Deep Down my Enemy is Good:
Thinking About the True Self Reduces Intergroup Bias

Julian De Freitas¹ & Mina Cikara²
Department of Psychology, Harvard University

Correspondence:

Julian De Freitas, William James Hall, 33 Kirkland St., Cambridge, MA, 02138

Intergroup bias — preference for one's in-group relative to out-groups — is one of the most robust phenomena in all of psychology. Here we investigate whether a positive bias that operates at the individual-level, belief in a good true self, may be leveraged to reduce intergroup bias. We find that even stereotypically threatening out-group agents are believed to have a good true self (Experiment 1). More importantly, consideration of in-group and out-group members' true self reduces intergroup bias, both in the form of explicit evaluative judgments (Experiment 2) and actual donation behavior (Experiment 3). Across studies, the palliative effects of thinking of an individual's true self generalize to that individual's entire group. In sum, a simple intervention — thinking about another's true self — reduces the gap in how people evaluate and treat out-group relative to in-group members. We discuss implications of these findings for conflict reduction strategies.

Keywords: Intergroup bias; true self; essentialism; lay theories

Deep Down my Enemy is Good:

Thinking About the True Self Reduces Intergroup Bias

“He’s not a bad guy, deep down,” I said.

My dad slipped the key into the door. “Deep down, no one is.” – Aaron Starmer, The Riverman

“We’re all Muslims deep down. We all yearn for peace.” – Boston Police Commissioner William B. Evans, speaking at the Islamic Society of Boston Cultural Center

Intergroup conflict is one of humanity’s greatest challenges. By at least one estimate, over 170 million civilians have perished as casualties of various forms of intergroup violence (Woolf & Hulsizer, 2004). As such, conflict reduction interventions have become a top priority for policy makers and researchers alike (Cohen & Insko, 2009). Psychologists have reported some success reducing prejudice and conflict with a variety of approaches, including: highlighting superordinate goals and identities (Gaertner & Dovidio, 2000), training regulation of negative emotions (Halperin, Porat, & Tamir, & Gross, 2013), fostering empathy across groups (Cikara, Bruneau, & Saxe, 2011), and initiating real as well as imagined contact between groups (Crisp & Turner, 2009; Pettigrew & Tropp, 2008). However, research examining the efficacy of these approaches reveals that positive effects may be short-lived (e.g., Bruneau & Saxe, 2012), may not generalize to entire groups (Brewer & Miller, 1984), and may backfire, particularly when parties are of unequal status or power (Dixon, Durrheim, & Tredoux, 2007; Dovidio, Gaertner, & Saguy, 2009; Vorauer & Sasaki, 2009). These ‘ironic’ effects of intervention come about often by bringing unsavory

information into focus (Zaki & Cikara, 2015) or by reinforcing an unjust status quo (Dixon et al., 2010).

All of these interventions start from the recognition that intergroup bias — the preferential evaluation and treatment of in-group relative to out-group members — is a fundamental facet of human psychology (Hewstone, Rubin, & Willis, 2002). This bias manifests across real and arbitrary groups in resource allocation (Tajfel & Turner, 1979), trait evaluations (Locksley, Ortiz, & Hepburn, 1980), implicit bias (Ashburn-Nardo, Voils, & Monteith, 2001), and physiological responses (Cikara & Van Bavel, 2014). In overt conflict, mere in-group preference is combined with out-group hostility, fostering anger, disgust, and aggressive behavior (Cuddy, Fiske, & Glick, 2007; Mackie & Hamilton, 1993). Informed by classic and contemporary theories of intergroup relations, most conflict interventions aim to close these gaps by targeting *group-related* cognitions and emotions (e.g., out-group empathy or anxiety, common identity, familiarity). Here we propose a novel approach in which we fight intractable intergroup bias with another robust, individual-level bias: the good true self bias (De Freitas, Cikara, Grossmann, & Schlegel, 2016a).¹

The Good True Self

People often explain the behavior of *individuals* by appealing to the concept of a good true self (Newman, Bloom, & Knobe, 2014; Knobe, 2015). That is, when evaluating a hypothetical agent, people reason as if there is something within the agent calling him or her to behave in a manner that is morally good. For example, if the agent changes from behaving badly to behaving virtuously then participants are more likely to report that this improvement reflects the emergence of the agent's true self; conversely, if the agent changes from behaving virtuously to behaving badly, the participants report that this deterioration reflects a movement away from the agent's true self (Newman et al., 2014). Perhaps the most

¹ We use the term 'bias' to refer to a response tendency, rather than to a normative departure from rationality.

convincing evidence that moral traits are seen as essential to the true self comes from experiments which find that various identity judgments (i.e., whether the person is still the same person) are consistently influenced by the removal of morally good traits more than the removal of morally bad traits or even a host of other mental faculties, including personality, memory, perception, and preferences (Prinz & Nichols, in press; Strohminger & Nichols, 2014; Tobia, 2016). In other words, when the morally good traits are removed, people are inclined to say that the person is no longer the same.

An emerging consensus is that this belief in a good true self is a form of psychological essentialism, whereby people tend to view morally good traits as an essential part of human identity (De Freitas et al., 2016b; De Freitas, Tobia, Newman, & Knobe, 2016c; Strohminger & Nichols, 2014). The belief in a good true self also has consequences for other judgments that psychologically derive from reasoning about agents in terms of good true selves. For instance, agents believed to be expressing their true selves are more likely to be judged as happy, strong-willed, and as purposeful than are agents who behave in a manner that conflicts with their true selves (Newman, De Freitas, & Knobe, 2015). These findings dovetail nicely with empirical work stemming from Self-Determination Theory, indicating that people do not associate their immoral behaviors with their broader core values (i.e., they compartmentalize bad behaviors; Ryan & Deci, 2000).

Finally, and most relevant to the current work, it is important to emphasize that belief in a good true self is perspective-independent; people regard both their own true selves (Bench et al., 2015; Molouki & Bartels, in press) and the true selves of others (Bench et al., 2015; Newman, Bloom, & Knobe, 2014) as fundamentally good. This stands in contrast to a large body of work on the self as a whole that shows robust perspective-dependent asymmetries in a variety of domains, such as fundamental attribution error (e.g., Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001; Taylor & Brown, 1994). The good true self bias is

also found in two unlikely boundary cases: 1) in interdependent cultures, where less emphasis is placed on the individual as a separate entity from others, and 2) in misanthropes, who have explicitly negative views about humanity (De Freitas et al., 2016b). The fact that belief in a good true self is robust across these boundary conditions provides support for the hypothesis that this belief is a fundamental aspect of people's commonsense understanding of others, and thus may have widespread consequences for other aspects of cognition, including intergroup bias. In particular, it is possible that even out-group members are viewed as having good true selves, rather than as being essentially bad.

Yet at first glance it would appear that the good true self bias has no chance of holding for out-group members, since aside from the strength of intergroup bias, it is well documented that people tend to think of the essence of an out-group as a negative category, e.g., 'the essence of Arab immigrants' (Haslam, Rothschild, & Ernst, 2000, 2002). Here we emphasize that the good true self bias is one that falls out of reasoning about the essence of an *individual person*. Extending this idea, it may be that when thinking about the essence of an individual out-group member, e.g. 'the essence of Alhadin', or 'the essence of Jafri', the same positive values normally associated with true selves are recruited. If so, a positive bias that falls out of thinking about the essence of an *individual person* can be leveraged to potentially reduce a negative bias that falls out of thinking about the essence of a disliked *out-group*.

Current Experiments and Hypotheses

Here we test whether even stereotypically threatening out-group members are believed to have good true selves (Experiment 1), and whether there is an asymmetry in attributions of a good true self to in-group and out-group targets as a function of whether the agent improves or deteriorates. If indeed out-group targets are equally likely to be attributed good true selves as in-group targets, we predict that asking participants to consider individual

out-group members' true selves *prior* to making group-level evaluative judgments (Experiment 2) and group-level resource allocation decisions (Experiment 3) will reduce intergroup bias (relative to when participants make these judgments and decisions without first considering targets' true selves). Experiment materials, data, and analysis code for all three experiments in the paper can be downloaded at:

https://osf.io/g6x7f/?view_only=a780cde9131745bebbe281fcd0029a9c

Experiment 1: Examining True Self Bias for In-Group Versus Out-Group Members

In this experiment, we test whether two stereotypically threatening out-group members — an Arab and an Arab immigrant in the U.S. — are equally likely as an in-group member — a white U.S. American — to be attributed a good true self. Specifically, we tested whether, for both agents, moral improvements are thought to reflect the true self more than moral deteriorations. Given the dominant phenomenon of intergroup bias, this would be a discovery in its own right, suggesting that intergroup bias does not extend to essentialist reasoning about individual members of the out-group.

It is worth noting that the following data were collected approximately two months after the San Bernadino mass shooting in the USA, which was widely attributed to “Islamic terrorism”. If historic context had any effect, it was to make our hypothesized results less likely.

Methods

Participants and exclusions. We recruited 1020 participants from the United States through Amazon's Mechanical Turk, an online labor crowdsourcing platform (see Buhrmester, Kwang, & Gosling, 2011; Goodman, Cryder, & Cheema, 2013; Ipeirotis, 2010; Paolacci, Chandler, & Ipeirotis, 2010; Berinsky, Huber, & Lenz, 2012), in order to obtain a final target $N = 600$ (100 per condition). This sample size provides 80% power to detect an effect $\eta_p^2 = .125$ (effect size is based on those observed in previous research on the true self;

see Newman et al., 2014). As planned a priori, we excluded from the analysis participants who incorrectly answered an attention check at the beginning of the study, or two comprehension questions at the end of the study. We also excluded participants who did not self-identify as being of White ethnicity (all questions used for exclusions are provided in the methods below), yielding a final sample of $N = 613$ ($M_{\text{age}} = 36.0$, 53.5% female).

Materials and procedure. Participants were randomly assigned to one of six conditions in a 2 (Condition: improvement vs. deterioration) $\times$ 3 (Ethnic group: Arab native vs. Arab Immigrant vs. White American) between-subjects design. As a robustness check, we used 6 different vignettes describing different agents (e.g. opposing/supporting terrorism, respecting/mistreating minorities, being a teetotaler/alcoholic), presented between-subjects (see link to archived materials and Newman et al., 2014). These vignettes were chosen because they span a variety of moral improvements and deteriorations, and have previously been shown to elicit a reliable good true self bias across cultures (i.e., participants consistently say that moral improvements reflect the agent's true self more than moral deteriorations; De Freitas et al., 2016b).

Participants were first introduced to the agent they would be learning about, "*Imagine an individual named [name]. [Name] is different from you in almost every way—he has a different occupation and prefers different things than you.*" This introduction is typically used in studies of the true self (De Freitas et al., 2016b; Newman, Bloom, & Knobe, 2014) in order to minimize the effect of perceived similarity between the participant and agent on more arbitrary factors (e.g. same job, lifestyle), since people tend to like those who are similar to them on just about any dimension (Ross, Greene, & House, 1977). For the Arab native and Arab immigrant conditions we translated the English names into Arab equivalents, e.g., Jeffrey to Jafri, or Chad to Chahid.

We then assessed participant's evaluations of the agent's group using a number of measures from the intergroup psychology literature (Cikara et al., 2014; Schmid, Ramiah, & Hewstone, 2014; Van Bavel & Cunningham, 2012). These measures included (i) perceived intergroup threat ("Please rate the extent to which you agree with the following statement: People from White [Arab] backgrounds like Jeffrey [Jafri] threaten the American way of life", 1 = strongly disagree, 5 = neither agree nor disagree, 9 = strongly agree), (ii) an attitude thermometer ("How do you feel about people from White [Arab] backgrounds like Jeffrey [Jafri]?" 1 = cold, 5 = neither warm nor cold, 9 = warm), and (iii) group identification, which consisted of three questions ("I [value/like/feel connected to] people from White [Arab] backgrounds like Jeffrey [Jafri]", 1 = not at all, 9 = very much). We created composite in-group and out-group identification scores by averaging the three items for each condition (American items $\alpha = .86$; Arab native items $\alpha = .88$, Arab immigrant items $\alpha = .90$).

Participants then read that the agent underwent a behavioral change from morally good to morally bad, or from morally bad to morally good (with the direction of moral change counterbalanced between subjects). For example:

Deterioration. "Al [Alhadin] used to be a very caring and involved father. In the past, he always showed real affection for his children and always expressed interest in his children's lives. Now, however, Al [Alhadin] is not a very caring and involved father and is not involved in his children's lives."

Improvement. "Al [Alhadin] used to be a "deadbeat" dad. In the past, he never showed any real affection for his children and never expressed any interest in his children's lives. Now, however, Al [Alhadin] is a very caring and involved father."

Following previous experiments on the good true self bias (Newman et al., 2014), participants then answered a forced-choice question about what caused the agent's behavior: "In your opinion, what aspect of Al's [Alhadin's] personality caused him to be a very caring

and involved father [not be a very caring father and not be involved in his children's lives]?" The answer choices were (a) His 'true self' (the deepest, most essential aspect of his being), (b) His 'surface self' (the things that he learned from society or others), and (c) None of the above. The third option included a space in which participants could explain their choice. A second question asked participants about the agent's behavior in relation to his/her true self, "Now that Al [Alhadin] is a very caring and involved father [is not a very caring father and is not involved in his children's lives], to what extent is he being true to the deepest, most essential aspects of his being?" (1 = not at all, 9 = very much so).

Participants then answered two comprehension questions, "What was the ethnicity of the man in the story?" (the answer options were: Black/ African American, Hispanic/ Latino, Asian/ Pacific Islander, Native American/ American Indian, Arab/ Middle Eastern, White/ Caucasian, and Other), and "Did the man's behavior improve, deteriorate, or neither (the answer options were: Improve, Deteriorate, Neither). We excluded participants who incorrectly answered these questions based on the condition to which they were assigned. Finally, participants reported their ethnicity, age, gender, education, and socioeconomic status, and were debriefed about the purpose of the experiment. We collected educational level and socioeconomic status for the sake of completeness only; we do not report any analyses including these measures.

Results

Where appropriate throughout this paper, we report statistical significance based on the Bonferroni-Holm corrected alpha level, α_{adj} , required to control for multiple comparisons in each multi-factor ANOVA (Cramer et al., 2015).² If $p > \alpha_{adj}$, then we

² The Bonferroni-Holm correction (aka sequential Bonferroni procedure) sorts all significant p -values of an ANOVA from the smallest to largest p -value, then computes adjusted α levels, α_{adj} , for each of these p -values. For the smallest p -value, α_{adj} equals α divided by the number of tests conducted in the ANOVA; for the next smallest p -value, α_{adj} equals α divided by the number of tests minus 1; and so forth until the largest significant p -value. The significance of each p -value is then evaluated against its respective adjusted alpha. Unlike many

conclude that the result is not significant. For all studies, we report all measures, manipulations and exclusions, and all sample sizes were determined before any data analysis.

Intergroup Bias. We do not include the condition (improvement vs. deterioration) or vignette factors in these analyses because in Experiment 1 the intergroup bias measures always preceded the vignette in which the agent's behavior improves or deteriorates.

Threat. A 3-way (Ethnicity: White American, Arab native, Arab immigrant) univariate ANOVA revealed the predicted main effect of ethnicity on threat judgments, $F(2, 610) = 6.88, p = .001, \eta_p^2 = .022$. The average threat rating for White Americans ($M = 2.59, SD = 1.92$) was lower than that for Arab natives ($M = 3.28, SD = 2.31$), mean difference = 0.69, 95% CI [0.28, 1.09], $t(411) = 3.29, p = .001, d = 0.32$, and Arab immigrants ($M = 3.24, SD = 2.14$), mean difference = 0.65, 95% CI [0.25, 1.03], $t(417) = 3.23, p = .001, d = 0.32$.

Attitude thermometer. A 3-way (Ethnicity: White American, Arab native, Arab immigrant) univariate ANOVA revealed the predicted main effect of ethnicity on attitudes, $F(2, 610) = 7.87, p = .0004, \eta_p^2 = .025$. The average attitude thermometer rating for White Americans ($M = 5.67, SD = 1.50$) was higher than that for Arab natives ($M = 5.12, SD = 1.72$), mean difference = 0.55, 95% CI [-0.86, -0.24], $t(411) = -3.46, p = .0006, d = -0.34$, and Arab immigrants ($M = 5.12, SD = 1.72$), mean difference = 0.55, 95% CI [-0.86, -0.24], $t(417) = -3.51, p = .0005, d = -0.34$.

Identification. A 3-way (Ethnicity: White American, Arab native, Arab immigrant) univariate ANOVA revealed a main effect of ethnicity on identification, $F(2, 610) = 20.02, p = 3.81 \times 10^{-9}, \eta_p^2 = .062$. The average identification rating for White Americans ($M = 5.72, SD = 1.59$) was higher than that for Arab natives ($M = 4.89, SD = 1.75$), mean difference = 0.83, 95% CI [-1.16, -0.51], $t(411) = -5.05, p = 6.7 \times 10^{-7}, d = -0.50$, and Arab immigrants ($M =$

other correction procedures, this procedure is not only strict but is also optimized for avoiding *both* Type 1 and 2 errors (Cramer et al., 2015).

4.76, $SD = 1.76$), mean difference = 0.96, 95% CI [-1.29, -0.64], $t(417) = -5.90$, $p = 7.68 \times 10^{-9}$, $d = -0.58$

Good True Self Bias. Forced-Choice. Following Newman et al. (2014), we recoded the forced-choice item as a binary response with “true self” response as “1” and “surface self” or “other” responses as “0”. A logistic regression including only condition found that participants were significantly more likely to attribute moral improvements (59%) than moral deteriorations (30%) to the true self, $z = -7.28$, $p = 3.24 \times 10^{-13}$. The AIC was not significantly lower for more complicated models which included target group membership or vignette as factors ($p = .398$ and $p = 0.014$ in favor of simplest model).³

Scaled-Response. Replicating previous work, a 2 (Condition: Deterioration, Improvement) $\times$ 2 (Ethnicity: White American, Arab native, Arab immigrant) $\times$ 6 (Vignette) univariate ANOVA found only a significant main effect of condition, $F(1, 577) = 112.76$, $p < 2 \times 10^{-16} < \alpha_{adj} = .0071$, $\eta_p^2 = .163$, but no significant main effects of ethnicity, $F(2, 577) = 0.90$, $p = .406$, $\eta_p^2 = .003$, or vignette, $F(5, 577) = 1.19$, $p = .314$, $\eta_p^2 = .010$. Collapsing across vignettes and ethnicities, the average true self attribution was higher in the improvement ($M = 6.21$, $SD = 1.93$) than the deterioration condition ($M = 4.40$, $SD = 2.23$), mean difference = 1.81, 95% CI [-2.14, -1.48], $t(611) = -10.73$, $p < 2.2 \times 10^{-16}$, $d = -0.87$. Thus, as predicted, this good true self bias was not moderated by target social group (see Figure 1).

³ For all forced-choice analyses in this paper, we run logistic regression models of increasing complexity. For instance, in Experiment 1 we ran models with (i) only the main effect of condition, (ii) the condition $\times$ ethnicity interaction, and (iii) the condition $\times$ ethnicity $\times$ vignette interaction. We then compare models in order to settle on the simplest model that provides the best quantitative fit to the data.

Figure 1. Mean true self attributions in Experiment 1. Error bars indicate standard error from the mean.

Discussion

Although participants exhibited the predicted evaluative intergroup bias on threat, attitude, and identification measures, even stereotypically threatening out-group members were equally subject to the good true self bias. In other words, in-group and threatening out-group members were judged as having equally good true selves. This effect was the same across groups irrespective of whether the agents improved (in which case their new behavior was viewed as more reflective of the true self) or deteriorated (in which case their new behavior was viewed as less reflective of the true self).

Experiment 2: Reducing Intergroup Bias in Explicit Evaluative Judgments

Experiment 2 used the same methods as Experiment 1, but also counterbalanced the order in which the intergroup judgments and true self measures were presented. Furthermore, since no significant difference emerged for judgments toward Arab natives versus Arab immigrants in Experiment 1, we used only the group that received the most negative judgments overall in that experiment: Arab immigrants. We predicted that participants who considered either in-group or out-group members' true selves prior to (as compared to after) making evaluative group judgments would exhibit reduced intergroup bias (i.e., relatively more neutral evaluations of both in-group and out-group).

Notice that in principle participants could be recruiting the notion of a true self when reasoning not only about moral improvements but also moral deteriorations, since for both conditions participants judge whether or not the behavior reflects the agent's true self. Yet given that the true self judgment task only asks participants to rate the extent to which a given behavioral change reflects the true self (and finds preferential attribution of moral improvements), this measure does not allow us to determine whether people are also activating the concept in the deterioration condition. If they are, then we should expect this activated concept to go on to reduce intergroup bias in *both* moral improvement and deterioration conditions (e.g., "he is simply acting in accordance with his true self in the improvement condition, or not in the deterioration condition"). In contrast, if the concept is only activated when reasoning about moral improvements, we should only expect to see a reduction in intergroup bias in the improvement condition.

Methods

Participants and exclusions. We recruited 1327 participants from the United States through Amazon's Mechanical Turk. We aimed for the same sample size per condition as in Experiment 1 ($N = 100$ after exclusions). As planned a priori, we excluded from the analysis participants who took part in Experiment 1, or who incorrectly answered the attention check at the beginning of the study or two comprehension questions at the end of the study. We also excluded participants who did not self-identify as being of White ethnicity, yielding a final sample of $N = 759$ ($M_{\text{age}} = 37.1$, 56.6% female).

Materials and procedure. Participants were randomly assigned to one of eight conditions in a 2 (Condition: Deterioration, Improvement) $\times$ 2 (Ethnicity: White American, Arab Immigrant) $\times$ 2 (Order: True Self Judgments After, Before Intergroup Judgments) between-subjects design. As a robustness check, we again used 6 different vignettes describing different agents, presented between-subjects (taken from Newman et al., 2014).

As in Experiment 2, participants were first introduced to the agent they would be learning about. They then received either the intergroup judgment measures followed by the true self measures (including the accompanying vignettes), or vice versa. Finally, they answered the same comprehension and demographic items as used in Experiment 1.

Results

Good True Self Bias. First we examined whether we replicated the main finding of Experiment 1: that participants were no less likely to attribute a good (i.e., improved) true self to a threatening out-group member than to an in-group member.

Forced-Choice. The forced-choice responses were recoded as in Experiment 1. A logistic regression including only condition found that participants were significantly more likely to attribute moral improvements (58%) than moral deteriorations (29%) to the true self, $z = -7.77$, $p = 7.96 \times 10^{-15}$, and the AIC was not significantly lower for more complicated models including the model that included order as a factor ($p = 0.091$ and $p = 0.042$ in favor

of simplest model). Therefore, we replicated the good true self bias for both in-group and out-group agents.

Scaled Item. A 2 (Condition: Deterioration, Improvement) $\times$ 2 (Ethnicity: White American, Arab immigrant) $\times$ 2 (Order: True Self Judgments After, Before) $\times$ 6 (Vignette) univariate ANOVA found only a significant main effect of condition, $F(1,711) = 201.55, p < 2 \times 10^{-16} < \alpha_{adj} = .0033, \eta_p^2 = .219$, and a marginal main effect of vignette, $F(5,711) = 3.59, p = .0032 < \alpha_{adj} = .0036, \eta_p^2 = .024$. All vignettes elicited higher true self ratings for the improvement than deterioration condition, but some elicited larger effects than others (see Supplementary Figure S1). There were no main effects of ethnicity, $F(1,711) = 4.86, p = .028 > \alpha_{adj} = .0038, \eta_p^2 = .007$, or order, $F(1,711) = 0.58, p = .445, \eta_p^2 = .0007$. We also did not find a significant ethnicity $\times$ order interaction, $F(1,711) = 2.90, p = .089, \eta_p^2 = .004$, a condition $\times$ ethnicity interaction $F(1,711) = 1.95, p = .163, \eta_p^2 = .003$, or any significant 3-way or 4-way interactions (all $ps > .210$). Collapsing across vignettes and ethnicities, the average true self attribution was higher in the improvement ($M = 6.51, SD = 1.89$) than the deterioration condition ($M = 4.32, SD = 2.38$), mean difference = 2.19, 95% CI [-2.50, -1.89], $t(757) = -14.09, p < 2.2 \times 10^{-16}, d = -1.02$. Therefore, mirroring the results from the forced-choice measure, we once again replicated the good true self bias for judgments of both in-group and out-group members.

Intergroup Bias. *Threat.* A 2 (Condition: Deterioration, Improvement) $\times$ 2 (Ethnicity: White American, Arab immigrant) $\times$ 2 (Order: True Self Judgments After, Before) $\times$ 6 (Vignette) univariate ANOVA revealed significant main effects of order, $F(1, 711) = 14.03, p = .0036 < \alpha_{adj} = 0.004, \eta_p^2 = .018$, and ethnicity, $F(1, 711) = 10.36, p = .001 < \alpha_{adj} = .0038, \eta_p^2 = .013$. These main effects were qualified by the predicted significant ethnicity $\times$ order interaction, $F(1, 711) = 19.87, p = 9.36 \times 10^{-6} < \alpha_{adj} = .0033, \eta_p^2 = .028$. There was no significant main effect of vignette, $F(5, 711) = 3.30, p = .006 > \alpha_{adj} = .0042$,

$\eta_p^2 = .024$, nor a significant vignette $\times$ order interaction, $F(5, 711) = 2.35, p = .039 > \alpha_{adj} = .0045, \eta_p^2 = .015$. Finally, there was also no significant main effect of condition, condition $\times$ order interaction, or 3-way or 4-way interactions (all $ps > 0.250$).

In order to unpack the significant ethnicity $\times$ order interaction, we first examined only those cases when the intergroup judgment measures were presented *prior* to the true self judgments (replicating the design of Experiment 1). The average threat rating for White Americans ($M = 2.38, SD = 1.81$) was significantly lower than that for Arab immigrants ($M = 3.54, SD = 2.29$), mean difference = 1.16, 95% CI [0.74, 1.56], $t(383) = 5.52, p = 6.28 \times 10^{-8}, d = 0.57$. When the true self measures were presented first, however, the average threat rating for White Americans ($M = 3.56, SD = 2.28$) did not differ significantly from that for Arab immigrants ($M = 3.35, SD = 2.21$), mean difference = -0.21, 95% CI [-0.67, 0.25], $t(372) = -0.89, p = .376, d = -0.09$. See Figure 2 for a summary of results.

Attitude thermometer. A 2 (Condition: Deterioration, Improvement) $\times$ 2 (Ethnicity: White American, Arab immigrant) $\times$ 2 (Order: True Self Judgments After, Before) $\times$ 6 (Vignette) univariate ANOVA revealed only the predicted, but marginally significant ethnicity $\times$ order interaction, $F(1, 711) = 8.09, p = .0046 > \alpha_{adj} = .0033, \eta_p^2 = .012$. Note that technically using an adjusted alpha here is overly conservative, since there was nothing exploratory about this predicted interaction (Cramer et al., 2015), and without this correction the interaction easily meets standard criteria for significance. Nonetheless, we report corrected alphas here for consistency.

None of the other tests was significant. There was no significant main effect of vignette, $F(5, 711) = 2.31, p = .043 > \alpha_{adj} = .0038, \eta_p^2 = .016$, nor significant main effects of ethnicity, condition, condition $\times$ order interaction, or 3-way or 4-way interactions (ethnicity $p = 0.137$; all other $ps > 0.250$). There was also no significant vignette $\times$ ethnicity $\times$ condition interaction, $F(5, 711) = 2.61, p = .024 > \alpha_{adj} = .0036, \eta_p^2 = .018$,

In order to unpack the ethnicity $\times$ order interaction, we first examined only those cases in which intergroup judgment measures were presented before true self judgments. The average attitude thermometer rating for White Americans ($M = 5.60$, $SD = 1.49$) was significantly higher than that for Arab immigrants ($M = 5.05$, $SD = 1.75$), mean difference = 0.55, 95% CI [-0.87, -0.22], $t(383) = -3.32$, $p = .0009$, $d = -0.34$. When the true self measures were presented first, however, the average attitude thermometer rating for White Americans ($M = 5.29$, $SD = 1.93$) did not differ significantly from that for Arab immigrants ($M = 5.49$, $SD = 1.77$), mean difference = -0.20, 95% CI [-0.19, 0.57], $t(372) = 0.98$, $p = .326$, $d = 0.10$. See Figure 2 for a summary of results.

Identification. As in Experiment 1, we first created composite in-group and out-group identification scores by averaging the three identification questions for each condition (American items $\alpha = .93$; Arab immigrant items $\alpha = .89$). A 2 (Condition: Deterioration, Improvement) $\times$ 2 (Ethnicity: White American, Arab immigrant) $\times$ 2 (Order: True Self Judgments After, Before) $\times$ 6 (Vignette) univariate ANOVA revealed a significant main effect of ethnicity, $F(1, 711) = 16.13$, $p = 6.55 \times 10^{-5} < \alpha_{adj} = .0033$, $\eta_p^2 = .022$, and, critically, the predicted ethnicity $\times$ order interaction, $F(1, 711) = 10.55$, $p = .001 < \alpha_{adj} = .0036$, $\eta_p^2 = .015$. There was no main effect of condition, significant condition $\times$ order interaction, or 3-way or 4-way interactions (vignette $\times$ ethnicity $\times$ condition $p = 0.062$; all other $ps > 0.250$).

Again, we first examined just those cases when the intergroup judgment measures were presented first. The average identification rating for White Americans ($M = 5.62$, $SD = 1.59$) was significantly higher than that for Arab immigrants ($M = 4.67$, $SD = 1.66$), mean difference = 0.95, 95% CI [-1.28, -0.62], $t(383) = -5.70$, $p = 2.37 \times 10^{-8}$, $d = -0.59$. When the true self measures were presented first, however, the average identification rating for White Americans ($M = 5.24$, $SD = 1.99$) did not differ significantly from that for Arab immigrants

($M = 5.16$, $SD = 1.88$), mean difference = 0.08, 95% CI [-0.48, 0.31], $t(372) = -0.42$, $p = .671$, $d = -0.04$. See Figure 2 for a summary of results.

Figure 2. Mean intergroup ratings for threat, attitude, and identification, in Experiment 2.

Error bars indicate standard error from the mean.

Discussion

The good true self bias replicated for both in-group and out-group targets: participants judged improvements as a better reflection of targets' true selves than deteriorations. More importantly, thinking first about the true self for in-group and out-group members (irrespective of whether it was a deterioration or improvement vignette) significantly reduced intergroup bias in the form of explicit evaluations of threat, attitudes, and identification. Further, note that the intergroup bias was attenuated by both reducing positive evaluations of the in-group and increasing positive-evaluations of the out-group. Thus, it appears that thinking about the true self has an equilibrating effect, simultaneously reducing overly positive attitudes toward us, and attenuating overly negative attitudes toward them.

Experiment 3: Reducing Intergroup Bias in Donation Behavior

Experiment 2 demonstrated that considering in-group and out-group members' true selves reduced intergroup bias in evaluation judgments. However, negative appraisals — both explicit attitudes and implicit associations — are not always highly correlated with discriminatory behavior (Ajzen & Cote, 2008; Talaska, Fiske, & Chaiken, 2008). Therefore in Experiment 3 we changed the outcome variable of interest to an incentive compatible behavioral measure. Participants had to decide how much of a bonus payment they wanted to allocate between an in-group charity versus an out-group charity. Aside from being a behavioral measure, this is an especially strong test of intergroup bias, since any amount given to the out-group charity is also an amount *not* given to the in-group charity. One side effect of this stronger measure, however, is that it removes our ability to directly measure whether a given change in donation behavior is driven by a reduction in in-group preference, an enhancement in out-group preference, or both. Yet, given the results of Experiment 2 (in which we found evidence for both), we expected that reasoning about *either* an in-group or out-group agent prior to making a donation would lead to a larger proportion of donations given to the out-group charity. Such a result would still be consistent with the interpretation that attitudes toward *both* in-group and out-group agents are shifted after thinking about the true self.

Because we did not find an effect of condition (improvement vs. deterioration) on the attenuation of the threat, attitude, and identification bias in Experiment 2, we include only improvement scenarios here. In sum, for Experiment 3 we predicted only a main effect of order (true self before vs. after donation measure) on donation behavior.

Methods

Participants and exclusions. We recruited 678 participants from the United States through Amazon's Mechanical Turk. We aimed for the same sample size per cell as in

Experiments 1–2 ($N = 100$ after exclusions). As planned a priori, we excluded from the analysis participants who took part in Experiments 1–2, and who incorrectly answered the attention check at the beginning of the study or two comprehension questions at the end of the study. We also excluded participants who did not self-identify as being of White ethnicity, yielding a final sample of $N = 374$ ($M_{\text{age}} = 36.7$, 52.0% female).

Materials and procedure. Participants were randomly assigned to one of four conditions in a 2 (Ethnicity: White American, Arab Immigrant) $\times$ 2 (Order: Charity Donation First, True Self First) between-subjects design. As a robustness check, we again used 6 different vignettes describing different agents, presented between-subjects (taken from Newman et al., 2014).

For the donation measure, participants were asked how much of a \$0.50 bonus they would like to allocate between an in-group charity (American Red Cross) and an out-group charity (Syrian Arab Red Crescent), both of which truly exist:

“We are giving you a 50c bonus that we would like you to split between the following two charities that are devoted to providing humanitarian relief during emergencies: the “American Red Cross” and the “Syrian Arab Red Crescent”. We will make this donation on your behalf.

How would you like to split the donation?” (the options were 100% to: American Red Cross, 90% to: ARC, 80% to: ARC, 70% to: ARC, 60% to: ARC, 50% to: EACH, 60% to SARC, 70% to: SARC, 80% to: SARC, 90% to: SARC, 100% to: Syrian Arab Red Crescent). For the purposes of analysis we recoded these choices from 1 (100% to in-group charity) to 11 (100% to out-group charity).

Participants either saw the donation measure first followed by the true self measure (including a vignette), or vice versa. Finally, they answered the same comprehension and

demographic items as used in Experiments 1–2. On behalf of the participants, we summed all of their donations and donated them to the intended charities.

Results

True Self Question. *Forced-Choice.* The forced-choice responses were coded as in Experiments 1–2. A logistic regression including only ethnicity found no significant difference depending on whether participants were judging a White American (57%) vs. an Arab immigrant (61%), $z = -0.71$, $p = .478$, and the AIC was not significantly lower for more complicated models ($p = 0.541$ and $p = 0.036$ in favor of simplest model). Participants did not judge improvements as more indicative of in-group true selves relative to out-group true selves.

Scaled Response. In line with the forced-choice response results, a 2 (Ethnicity: White American, Arab Immigrant) $\times$ 2 (Order: Donation First, True Self First) $\times$ 6 (Vignette) univariate ANOVA found no significant main effects of ethnicity, $F(1, 350) = 0.03$, $p = .870$, $\eta_p^2 = 2.78 \times 10^{-5}$, order, $F(1, 350) = 0.00$, $p = .989$, $\eta_p^2 = 1.87 \times 10^{-10}$, or vignette, $F(5, 350) = 1.54$, $p = .176$, $\eta_p^2 = 2.26 \times 10^{-2}$, and, critically, no significant ethnicity $\times$ order interaction, $F(1, 350) = 1.64$, $p = .201$, $\eta_p^2 = .003$.

Donation Behavior. A 2 (Ethnicity: White American, Arab immigrant) $\times$ 2 (Order: Donation First, True self first) $\times$ 6 (Vignette) univariate ANOVA revealed only the predicted significant main effect of order, $F(1, 350) = 12.68$, $p = .0004 < \alpha_{adj} = .0071$, $\eta_p^2 = .038$, conceptually replicating the results of Experiment 2. Collapsing across vignettes, the average donation amount to the Syrian Arab Red Crescent was significantly higher when the true self question was presented first ($M = 4.89$, $SD = 2.87$) than when the donation measure was presented first ($M = 3.81$, $SD = 2.87$), mean difference = 1.08, 95% CI [0.49, 1.66], $t(373) = 3.64$, $p = .0003$, $d = 0.38$.

Figure 3. Mean donations to the out-group charity (Syrian Arab Red Crescent) in Experiment 3. Error bars indicate standard error from the mean.

Discussion

As predicted, and in line with the results of Experiments 2, thinking about the true self of in-group and out-group members increased actual donations to an out-group charity.

General Discussion

We investigated whether beliefs in a good true self may be leveraged to reduce intergroup bias. Surprisingly, despite the strength and pervasiveness of intergroup bias, people reported that threatening out-group members have equally good true selves as do in-group members (Study 1). Furthermore, instructing participants to consider in-group and out-group true selves first reduced subsequent intergroup bias, both in the form of explicit

evaluative judgments (Study 2), and actual donation behavior (Study 3).

There are a number of aspects of this effect that make it unique. To our knowledge, belief in a good true self is one of very few social attributes that does not appear to be impacted by intergroup relations: it did not matter whether people were reasoning about in-group or threatening out-group agents, they consistently believed that the agent was good deep down. The intervention also does not originate at the group level. That is, we found that a positivity bias, which operates at the level of the *individual*, reduced a bias that falls out of reasoning about individuals in terms of their *group membership*. Thus, these findings are somewhat counterintuitive, suggesting that an effective intervention for reducing intergroup bias need not operate at the group level. Furthermore, the *way* in which the good true self bias reduced intergroup bias was also somewhat surprising. Effective interventions typically aim to reduce negative attitudes toward out-groups specifically (Cohen & Insko, 2008). Yet we found evidence that the good true self bias not only reduced negative attitudes toward out-group members, but also attenuated positive attitudes toward in-group members, leading to net equilibration. The reduction in in-group preferences makes sense in light of the literature on positive biases, which finds that people over-estimate the positive qualities of themselves and their in-group members to start with (e.g., Taylor & Brown, 1994). Thus, if the default state is to believe that one's in-group is 100% good, then thinking about the true self may ironically remind people that even in-group members are layered — yes, they are good, but they also have superficial aspects. The reverse would happen when thinking about the true self of an out-group member — they are not completely bad, since they are still good deep down. Thus, our findings show that intergroup interventions have legroom to operate within *both* out-group and in-group evaluations and behaviors, making both more neutral, and therefore more equitable.

In other words, we believe that these effects fall directly out of the concept of a good

true self, which is thought to arise from *psychological essentialism*, the tendency to understand entities in terms of a deeper, unobservable essence (Ahn et al., 2001; Gelman, 2003; Dar-Nimrod & Heine, 2011; Keil, 1989; Medin & Ortony, 1989; Xu, & Rhemtulla, 2005). Just as people show a robust, cross-cultural tendency to posit an unobservable essence for a variety of entities (e.g., Atran, 1993; Brown, 1991; Gil-White, 2001; Hirschfeld, 1998; Sousa, Atran, & Medin, 2002), they appear to posit an unobservable essence of the self, yielding the notion of a ‘true self’ (for a review, see De Freitas et al., 2016). Moreover, they view the true self as consisting of traits they morally value (Newman, Bloom, & Knobe, 2014; De Freitas et al., 2016), yielding the notion of a *good* true self. Thus, if the notion of a good true self is a fundamental way in which people intuitively make sense of the very fabric of the self, then we would expect it to generalize even in potential boundary cases. Here we find such generalization to intergroup reasoning.

We think the essentialist framework provides a more parsimonious account of these effects relative to other theories. For instance, the results are not consistent with the theory that thinking about good things in general reduces intergroup bias, since making true self judgments about *both* improvements and deteriorations reduced intergroup bias. Furthermore, it is unlikely that thinking about the true self leads one to think about feelings more than behaviors, thereby leading to more empathy (Andersen & Ross, 1984), since the good true self bias has previously been shown to apply to both feelings and overt behaviors (Newman, Bloom, & Knobe, 2014). In contrast to these alternatives, the results suggest that thinking about whether in-group and out-group members’ behaviors reflect their true self vs. surface self leads to a more nuanced representation of these individuals, enhancing positive views of an out-group member (which may be overly-negative to start with) and dampening positive views of an in-group member (which may be overly-positive to start with). Future studies can certainly probe this mechanism further, but the main objective of the current studies has been

to discover 1) that the good true self also applies to out-group members, and 2) it can thus be cheaply leveraged to reduce intergroup bias in both evaluations and behavior.

Critically, the intervention reported here is incredibly simple; it does not require the application of external, expensive, or elaborate methods and regimens. Instead, it would appear that the main ingredient for change comes conveniently pre-packaged in people's intuitive psychology: thinking of individuals in terms of morally good essences (De Freitas et al., 2016; Strohminger & Nichols, 2014). Of course, not all essentializing is good, as in the well-documented tendencies to think of the essence of an out-group as a category, e.g. 'the essence of Arab immigrants' (Haslam, Rothschild, & Ernst, 2000, 2002). In contrast, our studies show that asking people to think about the essence of an individual person, e.g. 'the essence of Alhadin', or 'the essence of Jafri', has immediate, powerful consequences for intergroup cognition and behavior.

Our approach is similar to other interventions that activate beliefs or mindsets which do not directly reference the relationship between groups but nevertheless have effects on intergroup dynamics. For example, inducing malleable beliefs about groups (e.g., that groups, in general, can change their basic characteristics) leads Israeli Jews and Palestinian citizens of Israel to express more positive attitudes towards the out-group, and a greater willingness to compromise with their out-groups (Halperin, Russell, Trzesniewski, Gross, & Dweck, 2011). Note that these malleable versus fixed prompts did not reference Israel, Palestine, or even the concept of out-groups, though the outcome measures referred specifically to the Israeli-Palestinian conflict (Halperin et al., 2011). Similar to these approaches, the true self bias may avoid the defensive reactance that is sometimes associated with directing participants to adopt out-group perspectives or to engage in a structured interaction because participants in the current experiments were never directly told that the true selves of in-group and out-group members were essentially good. Instead, they were simply asked to reflect on whether

a given behavior was a manifestation of the agent's true self.

In sum, although we tend to view in-group members as mostly good and out-group members as mostly bad, neither is viewed as rotten to the core. Directing people to consider this common, essential good may lead to more nuanced representations of us and them, fostering more equitable treatment across the group divide.

References

- Ahn, W., Kalish, C., Gelman, S. A., Medin, D. L., Luhmann, C., Atran, S., Coley, J. D., Shafto, P. (2001). Why essences are essential in the psychology of concepts. *Cognition*, 82, 59–69.
- Ashburn-Nardo, L., Voils, C. I., & Monteith, M. J. (2001). Implicit associations as the seeds of intergroup bias: How easily do they take root? *Journal of Personality and Social Psychology*, 81, 789–799.
- Atran, S. (1993). *Cognitive foundations of natural history: Towards an anthropology of science*. Cambridge University Press.
- Ajzen, I., & Cote, N. G. (2008). Attitudes and the prediction of behavior. In W.D. Crano and R. Preslin (Eds.). *Attitudes and Attitude Change*. New York: Psychology Press.
- Baumeister, R., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, 5, 323–370.
- Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating online labor markets for experimental research: Amazon.com’s Mechanical Turk. *Political Analysis*, 20, 351–368.
- 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.
- Brewer, M. B., & Miller, N. (Eds.). (1984). *Groups in contact: The psychology of desegregation*. Academic Press.
- Brown, D. E. (1991). *Human universals* (p. 118). New York: McGraw-Hill.
- Bruneau, E. G., & Saxe, R. (2012). The power of being heard: The benefits of “perspective-giving” in the context of intergroup conflict. *Journal of Experimental Social Psychology*, 48, 855–866.

- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6, 3–5.
- Cikara, M., Bruneau, E. G., & Saxe, R. R. (2011). Us and them intergroup failures of empathy. *Current Directions in Psychological Science*, 20, 149–153.
- Cikara, M., Bruneau, E., Van Bavel, J. J., & Saxe, R. (2014). Their pain gives us pleasure: How intergroup dynamics shape empathic failures and counter-empathic responses. *Journal of Experimental Social Psychology*, 55, 110–125.
- Cikara, M., & Van Bavel, J. J. (2014). The neuroscience of intergroup relations an integrative review. *Perspectives on Psychological Science*, 9, 245–274.
- Cohen, T. R., & Insko, C. A. (2008). War and peace: Possible approaches to reducing intergroup conflict. *Perspectives on Psychological Science*, 3, 87–93.
- Cramer, A. O., van Ravenzwaaij, D., Matzke, D., Steingroever, H., Wetzels, R., Grasman, R. P., Waldorp, L. J., & Wagenmakers, E. J. (2014). Hidden Multiplicity in Multiway ANOVA: Prevalence and Remedies. arXiv preprint arXiv:1412.3416.
- Crisp, R. J., & Turner, R. N. (2009). Can imagined interactions produce positive perceptions?: Reducing prejudice through simulated social contact. *American Psychologist*, 64, 231–240.
- Cuddy, A. J. C., Fiske, S. T., & Glick, P. (2007). The BIAS map: Behaviors from intergroup affect and stereotypes. *Journal of Personality and Social Psychology*, 92, 631–648.
- Dar-Nimrod, I., & Heine, S. J. (2011). Genetic essentialism: On the deceptive determinism of DNA. *Psychological Bulletin*, 137, 800–818.
- De Freitas, J., Cikara, M., Grossmann, I., & Schlegel, R. (2016a). *How deep is the good true self concept?* Manuscript under review
- De Freitas, J., Newman, G. E., Sarkissian, H., Grossmann, I., De Brigard, F., Luco, A.,

- & Knobe, J. (2016b). *Consistent belief in a good true self in misanthropes and three interdependent cultures*. Manuscript under review.
- De Freitas, J., Tobia, K., Newman, G. E., & Knobe, J. (2016c). Normative judgments and individual essence. *Cognitive Science*, 1–21.
- Dixon, J., Durrheim, K., & Tredoux, C. (2007). Intergroup contact and attitudes toward the principle and practice of racial equality. *Psychological Science*, 18, 867–872.
- Dovidio, J. F., Gaertner, S. L., & Saguy, T. (2009). Commonality and the complexity of “we”: Social attitudes and social change. *Personality and Social Psychology Review*, 13, 3–20.
- Dixon, J., Tropp, L. R., Durrheim, K., & Tredoux, C. (2010). “Let them eat harmony”: Prejudice-reduction strategies and attitudes of historically disadvantaged groups. *Current Directions in Psychological Science*, 19, 76–80.
- Gaertner, S. L., & Dovidio, J. F. (2000). *Reducing intergroup bias: The common ingroup identity model*. Philadelphia: Psychology Press.
- Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. Oxford, UK: Oxford University Press.
- Gil-White, F. J. (2001). Sorting is not categorization: A critique of the claim that Brazilians have fuzzy racial categories. *Journal of Cognition and Culture*, 1, 219–249.
- Goodman, J. K., Cryder, C. E., & Cheema, A. (2013). Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples. *Journal of Behavioral Decision Making*, 26, 213–224.
- Halperin, E., Porat, R., Tamir, M., & Gross, J. J. (2013). Can emotion regulation change political attitudes in intractable conflicts? From the laboratory to the field. *Psychological Science*, 24, 106–111.

- Halperin, E., Russell, A. G., Trzesniewski, K. H., Gross, J. J., & Dweck, C. S. (2011). Promoting the Middle East peace process by changing beliefs about group malleability. *Science*, 333, 1767–1769.
- Haslam, N., Rothschild, L., & Ernst, D. (2000). Essentialist beliefs about social categories. *British Journal of Social Psychology*, 39, 113–127.
- Haslam, N., Rothschild, L., & Ernst, D. (2002). Are essentialist beliefs associated with prejudice?. *British Journal of Social Psychology*, 41, 87–100.
- Hewstone, M., Rubin, M., & Willis, H. (2002). Intergroup bias. *Annual review of psychology*, 53, 575–604.
- Hirschfeld, L. A. (1998). *Race in the making: Cognition, culture, and the child's construction of human kinds*. MIT Press.
- Ipeirotis, P. (2010, March 9). The new demographics of Mechanical Turk. Retrieved from <http://www.behind-the-enemy-lines.com/2010/03/newdemographics-of-mechanical-turk.html>.
- Keil, F.C. (1989) *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.
- Locksley, A., Ortiz, V. & Hepburn, C. (1980). Social categorization and discriminatory behavior: Extinguishing the minimal intergroup discrimination effect. *Journal of Personality and Social Psychology*, 39, 773–783.
- Mackie, D., & Hamilton, D. (1993). *Affect, Cognition, and Stereotyping*. San Diego, CA: Academic.
- Medin, D. L., and Ortony, A. (1989). Psychological essentialism. In S. Vosniadou and A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 179–195). Cambridge: Cambridge University Press.
- Newman, G. E., Bloom, P., & Knobe, J. (2014). Value judgments and the true self.

- Personality and Social Psychology Bulletin*, 40, 203–216.
- Newman, G. E., De Freitas, J., & Knobe, J. (2015). Beliefs about the true self explain asymmetries based on moral judgment. *Cognitive Science*, 39, 96–125.
- Paolacci, G., Chandler, J., & Ipeirotis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5, 411–419.
- Pettigrew, T. F., & Tropp, L. R. (2008). How does intergroup contact reduce prejudice? Meta-analytic tests of three mediators. *European Journal of Social Psychology*, 38, 922–934.
- Prinz, J., & Nichols, S. (in press). Diachronic identity and the moral self. In J. Kiverstein (Ed.), *Handbook of the social mind*. London: Routledge.
- Ross, L., Greene, D., & House, P. (1977). The false consensus effect: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13, 279–301.
- Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American psychologist*, 55, 68–78.
- Schmid, K., Al Ramiah, A., & Hewstone, M. (2014). Neighborhood Ethnic Diversity and Trust The Role of Intergroup Contact and Perceived Threat. *Psychological Science*, 1–10.
- Sousa, P., Atran, S., & Medin, D. (2002). Essentialism and folkbiology: Evidence from Brazil. *Journal of Cognition and Culture*, 2, 195–223.
- Strohinger, N. and Nichols, S. (2014). The Essential Moral Self. *Cognition*, 31, 159–171.
- Tajfel, H., & Turner, J. (1979). An integrative theory of intergroup conflict. In W. G. Austin & S. Worschel (Eds.), *The Social Psychology of Intergroup Relations* (pp. 33–47). Pacific Grove, CA: Brooks/Cole Publishing.

- Talaska, C. A., Fiske, S. T., & Chaiken, S. (2008). Legitimizing racial discrimination: Emotions, not beliefs, best predict discrimination in a meta-analysis. *Social Justice Research, 21*, 263–296.
- Taylor, S. E., & Brown, J. D. (1994). Positive Illusions and Well-Being Revisited. *Psychological Bulletin, 116*, 21–27.
- Tobia, K. P. (2016). Personal identity, direction of change, and neuroethics. *Neuroethics, 1*, 1–7.
- Van Bavel, J. J., Packer, D. J., & Cunningham, W. A. (2011). Modulation of the Fusiform Face Area following minimal exposure to motivationally relevant faces: Evidence of in- group enhancement (not out-group disregard). *Journal of Cognitive Neuroscience, 23*, 3343–3354.
- Vorauer, J. D., & Sasaki, S. J. (2009). Helpful only in the abstract? Ironic effects of empathy in intergroup interaction. *Psychological Science, 20*, 191–197.
- Woolf, L. M., & Hulsizer, M. R. (2004). Hate groups for dummies: How to build a successful hate group. *Humanity and Society, 28*, 40–62.
- Xu, F., & Rhemtulla, M. (2005). In defense of psychological essentialism. In *Proceedings of the 27th Annual Conference of the Cognitive Science Society* (pp. 2377-2380). Erlbaum Mahwah, NJ.
- Zaki, J., & Cikara, M. (2015). Addressing empathic failures. *Current Directions in Psychological Science, 24*, 471–476.