*Empirical Research Paper*

# Perceived Power Polarizes Moral Evaluations

Personality and Social Psychology Bulletin
1–16

DOI: 10.1177/01461672241245181
journals.sagepub.com/home/pspb
Sage

**Russell Roberts and Alex Koch**

**Abstract**

We show an interactive effect of perceiver-target similarity in ideological beliefs and target power on impressions of target morality. Consistent with prior research, perceivers rated targets with dissimilar ideologies as less moral than targets with similar ideologies, but this difference in ratings was magnified for powerful targets relative to less powerful targets. We argue that these results emerged because perceivers expected similar-ideology, powerful (vs. powerless) targets to help the self more, and expected dissimilar-ideology, powerful (vs. powerless) targets to hurt the self more. We establish this effect when people evaluate politicians (Study 1), groups, and individuals (Studies 2a-2b); demonstrate its predictive power over other kinds of interpersonal similarity; and show that it affects morality judgments uniquely when compared with other consequential dimensions of social evaluation. Finally, we manipulated power experimentally and showed the interaction when the difference between high- and low-power manipulations was controlled over just $1 (Studies 3-4).

**Keywords**

social evaluation, belief similarity, power, morality

Received June 7, 2023; revision accepted March 15, 2024

People want to be seen as moral (Hauke & Abele, 2020). To this end, they tend to behave ethically, join groups they believe to be moral and contribute work and other resources to moral in-groups (Ellemers et al., 2013; Van Prooijen & Ellemers, 2015). Appearing moral has benefits. People prioritize judging others' morality (Leach et al., 2007), and they do so spontaneously (Koch, Imhoff, et al., 2020; Slepian & Koch, 2021), sooner, faster, and more often than judging others' sociability and competence (Nicolas et al., 2022; Willis & Todorov, 2006). They factor their impressions of morality most heavily in overall evaluations of groups (Brambilla et al., 2021) and other individuals (Brambilla et al., 2011; Goodwin et al., 2014; Wojciszke et al., 1998) and tend to share resources and cooperate with others they see as moral (Jenkins et al., 2018; Koch, Dorrough, et al., 2020). In contrast, impressions of immorality evoke powerful negative emotions such as moral outrage, which motivates people to dehumanize and punish transgressors, sometimes via collective action (Bastian et al., 2013; Crockett, 2017). In sum, morality matters more than other aspects of general evaluation (Brambilla et al., 2021; see also Koch et al., 2021), and both moral approbation and moral condemnation can powerfully influence how people behave toward one another.

## Ideological Belief Similarity Predicts Moral Evaluations

Given the spontaneous and consequential nature of moral impressions (Brambilla et al., 2021), it is important to understand their antecedents. What makes people think that others are moral? Although factors such as target characteristics (e.g., attractiveness, smiling) and perceiver traits (e.g., optimism, mood) contribute to positive moral evaluations, people do not always agree on who or what should be considered moral. One person's ally can be another person's foe, and some degree of moral evaluation hinges on who the perceiver and the target are in relation to one another (Koch, Imhoff, et al., 2020). That is, perceived morality is also a function of who evaluates whom, with greater similarity between them leading to more positive moral judgments (Montoya et al., 2008; Pinsof et al., 2023).

[1]The University of Chicago, IL, USA

**Corresponding Author:**
Russell Roberts, Booth School of Business, The University of Chicago, 5807 S Woodlawn Avenue, Chicago, IL 60637, USA.
Email: rrobert2@chicagobooth.edu

There are many ways in which two people can be similar (Medin et al., 1993). Although superficial similarities between people (e.g., mimicry) can affect moral evaluation (Bocian et al., 2018), some sources of similarity are more germane to this judgment than others (Imhoff et al., 2018; Koch et al., 2018; Zorn et al., 2022). Ideological beliefs encompass the shared core values, beliefs, and goals that people express through their political affiliation and opinions (Brandt & Crawford, 2020). For simplicity, we will refer to this construct as "belief similarity." Given that people's ideological beliefs include philosophically normative content (e.g., the goals people should [not] pursue, and the behaviors that are [not] permissible), belief similarity should affect moral evaluations to a greater extent than other forms of similarity. Indeed, people spontaneously evaluate others' ideological beliefs (Koch et al., 2016; Nicolas et al., 2022) and evaluate them as more moral if their beliefs appear more similar to the ideology of the self (Brandt & Crawford, 2020; Woitzel & Koch, 2023). Moreover, belief similarity predicts moral evaluation to a greater extent than, for example, shared group membership (Koch, Dorrough, et al., 2020), and it predicts perceived morality better than other aspects of general evaluation, including perceived sociability and competence (Bocian et al., 2018; Woitzel & Koch, 2023). The present research contextualizes this effect of belief similarity on moral evaluations.

## Power May Amplify the Effect of Belief Similarity on Moral Evaluations

Belief similarity matters for moral impressions because it signals a target's possible intentions toward the perceiver. That is, perceivers may expect more favorable treatment from others whose beliefs are similar to their own but may expect less favorable treatment from those whose beliefs differ from their own. Yet others' intentions only affect us to the extent that they can be carried out (Fiske et al., 2002). This research focuses on power as a measure of a target's ability to enact their intentions. Put simply, powerful people can carry out their intentions easily, whereas powerless people cannot. As a psychological construct, power has been defined in a variety of ways (Dépret & Fiske, 1993; Magee & Smith, 2013; Thibaut & Kelley, 1959). Our research does not focus on the effects of power *per se* but rather on the effects of perceiving another social actor as powerful (or not) in a way that is consistent with many of these definitions. To test our proposed mechanism, we experimentally manipulate power in terms of control over resources and, later, in terms of the effort required to exercise that control. Thus, for our purposes, we follow the common definition of power as "control over others' valued outcomes" (Dépret & Fiske, 1993).

We predict that power widens the gap in morality ratings between targets with beliefs that are similar versus different to those of the perceiver because people's impressions of others' morality depend on not just the anticipated direction of their actions (helping vs. harming the self) but also the anticipated magnitude of those actions (influencing the self a great deal vs. hardly at all; Cushman, 2008, 2015). We propose that people engage in a kind of consequentialist reasoning when rendering moral evaluations such that the expected impact of the target's enacted intentions—how much of a threat or opportunity they pose—is factored into their judgments (Brambilla et al., 2021; Carrier et al., 2019; Landy et al., 2016). As a result, perceivers' impressions of powerful targets' morality should be more extreme because powerful targets impact their fortunes to a greater extent. In contrast, their impressions of powerless targets' morality would be less extreme, consistent with their limited impact on their outcomes and well-being.

## How Does the Present Research Develop Previous Theory and Evidence?

Our hypothesis is rooted in the functionalist framework of social evaluation by which people's overall evaluation of others depends on whether they perceive their intentions as moral and whether they perceive them to be sufficiently able to carry out their intentions (e.g., Fiske et al., 2002; Wojciszke et al., 1998). More specifically, this hypothesis relates to the Moral Primacy Model (Brambilla et al., 2021) according to which morality judgments have a higher priority than judgments of competence and sociability, which function as personal and interpersonal proficiencies whose positivity depends on how moral the target is believed to be. For example, an article by Landy and colleagues (2016) finds that competence and sociability enhance overall impressions of moral targets to a greater extent than those of immoral targets. Similarly, work by Carrier and colleagues (2019) finds that perceivers expecting to cooperate with a target evaluate competent targets as warmer than incompetent targets but evaluate competent targets as less warm than incompetent targets when they expect to compete with them. Both findings highlight that a target's intentions affect how people factor the target's capabilities into their overall impression of the target. Our hypothesis complements these findings in that we posit another dimension that captures a target's ability to enact their intentions, power, as similarly moderating social impressions.

We extend the functionalist evaluation logic by hypothesizing that morality is not only a moderator of how other dimensions of social evaluation factor into general impressions, but that moral evaluations are themselves influenced by impressions of the target's ability to enact their intentions. By theorizing perceived belief-similarity as a key antecedent to moral impressions, we highlight the relational nature of these judgments (Bocian et al., 2018; Koch, Imhoff, et al., 2020; Pinsof et al., 2023) and provide a framework for

**Table 1.** Results of Study 1.

| Attitudes toward the target by . . . | *b* and 95% CI [LB, UB] | *t* | *p* |
|---|---|---|---|
| Perceiver-target similarity in beliefs | 0.86 [0.84, 0.87] | 95.31 | < .001 |
| Target power | −0.13 [−0.15, −0.10] | −14.21 | < .001 |
| Belief similarity * Power | 0.97 [0.94, 1.01] | 54.07 | < .001 |
| Power when Belief similarity = same | 0.36 [0.34, 0.38] | 29.78 | < .001 |
| Power when Belief similarity = different | −0.61 [-0.64, -0.59] | −46.18 | < .001 |

*Note.* *B* = estimate. 95% CI [UB, LB] = 95% confidence interval [lower bound, upper bound].

thinking about them not just as inputs into appraisals of the threat or opportunity posed by a social target but also as judgments that can themselves be influenced by the magnitude of that threat or opportunity.

### *Statistical Power*

We used the R package *simr* (Green & MacLeod, 2016) to run 100 simulation-based sensitivity power analyses for all central effects, namely moderations of the main effect of belief similarity on morality. When aiming for a small effect ($b = .2$ for standardized variables) for correlational analyses, and a medium-sized effect ($b = .5$) for experiments (when participants' judgments are likely to be influenced by manipulated variables to a greater degree), the paper's power of testing those moderations was always $1\text{-}\beta >= 0.95$.

### *Open Science*

Hypotheses and analyses for Studies 3 and 4 were preregistered (link for Study 3; and link for Study 4). All studies report all conditions and measures. None of the studies collected data after analyzing it. Studies 2a, and 2b excluded a few participants who recommended that their data not be analyzed. Studies 3 and 4 excluded a few participants who did not meet the eligibility criterion (identifying as liberal or conservative), as pre-registered. All studies standardized all variables to allow direct comparison of effect sizes across variables and studies. Some figures show unstandardized means for ease of interpretation. All study materials, data, code, and results are available here and on the Open Science Foundation website (link).

## Study 1

If power moderates the effect of belief-similarity on moral evaluations of others, people's ratings of targets with similar and different beliefs should show greater divergence for more powerful targets than for less powerful targets. Our initial test of this hypothesis examined the ratings of elected officials of different ranks and different ideological beliefs. Powerful (vs. less powerful) officials should engender more polarized ratings as a function of rater-target belief similarity.

### *Method*

The American National Election Studies (ANES, 2015) are nationally representative surveys of U.S. residents. ANES measured public opinions on political parties, candidates, and so on during most national election periods between 1948 and 2020. We analyzed the feeling thermometer ratings of presidents (high power; coded as 0.5) and respondents' incumbent district representatives in the House of Representatives (low power; coded as −0.5). These ratings were available for 19,004 people surveyed between 1978 and 2012. People rated on a 100-point scale how cold to warm they felt toward the targets, with 50 being neither cold nor warm. Feeling thermometer ratings are a standard measure of attitudes toward a social target (Krosnick et al., 2005), and attitudes toward targets are closely associated with moral evaluations of those targets (Bocian & Wojciszke, 2014; Brambilla & Leach, 2014; Leach et al., 2007). Thus, Study 1 took feeling thermometer ratings as a satisfactory approximation of target morality (ANES does not measure target morality for Congressional representatives). People could respond "don't know" if they were unfamiliar with a target; we excluded these responses. To determine perceiver–target similarity in beliefs, we matched their party affiliation (same, coded as 0.5 vs. different, coded as −0.5). If a person's affiliation was neither Republican nor Democrat, we excluded them.

### *Results*

In a linear mixed model with random intercepts for perceivers, we predicted attitudes toward the target from three fixed effects: target power (high: president, low: representative), perceiver–target similarity in beliefs (same party affiliation vs. one Republican, one Democrat), and their interaction. Table 1 and Figure 1 confirm the main effect of belief similarity and the interaction between power and belief similarity (Table 1, Figure 1). A contrast also showed that people rated high-power, same-beliefs targets more positively than low-power, same-beliefs targets; another contrast confirmed that people rated high-power, different-beliefs targets less positively than low-power, different-beliefs targets.

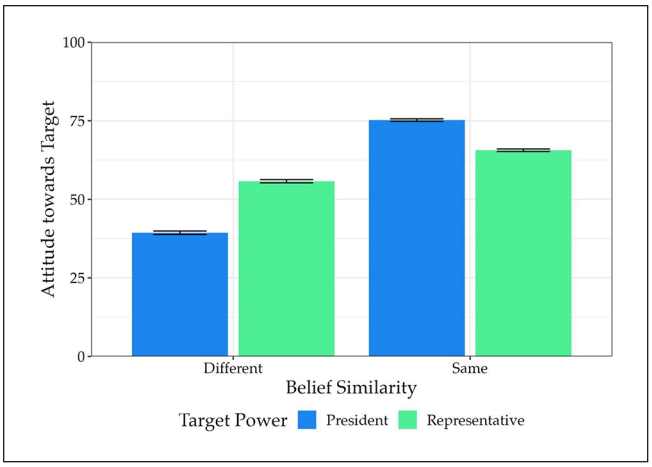


**Figure 1.** Feeling thermometer ratings (0: *very cold or unfavorable* – 100: *very warm or favorable*) (Study 1). Error bars indicate 95% confidence intervals.

### *Discussion*

Study 1 analyzed thousands of people who gave their opinions on thousands of politicians between 1978 and 2012. People rated high-power targets more positively than low-power targets when evaluating targets with similar beliefs, and they rated high-power targets less positively than low-power targets when evaluating targets with different beliefs.

Target power is likely confounded with target familiarity, and thus Study 1's effect of target power interacting with perceiver-target similarity in beliefs could actually be an effect of target familiarity interacting with perceiver-target similarity in beliefs. However, people in Study 1 could respond that they "don't know" a target. Study 1 excluded these responses, reducing this possibility to some extent.

According to our hypothesis, target power amplifies the effect of perceiver-target similarity in beliefs on target morality. However, Study 1 relied on measures of overall attitudes toward the target instead of the narrower construct of target morality specified in our hypothesis. Studies 2a-b address this limitation by measuring target morality directly rather than indirectly as in Study 1. In addition, Studies 2a-b examine the hypothesized interaction considering two other sources of similarity (shared group membership and living in the same U.S. state) as well as how this interaction affects two other consequential components of social evaluation (competence and sociability).

## Studies 2a and 2b

### *Method*

*Participants.* Study 2a recruited 658 U.S. residents on Prolific Academic. We excluded 28 participants who at the end of the study recommended to not analyze their data, leaving a final sample of 630 participants (383 men, 239 women, 8 Other/Prefer not to say; $M_{age}$ = 43.3).

Study 2b recruited 600 U.S. residents on Prolific Academic. We excluded 55 participants who at the end of the study recommended to not analyze their data, leaving a final sample of 545 participants (313 men, 226 women, 6 Other/Prefer not to say; $M_{age}$ = 44.0).

*Procedure in Study 2a.* Participants compared their own beliefs to the beliefs of 30 groups in society ("Compared to members of this group, my core values, beliefs, and goals are . . ."") using a slider scale from "Very Different" (0) to "Very Similar" (100). In addition, they rated whether or not they consider themselves a member of each group (no or yes, a binary measure). The order of these two measures was randomized. In a pilot study, roughly 600 people listed 20 types of people that they thought today's society (i.e., the U.S. in 2022) categorized into groups. People in Study 2a rated the 30 most frequently listed groups (i.e., an arguably society-representative sample of groups), which included Democrats, Republicans, Christians, rich people, LGBTQ+ people, poor people, students, Black people, young people, elderly people, White people, Hispanic people, blue-collar workers, Asian people, athletes, adults, middle-class people, Muslims, women, Jews, scientists, artists, men, atheists, parents, celebrities and influencers, teachers, politicians, immigrants, and military and veterans.

Next, participants used a slider scale from "Least Powerful" (0) to "Most Powerful" (100) to rate how powerful each group is with power defined as "how much control/influence over others this group has relative to other groups in society." Next and in randomized order, participants used three slider scales from "Not at all [. . .]" (0) to "Very [. . .]" (100) to rate how "Moral," "Competent," and "Sociable" each group is relative to others in society. Finally, participants provided demographic information, including their age and gender.

*Procedure in Study 2b.* Participants listed two adult individuals whom they saw as powerful ("they influence many others in society"), and two adult individuals whom they saw as powerless ("they influence hardly anyone in society"). For one powerful and one powerless individual, we asked the participants to list someone whose core values, beliefs, and goals are "very similar" to the participant's own core values, beliefs, and goals. For the other two individuals, participants listed someone whose core values, etc. are "very different" to the participant's own core values and so on. Participants listed four additional adult individuals that crossed a power manipulation (powerful vs. powerless as before) with a manipulation of similarity with respect to residence (i.e., living in the same vs. a different U.S. state). The order of listing the four individuals that crossed power and belief similarity

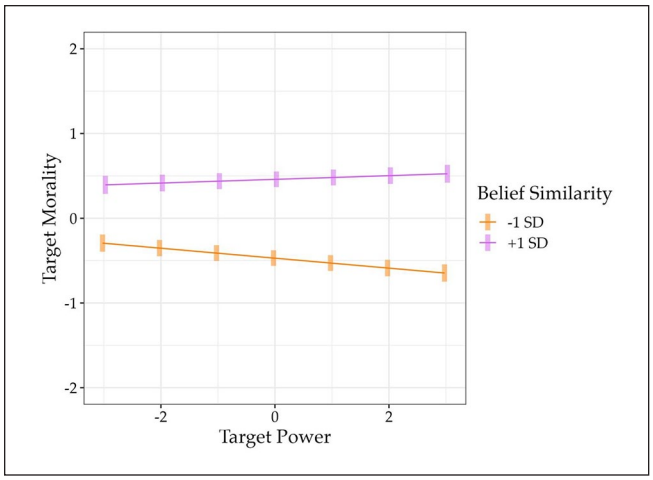


**Figure 2.** Ratings of target morality as a function of target power and perceiver-target belief similarity (Study 2a). Error bars indicate 95% confidence intervals.

versus that crossed power and residential similarity was randomized.

On the next three randomly ordered survey pages, participants used the same slider scales as in Study 2a to rate the morality, sociability, and competence of the eight individuals. On each survey page, the eight individuals appeared one below the other in random order. Finally, participants provided demographic information, including their age and gender.

## Results

In Study 2a, we ran a linear mixed model with random intercepts for both the perceivers and the targets. The model predicted the perceivers' ratings of the targets' morality from target power, perceiver-target belief similarity, and their interaction. This analysis revealed the predicted main effect of perceiver–target belief similarity ($b = 0.46$, 95% confidence interval [CI] = [0.45, 0.48], $t = 80.80$, $p < .001$) and the predicted interaction between belief similarity and target power ($b = 0.04$, 95% CI = [0.03, 0.05], $t = 8.61$, $p < .001$). Simple slope analysis (Aiken & West, 1991) probed this interaction. The coefficient of the effect of power on target morality was positive for targets at one standard deviation above the belief similarity mean ($b = 0.02$, 95% CI = [0.004, 0.039], $t = 2.43$, $p = .015$), compared to the negative coefficient for targets at one standard deviation below the belief similarity mean ($b = 0.06$, 95% CI = [−0.07, −0.04], $t = 7.94$, $p < .001$). Thus, target power amplified the effect of belief similarity on target morality (Figure 2). We also fit an identical model replacing perceiver–target belief similarity with shared group membership and found a similar pattern of results ($b_{membership} = .44$, 95% CI = [0.41, 0.47], $t = 30.94$, $p < .001$; $b_{power*membership} = .05$, 95% CI = [0.02, 0.07], $t = 3.55$, $p < .001$).

Next, we compared the main and interactive effects of shared group membership and belief similarity in the same model and when the dependent measure was either perceived morality, sociability, or competence. To this end, we fit three linear mixed models with random intercepts for the perceivers and targets in Study 2a. Each model included five fixed effects: group membership shared between the perceiver and the target (no = −0.5, yes = 0.5), belief similarity between the perceiver and the target, target power, and the interactions between target power and both other predictors. These models revealed significant interactive effects of target power and belief similarity on morality and, to a lesser extent, competence, but did not reveal significant interaction effects of target power and shared group membership (Table 2).

Next, we examined how people in Study 2b rendered these judgments when thinking of specific individuals rather than societal groups. At the same time, we pit belief similarity against residential similarity instead of shared group membership. We modeled Study 2b's data in three linear mixed models with random intercepts for the perceivers and the targets. Perceived morality was the dependent measure in one model, and perceived sociability and competence were the dependent measures in the other two models. The predictors in each model were all main effects, two-way interactions, and the three-way interaction between perceiver-target similarity (similar vs. different), target power (more vs. less powerful), and type of similarity (beliefs vs. residence).

In the model that predicted perceived morality, the three-way interaction was significant ($b = 0.45$, 95% CI = [0.25, 0.65], $t = 4.22$, $p < .001$). Contrasts revealed that target power depressed morality judgments when the beliefs of the perceiver and target were dissimilar ($b = 0.36$, 95% CI = [−0.46, −0.26], $t = 7.00$, $p < .001$). When their beliefs were similar, descriptive results showed the opposite effect, but this test failed to reach statistical significance ($b = 0.10$, 95% CI = [−0.003, 0.20], $t = 1.89$, $p = .058$). As a result, target power amplified the effect of belief similarity on perceived target morality (Figure 3). Target power influenced morality judgments neither when the perceiver and target shared state residency ($b = 0.03$, 95% CI = [−0.07, 0.13], $t = 0.51$, $p = .609$), nor when they resided in different states ($b = 0.02$, 95% CI = [−0.08, 0.12], $t = 0.46$, $p = .647$). So target power did not moderate the effect of residential similarity on perceived target morality. Table S1 in the Supplementary Materials shows all main effects, two-way interactions, and the three-way interaction in all three models.

In the two models that predicted perceived sociability and competence, the three-way interaction was significant as well (sociability: $b = 0.27$, 95% CI = [0.06, 0.47], $t = 2.52$, $p = .012$; competence: $b = 0.40$, 95% CI = [0.20, 0.60], $t = 3.97$, $p < .001$). However, the patterns of means in these two models differed from the pattern of means in the model that predicted perceived morality. Specifically, target power improved perceived sociability and competence regardless

**Table 2.** Results of Study 2a.

| Model | *b* and 95% CI [LB, UB] | *t* | *p* |
|---|---|---|---|
| DV = Morality | | | |
| Perceiver-target group membership | −0.05 [−0.08, −0.02] | 3.62 | < .001 |
| Perceiver-target belief similarity | 0.48 [0.46, 0.49] | 72.73 | < .001 |
| Target power | −0.02 [−0.03, −0.003] | 2.37 | .018 |
| Group membership * Target power | 0.01 [−0.01, 0.04] | 0.94 | .349 |
| Belief similarity * Target power | 0.04 [0.03, 0.05] | 7.25 | < .001 |
| DV = Sociability | | | |
| Perceiver-target group membership | −0.05 [−0.08, −0.01] | 2.82 | .005 |
| Perceiver-target belief similarity | 0.24 [0.23, 0.26] | 32.51 | < .001 |
| Target power | 0.14 [0.12, 0.15] | 16.45 | < .001 |
| Group membership * Target power | 0.01 [−0.02, 0.04] | 0.68 | .494 |
| Belief similarity * Target power | −0.01 [−0.02, 0.002] | 1.71 | .087 |
| DV = Competence | | | |
| Perceiver-target group membership | −0.02 [−0.05, 0.007] | 1.52 | .129 |
| Perceiver-target belief similarity | 0.40 [0.38, 0.41] | 58.27 | < .001 |
| Target power | 0.09 [0.07, 0.10] | 11.41 | < .001 |
| Group membership * Target power | −0.01 [−0.04, 0.01] | -1.04 | .300 |
| Belief similarity * Target power | 0.01 [0.001, 0.02] | 2.16 | .031 |

*Note.* *b* = estimate. 95% CI [UB, LB] = 95% confidence interval [lower bound, upper bound].

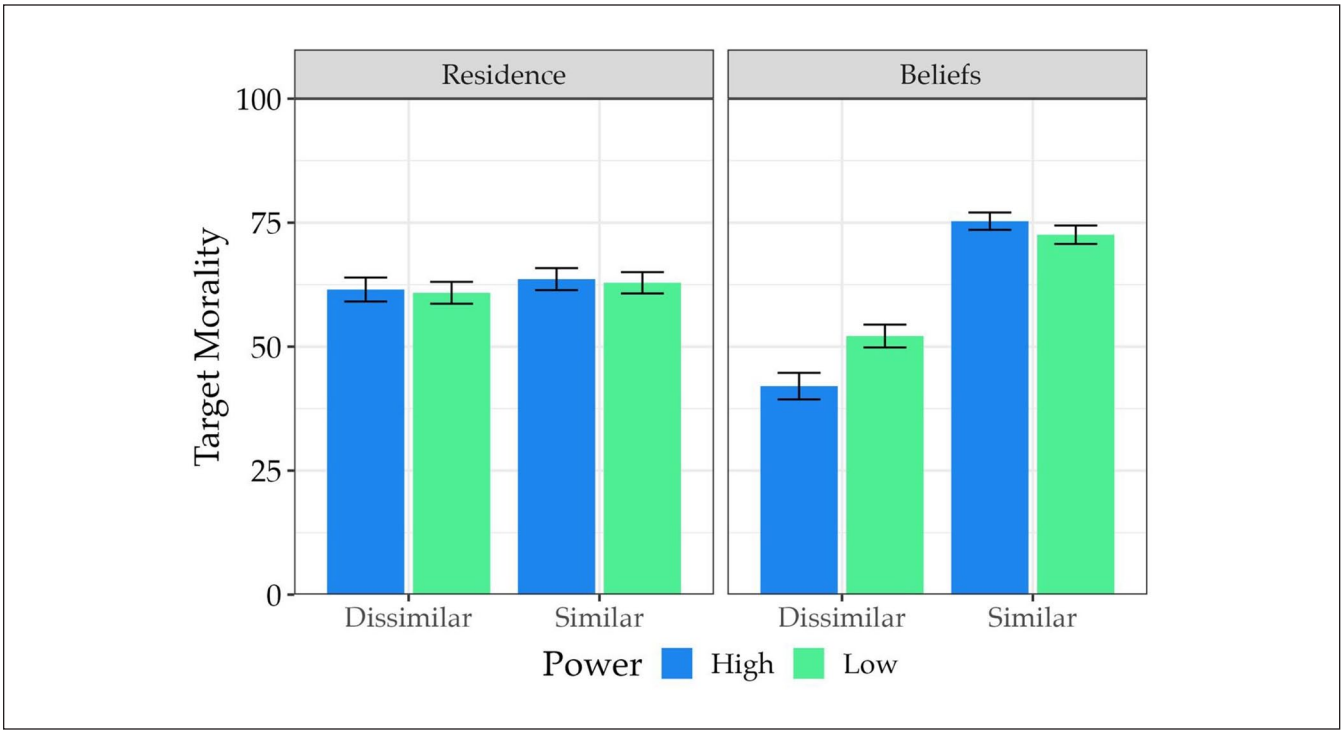


**Figure 3.** Ratings of target morality by target power and perceiver-target similarity (beliefs; state of residence) (Study 2b).
Error bars indicate 95% confidence intervals.

of whether the perceiver and target were similar or dissimilar, and regardless of whether their (dis)similarity was based on beliefs or U.S. state residence. Table S1 and Figures S1 and S2 in the Supplementary Materials visualize the two alternative models.

### *Discussion*

Studies 2a and 2b measured the hypothesized predictors (belief similarity between the perceiver and the target, and target power) directly and compared their interactive effect

**Table 3.** Distribution choice sets for high-power (left) and low-power (right) actors (Study 3).

| | | | | | |
|---|---|---|---|---|---|
| Option A | $0.70 for them | $0.80 for you | | $0.70 for them | $0.80 for you |
| Option B | $0.20 for them | $0.30 for you | Option B | $0.20 for them | $0.30 for you |
| Option C | $0.30 for them | $0.20 for you | Option C | $0.30 for them | $0.20 for you |
| Option D | $0.90 for them | $0.40 for you | | $0.90 for them | $0.40 for you |

to alternative interactive effects. Study 2a examined impressions of societal groups and found that when controlling for target power interacting with perceiver-target belief similarity, there was no interactive effect of target power and shared group membership on morality judgments. Study 2b examined impressions of adult individuals and found that when controlling for target power interacting with perceiver-target belief similarity, there was no interactive effect of target power and shared state residency on morality judgments.

In addition, in both studies, the interaction between target power and belief similarity had a smaller effect size when we replaced target morality with target competence or sociability as the dependent variable in our analysis. In fact, target competence and sociability were much better predicted (than target morality) through the main effect of target power. These results are consistent with functionalist models of social evaluation (e.g., the Moral Primacy Model, Brambilla et al., 2021; the Stereotype Content Model, Fiske et al., 2002), which conceive of competence, sociability, status, and power as personal and social proficiencies that moderate the impact of the intent of the target toward the perceiver (i.e., the self).

## Study 3

Study 3 manipulated both target power and belief similarity by having observers evaluate same-beliefs or different-beliefs actors who would choose between generous and selfish distributions of additional payments. The actor's choices were either limited to two relatively low-value options (low power) or these two options plus two high-value options (high power). We predicted that observers would expect to be treated favorably by similar–beliefs actors and unfavorably by different-beliefs actors, that they would expect more powerful actors to use that power to instantiate more extreme outcomes, and that morality judgments of high-power (vs. low-power) targets would therefore be more polarized.

### Method

*Participants.* Study 3 recruited 423 U.S. residents on Prolific Academic and excluded 23 people who identified as political moderates (and not as conservative or liberal), leaving a final sample of 400 participants ($M_{age}$ = 35.3; 220 female, 178 male, 2 other/prefer not to say; 210 liberal, 190 conservative). Similarity and dissimilarity were manipulated by recruiting participants who had been pre-screened as holding liberal and conservative ideological beliefs (which we verified with our own measure). Although this differs from the multifaceted continuous measure used in previous studies, it allowed for a simplified experimental design and ensured that roughly equal numbers of participants from each side of the ideological spectrum would be recruited. People received US$0.50 to complete the study.

*Procedure.* The observer began by providing demographic information, including their ideological beliefs (liberal or conservative; we excluded moderates at this point). Next, they completed all conditions of the 2 (high vs. low power) × 2 (similar vs. different beliefs) design in random order. In each condition, they learned that an actor, a real and unique participant, would decide which distribution of additional payment would be realized.

In the low-power conditions, the observer looked at the right panel of Table 3, which addressed them with "you" and the actor with "them." The observer learned that low-power actors would choose between a slightly generous distribution of additional payment (Option B) and a slightly selfish distribution (Option C). In the high-power conditions, the observer looked at the left panel of Table 3. They learned that high–power actors would choose between Option B, Option C, and two more options that bestowed them with more power because of greatly expanding the total amount of additional payment to be distributed. Option A was a slightly generous distribution. Option D was a selfish distribution. The observer also learned the actor's self-identified beliefs (liberal or conservative) and that before choosing the actor would learn theirs. Next, the observer predicted which distribution the actor would choose (i.e., the observer never learned which distribution the actor actually chose). Then, the observer used a 7-point scale to rate the morality of the predicted choice and then the morality of the actor (1 = "Very morally bad," 7 = "Very morally good"). About a week later, we bonused the observer according to the ice that most actors in a randomly selected condition had made.

There are two reasons why Options A and D were not numerically complementary as in the low-power conditions. The first is that dictator game players rarely give more resources to their co-players, compared with how much they keep for themselves (Engel, 2011). Thus, distributions more generous than Option A would have come across as unrealistic. Second, dictator game players typically keep more resources for themselves, compared to how much they give to their co-players (Engel, 2011). Thus, distributions less

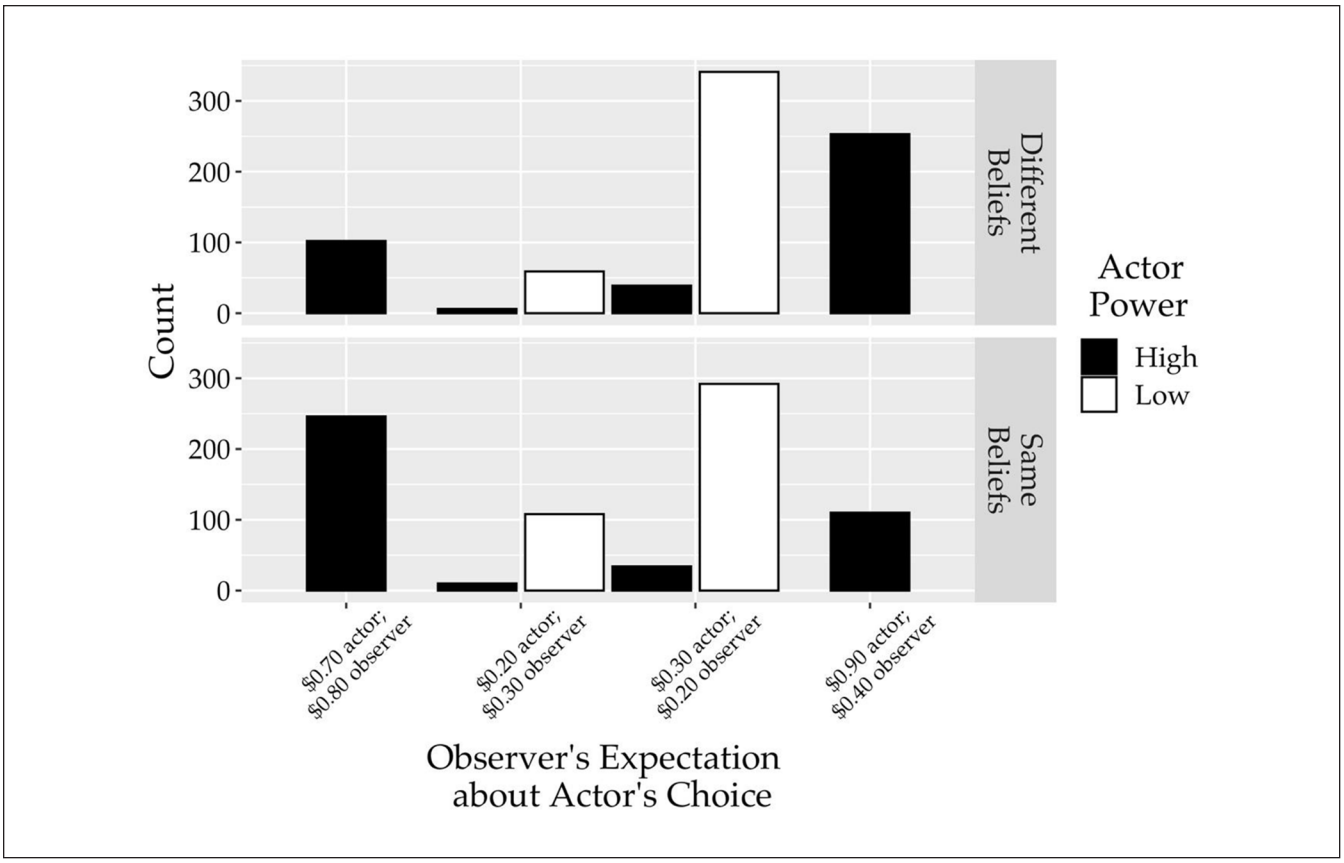


**Figure 4.** Observer Expectations of Actor's distribution choice by power and belief similarity (Study 3).

selfish than Option D would have come across as normal instead of particularly selfish.

## Results

*Predicted Choice.* In the high-power conditions, most observers expected actors to wield their power. Specifically, they predicted similar-beliefs actors to behave generously and choose Option A, whereas they predicted different-beliefs actors to behave selfishly and choose Option D. In the low-power conditions, most observers predicted both similar-beliefs actors and different-beliefs actors to behave selfishly and choose Option C, see Figure 4.

*Impressions of Morality.* In a linear mixed model with random intercepts for observers, we predicted impressions of morality from the perceiver-target similarity in beliefs (we coded different and similar beliefs as −0.5 and 0.5, respectively), target power (we coded low and high power as −0.5 and 0.5, respectively), and their interaction, which emerged as significant. As predicted and shown in Table 4, observers rated same-belief targets and their choices as more moral than different-belief targets, and this effect was moderated by target power. Probing this interaction, we found that different-belief actors rated their anticipated choice and the actor as more moral in the low-power condition than in the high-power condition (Figure 5 and Figure S3 in the Supplementary Materials). Observers evaluating same-belief actors rated their anticipated choice and the actor as more moral in the high-power condition than in the low–power condition.

We also performed mediation analyses to test whether the interactive effect of belief similarity and power on morality judgments is mediated by expectations about how the target will behave toward the perceiver. First, a mediator mixed model was fit predicting the target's expected choice (coded 1 = *most generous*, 4 = *most selfish*) from target power, perceiver-target belief similarity, their interaction, and perceiver-level random intercepts. Second, an outcome mixed model was fit predicting the perceived morality of the target's choice from target power, perceiver-target belief similarity, their interaction, expected target choice, and perceiver-level random intercepts. Mediation analysis was performed based on these two models (using the mediation R package; Tingley et al., 2014). Results indicated that of the total effect of the power/belief similarity interaction on choice morality ($b = 0.50$, 95% CI = [0.35, 0.64], $p < .001$), 86.7% was mediated by expected target choice ($b = 0.43$, 95% CI = [0.35, 0.52] $p < .001$). Thus, the interactive effect

**Table 4.** Results of Study 3.

| Dependent variables<br>fixed effects<br>simple effects | *b* and 95% CI [LB, UB] | *t* | *p* |
|---|---|---|---|
| Choice Morality | | | |
| Perceiver-target similarity in beliefs | 0.59 [0.52, 0.67] | 15.91 | < .001 |
| Target power | −0.03 [−0.10, 0.05] | -0.69 | .490 |
| Belief similarity * Power | 0.50 [0.35, 0.65] | 6.68 | < .001 |
| Power when Belief similarity = same | 0.22 [0.11, 0.32] | 4.24 | < .001 |
| Power when Belief similarity = different | −0.28 [−0.38, −0.18] | -5.21 | < .001 |
| Person Morality | | | |
| Perceiver-target similarity in beliefs | 0.60 [0.53, 0.67] | 15.92 | < .001 |
| Target power | 0.01 [-0.09, 0.06] | -0.31 | .757 |
| Belief similarity * Power | 0.42 [0.28, 0.57] | 5.64 | < .001 |
| Power when Belief similarity = same | 0.20 [0.10, 0.30] | 3.77 | < .001 |
| Power when Belief similarity = different | −0.22 [−0.32, −0.12] | -4.21 | < .001 |

*Note.* *b* = estimate. 95% CI [UB, LB] = 95% confidence interval [lower bound, upper bound].

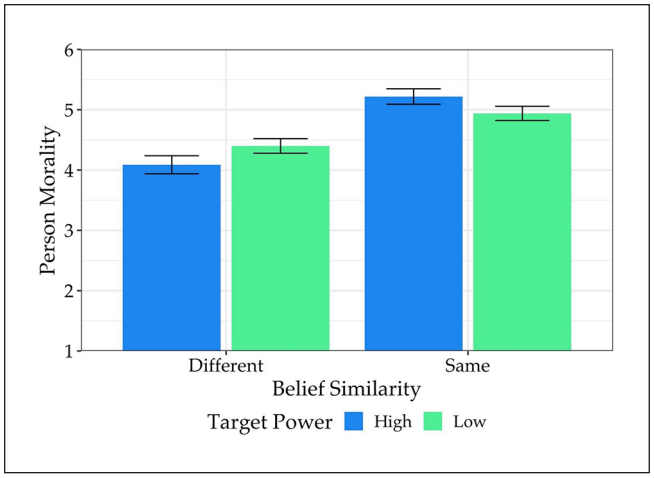


**Figure 5.** Ratings of person morality by target power and perceiver-target belief similarity (Study 3). Error bars indicate 95% confidence intervals.

of power/belief similarity on ratings of choice morality was mediated by the choice that the perceiver expected the target to make. The Supplementary Materials reports similar results for impressions of the target's morality (vs. the morality of their choice).

### *Discussion*

In Study 3, the observer predicted how the actor would distribute additional payment between them. Then, the observer evaluated both the morality of the anticipated distribution and the morality of the actor. The actor could choose both a more selfish distribution and a more generous distribution in the high-power (vs. low-power) condition. The actor's ideological beliefs were either different or similar to those of the observer. Results showed that on average, the observer predicted the high-power (vs. low-power), different-beliefs actor to wield their power and choose a more selfish distribution. Accordingly, the observer evaluated the choice and personal morality of the high-power (vs. low-power), different-beliefs actor as lower. On average, the observer predicted the high-power (vs. low-power), similar-beliefs actor to also wield their power and choose a more generous distribution. The observer evaluated the choice and person morality of the high-power (vs. low-power), similar-belief actor as higher.

Study 3 manipulated the actor's power and actor-observer similarity in beliefs, which interacted to cause changes in impressions of choice and person morality, addressing the inferential limitations of the correlational analyses reported in Studies 1-2b. Moreover, Study 3's power manipulation was minimal and fleeting; randomly designated high-power (vs. low-power) actors would distribute $1 more and make this choice only once. In everyday life, power differentials are typically neither arbitrary nor small and short-lived. Thus, power amplifying the effect of belief similarity on perceived morality likely has a considerably larger effect size in everyday life.

## Study 4

Study 4 manipulated the actor's power in a different and arguably more realistic way. Social power depends not just on the magnitude of a gain or loss that an actor can cause for an observer but also on the magnitude of the resources that the actor needs to pay to cause said gain or loss. The most powerful actors are the ones that can cause maximal gains or losses for an observer at minimal cost (e.g., the Roman emperor's thumb gesture determining whether to kill or spare the gladiator in the arena). Accordingly, Study 4 manipulated power in terms of both the magnitude of help or harm and the ease of helping or harming. In the low-power conditions, the observer encountered an actor who had three choices. First, the actor could do nothing. Second, they could

solve many puzzles for a small benefit to themselves at a small expense to the observer. And third, they could solve many puzzles for a small benefit to the observer (and none for themselves). In the high-power conditions, the actor had the same three choices, except that they would have to solve a few instead of many puzzles for a large instead of a small benefit. The beliefs of the actor and observer were similar or different as in Study 3, and the observer evaluated the morality of the anticipated choice of the actor and the morality of the actor per se as in Study 3. We predicted the same interaction effect as in the previous studies, namely that perceived power should amplify the effect of belief similarity on impressions of morality.

## Method

*Participants.* Study 4 recruited 854 U.S. residents on Prolific Academic and excluded 55 people who identified as political moderates (and not as conservative or liberal), leaving a final sample of 799 participants ($M_{age}$ = 35.28; 332 female, 449 male, 15 other/prefer not to say; 407 liberal, 392 conservative).[1] People received $0.50 to complete the study.

*Procedure.* The observer began by providing demographic information, including their ideology (liberal or conservative; we excluded moderates at this point). Next, they completed all conditions of the 2 (high vs. low power) × 2 (similar vs. different beliefs) design in random order. In each condition, they learned that an actor, a real and unique participant, would choose between three options. They could do nothing. In this case, the payment for participating would change for neither the actor nor the observer (neutral). Alternatively, the actor could solve puzzles to their own benefit at the expense of the observer (selfish), or to the observer's benefit at no cost to themselves (altruistic). The puzzles were 5-letter anagrams with one correct solution using all five letters. For example, the letters "WRNCO" can be rearranged to form the word "CROWN." Observers viewed example puzzles and learned that each puzzle takes 30 s to solve, on average. High-power actors could solve 4 puzzles for a $0.40 change in payment; low–power actors could solve 12 puzzles for a $0.05 change in payment. The observer also learned the actor's self-identified beliefs (liberal or conservative) and that before choosing the actor would learn theirs. Next, the observer predicted what the actor would choose. Then, the observer used a 7-point scale to rate the morality of the predicted choice and then the morality of the actor (1 = "Very morally bad," 7 = "Very morally good"). About a week later, we bonused the observer according to the choice that most actors in a randomly selected condition had made.

## Results

*Predicted Choice.* In the low-power conditions, the modal expectation for both similar-belief actors and different-belief actors was that they would do nothing, leaving payments for both parties unchanged. In the high-power conditions, most observers expected actors with different beliefs to selfishly solve puzzles for their own benefit at the expense of the observer, and most observers expected that similar-belief actors would altruistically solve puzzles for the benefit of the observer (Figure 6).

*Impressions of Morality.* In a linear mixed model with random intercepts for observers, we predicted impressions of morality from perceiver-target similarity in beliefs (we coded different and similar beliefs as −0.5 and 0.5, respectively), target power (we coded low and high power as -0.5 and 0.5, respectively), and their interaction. As predicted and shown in Table 5, observers rated same-belief targets and their choices as more moral than different-belief targets, and this effect was moderated by target power. As in Study 3, we examined the simple effects and found that observers evaluating different-belief actors rated their anticipated choice and the actor as more moral in the low-power condition than in the high-power condition (Figure 7 and Figure S6 in the Supplementary Materials). Observers evaluating same-belief actors rated their anticipated choice and the actor as more moral in the high-power (vs. low-power) condition.

We again conducted mediation analyses to test whether the interactive effect of belief similarity and power on morality judgments is mediated by expectations about how the target will behave toward the perceiver. These models were identical to those used to analyze the data in Study 3, except that in the mediator mixed model, expected actor choice was coded differently (*selfish* = −1, *neutral* = 0, and *generous* = 1). Results indicated that of the total effect of the power/belief-similarity interaction on choice morality ($b$ = 0.34, 95% CI [0.24, 0.45], $p < .001$), 58.0% was mediated by expected choice ($b$ = 0.20, 95% CI [0.13, 0.27], $p < .001$). Again, the interactive effect of power/belief-similarity on ratings of choice morality was mediated by the choice that the perceiver expected to target to make. The Supplementary Materials reports similar results for judgments of person (rather than choice) morality.

## Discussion

Most observers expected neither different-beliefs actors nor similar-beliefs actors to exert effort when they would have to solve twelve puzzles for little reward. However, when actors would only need to solve four puzzles for a much larger payoff, observers expected actors with different beliefs to work for their own benefit at the observer's expense but expected actors with similar beliefs to work for the observer's benefit. Consistent with these expectations, observers rated high-power (vs. low-power) actors with different beliefs as less moral, yet showed the opposite pattern when evaluating actors with similar ideological beliefs.

In Study 4, the payoff of an extra dollar made the difference between getting paid fairly versus generously,

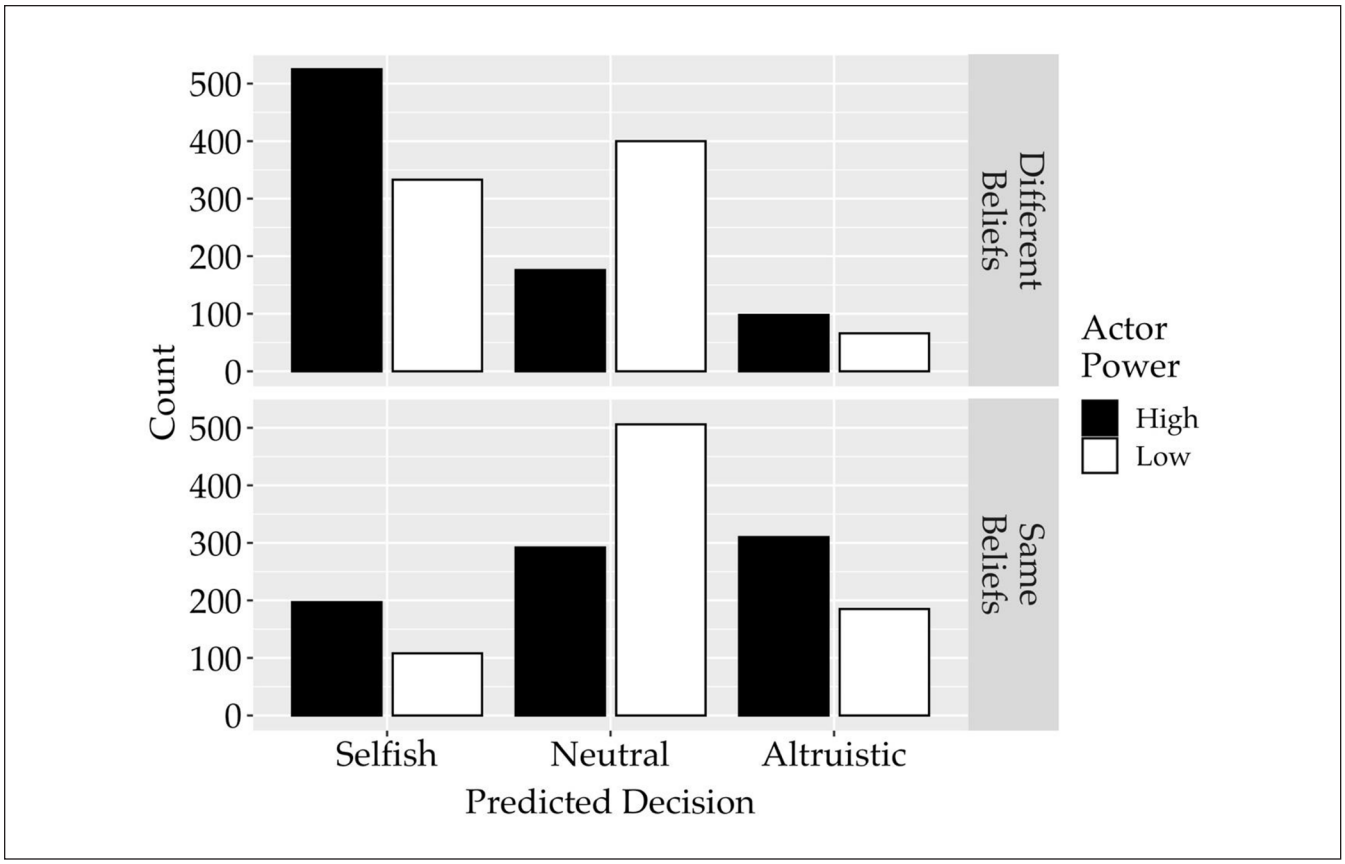


**Figure 6.** Predicted puzzle-solving choice by target power and perceiver-target belief similarity (Study 4).

**Table 5.** Results of Study 4.

| Dependent variables<br>fixed effects<br>simple effects | *b* and 95% CI [LB, UB] | *t* | *p* |
|---|---|---|---|
| Choice Morality | | | |
| Perceiver-target similarity in beliefs | 0.78 [0.73, 0.83] | 28.62 | < .001 |
| Target power | −0.07 [−0.13, −0.02] | −2.69 | .007 |
| Belief similarity * Power | 0.35 [0.24, 0.45] | 6.35 | < .001 |
| Power when Belief similarity = same | 0.10 [0.02, 0.18] | 2.59 | .010 |
| Power when Belief similarity = different | −0.25 [−0.33, −0.17] | −6.40 | < .001 |
| Person Morality | | | |
| Perceiver-target similarity in beliefs | 0.75 [0.70, 0.81] | 27.53 | < .001 |
| Target power | −0.07 [−0.12, −0.02] | −2.52 | .011 |
| Belief similarity * Power | 0.34 [0.23, 0.44] | 6.15 | < .001 |
| Power when Belief similarity = same | 0.10 [0.04, 0.26] | 2.56 | .011 |
| Power when Belief similarity = different | −0.24 [−0.47, −0.24] | −6.14 | < .001 |

*Note.* *b* = estimate. 95% CI [UB, LB] = 95% confidence interval [lower bound, upper bound].

according to the policies of the online worker platform Prolific Academic. That is, our manipulation of power was modest in absolute terms (one dollar) but nevertheless meaningful (relative to the normal reward for taking a brief survey). To be sure, a brief real-world encounter can involve a much greater difference in power, which translates to a stronger manipulation of power that may well shift perceived morality to a much greater extent than what we observed in Study 4. Future archival or field studies should examine this possibility.

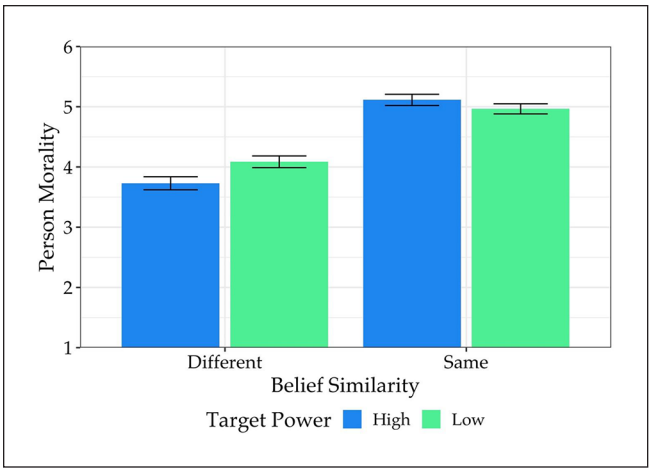


**Figure 7** Ratings of person morality by target power and perceiver-taregt belief similarity (Study 4). Error bars indicate 95% confidence intervals

## General Discussion

Our work revealed an interactive effect of target power and ideological belief similarity such that the gap in moral evaluations of targets with similar (vs. different) beliefs was greater for powerful targets than for less powerful targets. In general, when people evaluated targets whose beliefs appeared similar to their own, they rated high-power targets as *more* moral than low-power targets. In contrast, when perceivers evaluated targets with different beliefs, they rated high-power targets as *less* moral than low-power targets.

Study 1 found that U.S. residents rated same-party presidents more positively than same–party members of Congress but rated opposed-party presidents less positively than opposed party members of Congress. Studies 2a and 2b narrowed the focus to the hypothesized constructs of belief-similarity, power, and morality, and found a larger effect of perceiver-target belief similarity on moral evaluations when people rated groups and individuals that they perceived as more powerful. In addition, Studies 2a and 2b found that this interaction significantly influenced morality judgments over and above the effects of other forms of interpersonal similarity interacting with perceived power (shared group membership, living in the same geographic region), and found that two other basic dimensions of social evaluation, competence and sociability, were not similarly affected.

Study 3 experimentally manipulated both belief similarity and power, and found that observers, after anticipating how they would be treated in a real-money dictator game, rated high-power actors with similar beliefs as more moral than low-power, similar-belief actors, but reversed these judgments for high- and low-power actors with diverging beliefs. Study 4 generalized this pattern of results to an operationalization of power that considers the effort the actor needs to exert to influence the observer's outcomes. Observers expected actors with different beliefs to take money from them when this required little effort, but when the amount of money was reduced and required a lot of effort from the actor, observers expected different-belief actors to forego exerting that effort. In contrast, observers expected same-belief actors to readily exert little effort to benefit them somewhat but expected them to forego working hard to benefit them just a little. Importantly, observers' evaluations of actors' morality tracked whether they expected them to hurt them (different-belief, high-power actors), help them (same-belief, high-power actors), or not influence them (low-power actors whose beliefs were either similar or different to the beliefs of the observers).

Taken together, these studies empirically substantiated the hypothesized interaction effect between perceived belief similarity and perceived power on moral evaluations and suggest that it is rooted in expectations about how positively or negatively others will influence the outcomes of the self. People expect those with little power to have little impact on the self, and thus impressions of their morality are less extreme than those of powerful others whom they expect to help or harm the self to a greater and more meaningful extent, depending on whether their beliefs are aligned with or opposed to the those of the self.

It is noteworthy that in all studies, the size of the hypothesized interaction effect between belief similarity and target power was always considerably larger than the small negative main effect of target power shown in previous work (e.g., Wingen & Dohle, 2021). We too find that, all else equal, target power slightly reduces impressions of target morality, but additionally show that when belief similarity varies, target power enhances the perceived morality of similar-belief targets and reduces the perceived morality of different-belief targets.

### *Theoretical Implications*

The most direct theoretical implication of our work concerns functionalist models of social evaluation (Fiske et al., 2002; Wojciszke et al., 1998), including the Moral Primacy Model (Brambilla et al., 2021). These models posit that proficiency dimensions such as perceived competence, sociability, and power moderate the effect of perceived morality on liking (i.e., overall impression). We show that the proficiency dimension of perceived power also moderates the effect of a pivotal antecedent of perceived morality (belief similarity) on perceived morality.

Our findings further develop the Agency-Beliefs-Communion (ABC) model of social evaluation (Koch et al., 2016), which claims that impressions of communion (a composite of morality and friendliness) are an additive function of impressions of belief similarity and agency similarity (Koch, Imhoff, et al., 2020; agency is a composite of power, status, wealth, and assertiveness). The size of the effect of

agency similarity on communion is small. Thus, the present research suggests that agency, which includes power, is better modeled as a moderator of the considerably larger effect of belief similarity on communion, which includes morality. Predicting communion from belief similarity interacting with agency (one effect) instead of main effects of belief similarity and agency similarity (two effects) simplifies the ABC model.

Our findings are also consistent with the idea of coalitional cognition (Cikara, 2021), which emphasizes the mental representation of groups not simply as social categories, but as entities evaluated on the basis of context-dependent opportunity for cooperation or competition with the self. If coalition building is strategic, moral evaluations should be an interactive function of both belief similarity and power, with powerful allies being preferred to weak ones and nonthreatening (vs. threatening) out-groups being evaluated as more moral. Similarly, our findings resonate with recent theorizing that proposes that political belief systems are rooted in the alliances that people and groups form based on perceived similarity, transitivity, and interdependence (Pinsof et al., 2023). On this view, the content of people's political beliefs functions as a kind of ideological superstructure used to reinforce coalitional and competitive dynamics between groups. Within this framework, belief similarity constitutes an instant appraisal of whether the target is currently allied or opposed to the self, and power captures the degree and dynamic of interdependence between parties.

### *Practical Implications*

One potential implication of the present research is that as people attain higher power, in addition to being viewed more positively by their ideological allies, they may also expect to be viewed more negatively by ideological adversaries. For example, politicians who rise in the ranks of their own party should expect to become more polarizing to voters as their power and influence grow. They may become even more beloved by their supporters, but opponents will likely find them even more repellant because they represent an even greater potential threat. This is not surprising in the political realm, where rivals are vilified and where party systems make perceiver-target belief similarity highly salient. But it could also play out in other, less explicitly polarized contexts. For example, recently promoted managers might be viewed differently by various company factions to the extent that ideological competition rather than cooperation characterizes organizational culture. This could be particularly important in the context of corporate mergers, where distinctions between "us" and "them" are likely to be especially stark as powerful people from outside the organization move into management roles.

One additional insight from the present research is that people expect different-belief others to use their full power to act selfishly more often than they actually do (see Supplemental Studies 1 and 2). Although these expectations were, at an individual level, rational in the sense that selfishness was more common than generosity from different-belief actors, it also meant that chances for cooperation were missed. More generally, such overly cynical expectations of others with different beliefs may lead to preemptive competitive behavior that all but ensures hostility and conflict between parties rather than cooperation.

### *Limitations and Future Directions*

We theorized that perceived power amplifies the effect of perceived belief similarity on moral evaluations and that this interaction effect is rooted in people's expectations about how favorably others will behave toward them. To this end, our experiments asked observers to form expectations about actors' behavior toward the self. We cannot be sure that people spontaneously generate those expectations in the absence of prompts. Future research should clarify the role of prompts in the emergence of the hypothesized interaction effect. We suspect that it will emerge most reliably given strong and salient manipulations of both perceiver-target belief similarity and power.

A second limitation of our experiments is that we cannot accurately determine whether moral condemnation of powerful adversaries or moral approbation of powerful allies contribute equally to the interaction effect we observed. This is because actors' choice sets did not provide selfish and generous options of equal magnitude. While this was necessary given observers' baseline expectations about how others would behave in the economic games used in these experiments, it unfortunately does not allow us to compare the simple effect of power on morality given diverging beliefs with the same simple effect given aligned beliefs. Future research may clarify the conditions that produce the widest discrepancies in moral evaluation between high- and low-power co-believers, and whether and how they differ from those producing the widest discrepancies for targets that hold different beliefs.

Relatedly, our studies typically (but not exclusively) observed simple effects such that powerful targets with similar beliefs were rated as more moral than less powerful targets with similar beliefs, and that powerful targets with different beliefs were rated as less moral than less powerful targets with different beliefs. We argue that this pattern is a function of people's expectations about how such targets will treat them and obtain evidence consistent with this explanation. However, we expect that moral evaluations are not entirely driven by such consequentialist reasoning and that under some conditions either or both of these trends might be obfuscated or over-ridden by other, context-dependent factors. For instance, there could be circumstances under which all same-beliefs targets are rated as equally moral, or where the effect is reversed, perhaps in a situation where actual cooperation is more likely to come from a same–beliefs

target with a similar amount of power rather than one with a lot more power.

Looking ahead, while this analysis focuses specifically on the power of the target moderating the effect of belief similarity on moral evaluations, we expect that the power of the perceiver must also play some role in this process. Consistent with the finding (Russell & Fiske, 2010) that people with high status tend to ignore other people, particularly those with lower status, we predict that perceiver power should not amplify but attenuate the effect of perceiver–target belief similarity on moral evaluation of the target. This follows from the idea that powerful perceivers are likely less threatened by opposite-ideology targets and less likely to substantially benefit from same-beliefs targets. Future research could test this. Relatedly, future research could test whether impressions of morality are more polarized in societies with greater power imbalances (e.g., societies with greater income inequality as indicated by a larger GINI coefficient; Dorfman, 1979).

Finally, future research could revisit the question of the order of importance of different social-evaluative dimensions. Prior work suggests that perceivers prioritize targets' communion over their agency (Abele & Bruckmüller, 2011; Abele & Wojciszke, 2014). Yet from the standpoint of processing efficiency, learning a target's morality is wasted effort if they lack the power to enact it. For the same reason, power judgments may take priority over judgments of belief similarity, which predict judgments of target communion.

In conclusion, morality is not only a function of what someone *would* do if given the power to do it but also hinges on what people think they *will* do given the power they have.

## Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

## Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

## ORCID iDs

Russell Roberts https://orcid.org/0000-0001-8715-8932
Alex Koch https://orcid.org/0000-0002-6267-8066

## Supplemental Material

Supplemental material is available online with this article.

## Note

1. We recruited more participants than we had pre-registered in Study 4 because we aimed for at least 300 conservatives and 300 liberals, but there are more liberals than conservatives on online worker platforms, including Prolific Academic. Thus, recruiting at least 300 conservatives required recruiting a total number of participants that was higher than 600.

## References

Abele, A. E., & Bruckmüller, S. (2011). The bigger one of the "big two"? Preferential processing of communal information. *Journal of Experimental Social Psychology*, *47*(5), 935–948. https://doi.org/10.1016/j.jesp.2011.03.028

Abele, A. E., & Wojciszke, B. (2014). Communal and agentic content in social cognition: A dual perspective model. In J. M. Olson & M. P. Zanna (Eds.), *Advances in experimental social psychology* (Vol. 50, pp. 195–255). Academic Press. https://doi.org/10.1016/B978-0-12-800284-1.00004-7

Aiken, L. S., & West, S. G. (1991). *Testing and interpreting interactions in multiple regression*. Sage.

American National Election Studies. (2015). *ANES time series cumulative data file (1948–2012)*. Inter-University Consortium for Political and Social Research distributor. https://doi.org/10.3886/ICPSR08475.v15

Bastian, B., Denson, T. F., & Haslam, N. (2013). The roles of dehumanization and moral outrage in retributive justice. *PLOS ONE*, *8*(4), Article e61842. https://doi.org/10.1371/journal.pone.0061842

Bocian, K., Baryla, W., Kulesza, W. M., Schnall, S., & Wojciszke, B. (2018). The mere liking effect: Attitudinal influences on attributions of moral character. *Journal of Experimental Social Psychology*, *79*, 9–20. https://doi.org/10.1016/j.jesp.2018.06.007

Bocian, K., & Wojciszke, B. (2014). Self-interest bias in moral judgments of others' actions. *Personality and Social Psychology Bulletin*, *40*(7), 898–909. https://doi.org/10.1177/0146167214529800

Brambilla, M., & Leach, C. W. (2014). On the importance of being moral: The distinctive role of morality in social judgment. *Social Cognition*, *32*(4), 397–408. https://doi.org/10.1521/soco.2014.32.4.397

Brambilla, M., Rusconi, P., Sacchi, S., & Cherubini, P. (2011). Looking for honesty: The primary role of morality (vs. sociability and competence) in information gathering. *European Journal of Social Psychology*, *41*(2), 135–143. https://doi.org/10.1002/ejsp.744

Brambilla, M., Sacchi, S., Rusconi, P., & Goodwin, G. P. (2021). The primacy of morality in impression development: Theory, research, and future directions. In B. Gawronski (Ed.), *Advances in experimental social psychology* (Vol. 64, pp. 187–262). Academic Press. https://doi.org/10.1016/bs.aesp.2021.03.001

Brandt, M. J., & Crawford, J. T. (2020). Worldview conflict and prejudice. In B. Gawronski (Ed.), *Advances in experimental social psychology* (Vol. 61, pp. 1–66). Academic Press. https://doi.org/10.1016/bs.aesp.2019.09.002

Carrier, A., Dompnier, B., & Yzerbyt, V. (2019). Of nice and mean: The personal relevance of others' competence drives perceptions of warmth. *Personality and Social Psychology Bulletin*, *45*(11), 1549–1562. https://doi.org/10.1177%2F0146167219835213

Cikara, M. (2021). Causes and consequences of coalitional cognition. In B. Gawronski (Ed.), *Advances in experimental social psychology* (Vol. 64, pp. 65–128). Academic Press. https://doi.org/10.1016/bs.aesp.2021.04.002

Crockett, M. J. (2017). Moral outrage in the digital age. *Nature Human Behaviour*, *1*(11), 769–771. https://doi.org/10.1038/s41562-017-0213-3.3

Cushman, F. (2008). Crime and punishment: Distinguishing the roles of causal and intentional analyses in moral judgment. *Cognition*, *108*(2), 353–380. https://doi.org/10.1016/j.cognition.2008.03.006

Cushman, F. (2015). Deconstructing intent to reconstruct morality. *Current Opinion in Psychology*, *6*, 97–103. https://doi.org/10.1016/j.copsyc.2015.06.003

Dépret, E., & Fiske, S. T. (1993). Social cognition and power: Some cognitive consequences of social structure as a source of control deprivation. In G. Weary, F. Gleicher, & K. L. Marsh (Eds.), *Control motivation and social cognition* (pp. 176–202). Springer.

Dorfman, R. (1979). A formula for the Gini coefficient. *The Review of Economics and Statistics*, *61*(1), 146–149. https://doi.org/10.2307/1924845

Ellemers, N., Pagliaro, S., & Barreto, M. (2013). Morality and behavioural regulation in groups: A social identity approach. *European Review of Social Psychology*, *24*(1), 160–193. https://doi.org/10.1080/10463283.2013.841490

Engel, C. (2011). Dictator games: A meta study. *Experimental Economics*, *14*(4), 583–610. https://doi.org/10.1007/s10683-011-9283-7

Fiske, S. T., Cuddy, A. J., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology*, *82*(6), 878–902. https://doi.org/10.1037/0022-3514.82.6.878

Goodwin, G. P., Piazza, J., & Rozin, P. (2014). Moral character predominates in person perception and evaluation. *Journal of Personality and Social Psychology*, *106*(1), 148–168. https://psycnet.apa.org/doi/10.1037/a0034726

Green, P., & MacLeod, C. J. (2016). SIMR: An R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution*, *7*(4), 493–498. https://doi.org/10.1111/2041-210X.12504

Hauke, N., & Abele, A. E. (2020). The impact of negative gossip on target and receiver. A "big two" analysis. *Basic and Applied Social Psychology*, *42*(2), 115–132. https://doi.org/10.1080/01973533.2019.1702881

Imhoff, R., Koch, A., & Flade, F. (2018). (Pre)occupations: A data-driven model of jobs and its consequences for categorization and evaluation. *Journal of Experimental Social Psychology*, *77*, 76–88. https://doi.org/10.1016/j.jesp.2018.04.001

Jenkins, A. C., Karashchuk, P., Zhu, L., & Hsu, M. (2018). Predicting human behavior toward members of different social groups. *Proceedings of the National Academy of Sciences*, *115*(39), 9696–9701. https://doi.org/10.1073/pnas.1719452115

Koch, A., Dorrough, A., Glöckner, A., & Imhoff, R. (2020). The ABC of society: Perceived similarity in agency/socioeconomic success and conservative-progressive beliefs increases intergroup cooperation. *Journal of Experimental Social Psychology*, *90*, Article 103996. https://doi.org/10.1016/j.jesp.2020.103996

Koch, A., Imhoff, R., Dotsch, R., Unkelbach, C., & Alves, H. (2016). The ABC of stereotypes about groups: Agency/socioeconomic success, conservative–progressive beliefs, and communion. *Journal of Personality and Social Psychology*, *110*(5), 675–709. https://psycnet.apa.org/doi/10.1037/pspa0000046

Koch, A., Imhoff, R., Unkelbach, C., Nicolas, G., Fiske, S., Terache, J., Carrier, A., & Yzerbyt, V. (2020). Groups' warmth is a personal matter: Understanding consensus on stereotype dimensions reconciles adversarial models of social evaluation. *Journal of Experimental Social Psychology*, *89*, Article 103995. https://doi.org/10.1016/j.jesp.2020.103995

Koch, A., Kervyn, N., Kervyn, M., & Imhoff, R. (2018). Studying the cognitive map of the U.S. States: Ideology and prosperity stereotypes predict interstate prejudice. *Social Psychological and Personality Science*, *9*(5), 530–538. https://doi.org/10.1177/1948550617715070

Koch, A., Yzerbyt, V., Abele, A., Ellemers, N., & Fiske, S. T. (2021). Social evaluation: Comparing models across interpersonal, intragroup, intergroup, several-group, and many-group contexts. In B. Gawronski (Ed.), *Advances in experimental social psychology* (Vol. 63, pp. 1–68). Academic Press. https://doi.org/10.1016/bs.aesp.2020.11.001

Krosnick, J. A., Judd, C. M., & Wittenbrink, B. (2005). Attitude measurement. In D. Albarracin, B. T. Johnson, & M. P. Zanna (Eds.), *Handbook of attitudes and attitude change* (pp. 21–76). Erlbaum.

Landy, J. F., Piazza, J., & Goodwin, G. P. (2016). When it's bad to be friendly and smart: The desirability of sociability and competence depends on morality. *Personality and Social Psychology Bulletin*, *42*(9), 1272–1290. https://doi.org/10.1177%2F0146167216655984

Leach, C. W., Ellemers, N., & Barreto, M. (2007). Group virtue: the importance of morality (vs. competence and sociability) in the positive evaluation of in-groups. *Journal of Personality and Social Psychology*, *93*(2), 234–249. https://psycnet.apa.org/doi/10.1037/0022-3514.93.2.234

Magee, J. C., & Smith, P. K. (2013). The social distance theory of power. *Personality and Social Psychology Review*, *17*(2), 158–186. https://doi.org/10.1177/1088868312472732

Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, *100*(2), 254–278. https://doi.org/10.1037/0033-295X.100.2.254

Montoya, R. M., Horton, R. S., & Kirchner, J. (2008). Is actual similarity necessary for attraction? A meta-analysis of actual and perceived similarity. *Journal of Social and Personal Relationships*, *25*(6), 889–922. https://doi.org/10.1177%2F0265407508096700

Nicolas, G., Fiske, S. T., Koch, A., Imhoff, R., Unkelbach, C., Terache, J., Carrier, A., & Yzerbyt, V. (2022). Relational versus structural goals prioritize different social information. *Journal of Personality and Social Psychology*, *122*(4), 659–682. https://doi.org/10.1037/pspi0000366

Pinsof, D., Sears, D. O., & Haselton, M. G. (2023). Strange bedfellows: The alliance theory of political belief systems. *Psychological Inquiry*, *34*(3), 139–160. https://doi.org/10.1080/1047840X.2023.2274433

Russell, A. M., & Fiske, S. T. (2010). Power and social perception. In A. Guinote & T. K. Vescio (Eds.) *The social psychology of power* (pp. 231–250). Guilford Press.

Slepian, M. L., & Koch, A. (2021). Identifying the dimensions of secrets to reduce their harms. *Journal of Personality and Social Psychology*, *120*(6), 1431–1456. https://doi.org/10.1037/pspa0000202

Thibaut, J. W., & Kelley, H. H. (1959). *The social psychology of groups*. Wiley.

Tingley, D., Yamamoto, T., Hirose, K., Keele, L., & Imai, K. (2014). Mediation: R package for causal mediation analysis. http://hdl.handle.net/1721.1/91154

Van Prooijen, A. M., & Ellemers, N. (2015). Does it pay to be moral? How indicators of morality and competence enhance organizational and work team attractiveness. *British Journal of Management*, *26*(2), 225–236. https://doi.org/10.1111/1467-8551.12055

Willis, J., & Todorov, A. (2006). First Impressions: Making up your mind after a 100-Ms exposure to a face. *Psychological Science*, *17*(7), 592–598. https://doi.org/10.1111/j.1467-9280.2006.01750.x

Wingen, T., & Dohle, S. (2021). Exploring negative beliefs about power: The role of autonomy and influence. *Social Psychology*, *52*(4), 250–263. https://doi.org/10.1027/1864-9335/a000453

Woitzel, J., & Koch, A. (2023). Ideological prejudice is stronger in ideological extremists (vs. moderates). *Group Processes & Intergroup Relations*, *26*(8), 1685–1705. https://doi.org/10.1177/13684302221135083

Wojciszke, B., Bazinska, R., & Jaworski, M. (1998). On the dominance of moral categories in impression formation. *Personality and Social Psychology Bulletin*, *24*(12), 1251–1263. https://doi.org/10.1177/01461672982412001

Zorn, T. J., Mata, A., & Alves, H. (2022). Attitude similarity and interpersonal liking: A dominance of positive over negative attitudes. *Journal of Experimental Social Psychology*, *100*, Article 104281. https://doi.org/10.1016/j.jesp.2021.104281