. On the flip side, inter-observer agreement on moral character may have important consequences for coalition-building and group processes. Group decisions about whether to include or ostracize, hire or fire, promote or demote a particular individual might be based not only upon whether the target is seen as highly moral, but whether such perceptions are widely shared by different people. Additional research will be needed to understand the downstream effects of shared moral perception.

Conclusions

This study provides evidence that that different judges agree significantly on the structure and contents of a person's moral character. Using two measures of agreement and a diverse sample of targets and informants, we found ratings provided by friends, family, and acquaintances agreed with ratings furnished by the self, and, to a lesser extent, with ratings provided by other informants. Results suggest that moral behavior is consistent, such that ratings provided by different judges converge on similar assessments of moral character. This research also highlights the importance of a person-centered approach to moral psychology, and suggests exciting new questions for advancing current knowledge.

Footnotes

1. Due to pragmatic constraints, about 1/3 of community participants were asked for only four informants.
2. It is unclear why self-ratings of Honesty showed relatively poor reliability, while informant-ratings using the same items showed satisfactory reliability. Analysis revealed no specific items that were responsible for lowered reliability.
3. Across the six traits measured, informant ratings were in almost every case higher than self-ratings (the only exception was compassion, the highest-rated trait overall). In four out of six cases (fairness, honesty, temperance, and general morality) that difference was significant, $t_s > 5.07$, $p_s < .0001$.
4. We used a variation of Model 6 in Table 2 of Furr and Wood (2013). We modeled correlations between a single target variable and multiple informants instead of directional paths. To account for exchangeability among informants, means and standard deviations were held constant across informants. Analyses were conducted via AMOS 19. Model 6 in Table 3 of the online supplemental document accompanying Furr and Wood (2013) provides AMOS syntax that was subsequently adapted for the current analyses.
5. For inter-informant agreement, we used a variation of Model 1 in Table 2 of Furr and Wood (2013), adapted to include up to six informants. Model 1 in Table 3 of the online supplemental document accompanying Furr and Wood (2013) provides AMOS syntax that was subsequently adapted for the current analyses.
6. We confirmed that evaluativeness was not driving agreement among raters by analyzing self-other and inter-informant profile agreement while controlling for the desirability of each of the four traits. Even after controlling for evaluativeness, self-other ($\bar{r} = .73$, $p < .001$) and inter-informant ($\bar{r} = .61$, $p < .001$) profile agreement were robustly significant.
7. Alternatively, profile correlations can be calculated by computing the average self-other and inter-informant agreement for individual items from the Fairness, Honesty, Compassion, and Temperance scales. Overall item-level profile correlations were significant for both self-other ($\bar{r} = .35$, $p < .001$) and inter-informant ($\bar{r} = .28$, $p < .001$) profiles. Distinctive item-level profile correlations were also significant for both self-other ($\bar{r} = .12$, $p < .001$) and inter-informant ($\bar{r} = .19$, $p < .001$) agreement.

References

- Allport, G.W. (1937). *Personality: A psychological interpretation*. New York: Henry Holt & Co.
- Ashton, M.C., & Lee, K. (2010). Trait and source factors in HEXACO-PI-R self- and observer reports. *European Journal of Personality, 24*, 278-289.
- Biesanz, J.C. (2010). The social accuracy model of interpersonal perception: Assessing individual differences in perceptive and expressive accuracy. *Multivariate Behavioral Research, 45*, 853-885.
- Biesanz, J.C., West, S.G., & Millevoi, A. (2007). What do you learn about someone over time? The relationship between length of acquaintance and consensus and self-other agreement in judgments of personality. *Journal of Personality and Social Psychology, 92*, 119-135.
- Carré, J.M., McCormick, C. M., & Mondloch, C.J. (2009). Facial structure is a reliable cue of aggressive behavior. *Psychological Science, 20*, 1194-1198.
- Cloninger, R., Przybeck, T.R., Svrakic, D.M., & Wetzel, R. (1994). *The Temperament and Character Inventory (TCI): A guide to its development and use*. St. Louis, MO: Center for Psychobiology of Personality.
- Cohen, T.R., Panter, A. T., Turan, N., Morse, L., & Kim, Y. (2013). Agreement and similarity in self-other perceptions of moral character. *Journal of Research in Personality, 47*, 816-830.
- Darley, J.M., & Batson, C.D. (1973). From Jerusalem to Jericho: A study of situational and dispositional variables in helping behavior. *Journal of Personality and Social Psychology, 27*, 100-108.
- Doris, J.M. (2002). *Lack of character: Personality and moral behavior*. New York: Cambridge University Press.
- Fiske, S.T., Cuddy, A.J.C., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Science, 11*, 77-83.
- Fleeson, W., Furr, R.M., Jayawickreme, E., Meindl, P., & Helzer, E.G. (2014). Character: The prospects for a personality-based perspective on morality. *Personality and Social Psychology Compass, 8*, 178-191.
- Frimer, J.A., Walker, L.J., Lee, B.H., Riches, A., & Dunlop, W.L. (2012). Hierarchical integration of agency and communion: A study of influential moral figures. *Journal of Personality, 80*, 1117-1145.
- Frimer, J.A., Biesanz, J.C., Walker, L.J., & MacKinlay, C.W. (2013). Liberals and conservatives rely on common moral foundations when making moral judgments about influential people. *Journal of Personality and Social Psychology, 104*, 1040-1059.
- Funder, D.C. (1980). On seeing ourselves as others see us: Self-other agreement and discrepancy in

- personality ratings. *Journal of Personality*, 48, 473-493.
- Funder, D.C. (1991). Global traits: A neo-Allportian approach to personality. *Psychological Science*, 2, 31-39.
- Funder, D.C. (1995). On the accuracy of personality judgment: A realistic approach. *Psychological Review*, 102, 652-670.
- Funder, D.C., & Colvin, C.R. (1988). Friends and strangers: Acquaintanceship, agreement, and the accuracy of personality judgment. *Journal of Personality and Social Psychology*, 55, 149-158.
- Funder, D.C., & Colvin, C.R. (1997). Congruence of others' and self-judgments of personality. In R. Hogan, J. A. Johnson, & S.R. Briggs (Eds.), *Handbook of Personality Psychology* (pp. 617-647). San Diego, CA: Academic Press.
- Furr, R.M. (2008). A framework for profile similarity: Integrating similarity, normativeness, and distinctiveness. *Journal of Personality*, 76, 1267-1316.
- Furr, R.M. (2009). Profile analysis in person-situation integration. *Journal of Research in Personality*, 43, 196-207.
- Furr, R.M. (2010). The double-entry intraclass correlation as an index of profile similarity: Meaning, problems, and alternatives. *Journal of Personality Assessment*, 92, 1-15.
- Furr, R. M., Dougherty, D.M., Marsh, D.M., & Mathias, C.W. (2007). Personality judgment and personality pathology: Self-other agreement in adolescents with conduct disorder. *Journal of Personality*, 75, 629-662.
- Furr, R.M., & Wood, D. (2013). On the similarity between exchangeable profiles: A psychometric model, analytic strategy, and empirical illustration. *Journal of Research in Personality*, 47, 233-247.
- Goodwin, G.P., Piazza, J., & Rozin, P. (2014). Moral character predominates in person perception and evaluation. *Journal of Personality and Social Psychology*, 106, 148-168.
- Graham, J., Nosek, B.A., Haidt, J., Iyer, R., Koleva, S., & Ditto, P.H. (2011). Mapping the moral domain. *Journal of Personality and Social Psychology*, 101, 366-385.
- Haney, C., Banks, W., & Zimbardo, P. (1973). Interpersonal dynamics of a simulated prison. *International Journal of Criminology and Penology*, 1, 69-97.
- Hartshorne, H., & May, M.A. (1928). *Studies in the nature of character*. New York: Macmillan.
- Harman, G. (2000). The nonexistence of character traits. *Proceedings of the Aristotelian Society, New Series*, Vol. 100 (pp. 223-226). New York: Blackwell Publishing.
- Harman, G. (2003). No character or personality. *Business Ethics Quarterly*, 13, 87-94.

- Harman, G. (2009). Skepticism about character traits. *The Journal of Ethics*, 13, 235-242.
- Haselhuhn, M.P., & Wong, E.M. (2012). Bad to the bone: Facial structure predicts unethical behaviour. *Proceedings of the Royal Society B: Biological Sciences*, 279, 571-576.
- Jayawickreme, E., Meindl, P., Helzer, E.G., Furr, R.M., & Fleeson, W. (in press). Virtuous states and virtuous traits: How the empirical evidence regarding the existence of broad traits scaffolds virtue ethics. *Theory and Research in Education*.
- 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.
- Kammrath, L.K., Mendoza-Denton, R., & Mischel, W. (2005). Incorporating *If... Then...* personality signatures in person perception: Beyond the person-situation dichotomy. *Journal of Personality and Social Psychology*, 88, 605-618.
- Kendrick, D.T., & Funder, D.C. (1988). Profiting from controversy: Lessons from the person-situation debate. *American Psychologist*, 43, 23-34.
- Kenny, D.A., Kashy, D., & Cook, W.L. (2006). *Dyadic data analysis*. New York: Guilford Press.
- Kenny, D.A., & West, T.V. (2010). Similarity and agreement in self- and other perception: A meta-analysis. *Personality and Social Psychology Review*, 14, 196-213.
- Kruger, J., & Gilovich, T. (2004). Actions, intentions, and self-assessment: The road to self-enhancement is paved with good intentions. *Personality and Social Psychology Bulletin*, 30, 328-339.
- Lee, K., & Ashton, M.C. (2004). Psychometric properties of the HEXACO Personality Inventory. *Multivariate Behavioral Research*, 39, 329-358.
- Lee, K., Ashton, M.C., Pozzebon, J.A., Visser, B. A., Bourdage, J.S., & Ogunfowora, B. (2009). Similarity and assumed similarity in personality reports of well-acquainted persons. *Journal of Personality and Social Psychology*, 96, 460-472.
- Milgram, S. (1974). *Obedience to authority*. New York: Harper and Row.
- Miller, C.B. (2013). *Moral Character: An Empirical Theory*. Oxford: Oxford University Press.
- Paulhus, D.L., & John, O.P. (1998). Egoistic and moralistic bias in self-perceptions: The interplay of self-deceptive styles with basic traits and motives. *Journal of Personality*, 66, 2024-1060.
- Pizarro, D.A. & Tannenbaum, D. (2011). Bringing character back: How the motivation to evaluate character influences judgments of moral blame. In M. Mikulincer & Shaver, P. (Eds) *The Social psychology of morality: Exploring the causes of good and evil*. New York: APA Press.
- Stirrat, M., & Perrett, D.I. (2010). Valid facial cues to cooperation and trust: Male facial width and

trustworthiness. *Psychological Science*, 21, 349-354.

Strohminger, N., & Nichols, S. (2014). The essential moral self. *Cognition*, 131, 159-171.

Swann, W.B., Jr. (1983). Self-verification: Bringing social reality into harmony with the self. In J. Suls & A. G. Greenwald (Eds.), *Psychological perspectives on the self* (Vol. 2, pp. 33–66), Hillsdale, NJ: Erlbaum.

Tangney, J.P., Baumeister, R.F., & Boone, A.L. (2004). High self-control predicts good adjustment, less pathology, better grades, and interpersonal success. *Journal of personality*, 72, 271-324.

Tenney, E.R., Vazire, S., & Mehl, M.R. (2013). This examined life: The upside of self-knowledge for interpersonal relationships. *PLoS ONE*, 8, e69605. doi:10.1371/journal.pone.0069605.

van't Wout, M., & Snafey, A.G. (2008). Friend or foe: The effect of implicit trustworthiness judgments in social decision-making. *Cognition*, 108, 796-803.

Vazire, S. (2010). Who knows what about a person? The self-other knowledge asymmetry (SOKA) model. *Journal of Personality and Social Psychology*, 98, 281-300.

Vazire, S., & Carlson, E.N. (2010). Self-knowledge of personality: Do people know themselves? *Social and Personality Psychology Compass*, 4/8, 605-620.

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

Wood, D., Nye, C.D., & Saucier, G. (2010). Identification and measurement of a more comprehensive set of person-descriptive trait markers from the English lexicon. *Journal of Research in Personality*, 44, 258-272.

Table 1

Descriptive statistics for the Moral Character Questionnaire, *N*=173 Targets, 493 informants

Moral trait scale	Reliability		Mean(SD)	
	Target	Informant	Target	Informant
Fairness	.66	.78	3.92(.68)	4.08(.48)
Honesty	.52	.73	4.01(.43)	4.18(.34)
Compassion	.75	.78	4.39(.23)	4.39(.21)
Temperance	.76	.87	3.23(.22)	3.58(.21)
Moral concern	.71	.78	4.12(.43)	4.16(.28)
General morality	.86	.85	4.40(.12)	4.56(.09)
Global Moral Character	.84	.89	4.01(.43)	4.15(.33)

Table 2

Correlations among the six moral traits and extraction communality estimates, $n= 173$
(target-reports only)

	Fairness	Honesty	Compassion	Temperance	Moral Concern	Extraction communality
Fairness						.55
Honesty	.57***					.50
Compassion	.51***	.46***				.41
Temperance	.45***	.38***	.17*			.24
Moral Concern	.50***	.45***	.48***	.39***		.53
General Morality	.54***	.60***	.57***	.40***	.68***	.71

Table 3. Estimates of trait-level agreement, residual agreement, and agreement controlling for liking [95% CIs] for self-other and inter-informant agreement in the overall sample (173 Targets, 493 Informants)

	Global Moral Character	Facets of Character					
		Fairness	Compassion	Honesty	Temperance	Moral concern	General morality
Self-other	.36*** [.24,.48]	.30*** [.19,.41]	.15** [.03,.27]	.20*** [.09,.30]	.42*** [.28,.55]	.28*** [.16,.39]	.26*** [.15,.37]
Residual	--	.18** [.06,.29]	.07 [-.03,.18]	.14** [.04,.24]	.32*** [.20,.44]	.02 [-.08,.11]	.17*** [.07,.26]
w/o Liking	.36*** [.23,.48]	.29*** [.17,.41]	.15* [.03,.26]	.20*** [.10,.31]	.43*** [.29,.56]	.26*** [.14,.37]	.25*** [.13,.37]
Inter-informant	.24*** [.13,.36]	.20*** [.08,.31]	.22*** [.11,.34]	.09† [.01,.18]	.35*** [.22,.48]	.21*** [.09,.32]	.15** [.05,.25]
Residual	--	.15** [.04,.25]	.11* [.01,.22]	.04 [-.05,.13]	.23*** [.11,.35]	.02 [-.07,.10]	-.01 [-.09,.06]
w/o Liking	.28*** [.16,.41]	.19*** [.08,.31]	.25*** [.13,.37]	.14** [.04,.24]	.34*** [.21,.47]	.21*** [.10,.33]	.18** [.07,.28]

 Note. † $p < .10$, * $p < .05$, ** $p < .01$, *** $p < .001$

Table 4

Trait-level agreement [95% CI] for community and student subsamples.

Subsample	Moral Character Trait						
	Global Moral Character	Fairness	Compassion	Honesty	Temperance	Moral concern	General morality
<u>Community</u> (100 Targets, 283 Informants)							
Self-other	.45*** [.27,.62]	.37*** [.21,.54]	.21** [.05,.37]	.21** [.06,.36]	.47*** [.28,.66]	.32*** [.16,.48]	.28*** [.12,.44]
Inter-informant	.36*** [.19,.53]	.25*** [.10,.41]	.27*** [.10,.43]	.17* [.03,.31]	.51*** [.31,.72]	.26*** [.10,.42]	.22** [.07,.38]
<u>Student</u> (73 Targets, 210 Informants)							
Self-other	.23** [.07,.39]	.18* [.02,.34]	.07 [-.09,.23]	.13† [-.00,.27]	.30*** [.15,.46]	.25** [.09,.42]	.25** [.10,.41]
Inter-informant	.08 [-.06,.22]	.10 [-.05,.24]	.15† [-.00,.31]	.00 [-.11,.12]	.04 [-.08,.16]	.12 [-.03,.27]	.08 [-.06,.21]

Note. † $p < .10$, * $p < .05$, ** $p < .01$, *** $p < .001$

Table 5

Mean agreement on moral profiles [95% Confidence Interval]

Sample/subsample	Type of Profile Agreement	
	Overall	Distinctive
<u>Overall Sample</u> (173 Targets, 493 Informants)		
Self-other	.73***	.32***
95% CI	[.67,.78]	[.22,.42]
Inter-informant	.66***	.30***
95% CI	[.56,.74]	[.16,.43]
<u>Community</u> (100 Targets, 283 Informants)		
Self-other	.76***	.41***
95% CI	[.68,.83]	[.27,.53]
Inter-informant	.72***	.41***
95% CI	[.59,.82]	[.25,.56]
<u>Student</u> (73 Targets, 210 Informants)		
Self-other	.67***	.20*
95% CI	[.59,.75]	[.03,.35]
Inter-informant	.57***	.14
95% CI	[.40,.70]	[-.08,.36]

Note. * $p < .05$, ** $p < .01$, *** $p < .001$