Good for What? Social Rank, Character, and Moral Judgment

Seth Robertson*

Department of Philosophy

Harvard University

Alexander F. Danvers*

Department of Psychology

University of Arizona

*Both authors contributed equally to this manuscript

Date of Submission: April 23, 2019

Corresponding Author:
Alexander F. Danvers
Department of Psychology
University of Arizona
1503 E University Blvd, Building 68
Tucson, AZ, 85721
Tel: (520) 621-7447
Email: adanvers@email.arizona.edu

Abstract word count: 184
Main text word count: 5,514
References: 34

Abstract

The person-centered account of moral judgment suggests an evolutionary explanation for two puzzling findings from moral psychology: that people judge an act that had bad side effects as more intentional than an act that had good side effects (Knobe, 2003), and that people judge an accident as being more under the control of the target who caused it if that target has good values (Alicke, 1992). In both cases, people's intuitions are biased by their "fundamental need" to judge a person's character (Uhlmann, Pizarro, & Diermeier, 2015). People have evolved to judge people not just by specific actions, but by considering whether they would be good cooperative partners and members of a social group. We tested implications of this explanation: people alter their judgments of intention and causal control to benefit people with greater value to the group. In two studies, we examined the moderating role of social value on moral judgment using highly cited and often replicated experimental paradigms. Results do not corroborate this influential theory of moral judgment. High status individuals are not given "moral license." Instead, results suggest that dominant, but not prestigious, individuals are judged as more efficacious, which may have downstream consequences for moral judgment. Data and analysis scripts available at: <https://osf.io/u9pfr/>.

Keywords: Moral Judgment, Moral Character, Person Perception, Intention, Experimental Philosophy

In an experiment cited over 800 times in the last 15 years and replicated numerous times, Joshua Knobe (2003) found that people are more likely to attribute intention to a person who causes bad—as opposed to good—side effects (Knobe, 2003; 2004; 2006; Nadelhoffer, 2006; Knobe & Mendlow, 2004; Tannenbaum, Ditto, & Pizarro, 2007; Beebe & Buckwalter, 2010; Leslie, Knobe, & Cohen, 2006; Beebe & Kenson, 2012; Sripada, 2010; 2012). Another experiment cited over 400 times and replicated several times finds that judgments of a person’s control over a situation are influenced by perceptions of that person’s values (Alicke, 1992; Alicke, 2000; Alicke, Buckingham, Zell, & Davis, 2008; Alicke & Rose, 2012). People’s intuitive judgments about intention and control are susceptible to what appear to be irrelevant considerations: whether an action leads to a good or bad outcome, and whether a person has good or bad moral character. The “person-centered account” of moral judgment attempts to explain these results via evolutionary theory (Pizarro & Tannenbaum, 2011; Uhlmann, Pizarro, & Diermeier, 2015). People’s intuitive moral judgments aren’t designed to weigh *acts* alone: people also have a “fundamental human motivation to determine the moral character of others” (pp. 72, Uhlmann, Pizarro, & Diermeier, 2015). According to this theory, Knobe’s (2003) “Side Effect Effect” and Alicke’s (1992) “Culpable Causation Effect” are both instances where the ability to make impartial decisions about specific acts is driven by the need to judge a person’s underlying character.

The evolutionary logic of person-centered morality suggests that people are motivated to make character judgments because humans evolved to live in cooperative groups. In the context of a small group, where people would need to interact with each other repeatedly over time, identifying trustworthy cooperation partners was a fitness-relevant task. As Pizarro and Tannenbaum (2011) put it, the “ability and motivation to evaluate others on the basis of moral

character was likely of such fundamental importance during primate and human evolution that it is most likely a product of natural selection” (p. 92).

Yet if moral judgments are influenced by perceptions of how good a person would be as a cooperative partner and group member (a quality distinct from moral character), trustworthiness and other familiar moral character traits would not be the only consideration. Having ample resources and social power also makes one an attractive cooperative partner. This evolutionary explanation for the utilization of evaluations of character in judgements of acts also implies that *power* should matter in moral judgment. Further, if an evolved preference for identifying good cooperative partners underlies the seemingly illogical intuitions seen in the Knobe and Alicke effects, then judgments about the concepts of intention and causation should also be influenced by social power. Stated plainly, people should make allowances in morality, even at the level of attributions of *causation*, for the socially powerful.

Person-Centered Moral Judgment

Recently, supporters of “person-centered accounts” of moral judgment (Pizarro & Tannenbaum, 2011; Tannenbaum, Uhlmann, & Diermeier, 2011; Uhlmann et al., 2015) have claimed that research in empirical moral psychology, influenced as it has been by broadly deontological (Kohlberg, 1969; Piaget, 1965) or utilitarian (Greene, Morelli, Lowenberg, Nystrom, & Cohen, 2008; Greene, Sommerville, Nystrom, Darley, & Cohen, 2001) ethical theories, may have been overly focused on moral judgments of particular acts (Pizarro & Tannebaum, 2011, pp. 94–96). Researchers therefore have failed to appreciate ways in which our moral judgments about *people* affect our moral judgments about their *actions*. Supporters of person-centered accounts of moral judgment have argued that our inclinations and abilities to

evaluate other people's moral character may be, in some sense, more primary than our inclinations and abilities in judging their actions (Pizarro & Tannebaum 2011, pp. 92-93).

There are reasons to believe we may be evolutionarily predisposed to be attuned to character: the ability and inclination to accurately identify, via their "moral character," others' potential as especially helpful allies or especially dangerous enemies would have had a positive impact on fitness in relevant ancestral environments (Pizarro & Tannebaum, 2011, p. 93; Uhlmann et al., 2015, p. 74). Further, under the logic of cultural evolution and group selection, providing benefits to prestigious others should lead to increased fitness for the group overall (Henrich & Gil-White, 2001). If this is true, we should expect to see certain types of distortions in our moral judgment. For example, we should expect that some features of a person that are *not* typically associated with moral character but would contribute to their qualities as potential allies or enemies, such as their degree of social power, might influence our moral judgments.

Although prior theoretical statements about moral judgments being influenced by power treat it as roughly synonymous with status, the literature on social rank makes several useful distinctions between rank-related concepts. The evolutionary literature emphasizes a distinction between dominance and prestige as distinct ways to achieve social rank (Cheng, Tracy, Foulsham, Kingston, & Henrich, 2013; Cheng & Tracy, 2014; Henrich & Gil-White, 2001). Dominance involves the use of intimidation and coercion to gain and maintain status, and typically involves high rank individuals instilling fear in subordinates. Prestige is social rank freely granted by others as a result of demonstrating high levels of expertise and prosocial behavior. In a related literature, social rank is divided into two mutually reinforcing elements: status and power (Blader & Chen, 2012; 2014; Magee & Galinsky, 2008). Power is defined as control over valued resources, while status is defined as respect freely accorded by others. Since

both status and prestige involves respect and admiration for perceived social value, recent reviews have suggested that these (and the term “sociometric status”) are synonymous (Anderson, Hildreth, & Howland, 2015; Cheng et al., 2013). Other reviews suggest they are related, but that prestige is an antecedent—and not identical to—status, and may also be an antecedent to power (Blader & Chen, 2014).

These conceptual distinctions suggest an important point of clarification for the person-centered account of morality we will test. People should especially value individuals of high social rank who have the expertise and prosocial tendencies that would contribute to their—and their group’s—well-being, but they should also want to avoid “despotic” leaders who they perceive as likely to harm or exploit them¹. Moral judgments might be biased *against* an individual of high social rank who uses intimidation and coercive tactics against the individual or their allies. This distinction between leaders who value helping the group versus valuing helping themselves maps most cleanly onto the conceptual distinction between people who use prestige versus dominance to maintain status. Most theorists who use the dominance and prestige framework suggest that prestige requires a prosocial disposition (Cheng & Tracy, 2014), although some researchers have suggested that being moral—independent of having expertise—is a separate route to status (Bai, 2017). Limiting ourselves to the case where an individual expects to be the target of whatever strategy the high rank person uses, dominant leaders would be those who we expect to coerce and intimidate us into compliance, while prestigious leaders are those we naturally want to follow due to their skill and goodness. The dynamics may become more complicated if a dominant leader intimidates one faction of a group that we dislike (or a prestigious leader only shares knowledge with that disliked faction). However, the puzzles that

¹ We thank an anonymous reviewer for suggesting we consider the possibility of “harmful” leaders.

motivated our hypothesizing are moral dilemmas in which the acts described do not have clearly distinct effects on different factions within a group: environmental impact or a penalty for a car accident do not fall more on one coalition than another. We therefore hypothesize that, if there is an evolved bias in moral judgment towards people who provide benefits to the individual, in these moral dilemmas it should be towards individuals who have demonstrated valuable expertise and a willingness to help others in the past, i.e. those who high in prestige. It should not be for—and may act in opposition to—those high in dominance. The way in which an individual gained and maintains their high social rank should thus be a crucial theoretical moderator of this effect.

Defenders of person-centered accounts of moral judgment have not, to our knowledge, provided an explicit and detailed definition of character. Most often they talk about character in terms of personality traits such as integrity, empathy, trustworthiness, fair treatment of others, or more generally, whether a person is a “good” or “bad” person (Pizarro & Tannebaum, 2011, pp. 92–93; Uhlmann et al., 2015, p. 74). At times they also include a person’s perspectives, deeply-held moral principles, and values as part of a person-centered account – our moral judgments track features of actions that are “informative in that they indicate who the actor is and what he or she values and considers when performing moral actions” (Pizarro & Tannebaum, 2011, p. 104; Uhlmann et al., 2015, p. 77). On our reading, by “character” supporters of person-centered accounts of moral judgment are thinking of morally relevant traits of varying degrees of generality, long-term and deeply-held moral commitments such as values and principles, and, potentially, the moral perspectives that tend to be adopted.

Of course, there has been a great deal of research aimed at developing detailed, empirically-informed accounts of moral character (see Sinnott-Armstrong & Miller 2017; Jayawickreme &

Fleeson 2017; Miller 2015; Miller 2016; Snow 2010; Goodwin, Piazza, & Ronin 2014; Cohen & Morse 2014; Cohen, Panter, Turan, Morse & Kim (2014); Fleeson, Furr, Jayawickreme; Meindl, & Helzer 2014). But, there is a key distinction here that is essential to the rationale behind this study: the distinction between scholarly theories of moral character and natural or folk understandings of (something like) moral character. Person-centered accounts of moral judgment do not claim that our moral judgments of acts are surreptitiously influenced by our *explicit*, highly-abstract, philosophically and scientifically-informed notions of moral character, but instead that they influenced by our evolutionarily-shaped *implicit* notions of (something like) moral character. Thus, person-centered accounts of moral judgment should not be evaluated using familiar accounts of moral character (e.g., Cohen & Morse, 2014). This key distinction underlies the hypothesis we are interested in here: if the person-centered account of moral judgment is correct and our evolutionarily-influenced implicit notions of moral character influence our moral judgments about acts, we should expect to detect that evolutionary influence in what might otherwise be surprising ways.²

One key claim of person-centered accounts of moral judgment is that the inclination to evaluate a person's character and the ability to evaluate it well are beneficial, perhaps even beneficial enough to have been affected by natural selection (Pizarro & Tannebaum, 2011, p. 93). Pizarro and Tannenbaum liken this to attending to reputation in mathematical models of the evolution of cooperation (Pizarro & Tannebaum, 2011, p. 92). In these models, one way that cooperation can evolve in a population is to allow for *indirect reciprocity* (Nowak, 2006). Indirect reciprocity occurs when an individual who cooperates with or helps another person does not need to have their help reciprocated by the recipient. Instead, the helping individual develops

² Thanks to an anonymous reviewer for helping us to clarify this point and for suggesting several helpful references.

a positive reputation that makes it more likely anyone—including an outside party—will help or cooperate with her in the future. Theoretical biologists have long argued that reputation could allow for the development of morality (e.g., Alexander, 1987), and theoretical work on indirect reciprocity often discusses reputation broadly as incorporating “good” behavior (Ohtsuki & Iwasa, 2004).

This logic could be further extended to consider other features that would make individuals valuable as cooperators in a group. For example, individuals with access to many resources or specialized knowledge might have the capacity to provide greater benefit to others. Each instance of cooperation from these high value individuals would provide greater benefits to a recipient than help from an ordinary person. The person-centered account of morality implies that part of evaluating character is evaluating how likely a person will be to cooperate in the future, and how beneficial that cooperation will be. If an individual is especially bad at detecting defectors or incompetent cooperators, she or he will lose out on resources. But not all cooperation is equal. A highly valued person’s cooperation might be more valuable than the cooperation of a less valued individual, and therefore being too harsh in judging the valued individual will also lead the person judging to lose out on resources. Judgments of people as cooperators should be scaled by the amount of benefit their cooperation can provide; even if a high-value person is not a very reliable cooperator, their cooperation (when it happens) is highly valuable.

When we combine this insight with person-centered accounts of moral judgment, according to which our judgments of a person’s character (which are in part, judgments of successful future likelihood of cooperation) influence our judgments of that person’s acts, then we should, whether we are aware of it or not, judge defections of a highly valuable member of

society less harshly than we judge defections of a less valuable other. This evolutionary explanation for person-centered moral judgment suggests that judgments of individuals with the ability to contribute more to the well-being of others should be subtly influenced to account for this social value. Similarly, a person (even a person with good moral character) *without* many resources that might prove useful to others therefore would have judgments subtly biased against her or him.

Judging Intention and Culpability

We propose to test this model using two well-known effects in moral psychology that have been explained by considering the moral character of the people being evaluated: Knobe's (2003) "Side Effect Effect" and Alicke's (1992) culpable causation effect. These results are thought to be driven primarily by implicit considerations of character. If the importance of moral character is based on its function in increasing group-based fitness across repeated interactions, then these effects should be further moderated by other characteristics of the individual that would increase group-based fitness across repeated interactions.

In Knobe's secondary moral consequences experiments, an executive's decision to ignore environmental regulations is judged as intentional because it is seen as a behavior diagnostic of character (Knobe, 2003; 2006). But if the connection between power and moral judgment we've suggested holds, then the intentionality of the decision might also be influenced by features related to the company's ability to clean up the environment. If the company has the resources available to clean up after an environmental contamination and provide new jobs, then it might be considered a valuable "potential cooperator"—an agent who has the potential to provide large benefits to society. People might therefore judge the executive's decision as less intentional if he is acting on behalf of a rich, as opposed to poor, company. Similarly, in Alicke's culpable

causation experiment, people judged a cocaine user as more causal in his own traffic accident because it was seen as more in line with his character (Alicke, 1992). Yet if the individual were a cocaine user and an important member of the community, people might be willing to judge his actions as less causal, because he would be seen as providing value to society more broadly. Manipulating social value—indicated through prestige—should therefore change the size of the Knobe and Alicke effects. The theoretical model is illustrated in Figure 1.

Figure 1: Theoretical Model to be Tested

A:

B:

Note: The Chairman Scenario is the well-studied moral judgment task in Knobe's research. The Car Accident Scenario is the well-studied moral judgment task in Alicke's research. In both cases, judgment of an individual's ability to contribute to the success of a group, labeled "Social Value" here, is the mechanism explaining non-logical responses to the scenarios.

We used Amazon's Mechanical Turk (Mturk) to collect our research sample. The use of Mturk in social science research has recently come under scrutiny, so we here address several common concerns regarding the use of these samples. First, Mturk participants are more attentive to tasks than traditional undergraduate research populations, as indicated by recent research showing they are both more likely to respond correctly to an "attention check" question and to show larger effects in response to text manipulations than undergraduate students (Hauser & Schwarz, 2016). Additionally, results obtained using Mturk samples are comparable to those using other online platforms and student samples (Bartnek, Duenser, Moltchanova, & Zawieska, 2015). A recent commentary on the use of Mturk in industrial and organizational psychology addressed several other common concerns (Woo, Keith, & Thornton, 2015). One concern is that

participants complete multiple similar studies, which they conclude is problematic to the extent that the measure being used is susceptible to practice effects; taking a personality survey multiple times, for example, is not problematic, while taking an intelligence test multiple times is. To address this concern, we only recruit participants who have not previously completed any other Mturk tasks.

Another concern is that the people choosing to participate in the study are not representative. Woo, Keith, and Thornton (2015) explain that Mturk reaches a more diverse population than typical undergraduate samples and many sampling strategies used by industrial and organizational psychologists, and so they consider it a better option for reaching a representative sample than common alternatives. Finally, we argue that Mturk allows greater participation of researchers from outside of large, well-funded laboratories at major research universities. Scholars who may not have their own research labs at which to run studies, who may not be competitive in obtaining grants due to their teaching loads, or who may not have access to a large subject pool at their university are given the ability to contribute their research ideas through using Mturk samples.

Study 1

The first study examines the way social rank changes judgments of intentionality using Knobe's "Side Effect Effect" scenarios. Uhlmann, Pizarro, and Diermeier (2015) specifically discuss Knobe's (2003) studies of secondary moral consequences as case studies of how their approach can explain surprising findings in the field of moral judgment. To provide a fair test of the person-centered account, we used procedures matching those that have been well-established in the literature (Knobe, 2003; 2006; Sripada, 2012; Sripada & Konrath, 2011). Since we tested a proposed moderator of an often-replicated effect, we chose to alter as few details of the original

methods as possible. This allows us to closely compare our results with those of other researchers; if methodological changes beyond our targeted manipulation are introduced, any differences from prior findings could be attributed to these methodological changes and not the experimental manipulation of interest (Meehl, 1978). For example, if we were to attempt to stage a “live” version of the Knobe effect in lab, we might find that making decisions in person is different from making decisions while considering a situation in a removed setting. While this might be an interesting addition to the literature surrounding these effects, it would leave us uncertain as to whether a person’s value as a cooperative partner can explain the core effect, or the moderation was due to factors related to making the decision in person.

Modifications were made to vignettes used by Knobe, Sripada, and others to explore this effect. In addition to manipulating whether side effects are good or bad, we manipulated the dominance and prestige the individual being judged has. In two separate conditions, the individual was described as maintaining high social rank through dominance—by using intimidation and coercion, or through prestige—by demonstrating expertise and a prosocial orientation. These were compared to a control condition where the individual has low social rank. This will create a two-factor design, with type of side effects (good or bad) crossed with rank (high prestige, high dominance, or low rank).

Based on our theoretical model, we conducted two tests of the role of social rank on moral judgments. The first considered only the high prestige versus low rank conditions of the rank factor. We tested whether the “Knobe effect”—the difference between judgments of good and bad side effects—differs between the high prestige versus low rank conditions. If there is a significant interaction, with high prestige people being judged as having acted less intentionally in the “bad” condition, this would support the account of moral judgment as being biased by

considerations of how much an individual will benefit the group. If the interaction acts in the opposite direction, such that low rank people are seen as having acted less intentionally in the “bad” condition, this would suggest an account of morality in which people of high rank are held to a higher standard—contrary to the theoretical predictions of the person-centered account of morality. If there is an independent effect of prestige on moral judgments without an interaction, this would support a generalized halo effect, in which high prestige is thought to be related to goodness. If there is no effect of prestige, this would suggest moral judgments are not influenced by the social rank of the person being judged, conforming with an intuitive distinction between status and goodness.

Our second test considers only the high dominance versus low rank conditions of the rank factor. We will again test the interaction of these conditions with the good and bad side effect conditions, but in this case we would expect the opposite interaction. If the interaction is significant and high dominance people are judged as having acted more intentionally, this would support an account in which people are biased against those who have a net negative social value on them or their group. If the opposite reaction is present, this would suggest that people are biased in favor of high rank individuals in a group regardless of their degree of prosocial intentions—suggesting that strength can lead to moral leniency. If there is an independent effect of dominance, this would suggest rank does influence moral judgment, but as a separate, additive factor. If there is no effect of dominance, this would suggest a distinction between rank and goodness.

Methods

Participants. Participants were recruited from the crowdsourcing website Mechanical Turk (Mturk), specifying that they cannot have previously completed any tasks on Mturk.

Participants were paid \$1 for completing a 10 minute questionnaire. As part of our funding agreement, a secondary, unrelated survey was completed after our experiment.

In a study using three versions of the secondary consequences effect, Sripada (2012) reported differences in means corresponding to Cohen's d 's of 0.73, 0.82, and 1.24. Assuming we would find the smallest of these effect sizes in our study ($d = .73$), we would need 50 participants per cell to have 95% power replicate the established main effect. In his blog post describing power analyses for interaction effects, Roger Giner-Sorolla calculates that when finding power for an interaction effect where an established effect is "knocked out" by an interaction, the cell size should be twice that of the original study (Giner-Sorolla, 2018). We are adding one additional experimental condition, for a 2 x 3 ANOVA design. Including 100 participants per cell would yield a total N of 600. We therefore aimed to collect 600 participants.

Based on our preregistration, we excluded participants who did not answer our dependent variable question, and participants who complete the entire study in less than 60 seconds (on the basis that this would not give adequate time to read the full narrative and respond to questions). Excluding participants who took less than 60 seconds reduced the full sample size from 606 people who started the survey) to 520. The average age of participants was 38 years old ($SD = 11$). 37 % identified as female, 62% as male, and 0.4% as non-binary. The sample was majority white (64%), followed by black or African American (20%), Asian (5%), and Hispanic / Latino (4%). Less than 10 participants were included in each other category. Full details of ethnicity are available in our dataset, posted online (<https://osf.io/u9pfr/>).

Procedures. Participants were randomly assigned to one of six different conditions. Each of the conditions varied two factors: whether the side effect was helping or harming the

environment, and whether the company had high prestige, high dominance, or low rank. An example of the scenario is written below:

The vice-president of a company went to the chairman of the board and said, ‘We are thinking of starting a new program. It will help us increase profits, but it will also **help/harm** the environment.’

(high prestige) The company was very successful, because it was innovative and contributed to the community. It was admired by competitors.

(high dominance) The company was very successful, because it was uncompromising and played to win. It was feared by competitors.

(low rank) The company was not very successful. It did not have a lot of money or a large influence on its local community.

The chairman of the board answered, ‘I don’t care at all about **helping/harming** the environment. I just want to make as much profit as I can. Let’s start the new program.’

They started the new program. Sure enough, the environment was **helped/harmed**.

The social rank condition is indicated in bold, with the specific condition name in parentheses preceding the associated text. The helping versus harming manipulation is bolded and designated with a forward slash. In the helping condition only versions of the verb “help” are used; in the harming condition only versions of the verb “harm” are used.

Participants were asked a series of questions about the scenario that closely mirror prior work. These include the following questions used in the research of Sripada and Konrath (2011), and Sripada (2012):

1. How much do you agree with the statement “The Chairman intentionally harmed [helped] the environment?” (From Strongly Agree to Strongly Disagree)
2. What are the Chairman’s values and attitudes towards the environment? (From Very Pro-Environment to Very Anti-Environment)
3. In your opinion, what is the Chairman’s moral status? (From Very Moral to Very Immoral)
4. How much do you agree with the statement “The Chairman desires that the environment is harmed [helped]”? (From Strongly Agree to Strongly Disagree)

In earlier research, Question 1 from above was the final outcome variable in statistical models, Questions 2-4 were mediators, and the experimental manipulation was an exogenous predictor. We also asked the following question not included in the earlier materials:

5. How much do you agree with the statement “What the Chairman did is morally wrong”?

All questions were answered on a 7-point Likert scale.

Results

Our primary, preregistered hypothesis was that the strength of the “Side Effect Effect” would be moderated by the status of the company. This was tested through the interaction term in a 2 (help, harm) x 3 (dominance, prestige, low status) ANOVA. Results indicate that the “Side Effect Effect” was detected, as indicated by a main effect of the helping versus harming condition ($F(1, 514) = 56.89, p < .001, \text{partial } \epsilon^2 = .10, 95\% \text{ CI } [.06, .15]$). However, there was no interaction effect ($F(2, 514) = 0.43, p = .650, \text{partial } \epsilon^2 = .00, 95\% \text{ CI } [.00, .01]$), nor was

there a main effect of the status condition ($F(2, 514) = 0.05, p = .955$, partial $\epsilon^2 = .00$, 95% CI [.00, .01]). Our hypothesis was therefore *not supported*. These results are plotted in Figure 2A.

We also preregistered a secondary hypothesis, that judgments of the Chairman's morality would follow a similar pattern, with a significant interaction between the help versus harm condition and the status condition. Results of this study indicated that a main effect of the helping versus harming condition was again present ($F(1, 514) = 40.31, p < .001$, partial $\epsilon^2 = .07$, 95% CI [.04, .12]). However, there was no interaction effect ($F(2, 514) = 0.14, p = .869$, partial $\epsilon^2 = .00$, 95% CI [.00, .01]) nor a main effect of status ($F(2, 514) = 0.36, p = .695$, partial $\epsilon^2 = .00$, 95% CI [.00, .01]). This secondary hypothesis was therefore also *not supported*. These results are plotted in Figure 2B.

Figure 2: Results of Preregistered Hypotheses for Study 1

A: Primary Hypothesis, Regarding Judgment of Chairman's Intentionality

B: Secondary Hypothesis, Regarding Judgment of Chairman's Morality

Note: Results are presented as boxplots, with median point “cinched” to accentuate comparisons. Raw data is plotted over the boxplot with partial transparency and a slight jitter to aid visualization.

Preregistered Exploratory Mediation Analysis. We also preregistered an exploratory mediation analysis that would replicate the analysis performed by Sripada (2011). In this analysis, the effects of the experimental condition on perceptions of intentionality were mediated by three other perceptions of the Chairman: the Chairman’s values and attitudes towards the environment (from anti to pro); the Chairman’s moral status (from immoral to moral); and how well the Chairman’s desires lined up with the outcome of the company’s action. The model was estimated in the R package “lavaan,” with bootstrapping used to generate standard errors, as recommended by MacKinnon (2008).

Results of the direct effects indicate that, above-and-beyond the mediated effects, being in the helping condition significantly reduced judgments of intentionality ($B^* = -0.40$, $Z = -7.19$, $p < .001$)³, and being in the dominance condition (relative to the low status condition) reduced judgments of intentionality ($B^* = -0.12$, $Z = -2.40$, $p = .016$), but being in the prestige condition (relative to the low status condition) did not significantly alter judgments of intentionality ($B^* = -.02$, $Z = -0.34$, $p = .737$). The main effects of the help / harm condition and of the dominance condition were qualified by a statistically significant interaction ($B^* = 0.10$, $Z = 1.98$, $p = .048$). Participants in both the helping condition and the dominance condition were judged to be acting more intentionally than would be expected, based on the independent effects of these conditions. That is, the “Side Effect Effect” was approximately 25% smaller than would be expected if the company was dominant. No such interaction effect was statistically significant for the prestige effect ($B^* = .02$, $Z = 0.35$, $p = .726$).

Operating separately from these direct effects were the effects of the experimental conditions on the mediators. Perceptions of the Chairman’s attitudes towards the environment were only significantly influenced by the help / harm condition ($B^* = 0.22$, $Z = 2.97$, $p = .003$). If the Chairman made a decision that helped the environment, he was perceived as having more pro-environmental attitudes. No other effects of experimental conditions were statistically significant. Perceptions of the Chairman’s status from immoral to moral were also only significantly influenced by the help / harm condition ($B^* = .25$, $Z = 3.24$, $p = .001$). The Chairman was perceived as more moral if he helped the environment.

³ Note that B^* is used to indicate the standardized regression coefficient. Supplementary materials contain the unstandardized coefficients, for comparison.

Perceptions of the Chairman's desires lining up with the company's actions were significantly influenced by both being in the help versus harm condition ($B^* = -0.28$, $Z = -3.82$, $p < .001$) and being in the dominance versus low status condition ($B^* = .15$, $Z = 2.54$, $p = .011$). There were no other significant effects of the experimental conditions. Direction of the results indicate that people who were in the help condition were perceived to be acting *less* in line with their beliefs, as compared to those in the harm condition. People in the dominance condition, on the other hand, were perceived as acting *more* in line with their beliefs (as compared to those in the low status condition).

Only one of the three perceptions of the Chairman ultimately had a significant effect on judgments of intentionality: perceptions that the Chairman acted in line with his desires ($B^* = .61$, $Z = 14.32$, $p < .001$). Participants who perceived the outcome of the company policy as being more in line with the Chairman's desires also perceived the Chairman as acting more intentionally. Perceptions of the Chairman's attitudes towards the environment ($B^* = -0.01$, $Z = -0.13$, $p = .897$) and moral status ($B^* = .06$, $Z = 1.13$, $p = .260$) were not significantly related to judgments of intentionality.

Only two tests of overall mediated effects were statistically significant: the path from being in the help / harm condition, through perceptions of the Chairman's desires ($AB^* = -.17$, $Z = -3.65$, $p < .001$), and the path from being in the dominance (versus low status) condition, through perceptions of the Chairman's desires ($AB^* = 0.09$, $Z = 2.53$, $p = .011$). The first result indicates that being in the help condition reduced perceptions that the Chairman was acting in line with his own desires, and through this change in the perception of the Chairman's desire, judgments of intentionality were reduced. The second result indicates that being in the dominance condition (relative to the low status condition) increased perceptions that the

Chairman was acting in line with his desires, and through this change in the perception of the Chairman's desire, judgments of intentionality were increased. It is important to stress that these results are statistically consistent with such a causal interpretation, but that the model does not provide definitive proof of causal mediation. The model with all significant effects is illustrated in Figure 3.

Figure 3: Study 1 Mediation Model

Note: Only statistically significant effects are included in this model, for the sake of clarity of presentation. * indicates $p < .05$, ** indicates $p < .01$, *** indicates $p < .001$. Full results are available in online supplementary materials.

Discussion

Results of Study 1 did not confirm our hypothesis. In line with prior research, we did replicate the “Side Effect Effect.” Participants did judge the decision of the Chairman to harm

the environment as more intentional than the decision to help the environment. The Chairman's social status did not alter this effect. The "Side Effect Effect" was also observed for our secondary outcome, the morality of the Chairman. Like with the primary outcome, however, there was no interaction effect and therefore our hypothesis was not supported.

Our hypothesis was derived directly from the person-centered account of morality laid out by Uhlmann, Pizarro, and Diermeier (2015). These authors sought to explain the "Side Effect Effect" by positing that people give those who are likely to be valuable, cooperative group members the benefit of the doubt in moral judgment. In our experiment, the prestigious company represents a valuable cooperator, which should have led to its being given greater "moral license." If this explanation had been right, it should have manifested in a willingness to judge the prestigious company's decision to harm the environment as less intentional and the decision to help the environment as more intentional, thereby reducing the "Side Effect Effect." Instead, we found very little difference in judgments of prestigious people—in fact, the direction of the effects suggest that people held the prestigious to a slightly higher standard. This null result, by itself, does not suggest that the person-centered account of morality is wrong in all regards, but it does suggest that highly successful companies are not given moral license due to their status.

The exploratory statistical mediation analysis performed gives greater insight into what was going on in the data, and can be used to suggest further hypotheses. Sripada's (2011) earlier mediation analysis found that the values and attitudes of the Chairman towards the environment and the Chairman's desire were both statistically significant mediators of the effect of the help versus harm condition. In that study, the effect of the Chairman's morality were not a significant mediator. In this study, we replicated the mediation effect of the Chairman's desires: being in the helping condition was associated with a significant reduction in perceptions that the Chairman

was acting in line with his desires, and this in turn was associated with reduced judgments of intentionality. We also did not find that the morality (or immorality) of the Chairman's actions were a statistically significant mediator. However, Sripada's finding that the Chairman's values and attitudes mediated the effect of the help versus harm condition on judgments of intentionality was not replicated. Our study found that the Chairman was judged as being more pro-environment when the decision ended up helping the environment, but how anti to pro-environment the Chairman was did not have a statistically significant relationship to judgments of intentionality. Note also that the analysis we performed estimated the effects of the multiple mediators simultaneously, as is currently recommended in the literature (MacKinnon, 2008), as opposed to individually in separate models, as done by Sripada. This might influence the results, as in our analysis each mediation pathway captures effects above-and-beyond those of the other mediators.

A novel finding that emerged from this analysis was the mediated effect of the dominance condition (versus the low rank condition) on judgments of intentionality through perceptions of desire. When the company was described as dominant, this increased perceptions that the Chairman was acting in line with his desires, which in turn increased perceptions that the Chairman was acting intentionally. This suggests that people view dominant companies as acting more "willfully," and being less willing to take actions they don't want to. This judgment is in line with a theoretical understanding of the difference between dominance and prestige. Prestige is social status attained through other people's admiration, and so it must be earned through actions that others see as praiseworthy. This might lead a prestigious person (or organization) to not do exactly what they want, but instead to act in a way that benefits the broader group to earn praise. Dominance is social status attained through power and intimidation, and so it involves

taking actions that benefit the individual over the group. This would lead a dominant individual (or organization) to only take actions that they see as personally beneficial, and to be less likely to take actions that they didn't otherwise want to take. Theoretically, therefore, this study suggests that people perceive dominant individuals—as compared to low rank or prestigious individuals—as being less willing to compromise on their own desires. This is a useful direction for further follow-up.

Another set of novel findings emerged after statistically adjusting for the effects of perceptions of the Chairman. Although there was no overall interaction effect of the help versus harm condition and social status, there was a significant interaction between the help versus harm condition and the dominance condition above-and-beyond the mediation effects. The “Side Effect Effect” was reduced in the dominance condition, relative to the low status condition. It was not reduced for the prestige condition. Our prediction was the reverse of this: that the “Side Effect Effect” would be reduced in the prestige, but not the dominance condition. Our hypothesis had been based on the idea that prestigious groups, because they are successful and oriented towards the good of the group, would be given “moral license” to harm the environment. We thought that dominant groups, because they bully and intimidate, would not get such “moral license.” Yet these results suggest that the reduction in the “Side Effect Effect” was not due to moral licensing—which would be seen in viewing environmental harm as less intentional—but due to perceptions that dominant people were acting more intentionally when helping the environment. This is consistent with the explanation posited above: dominant people were seen as acting more “willfully,” or doing what they wanted independent of what others believe.

Overall, the results of the mediation analysis suggest a nuanced story. When the company is presented as dominant, the results of its actions are judged to be more in line with its desires.

This suggests that people believe dominant companies *do what they want* regardless of the consequences. Indeed, this seems to be a central theme of dominance: taking what is wanted, as opposed to working to have others grant it to you. Our pattern of results suggests that the actions of dominant companies are perceived as more in line with their desires, and that they are given more credit for causing positive outcomes. It is important to note that these results were not the core, hypothesized effects of the study. Comparing the results of the preregistered analysis with these preregistered exploratory analyses demonstrate that these effects were not due simply to the effects of the experimental conditions, but to the effects of the experimental conditions adjusted for other judgments and perceptions.

Study 2

In the second study, we tested whether social rank biases people's judgment of causal control using Alicke's (1992) "culpable causation" paradigm. The vignettes Alicke and colleagues have used to examine culpability was the other case study discussed at length by Uhlmann, Pizarro, and Diermeier (2015). The person-centered account was specified as an explanation of this effect. As in Study 1, we have chosen to remain as faithful to the original methodology as possible, to provide a fair test of theory.

Our experiment added a manipulation of the social rank of the person being judged in the scenario. Three conditions were included: a high prestige condition, a high dominance condition, and a low rank condition. As in Study 1, the critical test of our hypothesis is the significance of the interaction between the moral character of the individual being judged (good vs bad) and that person's rank. If the effect of moral character on judgments of culpability is dampened when the

person has high prestige and amplified when the person has high dominance, this would support the evolutionary account. It would demonstrate that character-based moral judgments are influenced by practical considerations of the benefits an individual provides to the group.

Methods

Participants. As in Study 1, participants were recruited from Mturk. Participants were paid \$1 for completing the 10-minute questionnaire. In Alicke's (1992) Study 1, the effect size was $d = .45$. To detect this effect with 95% power, we would need 130 participants per cell. To detect a "knock out" of the effect with our interaction we would need 1560 participants. We therefore proposed to collect 1560 participants. Based on our preregistered exclusion criteria, we excluded participants who did not answer our dependent variable question, and participants who completed the entire study in less than 60 seconds (on the basis that this would not give adequate time to read the full narrative and respond to questions). Some dropout was anticipated in this study, so we allowed enrollment in the study to go a bit above 1560, so that 1633 people ultimately started the survey. After excluding participants, we had 1591 participants remaining. Given that this was only 2% more than our planned sample size, we decided to analyze the entire sample together. Average age of participants was 39 years old ($SD = 13$). Participants were 44% female, 55% male, and 0.5% non-binary. Participants were 66% white, 14% black or African American, 10% Asian, 5% Hispanic or Latino, and 1% Hispanic or Latino and white. There were less than 10 participants in each of the other ethnic categories. Full responses are available online (<https://osf.io/u9pfr/>).

Procedures. Participants were randomly assigned to one of four different conditions. Each of the conditions varies two factors: whether the driver was pursuing a socially desirable or

undesirable motive, and whether the driver has low social rank, high prestige, or high dominance. An example of the scenario is written below:

(High prestige) You have a new boss at work named John, who was promoted because of his expert training and cooperative attitude. Other people in the office admire him.

(High dominance) You have a new boss at work named John, who was promoted because of his dominant personality and willingness to push others to get work done. Other people in the office are intimidated by him.

(Low rank) You have a new coworker named John. He just started and has no influence on your day to day work.

John was driving over the speed limit (about 40 mph in a 30-mph zone) in order to get home in time to ...

(Socially desirable) hide an anniversary present for wife that he had left out in the open before his family could see it.

(Socially undesirable) hide a vial of cocaine he had left out in the open before family could see it.

As John came to an intersection, he applied his brakes, but he was unable to stop as quickly as usual because of some oil that had spilled on the road. As a result, John hit a car that was coming from the other direction.

John hit the driver on the driver's side, causing him multiple lacerations, a broken collar bone, and a fractured arm. John was uninjured in the accident.

The bold text indicates the portions of the text that are manipulated. After reading the narrative, participants rated:

1. The extent to which John was responsible for the accident on a scale of 0 (not at all responsible) to 10 (very responsible)
2. The extent to which John was the cause of the accident on a scale of 0 (not at all the cause) to 10 (very much a cause).
3. How much money, if any, they think should be awarded to the driver that John hit.
Participants will be told that the other driver could be awarded between \$0 and \$100,000, and that the average amount in cases such as this, if John were believed to be at fault, would be about \$25,000.

These questions exactly mimic the outcomes used in Alicke's (1992) earlier research on this scenario. Alicke also included an open-ended question about the cause of the accident, but we did not include this question because we are basing our study on the quantitative results he reported. We will add one new question:

4. How much do you agree with the statement "What John did was morally wrong"?
rated on a scale from 1 to 7.

Additionally, Alicke manipulated the situational factors that influenced the accident, so that in one condition it was an oil spill (described above), in another it was a tree branch obstructing his view, and in a third it was the other car running a stop sign. We only use the oil spill condition because the effect of the socially desirable versus undesirable motive was largest in this condition. Our research question is specifically focused on the moral character of the individual, so we chose a larger effect that would be easier to replicate. We hypothesized that

there would be a significant interaction between John's social rank and the social desirability of his motives in predicting John's responsibility for the accident. We treated judgment of how much John caused the accident and how much John should pay the driver he hit as secondary outcomes. In Alicke's (1992) prior research, there was an effect of the social desirability manipulation on all of these outcomes. We will specify the response to Question 1 as our primary outcome. However, if our manipulation does not influence all of the outcomes in the same way, this would suggest that participants are interpreting these responses in differently from the way they were in the original study.

Results

Our main preregistered hypothesis was that there would be an interaction between social status and desirability of the individual in predicting judgments of that individual's responsibility. The test of the interaction was non-significant ($F(2, 1485) = 1.95, p = .142$, partial $\epsilon^2 = .00$, 95% CI [.00, .01]), indicating that our hypothesis was not supported. Replicating prior research, the Culpable Causation effect did appear, as indicated by a main effect of the desirability of the individual's traits ($F(1, 1485) = 6.14, p = .013$, partial $\epsilon^2 = .00$, 95% CI [.00, .01]). There was a main effect of social status that was just significant ($F(2, 1485) = 3.02, p = .049$, partial $\epsilon^2 = .00$, 95% CI [.00, .01]). A follow-up Tukey's Honestly Significant Difference test comparing the social status groups found no significant differences after adjusting for multiple comparisons. Results are visualized in Figure 4A.

We preregistered a plan to analyze two secondary outcomes previously used in the literature: whether the individual was perceived as having caused the accident, and the dollar

value of damages participants said should be awarded for the accident. The interaction effect on perceptions of causation was statistically significant ($F(2, 1485) = 4.37, p = .013$, partial $\epsilon^2 = .01$, 95% CI [.00, .02]). The main effect of desirability was not significant ($F(1, 1485) = 1.58, p = .210$, partial $\epsilon^2 = .00$, 95% CI [.00, .01]), but the main effect of social status was ($F(2, 1485) = 4.00, p = .018$, partial $\epsilon^2 = .01$, 95% CI [.00, .01]). Pairwise comparisons of all cells were conducted as a follow-up, using Tukey's Honestly Significant Difference test. Results indicate that there was a "Culpable Causation Effect" for the low rank condition ($difference = 0.75, p < .001$) and for the prestige condition ($difference = 0.89, p < .001$), but not for the dominance condition ($difference = 0.22, p = .809$). This was because the dominant individuals were perceived as having caused the outcome at high rates *whether or not* they were in the socially desirable condition or not. This effect can be seen in Figure 4B.

The interaction effect on judgments of damages to be awarded was not statistically significant ($F(2, 1471) = 1.57, p = .209$, partial $\epsilon^2 = .00$, 95% CI [.00, .01]). Neither was the main effect of social status ($F(2, 1471) = 0.68, p = .505$, partial $\epsilon^2 = .00$, 95% CI [.00, .01]). The main effect of desirability was statistically significant ($F(2, 1471) = 18.13, p < .001$, partial $\epsilon^2 = .01$, 95% CI [.00, .03]). The main effect of desirability indicates that the Culpable Causation Effect was replicated using this outcome variable (See Figure 4C).

One other exploratory item was added to this study, which asked participants to judge whether the target in the story was moral. There was no interaction effect predicting this effect ($F(2, 1485) = 1.37, p = .254$, partial $\epsilon^2 = .00$, 95% CI [.00, .01]), but there were main effects of social status ($F(2, 1485) = 3.03, p = .037$, partial $\epsilon^2 = .00$, 95% CI [.00, .01]) and of the desirability condition ($F(1, 1485) = 52.03, p < .001$, partial $\epsilon^2 = .03$, 95% CI [.02, .05]). The main effect of the desirability condition mimics the overall Culpable Causation Effect, with the

acts of socially desirable individuals being judged as less wrong (see Figure 4D). The difference between the social status conditions was further analyzed using a Tukey's Honestly Significant Difference test. After adjusting for the number of tests being performed, Tukey's test did not find a statistically significant difference between any of the status conditions. The closest was the comparison of the dominance and prestige conditions ($\text{difference} = -0.18, p = .115$), which indicated that the actions of prestigious people were perceived as slightly less wrong overall.

Figure 4: Effects of the Experimental Conditions in Study 2

A: Primary Outcome, Judgments of Responsibility

B: Secondary Outcome, Judgments of Causality

C: Secondary Outcome, Damages Awarded

D: Secondary Outcome, Morality of Action

Note: Results are presented as boxplots, with median point “cinched” to accentuate comparisons. Raw data is plotted over the boxplot with partial transparency and a slight jitter to aid visualization.

Discussion

Study 2 did not support our preregistered hypotheses. The Culpable Causation Effect was found for the primary outcome, but social status did not moderate this effect. The same pattern of results was found for one of our secondary outcomes, the amount of damages to be awarded for the accident. The original Culpable Causation Effect was found, but it was not moderated by social status. For our second and third variables, there were statistically significant interaction effect, indicating that the Culpable Causation Effect was moderated by social status. However, our prediction was that participants in the prestige condition would be granted a “moral license” due to their ability to contribute to the group’s outcome and prosocial orientation. Instead, we

found “moral penalties” for the dominant individual. Even if he was rushing home for a prosocial reason—to hide an anniversary gift for his parents—he was seen as more culpable and forced to pay more damages if he was dominant. These results, like those of Study 1, should not be taken as invalidating the person-centered account of morality. Instead, they should be taken as evidence that the person-centered account cannot explain why the “Side Effect Effect” and Culpable Causation Effect occur.

One unexpected effect does stand out from this experiment: individuals who were dominant were perceived as having been more responsible for the car accident, having caused the car accident to a greater degree, and having performed a morally worse action. This effect was driven by being less lenient on dominant individuals when they were rushing home to hide an anniversary gift. This effect was not present for damages awarded. This pattern of results suggests that participants saw dominant individuals as particularly efficacious in causing outcomes. This effect might be interpreted as dominance superseding social desirability when making ratings of causality. Even though people pursuing “good” ends are generally seen as less likely to cause bad outcomes in general, that moral license goes away if the individual is dominant.

General Discussion

This manuscript presents the results of two studies designed to test a theoretical explanation for two unusual findings in the moral psychology literature: the “Side Effect Effect” and the Culpable Causation Effect. The person-centered account of moral judgment suggests that, contrary to dominant moral theories in psychology and experimental philosophy that focus

on the morality of specific acts, moral judgments are driven to a significant degree by attempts to uncover the underlying character of moral actors. Drawing on evolutionary theory, the person-centered account suggests that judgments about the morality of actions are influenced by broader considerations about whether an individual can contribute to the general success of the perceiver's group. The overall power of the individual to contribute to the group should be directly relevant to moral judgment, even though, from the "local" point of view of the specific act, such information is not relevant. Manuscripts laying out the theory of person-centered moral judgments specifically invoke it to explain the "Side Effect Effect" and the Culpable Causation Effect (Pizarro & Tannenbaum, 2011; Uhlmann, Pizarro, & Diermeier, 2015). Our tests of two predictions derived from the person-centered account of morality failed to support the theory. Companies and people who are likely to be particularly valuable allies and contributors to the broader group's success are not given moral license.

Instead, our two studies uncovered another unusual consistency. Participants who achieved status via dominance, as opposed to prestige, were judged as acting more in line with their own desires, and were judged as having more causal power on the outcome of a car crash. Together, these findings suggest that participants viewed dominant individuals as having greater agency. As mentioned in the discussion of Study 1, this fits with the theoretical distinction between dominance and prestige. Prestigious individuals rely on others to grant them status, and so are more likely to alter or moderate their behavior to better account for the needs or goals of others. Dominant individuals, on the other hand, claim status through intimidation and power. People who achieve dominance should therefore be thought of as being particularly effective in achieving their ends. They pursue what they want, and they are likely to get it.

In the context of moral decision-making, these results suggest that social status is likely to play a greater role in influencing perceptions of overall causal efficacy. If there is a dilemma that hinges on determining if an individual wanted to or was able to affect an outcome, perceiving that individual as dominant (versus not) is likely to sway the moral judgment. Prestige, which we thought would be associated with a sort of halo of moral licensing, did not have any effects. So we would not expect more prestigious people to be judged any differently from low status individuals. Of course, our new hypotheses regarding dominance are generated from these results, and so require new data to be adequately tested. However, they do demonstrate the importance of failed hypothesis tests for generating new hypotheses.

These studies were carefully designed to replicate and extend previous findings. They therefore used a very specific experimental paradigm: judgments made online, by U.S. adults reading short vignettes, and answering questions about them. There are certainly other ways to approach this domain area, such as analyzing decisions made by individuals regarding real-life scenarios with moral stakes, or by analyzing changing perceptions of individuals as they gain status over time. Our findings may inform hypotheses using other methods, but they do not preclude finding different outcomes using alternate paradigms. Our sample was also drawn from the same pool as well-known studies on this topic, but this pool does not represent all individuals for whom this hypothesis is relevant. It might be that among individuals from different socio-economic statuses or from different cultures, results might differ. Research using different paradigms and samples would therefore be a welcome addition to the literature, either for establishing the original effects or for examining how these judgments interact with social status.

A further consideration for future research is how low a low-status individual or group should be. Our low status conditions include a less successful company, and a new coworker

without much power at his organization. However, these are still a company and person in good standing with the group. Results might have been different for individuals who are seen as detriments to the group's success, or are very low status. More broadly, an experimental (or observational) paradigm that samples more fully from individuals with a range of social statuses—from very low to very high—might provide more insight into broader effects of status on moral judgment.

Overall, this manuscript suggests a limitation on the scope of person-centered morality. People may allow their moral decision-making to be swayed by perceptions of the underlying character of individuals, but our evidence suggests that this character does not include broader considerations of the person's status and ability to contribute materially to the group's success. Instead, it may be that moral character is judged separately from considerations of status and power. Moral license may be granted to people who have done good deeds previously, even though it is not granted to people who have status.

Acknowledgements. Special thanks to Simine Vazire, for contributing her discretionary funds to the completion of this research project. Thanks to Sarah Schiavone for assistance with online data collection.

References

- Alicke, M. D. (1992). Culpable causation. *Journal of Personality and Social Psychology*, 63(3), 368-378.
- Bartneck, C., Duenser, A., Moltchanova, E., & Zawieska, K. (2015). Comparing the similarity of responses received from studies in amazon's mechanical turk to studies conducted online and with direct recruitment. *PloS one*, 10(4), e0121595.
- Cohen, T. R., & Morse, L. (2014). Moral character: What it is and what it does. *Research in Organizational Behavior*, 34, 43-61.
- Cohen, T. R., Panter, A. T., Turan, N., Morse, L., & Kim, Y. (2014). Moral character in the workplace. *Journal of Personality and Social Psychology*, 107, 943-963.
- 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, 171-191.
- Giner-Sorolla. (2018, January 24). Powering Your Interaction. [blog post]. Retrieved from <https://approachingblog.wordpress.com/2018/01/24/powering-your-interaction-2/>.
- 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.
- Greene, J. D., Morelli, S. A., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2008). Cognitive load selectively interferes with utilitarian moral judgment. *Cognition*, 107(3), 1144–1154.
- Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, 293(5537), 2105–2108.

- Hauser, D. J., & Schwarz, N. (2016). Attentive Turkers: MTurk participants perform better on online attention checks than do subject pool participants. *Behavior research methods*, 48(1), 400-407.
- Henrich, J., & Gil-White, F. J. (2001). The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior*, 22(3), 165–196.
- Jayawickreme, E., & Fleeson, W. (2017). Does whole trait theory work for the virtues? In W. Sinnott-Armstrong & C. Miller (Eds.), *Moral Psychology, Volume 5: Virtue and Character* (75–104). Cambridge, MA: MIT Press.
- Knobe, J. (2003). Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology*, 16(2), 309–324.
- Knobe, J. (2006). The concept of intentional action: A case study in the uses of folk psychology. *Philosophical Studies*, 130(2), 203–231.
- Kohlberg, L. (1969). *Stages in the development of moral thought and action*. New York: Holt, Rinehart & Winston.
- Miller, C. (2013). *Moral character: An empirical theory*. Oxford: Oxford University Press.
- Miller, C. (2016). Empirical approaches to moral character. Stanford Encyclopedia of Philosophy. Retrieved from <https://plato.stanford.edu/entries/moral-character-empirical/>
- Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science*, 314(5805), 1560–1563.
- Ohtsuki, H., & Iwasa, Y. (2004). How should we define goodness?—reputation dynamics in indirect reciprocity. *Journal of Theoretical Biology*, 231(1), 107–120.
- Piaget, J. (1965). The moral judgment of the child (1932). *New York: The Free.*

- Pizarro, D. A., & Tannenbaum, D. (2011). Bringing character back: How the motivation to evaluate character influences judgments of moral blame. *The Social Psychology of Morality: Exploring the Causes of Good and Evil*, 91–108.
- Sinnott-Armstrong W., & Miller, C. (2017). *Moral psychology, Volume 5: Virtue and character*. Cambridge: Cambridge, MA: MIT Press.
- Sripada, C. S. (2012). Mental state attributions and the side-effect effect. *Journal of Experimental Social Psychology*, 48(1), 232–238.
- Sripada, C. S., & Konrath, S. (2011). Telling more than we can know about intentional action. *Mind & Language*, 26(3), 353–380.
- Tannenbaum, D., Uhlmann, E. L., & Diermeier, D. (2011). Moral signals, public outrage, and immaterial harms. *Journal of Experimental Social Psychology*, 47(6), 1249–1254.
- Uhlmann, E. L., Pizarro, D. A., & Diermeier, D. (2015). A Person-Centered Approach to Moral Judgment. *Perspectives on Psychological Science*, 10(1), 72–81.
<https://doi.org/10.1177/1745691614556679>
- Woo, S. E., Keith, M., & Thornton, M. A. (2015). Amazon Mechanical Turk for industrial and organizational psychology: Advantages, challenges, and practical recommendations. *Industrial and Organizational Psychology*, 8(2), 171-179.