

Do People Want to Be More Moral?



Jessie Sun¹  and Geoffrey P. Goodwin²

¹Department of Psychology, University of California, Davis, and ²Department of Psychology, University of Pennsylvania

Psychological Science
2020, Vol. 31(3) 243–257
© The Author(s) 2019
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/0956797619893078
www.psychologicalscience.org/PS



Abstract

Most people want to change some aspects of their personality, but does this phenomenon extend to moral character and to close others? Targets ($n = 800$) rated their personality traits and reported how much they wanted to change on each trait; well-acquainted informants ($n = 958$) rated targets' personality traits and how much they wanted the targets to change on those same traits. Targets and informants reported a lower desire to change more morally relevant traits (e.g., honesty, compassion, fairness) compared with less morally relevant traits (e.g., anxiety, sociability, productiveness)—even after we controlled for current trait levels. Moreover, although targets generally wanted to improve more on traits that they had less desirable levels of, and informants wanted their targets to improve more on those traits as well, targets' moral change goals were less calibrated to their current trait levels. Finally, informants wanted targets to change in similar ways, but to a lesser extent, than targets themselves did. These findings suggest that moral considerations take a back seat when it comes to self-improvement.

Keywords

volitional personality change, moral character, self-improvement, moral motivations, self–other agreement, open data, open materials, preregistered

Received 5/10/19; Revision accepted 10/16/19

Most people want to change some aspects of their personality—across adulthood, more than 78% of people express desires to be more extraverted, emotionally stable, conscientious, agreeable, or open to experience (Hudson & Fraley, 2016b). To date, personality change goals have primarily been examined through the lens of the Big Five framework and the goals that people have for themselves (Baranski, Morse, & Dunlop, 2017; Hudson & Fraley, 2015; Hudson & Roberts, 2014). Here, we examined whether the common desire for personality change extends to moral character and to close others. In other words, do people want to be more moral, and do people who know us well want us to change in similar ways?

Personality traits describe relatively stable, enduring patterns of thoughts, feelings, and behaviors that make people different from one another (DeYoung, 2015). Moral character traits are personality traits that capture individual differences in the tendency to follow relevant moral standards—essentially, the extent to which someone is a “good” or “bad” person (Fleeson, Furr, Jayawickreme, Meindl, & Helzer, 2014). Traits differ in the degree to which

they are seen as being relevant to moral character (but note that such judgments can vary across cultures and individuals; Meindl & Graham, 2014). For example, at least in Western contexts, extraversion is considered a *nonmoral trait* (i.e., not particularly morally relevant), whereas honesty is considered a *moral trait* (i.e., highly morally relevant; Meindl & Graham, 2014). Thus, being generally introverted does not make someone a “bad” person, whereas being generally dishonest might.

People place great value on morality in themselves and in others; indeed, some research suggests that feeling moral might be a basic psychological need (Prentice, Jayawickreme, et al., 2019). Moral character traits are also among the most powerful determinants of the overall impressions we form of other people (Goodwin, 2015; Goodwin, Piazza, & Rozin, 2014; Hartley et al., 2016). Thus, if people recognize that their moral

Corresponding Author:

Jessie Sun, University of California, Davis, Department of Psychology,
200 East Quad, Davis, CA 95616
E-mail: jesun@ucdavis.edu

character influences how much other people like and respect them, they might want to become more moral not only to reach their own moral standards but also to attain the reputational benefits of being viewed as a moral person.

However, there are equally compelling reasons to believe that people might be less motivated to change their moral traits (compared with their nonmoral traits). First, because people appear to inaccurately inflate their positive moral qualities to a greater extent compared with their nonmoral qualities (Tappin & McKay, 2017), they might see no need to further enhance their moral traits. Second, people might typically be more inclined to make changes that will assuage their dissatisfaction with aspects of their lives (Baumeister, 1994; Hudson & Fraley, 2016a). For example, people who are dissatisfied with their friendships might strive to become more extraverted to overcome these social difficulties (Hudson & Roberts, 2014). In contrast, insofar as morality involves overcoming selfish impulses for the benefit of other people (Baumeister & Exline, 1999), the perceived costs of becoming more moral might outweigh its reputational benefits, dissuading people from striving for moral self-improvement.

Further complicating the picture is the fact that moral traits are seen as especially central to personal identity (for a review, see Strohminger, Knobe, & Newman, 2017). Whereas some research suggests that people are (hypothetically) less willing to take pharmaceuticals to enhance traits that are more fundamental to identity (e.g., empathy, kindness; Riis, Simmons, & Goodwin, 2008), other studies suggest that hypothetical improvements to moral traits might be seen as bringing people closer to their “true selves” (Bench, Schlegel, Davis, & Vess, 2015; Christy, Kim, Vess, Schlegel, & Hicks, 2017).

Little empirical evidence addresses the question of whether people prefer to change moral or nonmoral traits. In one study, in which participants selected five characteristics that they most wanted to improve, the probability of wanting a moral trait to improve was 21% (vs. 47% for nonmoral traits; Molouki & Bartels, 2017). However, because the researchers did not design the study to compare moral and nonmoral personality change goals, they did not comprehensively sample moral and nonmoral traits. Therefore, our primary aim was to comprehensively examine the extent to which people want to change more or less morally relevant traits.

Our second goal was to document the ways in which other people want us to change our personality traits (*interpersonal* change goals). Because people who know us well have a unique perspective on our personalities (Vazire & Carlson, 2011), these close others could shape our personality change goals—and ultimately our

personalities—by pointing out areas for improvement that may or may not have occurred to us (Bollich, Johannet, & Vazire, 2011). Consistent with this idea, past research found that when informants agreed with targets’ assessments of themselves as having low levels of agreeableness and extraversion, targets reported stronger desires to improve these traits, compared with when informants disagreed with targets’ assessments of their current levels of these traits (Quintus, Egloff, & Wrzus, 2017). However, to our knowledge, no study has asked informants about how they would like targets to *change*. Doing so would shed light on whether the people who know us well want us to change in ways that are similar to how we want ourselves to change. Evidence of self–other agreement on personality change goals would show that close others can corroborate our own assessments of which traits are most in need of improvement, whereas self–other disagreement might suggest that close others either have unique insight into our shortcomings or want us to improve in ways that would benefit them.

Our final goal was to explore individual differences in the desire for moral self-improvement. People who have less desirable levels of a given Big Five trait generally want to improve more on that trait (Baranski et al., 2017; Hudson & Fraley, 2015; Hudson & Roberts, 2014), but these correlations have been more consistent for extraversion, conscientiousness, and emotional stability than for agreeableness and openness. Here, we examined whether less desirable trait levels are associated with a greater desire to improve moral traits or whether, alternatively, the most moral people are the ones who most want to improve their moral traits (i.e., whether virtue begets virtue in the moral domain).

Method

We recruited two samples of participants. We preregistered stopping rules for both samples and an initial exploratory analysis plan for Sample 1 on the Open Science Framework (OSF) at <https://osf.io/qjfnw/>. After freely exploring the data in Sample 1, we preregistered our analysis plan for Sample 2 at <https://osf.io/6mdwp/> and directly replicated our Sample 1 analyses. We determined our sample sizes (300 targets in Sample 1; 500 targets in Sample 2) on the basis of time and budget constraints (and however many informants would respond); the resulting sample sizes exceeded recommendations for obtaining stable estimates of the average published effect in personality and social psychology ($N > 250$; Schönbrodt & Perugini, 2013). We report all data exclusions and all measures (measures not included in the current article are reported in our OSF repository). All materials and the data and scripts needed to

reproduce the analyses are available at <https://osf.io/cbxjh>.

Participants and procedure

We recruited two samples of target participants from the undergraduate psychology research pools at the University of Pennsylvania (Sample 1) and the University of California, Davis (Sample 2). Apart from minor differences described below, we used the same data-collection procedures for the two samples. Targets completed an online questionnaire (implemented using Qualtrics software; <https://www.qualtrics.com/>) in which they nominated up to four informants who knew them well, then self-reported their personality traits and change goals (described below). When nominating informants, targets were asked to choose at least two people who were not current romantic partners or family members. We invited informants to participate in the study by e-mailing them a unique link to an online questionnaire. Informants rated their target's personality traits and reported the personality changes they wanted to see in their target. We sent informants who had not yet completed the survey three (Sample 1) to six (Sample 2) e-mail reminders, spaced approximately 1 week apart. In both samples, we used all available data for analyses involving only target self-reports (even if the target had no informant reports).

Targets participated in exchange for course credit. Informants who completed the survey were entered into a prize drawing for a 1 in 10 chance of winning a \$20 Amazon.com gift card. In Sample 2, informants could also ask to be entered into the prize drawing even if they did not participate (in compliance with California law). Data-collection procedures were approved by the institutional review board (IRB) at the University of Pennsylvania (Sample 1; IRB ID: 831767) and the University of California, Davis (Sample 2; IRB ID: 1328211-2).

Sample 1. In Sample 1, 300 targets (224 female, 74 male, 2 other or not disclosed) between the ages of 18 and 29 years (age: $M = 19.57$, $SD = 1.29$) completed the study and met our preregistered inclusion criteria (we excluded 4 participants who were under the age of 18). Targets identified as White/Caucasian ($n = 143$), Asian ($n = 85$), Hispanic/Latino ($n = 24$), Black/African American ($n = 20$), Pacific Islander ($n = 1$), other/multiple ($n = 25$), or did not disclose their ethnicity ($n = 2$). Targets successfully nominated 1,023 informants, of whom 417 (288 female, 124 male, 5 other or not disclosed; age: range = 18–98 years, $M = 28.54$, $SD = 15.22$) completed the study (response rate = 41%). Of the 221 targets who had informants, 92 had 1 informant and 129 had 2 or more informants.

Informants reported being friends of the target ($n = 234$), parents ($n = 97$), current romantic partners ($n = 37$), siblings ($n = 32$), other family members ($n = 8$), former romantic partners ($n = 3$), or coworkers ($n = 1$), and 5 reported their relationship with the target as “other.” On average, informants reported having known their targets for 9.39 years ($SD = 8.17$). Following our preregistered stopping rules, we ended data collection for targets when 300 targets (who met our inclusion criteria) completed the study and ended data collection for informants 1 week after the third e-mail reminder had been sent to the last informants who were nominated.

Sample 2. Because informant response rates were lower in Sample 2, we deviated from our preregistered sampling plan by (a) sending three more e-mail reminders to informants nominated by the first 300 targets and (b) recruiting 200 more targets (prior to analyzing the Sample 2 data). Thus, Sample 2 comprised 500 targets (404 female, 93 male, 3 other or not disclosed; age: range = 18–47, $M = 19.85$, $SD = 2.55$) who completed the study and met our preregistered inclusion criteria (none were excluded). Targets identified as White/Caucasian ($n = 102$), Asian ($n = 236$), Hispanic/Latino ($n = 89$), Black/African American ($n = 3$), Pacific Islander ($n = 3$), other/multiple ($n = 64$), or did not disclose their ethnicity ($n = 3$). Targets successfully nominated 1,464 informants, of whom 541 (396 female, 135 male, 10 other or not disclosed; age: range = 18–81, $M = 25.63$, $SD = 12.56$) completed the study (response rate = 37%). Of the 288 targets who had informants, 122 had 1 informant and 166 had 2 or more informants. Informants reported being friends of the target ($n = 334$), parents ($n = 80$), current romantic partners ($n = 30$), siblings ($n = 57$), other family members ($n = 20$), former romantic partners ($n = 3$), or coworkers ($n = 1$), and 8 reported their relationship with the target as “other.” On average, informants reported having known their targets for 8.95 years ($SD = 7.85$). We ended data collection 1 week after we sent the third e-mail reminder to the informants who were nominated by the last 200 targets.

Measures

Targets rated their current standing on 21 personality traits and reported the extent to which they wanted to change on each of these traits, using the measures described below. Informants rated their target's personality traits and reported the extent to which they wanted their target to change on each of these traits, using the same measures.

Personality traits.

Big Five facets. The Big Five Inventory-2 (BFI-2; Soto & John, 2017) measures 15 facets—3 facets for each of the

Big Five personality domains (extraversion: sociability, assertiveness, energy level; agreeableness: compassion, respectfulness, trust; conscientiousness: organization, productiveness, responsibility; negative emotionality: anxiety, depression, emotional volatility; and open-mindedness: intellectual curiosity, aesthetic sensitivity, creative imagination). Each facet is measured with four items (60 items total). We also assessed an additional facet of negative emotionality (anger), using two items from the Big Five Aspects Scale (DeYoung, Quilty, & Peterson, 2007): “Gets angry easily” and “Is not easily annoyed” (the latter was reverse scored). Targets rated the extent to which they agreed with 62 statements (e.g., “I am someone who is outgoing, sociable”) on a 5-point scale (1 = *disagree strongly*, 5 = *agree strongly*). Informants rated the extent to which they agreed with the same statements about the target (e.g., “[target’s name] is someone who is outgoing, sociable”). We dropped one item from the aesthetic-sensitivity measure (“Values art and beauty”) for reasons described in the next section.

Moral Characteristics Questionnaire (MCQ). The BFI-2 contains some morally relevant content, especially for the agreeableness (e.g., “Is helpful and unselfish with others”) and conscientiousness (e.g., “Is reliable, can always be counted on”) domains. However, other aspects of moral character are not well captured by the BFI-2. Thus, we used the MCQ (Prentice, Furr, & Hawkins, 2019) to measure general morality (e.g., “I am a person of strong moral character”; four items) and the specific domains of honesty (e.g., “I consistently tell the truth”), fairness (e.g., “I treat people fairly”), loyalty (e.g., “I shift my loyalties easily”; reverse scored), and purity (“I would say that I’m a wholesome person, relatively ‘pure’”; two items per domain). Targets rated the extent to which they agreed with such statements about themselves, and informants rated the extent to which they agreed with such statements about the target, on a 5-point scale (1 = *strongly disagree*, 5 = *strongly agree*).

The original measures included two additional items for each domain (i.e., six items for general morality and four items for each of the other domains). However, whereas all but one of the BFI-2 items describe current tendencies, two items from each of the MCQ scales captured values and moral strivings (e.g., “I don’t believe that honesty is that important,” “I want to be honest even when it’s hard”). To avoid a confound when examining the correlations between traits and change goals (which, by definition, capture strivings), we dropped the MCQ items that captured values and strivings (and the “values art and beauty” BFI-2 item) and included only the items that described current tendencies (e.g., “I tend to act morally”) and overall self-perceptions (e.g., “I am an honest person”), as described in the Sample 2 preregistration.

Personality change goals.

Change-goals scale. We measured change goals by modifying two items for each of the 21 personality traits (16 Big Five facets and 5 MCQ domains) described above. Following Hudson and Roberts (2014), we reworded the instructions, items, and response scales in terms of how much targets wanted to change, and how much informants wanted targets to change, on each personality trait (see the codebooks at <https://osf.io/rbeuw/> and <https://osf.io/87dkh/> for full item wordings). For example, “I am someone who is helpful and unselfish with others” was reworded to “I want to be helpful and unselfish with others” (self-report) or “I want [target’s name] to be helpful and unselfish with others” (informant report), with response options indicating the magnitude and direction of the desired changes (−2 = *much less than I currently am*, −1 = *less than I currently am*, 0 = *I do not want to change in this trait*, 1 = *more than I currently am*, 2 = *much more than I currently am*; we adjusted the pronouns and grammar as needed for different items and for the informant reports).

We computed the mean of the two change-goal items for each trait, separately for targets and informants. We also extracted an index of the overall desire for change by taking the average of the absolute scores across the 42 change-goal items, separately for targets and informants (resulting in a continuous measure that had a possible range of 0, *no change desired on any of the 42 items*, to 2, *“much more” or “much less” on all 42 items*).

Change-goal priorities. To get a better sense of which personality changes people would most prioritize, we showed targets and informants all of the desired changes they had selected from the full set of change goals (i.e., excluding the items for which they selected the response, “I do not want [target’s name] to change in this trait”) and asked them to select the top three most desired changes (“Which three changes would you most like to see in yourself/[target’s name]?”). On average, targets selected priorities out of a list of 26.44 desired changes (Sample 1; $SD = 8.44$) and 27.63 desired changes (Sample 2; $SD = 9.32$), and informants selected priorities out of a list of 12.40 desired changes (Sample 1; $SD = 7.95$) and 12.86 desired changes (Sample 2; $SD = 9$). See the Supplemental Material (Section 1) for a description of the order in which this list was presented.

From the top three priorities, we computed 42 binary variables that represented whether or not a target or informant (respectively) prioritized a goal to improve each of the 21 traits as one of their top three priorities. Each of these 21 traits has an unambiguous “positive pole.” For instance, previous research shows that most people want to be more extraverted, conscientious, agreeable, and open to experience and to be less

neurotic—and fewer than 3% want to change in the opposite direction (Hudson & Roberts, 2014). This indicates that most people think that low neuroticism and high levels of the other Big Five traits are desirable and that changes toward these positive poles are improvements. Similarly, our moral-valence norms (described in the next section) showed that people generally consider general morality, honesty, fairness, loyalty, and purity to be morally positive. Thus, we operationalized “improvements” as change goals in the desirable direction (i.e., decreases on facets of negative emotionality and increases on all other traits), which were coded as 1. Not prioritizing a goal was coded as 0, as were goals to change in the undesirable direction (because very few participants prioritized goals to change in the undesirable direction). Out of these top three priorities, we also asked targets and informants to select and provide an open-ended justification for their top change goal (which we used in supplemental analyses; see the Supplemental Material, Section 7).

Moral relevance. To examine whether change goals and other effects varied depending on the trait’s moral relevance, we obtained norming data for the 42 items included in our change-goals scale. We recruited two separate samples of trait raters from the same subject pools that the targets were drawn from (and prevented target participants from signing up for the trait-norming study). Each trait rater was randomly assigned to one of two versions of the norming task (for a description of additional exploratory dimensions, see <https://osf.io/me9yp/>). Only the first version of the task involved rating moral relevance. Thus, the trait raters included in this study were 114 University of Pennsylvania undergraduates (81 female, 32 male, 1 not disclosed; age: $M = 19.53$ years, $SD = 1.20$) and 203 University of California, Davis, undergraduates (165 female, 36 male, 2 not disclosed; age: $M = 20.30$ years, $SD = 3.32$).

Trait raters rated the moral valence of the 42 items in our change-goals scale. Specifically, after reading a brief explanation of moral traits, participants were asked, “How morally good or morally bad is it to be high on each of the following traits?” (e.g., “Having a forgiving nature”; $-3 = \textit{very morally bad}$, $-2 = \textit{moderately morally bad}$, $-1 = \textit{slightly morally bad}$, $0 = \textit{neither morally good nor morally bad}$, $1 = \textit{slightly morally good}$, $2 = \textit{moderately morally good}$, $3 = \textit{very morally good}$). To index the moral relevance of each of the 21 traits, we computed the absolute value of the moral-valence rating for each item for each rater ($0 = \textit{not morally relevant}$, $3 = \textit{very morally relevant}$). We then averaged the two items for each trait before computing the average moral-relevance rating for each trait across all raters. The raters also rated the perceived changeability of each trait, which we used for supplemental analyses

(described in the Supplemental Material, Section 3). Descriptive statistics for the norming data are reported in Table S1 in the Supplemental Material.

We conducted analyses for each sample using the norms for that sample. The between-traits analyses (e.g., correlations between moral relevance and average absolute change goals) involved a relatively small set of 21 traits. This provided 80% power to detect correlations greater than $|.57|$ but limited our ability to detect smaller correlations (and therefore to draw conclusions about null effects for the between-traits analyses).

Data analyses

Most analyses were conducted in the R programming environment (R Core Team, 2018); for some supplemental analyses, we used Mplus Version 8.3 (Muthén & Muthén, 2017). We used two complementary approaches: (a) the full change-goals scale and (b) the top three priorities. The first method captures the direction and amount of change that participants desired for each trait, whereas the second method summarizes the improvements that participants most prioritized. For each approach, we report descriptive statistics (means for the scale, frequencies for the priorities) and examine the associations between trait levels and change goals (correlations for the full scale are presented below; odds ratios for the priorities are in the Supplemental Material, Section 5).

For analyses involving informant reports, we computed an aggregate score across all informants for a given target, with two exceptions for our supplemental analyses: (a) We used all open-ended responses when examining reasons for the personality change that informants most prioritized, and (b) we randomly selected one informant per target for the logistic regression analyses predicting informant-reported priorities from trait levels. For scale-reliability estimates, we computed the ω coefficient for scales that had three or more items and α for scales that had two items (using the *MBESS* package for R; Kelley, 2018). We repeated this procedure to compute scale reliabilities for the informant reports after aggregating scores for each item across all informants for a given target. For the informant-reported measures, we computed the intraclass correlation coefficient, ICC(1), from a random-effects model for targets who had two or more informants (using the *lme4* package; Bates, Mächler, Bolker, & Walker, 2015). This represents the proportion of variation in the informant reports due to variation between targets.

For most of the key analyses (described below), we estimated latent correlations using structural equation models, implemented via *lavaan* (Rosseel, 2012), to ensure that any differences in effect sizes across traits were not due to differences in measurement reliability

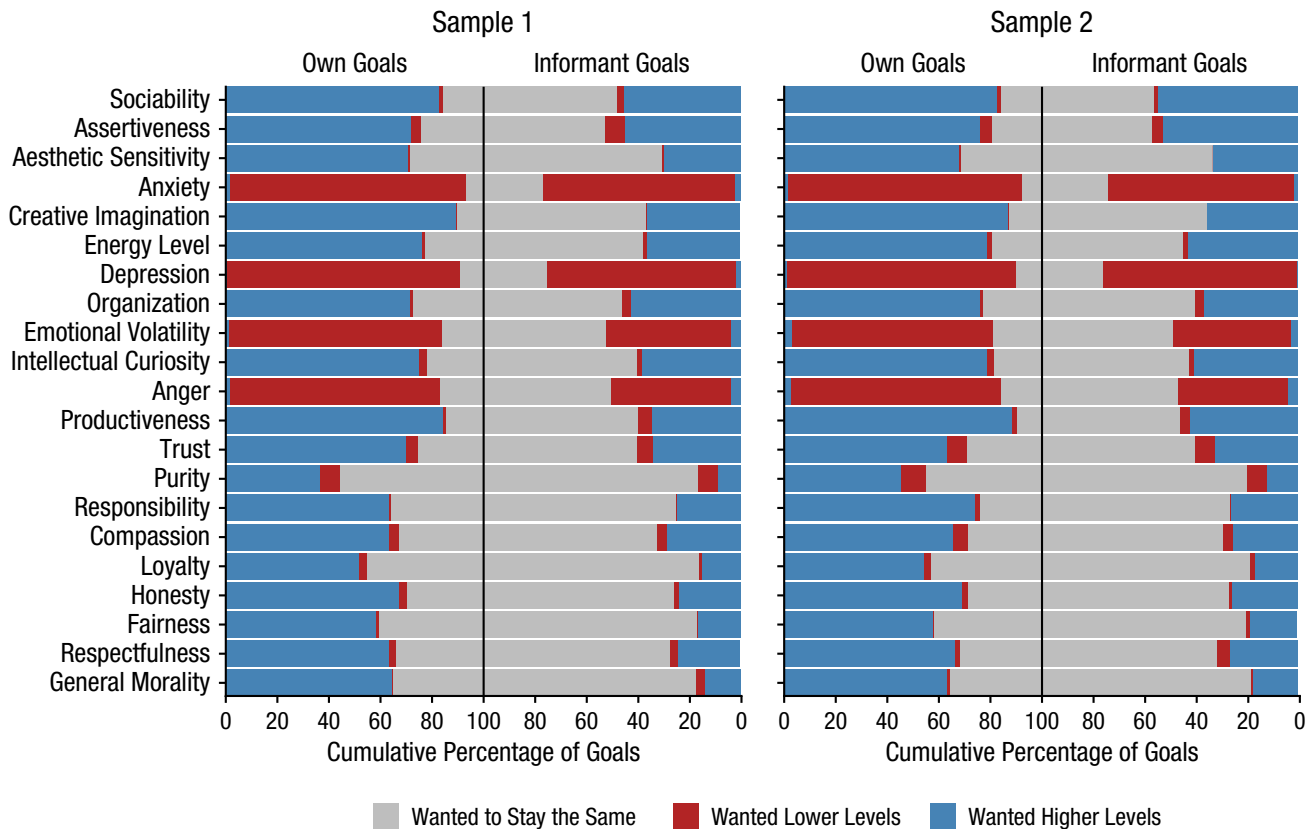


Fig. 1. Categorical summary of ratings on the change-goals scale, separately for Samples 1 and 2. Stacked bars on the left of each graph show the percentage of targets who wanted higher levels, lower levels, or the same level of each trait. Stacked bars on the right of each graph show the percentage of individual (i.e., not aggregated) informants who wanted their target to have higher levels, lower levels, or the same level of each trait. To facilitate visual comparison, we have ordered the traits from least to most morally relevant on the basis of the average of the moral-relevance norms across the two samples (weighted equally). However, in the analyses for each sample, we used the moral-relevance norms for the respective sample.

(because some traits were measured with two items and others with four items). For these structural equation models, when there were only two indicators, we constrained their factor loadings to be equal. Average effect sizes and comparisons of effect sizes were computed after appropriate transformations (described in the Supplemental Material, Section 1).

Results

Descriptive statistics and self–other agreement correlations for the key measures are shown in Table 1 (for personality traits) and Table 2 (for change goals). The full correlation matrices for personality traits (Tables S6–S7) and change goals (Tables S8–S9) are available in our OSF repository (<https://osf.io/cybtX/>).

Do people want to be more moral?

As shown in Table 2 and Figures 1 and 2, targets showed the strongest desires to be less anxious, depressed, emotionally volatile, and angry (i.e., all the facets of negative

emotionality that we measured) and to be more creative, productive, and sociable. In contrast, targets reported weaker desires to change on traits that were more morally relevant (e.g., honesty, general morality, compassion, fairness). To further understand this pattern, we converted each participant's change goal for each trait into an absolute value, then averaged those values across participants to compute the average amount of desired change for each trait. This procedure placed equal weight on goals to change in either direction (e.g., goals to become more or less compassionate both contributed to higher average absolute change goals for this analysis). We then correlated the average absolute change goal for each trait with the traits' moral-relevance scores (rated by a separate sample of participants). This showed that on average, targets showed a weaker desire to change on traits that were more morally relevant, Sample 1: $r(19) = -.69$, 95% confidence interval (CI) = $[-.86, -.37]$, $p < .001$; Sample 2: $r(19) = -.62$, 95% CI = $[-.83, -.26]$, $p = .003$.

Figure 2, which summarizes participants' top three most desired changes for the target, makes the targets' personality-change priorities even clearer. For these

Table 1. Descriptive Statistics for Personality Traits

Trait	Sample 1						Sample 2									
	Self-report			Informant report			Self-report			Informant report						
	<i>M</i>	<i>SD</i>	Rel.	<i>M</i>	<i>SD</i>	Rel.	<i>M</i>	<i>SD</i>	Rel.	<i>M</i>	<i>SD</i>	Rel.	<i>r</i>			
Sociability	3.09	1.01	.86	3.49	0.94	.88	.61	.74***	2.82	0.91	.82	3.18	0.86	.84	.46	.63***
Assertiveness	3.36	0.88	.80	3.52	0.80	.83	.40	.55***	3.03	0.90	.79	3.23	0.74	.79	.43	.51***
Energy level	3.70	0.78	.72	4.04	0.69	.73	.43	.58***	3.43	0.79	.66	3.79	0.72	.75	.35	.50***
Compassion	3.93	0.72	.57	4.10	0.70	.72	.16	.54***	3.78	0.69	.47	4.00	0.61	.61	.12	.39***
Respectfulness	3.98	0.67	.62	4.26	0.62	.75	.28	.56***	3.97	0.65	.59	4.22	0.64	.76	.26	.18*
Trust	3.18	0.79	.70	3.61	0.73	.77	.18	.48***	3.27	0.75	.63	3.55	0.72	.75	.28	.44***
Organization	3.80	0.90	.82	3.79	0.95	.90	.51	.66***	3.83	0.85	.81	3.83	0.90	.89	.40	.57***
Productiveness	3.51	0.86	.76	4.12	0.71	.78	.40	.59***	3.31	0.79	.71	3.90	0.72	.77	.35	.56***
Responsibility	3.56	0.69	.59	4.00	0.67	.74	.25	.24**	3.52	0.66	.51	3.91	0.69	.70	.26	.33***
Anxiety	3.66	0.86	.79	3.35	0.87	.84	.31	.57***	3.64	0.84	.74	3.29	0.75	.78	.26	.47***
Depression	2.73	0.93	.80	2.39	0.91	.88	.45	.62***	2.86	0.93	.79	2.50	0.82	.84	.27	.40***
Emotional volatility	2.91	0.98	.84	2.63	0.96	.87	.41	.55***	2.97	0.95	.81	2.58	0.94	.89	.25	.43***
Anger	2.94	1.02	.62	2.60	0.94	.73	.31	.60***	2.94	1.01	.56	2.68	0.91	.71	.28	.44***
Intellectual curiosity	4.08	0.66	.68	4.13	0.65	.76	.25	.52***	3.79	0.73	.72	3.90	0.66	.75	.33	.62***
Aesthetic sensitivity	3.61	1.00	.73	3.69	0.88	.82	.44	.71***	3.44	0.94	.62	3.57	0.81	.70	.39	.65***
Creative imagination	3.76	0.76	.77	3.98	0.60	.78	.20	.44***	3.56	0.76	.68	3.91	0.62	.73	.18	.20**
General morality	4.06	0.60	.77	4.47	0.46	.73	.05	.15	3.99	0.57	.69	4.35	0.47	.70	.12	.10
Honesty	3.94	0.69	.75	4.37	0.58	.76	.05	.29***	3.89	0.74	.77	4.31	0.59	.74	.13	.13
Fairness	4.16	0.54	.67	4.47	0.50	.75	.17	.12	4.18	0.59	.68	4.44	0.50	.74	.13	.10
Loyalty	4.25	0.62	.49	4.41	0.55	.68	.06	.24*	4.15	0.67	.48	4.27	0.60	.60	.11	-.02
Purity	2.87	0.81	.40	3.85	0.84	.70	.19	.31*	2.96	0.75	.16	3.84	0.74	.64	.14	.39*

Note: "Rel." denotes the reliability estimate for each trait. We used α for the two-item honesty, fairness, loyalty, and purity scales and ω for the remaining three- to four-item scales. The intraclass correlation coefficient, ICC(1), shows the proportion of variability in informant ratings due to variability between targets (computed across informant ratings for targets who had two or more informants). Latent self-other agreement correlations (*r*s) were estimated using structural equation models.

* $p < .05$. ** $p < .01$. *** $p < .001$. (uncorrected for multiple comparisons)

Table 2. Descriptive Statistics for the Change-Goals Scale

Trait	Sample 1						Sample 2									
	Self-report			Informant report			Self-report			Informant report						
	<i>M</i>	<i>SD</i>	α	<i>M</i>	<i>SD</i>	α	<i>M</i>	<i>SD</i>	α	<i>M</i>	<i>SD</i>	α	<i>r</i>			
Anxiety	-1.18	0.63	.56	-0.79	0.61	.79	.30	.59***	0.67	0.67	.65	-0.78	0.58	.70	.16	.46***
Depression	-1.12	0.61	.53	-0.77	0.59	.60	.29	.61***	0.67	0.67	.58	-0.83	0.59	.62	.22	.52***
Productiveness	0.99	0.64	.68	0.24	0.43	.60	.30	.30**	0.67	0.67	.50	0.35	0.45	.59	.09	.25**
Creative imagination	1.04	0.57	.78	0.30	0.41	.82	.02	.13	0.63	0.63	.75	0.32	0.39	.74	.10	.12
Anger	-0.90	0.66	.69	-0.40	0.49	.79	.18	.33***	0.74	0.74	.75	-0.42	0.52	.80	.13	.31***
Emotional volatility	-0.96	0.64	.56	-0.47	0.54	.70	.29	.59***	0.73	0.73	.52	-0.46	0.53	.62	.06	.46***
Sociability	0.88	0.61	.69	0.41	0.52	.80	.34	.37***	0.67	0.67	.63	0.55	0.50	.73	.15	.23**
Organization	0.79	0.65	.84	0.33	0.47	.72	.20	.39***	0.66	0.66	.81	0.35	0.52	.89	.27	.33***
Energy level	0.74	0.55	.66	0.30	0.43	.77	.25	.40***	0.64	0.64	.67	0.38	0.45	.73	.18	.29***
Intellectual curiosity	0.69	0.60	.59	0.30	0.41	.68	.08	.10	0.62	0.62	.46	0.33	0.42	.57	.11	.31**
Assertiveness	0.67	0.61	.62	0.34	0.51	.75	.22	.41***	0.64	0.64	.53	0.44	0.52	.71	.16	.18*
Responsibility	0.62	0.60	.79	0.21	0.38	.86	.13	.08	0.64	0.64	.64	0.23	0.38	.75	.13	.18*
Honesty	0.60	0.58	.68	0.17	0.36	.78	.01	.07	0.65	0.65	.75	0.21	0.37	.79	.13	.29***
Aesthetic sensitivity	0.65	0.53	.59	0.27	0.40	.84	.01	.13	0.63	0.63	.69	0.32	0.42	.81	.16	.16*
Trust	0.60	0.60	.43	0.27	0.46	.69	.12	.43***	0.67	0.67	.36	0.24	0.43	.60	.07	.22*
Respectfulness	0.57	0.60	.64	0.16	0.33	.61	.10	.22*	0.67	0.67	.59	0.20	0.40	.66	.13	.30***
General morality	0.59	0.55	.75	0.08	0.30	.79	.17	.03	0.61	0.61	.69	0.15	0.34	.76	.17	.20**
Compassion	0.53	0.60	.62	0.22	0.40	.67	.15	.28**	0.64	0.64	.54	0.18	0.37	.65	.16	.27**
Fairness	0.52	0.55	.78	0.12	0.26	.71	.07	-.02	0.60	0.60	.77	0.15	0.33	.81	.11	.26***
Loyalty	0.40	0.54	.66	0.12	0.27	.84	.01	-.00	0.63	0.63	.75	0.16	0.36	.84	.12	.16*
Purity	0.26	0.50	.24	0.01	0.25	.21	.08	-.26	0.57	0.57	.08	0.03	0.31	.31	.10	.23

Note: The intraclass correlation coefficient, ICC(1), shows the proportion of variability in informant ratings due to variability between targets (computed across informant ratings for targets who had two or more informants). Latent self-other agreement correlations (*r*s) were estimated using structural equation models. Traits are ordered from the highest to lowest mean absolute desired change (averaged across both samples). The means reported in this table were computed from the raw change goals, but we used mean absolute change goals for the correlations with moral relevance (reported in the text) and perceived changeability (reported in the Supplemental Material, Section 3).

* $p < .05$. ** $p < .01$. *** $p < .001$. (uncorrected for multiple comparisons)

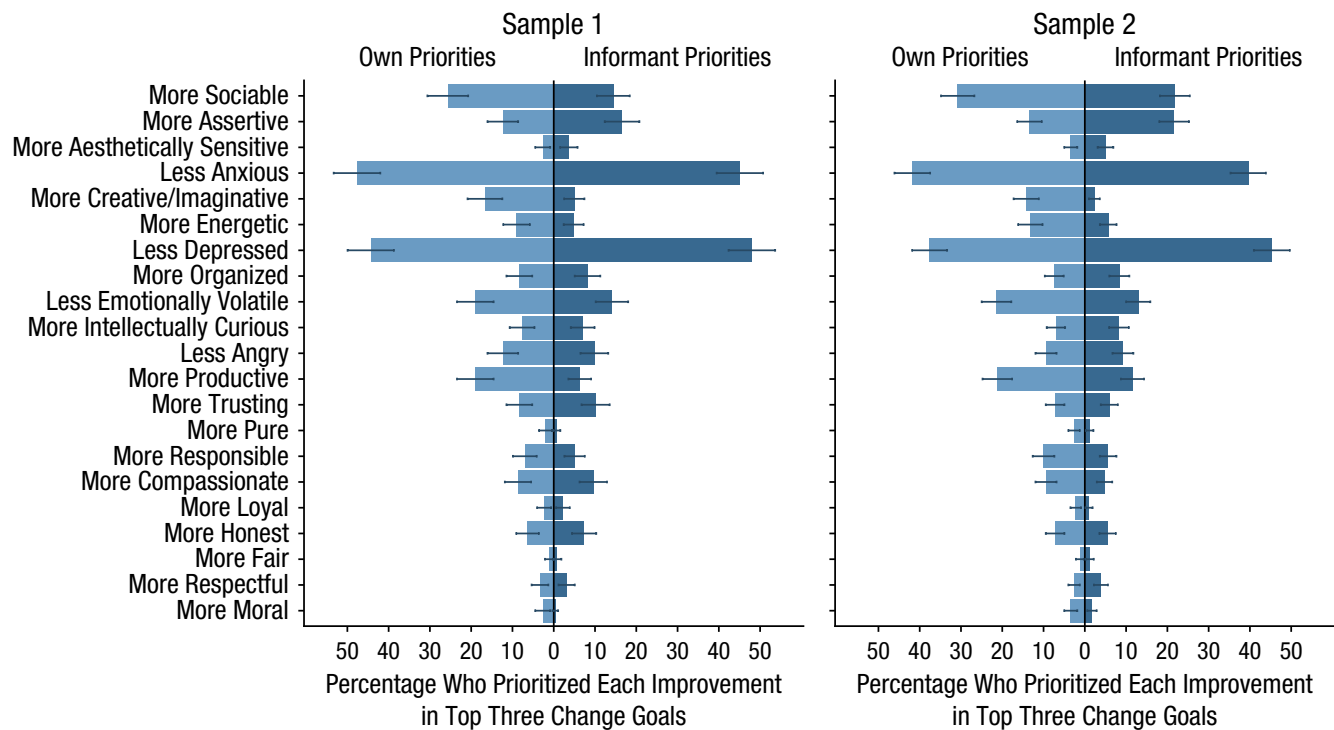


Fig. 2. Percentage of targets and informants who prioritized an improvement on each trait as one of their top three most desired changes in themselves or in their targets, respectively. Results are shown separately for Samples 1 and 2. To facilitate visual comparison, we have ordered the change goals from least to most morally relevant on the basis of the average of the moral-relevance norms across the two samples (weighted equally). However, in the analyses for each sample, we used the moral-relevance norms for the respective sample. Error bars depict 95% confidence intervals.

analyses, change goals that reflected a “worsening” of the trait (which fewer than 2% of participants prioritized for each trait) were not included. By focusing on participants’ desires to improve each trait, these analyses directly addressed the main question of whether people want to be *more* moral. We found that targets were less inclined to prioritize more morally relevant improvements, Sample 1: $r(19) = -.56$, 95% CI = $[-.80, -.17]$, $p = .009$; Sample 2: $r(19) = -.57$, 95% CI = $[-.80, -.18]$, $p = .007$. Instead, they focused on reducing negative emotionality. For example, a large proportion of the targets reported that becoming less anxious (Sample 1: 47.67%; Sample 2: 41.80%) or less depressed (Sample 1: 44.33%; Sample 2: 37.60%) were among their top three personality change goals. These were followed by goals to become more sociable (Sample 1: 25.67%; Sample 2: 30.80%), less emotionally volatile (Sample 1: 19%; Sample 2: 21.40%), and more productive (Sample 1: 19%; Sample 2: 21.20%). Moral improvements were rarely prioritized; for example, only about 9% and about 3% of targets prioritized a goal to become more compassionate or more generally moral, respectively.

Do close others want us to be more moral?

Next, we examined whether close others wanted their targets to change to a similar extent and in similar ways as the targets themselves did. First, we compared targets’ and informants’ average absolute change goals across the 42 items (scores on this index ranged from 0 to 2). Paired-samples t tests showed that on average, targets wanted to change themselves more (Sample 1: $M = 0.79$, $SD = 0.33$; Sample 2: $M = 0.86$, $SD = 0.38$) than their close others wanted them to change (Sample 1: $M = 0.35$, $SD = 0.22$, $g = 1.55$, 95% CI = $[1.33, 1.77]$; Sample 2: $M = 0.38$, $SD = 0.25$, $g = 1.50$, 95% CI = $[1.31, 1.68]$). This pattern can be seen most clearly in Figure 1: The gray bars (depicting the percentage of targets who did not want to change on each trait and the percentage of informants who did not want the target to change on each trait) are much larger for the informants across all traits. In addition, the average extent to which close others wanted the target to change was descriptively very similar for both friends (Sample 1: $M = 0.33$, $SD = 0.23$; Sample 2: $M = 0.36$, $SD = 0.27$) and parents

(Sample 1: $M = 0.34$, $SD = 0.23$; Sample 2: $M = 0.30$, $SD = 0.23$).

Although informants typically wanted the targets to change less than the targets themselves did, they reported a similar pattern of change goals across traits, across targets, and within targets. Across traits, targets and informants showed a similar pattern of personality-change priorities: The two sets of percentages (which included goals to change in the undesirable direction) were almost perfectly correlated, Sample 1: $r(40) = .95$, 95% CI = [.90, .97], $p < .001$; Sample 2: $r(40) = .93$, 95% CI = [.87, .96], $p < .001$. Crucially, informants also had lower desires to change more morally relevant traits, Sample 1: $r(19) = -.68$, 95% CI = [-.86, -.35], $p < .001$; Sample 2: $r(19) = -.71$, 95% CI = [-.87, -.40], $p < .001$, and were less likely to prioritize more morally relevant improvements for their targets, Sample 1: $r(19) = -.46$, 95% CI = [-.74, -.03], $p = .037$; Sample 2: $r(19) = -.53$, 95% CI = [-.78, -.13], $p = .013$.

We used two additional methods to examine the similarity between self- and informant-reported change goals. First, for each trait, we computed the latent correlation between self-reported change goals and informant-reported change goals (i.e., self-other agreement; see Table 2). These estimates were generally moderately positive (Sample 1: mean $r = .26$; Sample 2: mean $r = .28$), suggesting that targets and informants agreed to some extent on how much they wanted the target to be higher or lower on each trait. For example, if targets reported that they wanted to become much less anxious, their informants also tended to report greater goals to reduce the targets' anxiety (compared with targets who reported not wanting to reduce their anxiety).

Second, we computed profile correlations. Profile correlations describe the similarity between two sets of goals (i.e., across all 21 traits) as opposed to the between-persons agreement for a given trait. They therefore allow us to examine the extent of similarity between targets' idiosyncratic profiles of change goals and their informants' profiles of change goals for them. However, profile correlations can be positive simply because of normativeness effects (e.g., because the average person wants to be less depressed and more sociable; Furr, 2008). Because of such effects, it is theoretically possible that an "informant" who had never met a given target could report change goals that substantially overlap with that target's goals. Therefore, we compared the average overall profile correlation with a baseline based on many pseudosamples in which we randomly paired up each target's profile of 21 change goals with the profile of 21 change goals reported by a different target's informant (or informants). We then recomputed the profile correlations on the basis of 1,000 such pseudosamples (using the *multicon*

package; Sherman & Serfass, 2015). From this, we found that the mean overall correlation (r) for both samples was $.65$, $p < .001$. After removing normativeness effects (i.e., agreement due to the average change-goal profile), we still found a small amount of distinctive profile agreement (Sample 1: $r = .20$, $p < .001$; Sample 2: $r = .12$, $p < .001$). That is, targets and informants showed some agreement on the profile of changes that they wanted to see, over and above mere normativeness effects.

Who wants to be more moral?

Although people were generally less inclined to change on more morally relevant traits, some people showed a greater desire to change on more morally relevant traits (relative to other people). Thus, we explored the correlates of these individual differences in moral change goals.

Associations between traits and change goals. First, we examined whether change goals were generally calibrated to the targets' current traits. Figure 3 shows the latent correlations between current levels and change goals for each trait (see also Table S10 at <https://osf.io/cybtX/>). The correlations replicated the pattern observed in previous studies (e.g., Hudson & Roberts, 2014): In general, targets who reported having lower levels of a given trait wanted to increase more on that trait (Sample 1: mean $r = -.52$; Sample 2: mean $r = -.35$). However, the negative association between traits and change goals was smaller for more morally relevant traits, as shown by a strong positive association between moral-relevance scores and (Fisher r -to- z transformed) correlations between traits and change goals, Sample 1: $r(19) = .55$, 95% CI = [.16, .79], $p = .010$; Sample 2: $r(19) = .49$, 95% CI = [.07, .76], $p = .026$. That is, knowing someone's current standing on a more morally relevant trait provides relatively little information about whether they want to have higher or lower levels of that trait.

We then examined the associations between informant-reported current traits and change goals. A paired-samples t test showed that the negative association between traits and change goals was even stronger when both were informant-reported, compared with when both were self-reported, Sample 1: mean $r = -.79$, $t(20) = 5.97$, $p < .001$; Sample 2: mean $r = -.69$, $t(20) = 11.39$, $p < .001$. In addition, there was no evidence that the association between informant-reported traits and change goals was weaker for more morally relevant traits, Sample 1: $r(19) = .29$, 95% CI = [-.16, .64], $p = .198$; Sample 2: $r(19) = .33$, 95% CI = [-.12, .67], $p = .142$ (but note that we had relatively low power to detect these effects). However, there was no significant




Fig. 3. Associations between current traits and change goals, separately for targets and informants and for Samples 1 and 2. The traits are ordered from least to most morally relevant on the basis of the moral-relevance norms for each respective sample. Error bars depict 95% confidence intervals.

interaction between moral relevance and whether traits and change goals were both self-reported (the reference category) or both informant-reported (Sample 1: $b = -0.17$, 95% CI = $[-0.34, 0.01]$, $p = .083$; Sample 2: $b = -0.05$, 95% CI = $[-0.18, 0.09]$, $p = .507$).

We conducted conceptually similar logistic regression analyses, predicting change-goal priorities from current levels. This alternative approach (see the Supplemental Material, Section 5) showed that targets who had less desirable levels of a trait were more likely to prioritize improving that trait as one of their top three change goals and that informants' change goals remained relatively calibrated to their perceptions of targets' deficits, even for morally relevant traits.

Associations with additional individual differences. The above results showed that self-reported moral change goals were, at best, only weakly negatively correlated with self-reported levels of the respective traits. To find out whether moral change goals are related to other aspects of moral character and values, we used semipartial correlations to examine the associations that each self-reported moral change goal had with a number of additional individual differences relating to moral character and values, controlling for self-reported levels of the trait in question. Two replicable and theoretically noteworthy findings were that more religious targets tended to report greater desires to be more compassionate and loyal (when analyses controlled for self-reported compassion and loyalty, respectively) and that targets who valued impartially maximizing the greater good (Kahane et al., 2018) reported greater desires to become fairer and more compassionate (when analyses controlled for self-reported fairness and compassion, respectively; see the Supplemental Material, Section 6, and Tables S3 and S4 in the Supplemental Material for full details).

Do high trait levels explain why people do not want to be more moral?

Finally, because ratings tended to be quite high for most of the morally relevant traits, we explored the possibility that people are less motivated to improve moral traits because they already see themselves or their close others as having relatively high levels of such traits (compared with less morally relevant traits). To test this idea, we conducted supplemental within-persons analyses using multilevel models. Across both samples, person-mean-centered trait levels and moral relevance independently predicted lowered change goals for both the continuous and priority-based measures and for both self- and informant perspectives (these analyses were not preregistered; for details, see the Supplemental Material, Section 3 and Table S2). In other words,

people were less inclined to change on the traits that they saw themselves as having relatively more desirable levels of (compared with their other traits)—but even after we controlled for targets' perceptions of their current levels, they were still less inclined to change on more morally relevant traits (and the same was true from the perspective of informants).

Discussion

We examined whether people want to be more moral and whether close others also want our moral qualities to improve. Two findings stand out. First, people were less interested in changing the levels of moral traits (e.g., honesty, fairness, compassion), compared with nonmoral traits (e.g., anxiety, sociability, productivity), in themselves and in close others. Second, targets and informants showed similar patterns of change goals, but targets wanted to change themselves to a much greater extent than their informants wanted them to change.

Why do people not particularly want to be more moral? Although self-ratings tended to be high for moral traits, ceiling effects cannot explain our main result, because our measure of change goals allowed participants to report how much they wanted to increase, decrease, or stay the same on each trait, independently of how they currently rated themselves. A more psychologically interesting possibility is that people see less room for improvement on moral traits. Because we did not measure where people thought they stood relative to the extremes of each trait, our data cannot speak directly to this idea (i.e., people might not have been claiming that they had the lowest or highest possible levels of each trait when they "strongly disagreed" or "strongly agreed" with each trait description; Blanton & Jaccard, 2006). A related possibility is that people are less motivated to improve on moral traits because they already see themselves as having quite high levels of such traits and therefore morally "good enough"—even if they think they *could* be morally better (see Schwitzgebel, 2019). However, even after controlling for current levels, we found that people were still less inclined to change more morally relevant traits, suggesting that additional psychological factors might reduce people's desires to change morally relevant traits.

One such possibility is that people are typically motivated to change in ways that will improve their own well-being (Hudson & Fraley, 2016a). Whereas becoming less anxious has obvious personal benefits, people might believe that becoming more moral would result in few personal benefits (or even some costs). Supporting this idea, targets' and informants' justifications for their top change goal suggest that nonmoral improvements would primarily benefit the target, whereas moral

improvements would primarily benefit other people (see the Supplemental Material, Section 7). Our findings are also consistent with the idea that people are reluctant to change moral traits because those traits are fundamental to their identity (e.g., Riis et al., 2008).

Considering how much people value morality in others (Goodwin, 2015; Goodwin et al., 2014; Hartley et al., 2016), it is perhaps more surprising that people do not want their close others to improve their moral qualities. Instead, like targets, informants prioritized wanting targets to become less anxious and depressed—and their open-ended justifications for these goals overwhelmingly reflected a concern for the targets' well-being (see the Supplemental Material, Section 7). Similar mechanisms might explain why people do not want themselves and close others to become more moral. An additional possible explanation—specific to interpersonal change goals—is that people are less likely to become or stay close with social partners who have very different moral values in the first place (Haidt, Rosenberg, & Hom, 2003; Skitka, 2010).

Finally, targets who had less desirable levels of a given trait wanted to improve more on that trait, but this was less true for more morally relevant traits. For example, self-perceived deficiencies in compassion and general morality were not particularly indicative of how much targets wanted to improve on these traits. Interestingly, however, informants' change goals were generally well calibrated to their perceptions of targets' deficits—including moral deficits. This suggests that close others might have unique insight into not only our trait levels (Quintus et al., 2017; Vazire & Carlson, 2011) but also which of our traits are most in need of improvement.

Constraints on generality

Our goal in asking targets to self-nominate well-acquainted informants was to understand interpersonal change goals in the context of real-world relationships. Thus, we do not expect our findings regarding interpersonal change goals to generalize beyond people who already like and are close to their targets. We would likely see a greater overall desire to change targets—and perhaps a greater desire to change moral traits specifically—if we recruited informants who disliked the targets (Leising, Erbs, & Fritz, 2010). Likewise, our results do not speak to the moral change goals that people might have for abstract entities, such as outgroups or people in general.

Does our main finding—that people are relatively unenthusiastic about moral self-improvement—generalize beyond the unique developmental context of young adulthood? For example, given that concerns for

nurturing and guiding the next generation tend to peak in midlife (McAdams, de St. Aubin, & Logan, 1993), midlife adults might be more likely to prioritize moral improvements that serve these goals. We conducted another preregistered replication, which showed that relatively older adults (mean age = 45.4 years) are also less inclined to improve more morally relevant traits (see Supplemental Material, Section 8). Thus, we conclude that across the adult life span, people in Western cultures deprioritize moral self-improvements. Future research should investigate how other people want targets of different ages to change. For example, although parents of our young adult targets wanted them to change less than the targets themselves did, we speculate that parents might want to see greater moral development in younger children, even before children start thinking about changing their own personalities (but we have no idea what age range this might apply to).

Conclusion

People care a lot about morality, but does that mean that they want themselves and close others to be more moral? We found that North American college students might not value moral improvements as much as nonmoral improvements in themselves and that their close others typically feel the same way. This suggests that personality change goals might be primarily motivated by the desire to improve one's own life rather than by more noble considerations. Still, on an uplifting note, close others generally accept us for who we are but overwhelmingly want us to change in ways that are in our best interests.

Action Editor

Michael Inzlicht served as action editor for this article.

Author Contributions

J. Sun conceptualized the study, collected and analyzed the data, and drafted the manuscript. G. P. Goodwin provided extensive feedback on the study design, analyses, and interpretations of the results. Both authors revised the manuscript and approved the final version for submission.

ORCID iD

Jessie Sun  <https://orcid.org/0000-0001-6764-0721>

Acknowledgments

We are grateful to Ted Schwaba, Luke Smillie, and Joshua Wilt for comments on an earlier draft of this article; to colleagues at the University of California, Davis, and the University of Pennsylvania for insightful conversations; and to Rebecca Neufeld and Nicholas Hunt for assistance with data collection and coding.

Declaration of Conflicting Interests

The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.

Supplemental Material

Additional supporting information can be found at <http://journals.sagepub.com/doi/suppl/10.1177/0956797619893078>

Open Practices



All measures, data, R scripts, and Mplus input and output files required to reproduce the results are publicly available on the Open Science Framework (OSF) and can be accessed at <https://osf.io/cbxjh>. We preregistered the analyses for Sample 2 at <https://osf.io/6mdwp/>. The analyses for Sample 1 were exploratory, but we preregistered stopping rules and an initial exploratory analysis plan at <https://osf.io/qjfnw/>. The complete Open Practices Disclosure for this article can be found at <http://journals.sagepub.com/doi/suppl/10.1177/0956797619893078>. This article has received the badges for Open Data, Open Materials, and Preregistration. More information about the Open Practices badges can be found at <http://www.psychologicalscience.org/publications/badges>.

References

- Baranski, E. N., Morse, P. J., & Dunlop, W. L. (2017). Lay conceptions of volitional personality change: From strategies pursued to stories told. *Journal of Personality, 85*, 285–299. doi:10.1111/jopy.12240
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*(1). doi:10.18637/jss.v067.i01
- Baumeister, R. F. (1994). The crystallization of discontent in the process of major life change. In T. F. Heatherton & J. L. Weinberger (Eds.), *Can personality change?* (pp. 281–297). Washington, DC: American Psychological Association. doi:10.1037/10143-012
- Baumeister, R. F., & Exline, J. J. (1999). Virtue, personality, and social relations: Self-control as the moral muscle. *Journal of Personality, 67*, 1165–1194. doi:10.1111/1467-6494.00086
- Bench, S. W., Schlegel, R. J., Davis, W. E., & Vess, M. (2015). Thinking about change in the self and others: The role of self-discovery metaphors and the true self. *Social Cognition, 33*, 169–185. doi:10.1521/soco.2015.33.3.2
- Blanton, H., & Jaccard, J. (2006). Arbitrary metrics in psychology. *The American Psychologist, 61*(1), 27–41. doi:10.1037/0003-066X.61.1.62
- Bollich, K. L., Johannet, P. M., & Vazire, S. (2011). In search of our true selves: Feedback as a path to self-knowledge. *Frontiers in Psychology, 2*, Article 312. doi:10.3389/fpsyg.2011.00312
- Christy, A. G., Kim, J., Vess, M., Schlegel, R. J., & Hicks, J. A. (2017). The reciprocal relationship between perceptions of moral goodness and knowledge of others' true selves. *Social Psychological and Personality Science, 8*, 910–917. doi:10.1177/1948550617693061
- DeYoung, C. G. (2015). Cybernetic Big Five Theory. *Journal of Research in Personality, 56*, 33–58. doi:10.1016/j.jrp.2014.07.004
- DeYoung, C. G., Quilty, L. C., & Peterson, J. B. (2007). Between facets and domains: 10 aspects of the Big Five. *Journal of Personality and Social Psychology, 93*, 880–896. doi:10.1037/0022-3514.93.5.880
- Fleeson, W., Furr, R. M., Jayawickreme, E., Meindl, P., & Helzer, E. G. (2014). Character: The prospects for a personality-based perspective on morality. *Social and Personality Psychology Compass, 8*, 178–191. doi:10.1111/spc3.12094
- Furr, R. M. (2008). A framework for profile similarity: Integrating similarity, normativeness, and distinctiveness. *Journal of Personality, 76*, 1267–1316. doi:10.1111/j.1467-6494.2008.00521.x
- Goodwin, G. P. (2015). Moral character in person perception. *Current Directions in Psychological Science, 24*, 38–44. doi:10.1177/0963721414550709
- 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. doi:10.1037/a0034726
- Haidt, J., Rosenberg, E., & Hom, H. (2003). Differentiating diversities: Moral diversity is not like other kinds. *Journal of Applied Social Psychology, 33*, 1–36. doi:10.1111/j.1559-1816.2003.tb02071.x
- Hartley, A. G., Furr, R. M., Helzer, E. G., Jayawickreme, E., Velasquez, K. R., & Fleeson, W. (2016). Morality's centrality to liking, respecting, and understanding others. *Social Psychological and Personality Science, 7*, 648–657. doi:10.1177/1948550616655359
- Hudson, N. W., & Fraley, R. C. (2015). Volitional personality trait change: Can people choose to change their personality traits? *Journal of Personality and Social Psychology, 109*, 490–507. doi:10.1037/pspp0000021
- Hudson, N. W., & Fraley, R. C. (2016a). Changing for the better? Longitudinal associations between volitional personality change and psychological well-being. *Personality and Social Psychology Bulletin, 42*, 603–615. doi:10.1177/0146167216637840
- Hudson, N. W., & Fraley, R. C. (2016b). Do people's desires to change their personality traits vary with age? An examination of trait change goals across adulthood. *Social Psychological and Personality Science, 7*, 847–856. doi:10.1177/1948550616657598
- Hudson, N. W., & Roberts, B. W. (2014). Goals to change personality traits: Concurrent links between personality traits, daily behavior, and goals to change oneself. *Journal of Research in Personality, 53*, 68–83. doi:10.1016/j.jrp.2014.08.008
- Kahane, G., Everett, J. A. C., Earp, B. D., Caviola, L., Faber, N. S., Crockett, M. J., & Savulescu, J. (2018). Beyond sacrificial harm: A two-dimensional model of utilitarian psychology. *Psychological Review, 125*, 131–164. doi:10.1037/rev0000093

- Kelley, K. (2018). MBESS (Version 4.4.3) [R package]. Retrieved from <https://cran.r-project.org/web/packages/MBESS/index.html>
- Leising, D., Erbs, J., & Fritz, U. (2010). The letter of recommendation effect in informant ratings of personality. *Journal of Personality and Social Psychology, 98*, 668–682. doi:10.1037/a0018771
- McAdams, D. P., de St. Aubin, E., & Logan, R. L. (1993). Generativity among young, midlife, and older adults. *Psychology and Aging, 8*, 221–230. doi:10.1037/0882-7974.8.2.221
- Meindl, P., & Graham, J. (2014). Know thy participant: The trouble with nomothetic assumptions in moral psychology. In H. Sarkissian & J. C. Wright (Eds.), *Advances in experimental moral psychology* (pp. 233–252). London, England: Bloomsbury. doi:10.5040/9781472594150.ch-012
- Molouki, S., & Bartels, D. M. (2017). Personal change and the continuity of the self. *Cognitive Psychology, 93*, 1–17. doi:10.1016/j.cogpsych.2016.11.006
- Muthén, L., & Muthén, B. O. (2017). *Mplus user's guide* (8th ed.). Los Angeles, CA: Author.
- Prentice, M., Furr, R. M., & Hawkins, A. (2019). *Introducing the Moral Characteristics Questionnaire*. Manuscript in preparation.
- Prentice, M., Jayawickreme, E., Hawkins, A., Hartley, A., Furr, R. M., & Fleeson, W. (2019). Morality as a basic psychological need. *Social Psychological and Personality Science, 10*, 449–460. doi:10.1177/1948550618772011
- Quintus, M., Eglhoff, B., & Wrzus, C. (2017). Predictors of volitional personality change in younger and older adults: Response surface analyses signify the complementary perspectives of the self and knowledgeable others. *Journal of Research in Personality, 70*, 214–228. doi:10.1016/j.jrp.2017.08.001
- R Core Team. (2018). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Riis, J., Simmons, J. P., & Goodwin, G. P. (2008). Preferences for enhancement pharmaceuticals: The reluctance to enhance fundamental traits. *Journal of Consumer Research, 35*, 495–508. doi:10.1086/588746
- Rosseel, Y. (2012). lavaan: An R package for structural equation modeling. *Journal of Statistical Software, 48*(2). doi:10.18637/jss.v048.i02
- Schönbrodt, F. D., & Perugini, M. (2013). At what sample size do correlations stabilize? *Journal of Research in Personality, 47*, 609–612. doi:10.1016/j.jrp.2013.05.009
- Schwitzgebel, E. (2019). Aiming for moral mediocrity. *Res Philosophica, 96*, 347–368.
- Sherman, R. A., & Serfass, D. G. (2015). The comprehensive approach to analyzing multivariate constructs. *Journal of Research in Personality, 54*, 40–50. doi:10.1016/j.jrp.2014.05.002
- Skitka, L. J. (2010). The psychology of moral conviction. *Social and Personality Psychology Compass, 4*, 267–281. doi:10.1111/j.1751-9004.2010.00254.x
- Soto, C. J., & John, O. P. (2017). The next Big Five Inventory (BFI-2): Developing and assessing a hierarchical model with 15 facets to enhance bandwidth, fidelity, and predictive power. *Journal of Personality and Social Psychology, 113*, 117–143. doi:10.1037/pspp0000096
- Strohming, N., Knobe, J., & Newman, G. (2017). The true self: A psychological concept distinct from the self. *Perspectives on Psychological Science, 12*, 551–560. doi:10.1177/1745691616689495
- Tappin, B. M., & McKay, R. T. (2017). The illusion of moral superiority. *Social Psychological and Personality Science, 8*, 623–631. doi:10.1177/1948550616673878
- Vazire, S., & Carlson, E. N. (2011). Others sometimes know us better than we know ourselves. *Current Directions in Psychological Science, 20*, 104–108. doi:10.1177/0963721411402478