



LBS Research Online

M E Inesi and K Rios

How prosocial actors use power hierarchies to build moral reputation

Article

This version is available in the LBS Research Online repository: <https://lbsresearch.london.edu/id/eprint/2771/>

Inesi, M E and Rios, K

(2023)

How prosocial actors use power hierarchies to build moral reputation.

Journal of Experimental Social Psychology, 106.

ISSN 0022-1031

(In Press)

DOI: <https://doi.org/10.1016/j.jesp.2022.104441>

Elsevier

<https://www.sciencedirect-com.lbs.idm.oclc.org/sci...>

Users may download and/or print one copy of any article(s) in LBS Research Online for purposes of research and/or private study. Further distribution of the material, or use for any commercial gain, is not permitted.



Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jespHow prosocial actors use power hierarchies to build moral reputation[☆]M. Ena Inesi^{a,*}, Kimberly Rios^b^a Organizational Behaviour, London Business School, Park Road, London NW1 4SA, UK^b Ohio University, Department of Psychology, 219 Porter Hall, Athens, OH 45701, USA

ARTICLE INFO

Keywords:

Power
Hierarchy
Beneficiary choice
Moral character
Reputation
Prosocial behavior

ABSTRACT

Power hierarchies are ubiquitous, emerging formally and informally, in both personal and professional contexts. When prosocial acts are offered within power hierarchies, there is a widespread belief that people who choose lower-power beneficiaries are altruistically motivated, and that those who choose higher-power beneficiaries hold a self-interested motive to ingratiate. In contrast, the current research empirically demonstrates that people can also choose lower-power beneficiaries for self-interested reasons – namely, to bolster their own moral reputation in the group. Across three pre-registered studies, involving different contexts and types of prosocial behavior, and including real financial incentives, we demonstrate that people are more likely to choose lower-power beneficiaries when reputation concerns are more salient. We also provide evidence of the mechanism underlying this pattern: people believe that choosing a lower-power beneficiary more effectively signals their own moral character.

Whether at work, in social groups or even at home, people are often part of social hierarchies. Usually, their position in the hierarchy is somewhere between the top and the bottom. A person's position in the power hierarchy can depend on many factors, such as experience, expertise or assignment to a role (French & Raven, 1959; Magee & Galinsky, 2008). Wherever people find themselves in a hierarchy, however, two things are true: They will depend on those above them to provide access to desired resources, and they will control their subordinates' access to desired resources (Emerson, 1962).

Research suggests that this asymmetrical access to resources shapes prosocial behavior, including the motives that give rise to it. By definition, higher-power individuals control more valuable resources, such as deciding who gets promoted or who is invited to an exclusive event (Emerson, 1962; Keltner, Gruenfeld, & Anderson, 2003). This creates an incentive to offer acts of kindness to individuals with more power, albeit for self-interested reasons. For example, by ingratiating the boss, an employee may be more likely to gain access to the valued resources the boss controls (Jones, 1964). Such self-interested motives are less apparent when an employee chooses to offer kindness to a lower-power beneficiary. Considerable research has demonstrated that observers form inferences consistent with the motives outlined above. That is, observers view people who offer kindness to higher-power individuals as more likely to be motivated by self-interest (Inesi, Gruenfeld, &

Galinsky, 2012; Kunstman, Fitzpatrick, & Smith, 2018). In addition, they tend to judge those who select lower-power beneficiaries as particularly altruistic (Inesi et al., 2012; Milinski, Semmann, & Krambeck, 2002).

In contrast to outside observers' beliefs, however, we propose that people may also select lower-power beneficiaries for self-interested reasons. We theorize that, in addition to considering the resources and rewards that could emerge directly from beneficiaries (leading people to choose higher-power beneficiaries), prosocial actors also consider the reputational benefits their actions may garner in the wider social group. People are keen to develop a reputation for high moral character, and performing kind acts – especially those that are costly – can serve this purpose (Hardy & Van Vugt, 2006; Zahavi & Zahavi, 1997). Therefore, we predict that, due to the widely held belief that those who select lower-power beneficiaries have higher moral character, people may *strategically* select lower-power beneficiaries to bolster their moral image. Indeed, we predict that the tendency to use lower-power beneficiaries to signal high moral character will be stronger than the tendency to use higher-power beneficiaries, because the former sends a clearer signal of the prosocial actor's ostensible altruism.

By introducing this possibility, we shed greater light on the dark side of prosocial behavior: that people sometimes perform acts of ostensible altruism for self-interested reasons. At the same time, this “dark”

[☆] This paper has been recommended for acceptance by Dr. Joris Lammers.

* Corresponding author.

E-mail address: einesi@london.edu (M.E. Inesi).

<https://doi.org/10.1016/j.jesp.2022.104441>

Received 8 October 2022; Received in revised form 9 December 2022; Accepted 20 December 2022

Available online 17 January 2023

0022-1031/© 2023 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

dynamic may actually serve to re-distribute resources to lower-power individuals. That is, in contrast to historical accounts that motivate offering kindness to higher-power individuals (Jones, 1964), we propose a complementary dynamic that provides an incentive to offer time, favors, and resources to those at lower levels of the hierarchy.

1. Power hierarchies and ingratiation

Power hierarchies are ubiquitous in everyday life (Gruenfeld & Tiedens, 2010; Magee & Galinsky, 2008). Defined as asymmetrical control over valued resources, power is a relative construct: to the extent that a person has resources (e.g., money, information, ability to promote) on which others depend, the former person is more powerful than the latter (Emerson, 1962). In many contexts, power hierarchies are explicitly described in charts (e.g., diagrams indicating to whom one reports, and who reports to oneself) or embedded in roles. They can also emerge in more subtle forms, however. For example, if one person in a group of friends belongs to an exclusive membership club that others want to visit, then that person will have more power than the others - at least in that domain. Research suggests that people naturally form and even prefer hierarchies. For example, in dyadic interactions, people adjust their postures so that they complement rather than mimic their partner (e.g., adopt a more dominant posture if their partner is submissive, and vice versa; Tiedens & Fragale, 2003). Indeed, hierarchies have been shown to satisfy fundamental psychological needs for control and structure (because they provide clarity as to one's position; Friesen, Kay, Eibach, & Galinsky, 2014; Zitek & Phillips, 2020), as well as for ease and fluency of processing (because they are relatively easy to remember; Zitek & Tiedens, 2012; Zitek & Phillips, 2020).

In his classic 1964 book, Edward Jones surmised that power differences - that is, the dependence on another person for valued resources - can fundamentally alter the motives driving acts of kindness: the more a person depends on someone else to get what they want, the more likely their acts of kindness are motivated by self-interest. Acts of kindness offered to higher-power beneficiaries, he reasoned, become tainted by a self-interested goal of ingratiating the targets, and ultimately gaining access to the valued rewards that they control. By ingratiating their boss, for example, an employee may be more likely to secure a promotion.

Numerous studies support the notion that employees seek to ingratiate their bosses for self-interested reasons. In content-coded essays, employees were more likely to cite personal benefits as a reason for attempting to influence their superiors, relative to their coworkers or subordinates (Kipnis, Schmidt, & Wilkinson, 1980). Male undergraduate participants in another study reported being more likely to use ingratiation tactics to obtain rewards from a higher-power individual than from a stranger or a friend (Bohra & Pandey, 1984). Employees' tendencies to ingratiate their bosses more so than coworkers or subordinates are most pronounced in organizations with clear power structures (e.g., specialized roles, centralized authority; Drory & Zaidman, 2007). Taken together, these findings demonstrate the proclivity of people within hierarchies to perform ostensibly altruistic acts for higher-power beneficiaries in the hopes of gaining personal rewards.

2. Observer responses to prosocial behavior in hierarchies

People's tendencies to act kindly toward higher-power others for self-interested reasons do not go undetected by third parties. When observing prosocial acts, observers are keen to identify why the person acted in a generous way (Carlson & Zaki, 2018; De Freitas, DeScioli, Thomas, & Pinker, 2019), and they use power hierarchies - among other cues (e.g., whether the person bragged about their actions; Berman, Levine, Barasch, & Small, 2015; whether the person acted quickly or slowly; Critcher, Inbar, & Pizarro, 2013) - to guide their understanding

of the person's motives.

When observing kind acts offered within a power hierarchy, two related patterns emerge. First, people believe that those who offer kind acts to higher-power beneficiaries are more selfishly motivated than those who show the same generosity to peers. For example, people are more likely to attribute favors to selfish motives when offered to a higher-power individual versus to a peer (Inesi et al., 2012). They are also more likely to discount the authenticity of compliments that are directed at a higher-power individual versus at a peer (Kunstman et al., 2018). Furthermore, observers use these attributions to form a belief about the prosocial actor's moral character. For example, participants in one study believed individuals who chose a higher-power beneficiary had lower moral character (Inesi, Adams, & Gupta, 2021).

A second pattern that emerges in hierarchical contexts is that people believe those who offer kind acts to lower-power individuals are more altruistically motivated,¹ even more so than those who offer kind acts to peers. For example, when managers accept job candidate referrals from lower-power referrers, observers judge them as less selfishly motivated than when the referrer has equal power to the manager (Derfler-Rozin, Baker, & Gino, 2018). In another study, participants judged an employee who offered consistently kind acts to a subordinate as less "slimy" than someone who was both kind and unkind to individuals across the hierarchy (Vonk, 1998; see also Inesi et al., 2012; Milinski et al., 2002). Taken together, this research demonstrates that not only are people more skeptical about the motives driving acts of kindness offered to higher-power individuals, but they are also particularly generous in attributing altruistic motives to those who offer kindness to lower-power individuals.

3. Prosocial behavior and reputation-building

In contrast to these existing findings, we suggest that people may select lower-power beneficiaries for self-interested reasons linked to reputation. Specifically, people may choose lower-power beneficiaries because they believe doing so will most effectively communicate their own positive image to others. Research has consistently demonstrated that people are motivated to act in ways that improve their moral reputation, defined as judgments others make about one's personal character (Emler, 1990; Wu, Balliet, & Van Lange, 2016). Specifically, individuals seek to develop a reputation for being moral and cooperative because moral reputations are rewarded by others. People who have a reputation for high moral character are granted status, respect and resources - even from people they have never interacted with personally (Ariely, Bracha, & Meier, 2009; Bénabou & Tirole, 2006; Bereczkei, Birkas, & Kerekes, 2007; Berman & Silver, 2022; Cottrell, Neuberg, & Li, 2007; Hardy & Van Vugt, 2006; Nowak & Sigmund, 2005; Simpson & Willer, 2008; Wu et al., 2016; Zahavi & Zahavi, 1997).

Therefore, when opportunities arise to boost one's reputation, such as when one's actions are publicly known to others, people tend to act in more generous and cooperative ways. For example, people tend to express more interest in purchasing environmentally-friendly products when they are buying in public (i.e., at a shop) versus in private (i.e., online; Griskevicius, Tybur, & Van den Bergh, 2010) because such products signal the purchaser's motivation to benefit the common good. This tendency extends to the enactment of prosocial behavior. Participants in one study showed greater generosity to others when third parties would learn about their behavior versus when they would not (Simpson & Willer, 2008). Relatedly, members of distrusted groups (i.e., atheists) are more generous toward outgroup partners (i.e., Christians) in economic games when they believe their group identity has been

¹ For this review, we consider attributions to altruistic and selfish/self-interested motives to be two ends of a single dimension, consistent with existing research (Derfler-Rozin et al., 2018; Inesi et al., 2021; Kunstman et al., 2018).

made public rather than kept anonymous. That is, when in public, atheists are motivated to improve their group's negative reputation and act more generously as a result (Cowgill, Rios, & Simpson, 2017).

The research reviewed thus far demonstrates that reputational concerns affect the decision of *whether or not* to act generously. We know little, however, about how such choices unfold. Incipient research in this area shows that, when reputation concerns are salient, people make the decision to cooperate faster (Jordan, Hoffman, Nowak, & Rand, 2016), presumably because they believe that faster decision-making signals a more altruistic character to observers (Cricher et al., 2013). Here, we investigate whether reputational concerns affect *whom* prosocial actors choose as their beneficiaries. We ask: would they select lower-power beneficiaries over higher-power ones, not so much to benefit these individuals, but rather to maximize their personal reputational gain? We suggest that they would.

As reviewed previously, research has demonstrated that observers offer greater social and material rewards to prosocial actors who choose lower-power beneficiaries versus higher-power ones (Inesi et al., 2021). Choosing a lower-power beneficiary ostensibly sends a clearer signal of the prosocial actor's moral character because it is seen as more costly. That is, there are no obvious personal benefits involved, whereas this is not true for those choosing a higher-power beneficiary. We propose that prosocial actors are aware of this dynamic, leading to a lay theory that their reputation will benefit most from choosing lower-power rather than higher-power beneficiaries. Due to this lay theory, prosocial actors may be especially likely to select lower-power beneficiaries in contexts where reputational concerns are high - such as in public contexts or where gossip is common - because other people will learn about their actions and form judgments about their moral character (Feinberg, Willer, & Schultz, 2014). Therefore, we predict that, as reputational concerns increase, people will be more likely to choose lower-power beneficiaries, and that this pattern is driven by their belief that choosing lower-power beneficiaries will improve their reputation in the group. Furthermore, we predict that the increased tendency to choose lower-power beneficiaries when reputational concerns are salient will be larger than any change in the tendency to choose higher-power beneficiaries. In other words, people are not just more likely to select any beneficiary in a power hierarchy when reputational concerns increase. Rather, they are particularly likely to select lower-power beneficiaries, because they believe doing so will benefit their reputation most.

Notably, we do not believe that choosing a lower-power beneficiary is always driven by self-interest. It has been well documented that people do display altruism (Batson, 1998), and therefore choosing a lower-power beneficiary is likely often driven by a desire to help. Our contention here is that in addition to those who choose lower-power beneficiaries for altruistic reasons, there are others who hold more self-interested motives. Furthermore, this can be illustrated by changes in beneficiary choice: to the extent that some people choose lower-power beneficiaries for self-interested motives linked to reputation-building, then we should see an increase in the propensity to select such beneficiaries in public (versus private) contexts, where reputational concerns are more salient.

4. Contribution

The present research offers several theoretical contributions to existing literature. First, it proposes a new and counterintuitive dynamic within the realm of hierarchy and strategic generosity. In contrast to the dominant narrative that strategic prosocial behavior (i.e., enacted for personal gain) tends to be directed at higher-power beneficiaries (Inesi et al., 2021; Jones, 1964; Vonk, 1998), we suggest that prosocial actors may also strategically choose lower-power beneficiaries to boost their moral reputation. Indeed, a second contribution of this work is to fundamentally reframe our understanding of the benefits offered by power hierarchies. Researchers typically assume that the primary driver

of strategic prosocial behavior in hierarchy is a desire to gain access to the valued resources that are unequally distributed in power hierarchies (e.g., money, network access). Here, we reveal that the ability to build a positive reputation represents a second important currency that drives strategic prosocial behavior in hierarchies. Finally, our work makes a significant contribution to the literature on prosocial behavior by demonstrating that people are not only more likely to strategically act generously when in public, but also more likely to select certain individuals as their beneficiaries. In doing so, we demonstrate the sophisticated ways in which strategic generosity works: that individuals will select a beneficiary not to benefit the person, but rather to use this person's characteristics (i.e., position in the power hierarchy) to further their own reputation.

5. Overview of studies

We test our predictions across three studies. In Study 1, we test which of two past experiences working adults would prefer to share with a new colleague to create a good impression: a time they were kind and generous to a higher-power person or a time they were kind and generous to a lower-power person at work. In Study 2, we endow participants with real money and test who within a power hierarchy (if anyone) they choose as beneficiaries when their actions are public versus private. We also test whether participants choose lower-power beneficiaries to bolster their reputation. Finally, Study 3 focuses on a different type of prosocial behavior - volunteering. Participants first report their lay theories about the effects of choosing beneficiaries at different levels of the hierarchy (i.e., to what extent these choices would promote reputation-building, ingratiation and altruism). Then they actually choose a beneficiary. This procedure enables us to assess a variety of motives that may operate in the context of prosocial behavior in hierarchies, and to test the causal relationship between lay theories about reputation-building and beneficiary choice.

We report all measures, manipulations, and data exclusions. Target sample sizes were determined a priori, and we did not collect further data after analysis. Informed consent was obtained from all participants and all studies received institutional ethics approval. All studies were pre-registered and all study materials have been publicly posted.²

6. Study 1

In Study 1, we investigated participants' choices of whether they would rather publicize a kind act that they performed for a higher-power or lower-power beneficiary. Specifically, we invited individuals with work experience to recall separate instances in which they did something kind for a higher-power and a lower-power colleague. Then, we asked them to imagine they wanted to make a good impression on a new co-worker, thereby activating a reputational concern. Participants then reported how much they wanted to share each story with the new co-worker - in other words, which story they would prefer to be made public. Based on the theories outlined above, we predicted that participants would prefer to share a story about offering kindness to a lower-power colleague versus a higher-power colleague, because they believe showing kindness to a lower-power colleague more effectively signals their own high moral character.

This study was pre-registered.³

6.1. Method

6.1.1. Participants

A small pilot study ($N = 8$) indicated an effect size of $d_z = 0.5$ for the

² https://osf.io/mzkdx/?view_only=9564f554ef2a48ce903b4e35ae27f823

³ https://aspredicted.org/NVH_RRX

paired samples *t*-test on the primary dependent measure (desire for others to hear about a prosocial act directed at higher- vs. lower-power beneficiary). We used this to run a power analysis on G*Power (Faul, Erdfelder, Buchner, & Lang, 2009). To achieve power of 0.8, this required a total of $N = 34$. To be conservative, and given the possibility of exclusions, 50 participants were recruited.

The final dataset included 42 online participants (43% women; $M_{age} = 29.76$ years, $SD = 10.01$) with work experience. Using a sensitivity power analysis in G*Power, we found that with 42 participants, the smallest effect size we could detect at 80% power ($\alpha = 0.05$) would be $d_z = 0.44$. Participants were recruited from prolific.co and reported residing in the United Kingdom or Ireland. 73.8% identified as White, 2.4% as Black or African, 7.1% as Asian and 16.7% as Hispanic/Latino/a. We selected prolific.co as an online participant provider because we wanted working adults who speak English and who are familiar with the British currency system. Prolific.co allows researchers to select respondents within a certain geographic area, and who have different types of work experience. Because the study asked participants to reflect on prior work experiences, participants who had previously indicated to prolific.co that they are “Not in paid work (e.g. ‘homemaker’, ‘retired’, or ‘disabled’)” were not allowed to participate. There was one duplicate IP address in this study, which we specified a priori would be excluded from the analysis. Also as pre-registered, participants who failed an attention check were directed out of the study before they could see any study materials.

The design was a one-factor (relative power of beneficiary: higher vs. lower) within-participant design. Participants were paid £0.30.

6.1.2. Procedure

All participants recalled a time from their past work experience in which they had done something kind and generous for someone with less power in the company, as well as a time in which they had done something kind and generous for someone with more power in the company. The order of which came first (higher-power vs. lower-power) was randomized. Power hierarchies tend to be clearly defined at work, and therefore we assumed that examples would be relatively easy for participants to generate. After each story, participants were asked to describe their relationship to the person, what they did that was kind and generous, and how costly the act was to them (5-point scale: 1 = *Not at all costly*, 5 = *Extremely costly*). We asked this last question to ensure that the cost of the act was not systematically different across the two types of stories. It was not, $t(41) < 0.001$, $p > .999$, $d_z < 0.001$ ($M_{lower-power\ beneficiary} = 2.33$, $SD = 1.00$; $M_{higher-power\ beneficiary} = 2.33$, $SD = 0.95$).

Next participants were asked to imagine that a new person had joined the company, would be on their team, and that they wanted to make a good impression on this person. They answered how much they would want the new hire to hear about each of the two stories that they had generated earlier (5-point scale: 1 = *Would NOT want them to hear about it*, 5 = *Would DEFINITELY want them to hear about it*). Finally, participants answered a forced choice question: “If you had to choose one story for this new hire to hear about, which one would it be?” (1 = *The story about me doing something kind and generous to a higher power person*, 2 = *The story about me doing something kind and generous to a lower power person*). The order in which these items were presented matched the order in which the stories were previously generated by participants.

Finally, participants reported their age, ethnicity and gender.

6.2. Results and discussion

We predicted that participants would prefer a new colleague to hear about a story in which they were kind and generous to a lower-power beneficiary over one in which they were generous to a higher-power beneficiary. A paired-samples *t*-test on desire for a new colleague to hear about the story, with power of beneficiary (higher- vs. lower-) as a

within-participants factor, revealed a significant effect, $t(41) = -2.47$, $p = .018$, $d_z = 0.39$, with means in the predicted direction ($M_{lower-power\ beneficiary} = 3.10$, $SD = 1.12$; $M_{higher-power\ beneficiary} = 2.76$, $SD = 0.98$).⁴

A chi-square analysis on the forced-choice item - which tests the null hypothesis that participants are equally likely to choose the low-power and the high-power beneficiary stories - revealed a significant effect, $\chi^2(1, N = 42) = 16.10$, $p < .001$. When asked to choose which story to share with a new colleague to create a good impression, 81.1% of participants chose the story in which they were generous to a lower-power beneficiary and 18.9% chose the story in which they were generous to a higher-power individual.

The results of Study 1 support the notion that people believe choosing a lower-power beneficiary signals better reputational qualities than choosing a higher-power beneficiary does. This effect emerged even though we did not specify what type of impression participants wanted to make (e.g., related to morality, competence, etc.). Nevertheless, in a context in which reputation concerns are salient, participants chose to have a new colleague hear about a personal story in which they showed kindness to a lower-power co-worker over one in which they showed kindness to a higher-power co-worker. In addition to confirming our predictions, this finding is notable given past findings demonstrating that people who want to signal status tend to invest in products that are visibly associated with wealth and power (Han, Nunes, & Drèze, 2010; Rucker & Galinsky, 2008). Despite this, participants who wanted to create a good impression on a new colleague still preferred to share stories in which they associated with a lower-power individual rather than a higher-power individual. It may be that individuals consider context when demonstrating proximity to high and low power. Our findings show that, when sharing personal acts of kindness, they prefer to demonstrate proximity to lower-power individuals. When telling colleagues about other sorts of social interactions, such as collaborative work projects or after-work socializing, people may prefer to share experiences involving higher-power others. We return to this possibility in the General Discussion.

7. Study 2

Study 2 pushes the findings of Study 1 forward to understand whether and under what circumstances prosocial actors actually *choose* to benefit lower-power individuals in real-life situations. We tested this by endowing participants with real money that they could keep or give to either a higher- or lower-power beneficiary as a gift. The decision to give the money away constitutes a prosocial act, since it benefits the target and comes at a cost to the participant. Our key measure was to whom they would choose to give the money, if anyone. We further manipulated whether or not this choice would be shared with a third party, who in turn would divide a pot of money between themselves and the participant. By sharing this information, we created an incentive for participants to act in ways that the third party would judge positively.

We expected that in the private condition, a certain proportion of individuals would choose to keep the money, a certain proportion would choose the lower-power beneficiary (likely motivated by altruism), and a certain proportion would choose the higher-power beneficiary (likely motivated by ingratiation). We assumed that these motives would still be active and drive choices in the public condition. Therefore, any changes in choice behavior in public versus private must be attributable to reputation concerns. We predicted that a greater portion of participants would choose the lower-power beneficiary in public because they seek to strategically signal positive moral character. We further predicted that the increased tendency to select lower-power beneficiaries in public versus private contexts should be greater than any change in

⁴ Although the minimum detectable effect size (according to our power analysis) was larger than the actual effect size, which suggests this study was likely underpowered, the results supported our hypothesis.

tendencies to select the higher-power beneficiary.

Support for these predictions could manifest in several different choice patterns. For example, one possibility is that social visibility increases people's propensity to choose *both* high- and low-power beneficiaries, but the increase in proportion of those selecting lower-power beneficiaries is greater than the increase in proportions of those selecting higher-power beneficiaries. This pattern could emerge if people believe that there are reputational gains from choosing a higher-power beneficiary (albeit smaller ones than for choosing a lower-power beneficiary; see Carlson & Zaki, 2018; Inesi et al., 2012 for supportive data). Another possible pattern that could emerge is that more people choose lower-power beneficiaries in public contexts compared to private ones, whereas *fewer* people select higher-power beneficiaries. However, above and beyond the specific pattern that emerges, we predict that our key predictions will receive support.

A second goal of Study 2 was to empirically demonstrate the underlying mechanism: that participants are more likely to choose the lower-power beneficiary in public (versus private) *because* they seek to strategically improve their moral reputation. Thus, after making their choice of what to do with the money we endowed them, we asked participants to report to what extent a desire to signal moral character affected their choice. We predicted that a desire to signal moral character would more strongly drive participants' choice of the lower-power beneficiary in the public condition compared to the private one.

This study was pre-registered.⁵

7.1. Method

7.1.1. Participants

G*power (Faul et al., 2009) was used to determine sample size in advance. We used the results from another pre-registered study (see Supplemental Materials B) to determine sample size. At the time we ran this study, we believed that the best statistical test of our predictions was a multinomial logistic regression. Although we have since changed this analysis strategy, we computed our a priori power analyses using this statistical test. G*power can only compute sample sizes for binomial (not multinomial) logistic regression; therefore, we used a dichotomous version of the choice variable (0 = keep bonus or give to VP, 1 = give to Assistant) and ran a logistic regression analysis. Using the z test menu for logistic regression (binary), we entered the odds ratio results (3.27) and $\Pr(H1) = 0.45$, and assumed two-tailed, $\alpha = 0.05$ and equal cell sizes. To achieve power of 0.8, this suggested a minimum sample size of 112. We increased this to 150, since we also included a new, untested variable (i.e., the proposed mediator). We posted a request for 214 participants, which provides a further 30% buffer to leave room for exclusions based on the pre-registered criteria (e.g., fail attention checks and therefore do not provide data, come from duplicate IP addresses; Thomas & Clifford, 2017).

The final dataset included 176 online participants (73.9% women; $M_{age} = 32.60$ years, $SD = 10.78$) who were recruited through prolific.co. Using a sensitivity power analysis in G*Power, we found that with 176 participants, the smallest effect size we could detect at 80% power ($\alpha = 0.05$) would be $f = 0.21$. Because our study materials include reference to British currency, we recruited only participants residing in the UK or Ireland. 83.5% identified as White, 4.5% as Black, 6.8% as Asian, 1.7% as Latino/a and 3.4% identified as Other. As pre-registered, participants from duplicate IP addresses and with invalid prolific IDs were excluded from the final dataset in an effort to remove bots. For the same reason, attention checks were included in the study, one before any materials were shown, and others after key pieces of information were shared. Participants whose responses indicated that they were not paying attention were directed out of the study without responding to the dependent measures. The design was a one-factor (context: private vs.

public) between-participants design. The study was advertised as paying £0.35, with the possibility of earning additional bonus money. All participants earned a further £0.10 bonus.

7.1.2. Procedure

Upon entering the study, participants were told that they would complete two tasks with other online Prolific workers. The first task was described as placing them in a hierarchy with two other workers. The second task was described as involving one other participant (i.e., not one of the workers from the first task). In the *private* condition, the second task was described as a shapes puzzle that they would complete with another participant. In the *public* condition, participants were told that they would interact with another person who would be given £2. This person would decide how to split the money between themselves and the participant.⁶ We endowed the ostensible observer with money so that participants in the public condition would be financially incentivized to signal positive moral character. In real world contexts, people seek to build a reputation for high moral character because group members have resources (e.g., help, money) that they can choose to offer or withhold. Since our participants were anonymous to one another and would not interact again, they may have been less motivated to signal moral character if the financial incentive had not been introduced. Indeed, existing studies of cooperation under private versus public conditions - in which identities are not known (Jordan et al., 2016; Simpson & Willer, 2008) - endow third parties with resources that they can offer or withhold.

At this point, participants began the first task. They were told that they would work in a hierarchy with two other individuals. Participants were assigned to the Manager role, and they had a Vice President (VP) with more power than them, as well as an Assistant with less power than them. We intentionally did not specify the nature of the task, so as not to create expectations around normative behavior linked to certain jobs or industries (Kay, Wheeler, Bargh, & Ross, 2004). The power hierarchy was also displayed in an organizational chart. On the next page, they were told that their work would be evaluated by the VP at the end of the study, and that they would evaluate their Assistant's work. Furthermore, the evaluations would be directly linked to bonus money, with a high evaluation earning £0.20, a moderate evaluation earning £0.10 and a low evaluation earning nothing. Thus, the power hierarchy was communicated through title and responsibility, as well as control over desired resources (i.e., money; French & Raven, 1959; Emerson, 1962).

On the next page, participants were told that, before starting the first task, we were giving them a starting bonus of £0.05. They had three options of what to do with it: keep the bonus (the VP and Assistant would not be told about the bonus), give the bonus to the VP (the Assistant would not be told about the bonus), or give the bonus to the Assistant (the VP would not be told about the bonus). We reasoned that this procedure could activate two possible motives that could be satisfied through their choice of beneficiary. First, participants could choose the higher-power VP in the hopes of ingratiating and securing a higher evaluation (and therefore compensation). Second, participants could choose the lower-power Assistant, consistent with an altruistic desire to benefit others, especially those who are relatively disadvantaged and may appreciate it most.

Participants in the *private* condition then chose what to do with the bonus. Participants in the *public* condition, however, were first told that the other participant in the second task would be informed of their bonus choice before dividing the £2. Then they decided what to do with

⁶ Since the public condition involves allocating monetary units, we used puzzle shapes as a non-monetary proxy in the private condition. Supplemental Materials B contains a related study that describes the second task as a Dictator Game in both the private and public condition, but then later describes whether or not the Dictator would know their bonus choice in the first task. Similar findings emerge across these different manipulations.

⁵ https://aspredicted.org/UAS_RQA

Table 1
Count data and percentages by context condition (Study 2).

Bonus Choice:	Keep bonus	Choose HP beneficiary	Choose LP beneficiary	Total
Condition:				
Private Context	65 (69.1%)	10 (10.6%)	19 (20.2%)	94
Public Context	25 (30.5%)	17 (20.7%)	40 (48.8%)	82

(HP = Higher-power; LP = Lower-power)

the bonus. Thus, participants in the public condition had a third possible motive driving their choice of what to do with the bonus: to signal personal characteristics that the third party would financially reward.

Next, all participants answered four items that measured to what extent their bonus choice was driven by a desire to signal moral character, as well as three items that checked the power manipulation. Finally, participants provided demographic information before being told that they would not actually complete any tasks.

7.1.3. Measures

The choice measure was: What do you want to do with the £0.05 starting bonus? (1 = keep it all, 2 = Give it all to my Vice President, 3 = Give it all to my Assistant).

We measured moral character signaling motives by asking: “To what extent was your choice about the starting bonus driven by a desire...”: “...to be seen as trustworthy”, “...to be seen as kind”, “...not to be seen as strategic” and “...not to be seen as manipulative”. Participants responded on five-point scales (1 = *Not at all*, 5 = *A great deal*). Thus, two items were focused on signaling positive moral character (kind and trustworthy) and two items were focused on not signaling negative moral character (strategic and manipulative). These items were based on traits that are considered valuable (kind, trustworthy) and problematic (strategic, manipulative) from a reputational signaling perspective (Jordan et al., 2016; Wu et al., 2016). A principal components factor analysis with varimax rotation revealed two factors corresponding to positive (the correlation between the two items: $r(174) = 0.71, p < .001$) and negative (the correlation between the two items: $r(174) = 0.69, p < .001$) moral character signaling concerns (the correlation between the positive and negative moral character signaling concerns: $r(174) = 0.19, p = .014$). All loadings on each factor exceeded 0.9.

The power manipulation check questions asked participants to rate how much power each of the following roles had in the first task: Manager (your role), VP and Assistant. Each of these was responded to on a five-point scale (1 = *Very little power*, 5 = *A great deal of power*).

Finally, participant reported their age, ethnicity and gender.

7.2. Results and discussion

7.2.1. Power manipulation check

A repeated-measures analysis on the three power manipulation check questions revealed a significant omnibus effect, $F(1, 174) = 493.63, p < .001, \eta_p^2 = 0.74$. Follow-up contrasts showed that participants reported that they ($M = 3.47, SD = 0.76$) had more power than their Assistant ($M = 1.68, SD = 0.88$), $t(175) = 22.25, p < .001, d_z = 1.68$, and that their VP ($M = 4.26, SD = 1.12$) had more power than they did, $t(174) = -6.87, p < .001, d_z = 0.52$.

7.2.2. Bonus choice

A chi-square test on the 2×3 table comprised of (context: public vs. private) and (bonus choice: keep vs. share with VP vs. share with Assistant) revealed a significant effect, $\chi^2(2, N = 176) = 26.37, p < .001$ (Pearson test). Participants were more likely to choose the lower-power beneficiary in the public condition (48.8%) compared to the private condition (20.2%). Participants were also more likely to choose the higher-power beneficiary in the public (20.7%) versus the private condition (10.6%). See Table 1.

We predicted that the proportion of individuals choosing a lower-

power beneficiary would be larger in the public condition (48.8%) compared to the private condition (20.2%). The difference was statistically significant, $z = 4.14, p < .001$. We also predicted that the increase in the proportion of participants choosing the lower-power beneficiary in public versus private would be larger than the increase in the proportion of participants choosing the higher-power beneficiary. 28.6% more participants chose the lower-power beneficiary in public compared to private; 10.1% more participants chose the higher-power beneficiary in public compared to private. The difference did not reach conventional levels of statistical significance, $z = 1.79, p = .074$, although the pattern trended in the right direction.^{7,8}

7.2.3. Moral character signaling motives

We predicted that context would affect moral character signaling motives, such that participants in the public condition would have a stronger motive than those in the private condition. We ran two ANOVAs to test this prediction, first on the motive to signal positive moral character, and second on the desire to avoid signaling negative moral character.

A one-way ANOVA on positive moral character signaling motives revealed a significant effect of context, $F(1, 173) = 36.91, p < .001, \eta_p^2 = 0.18$. Participants in the public condition were significantly more likely to report that their choice was driven by a desire to signal positive moral character traits ($M_{\text{private}} = 2.24, SD = 1.19; M_{\text{public}} = 3.38, SD = 1.29$).

Similarly, an ANOVA on motive to avoid signaling negative moral character also revealed a significant effect of context, $F(1, 173) = 4.02, p = .047, \eta_p^2 = 0.02$ ($M_{\text{private}} = 2.44, SD = 1.35; M_{\text{public}} = 2.82, SD = 1.14$). Participants in the public condition were significantly more likely to report that their choice was driven by a desire to avoid signaling negative moral character traits.

7.2.4. Indirect effects

We predicted that participants would be more likely to select the lower-power beneficiary (over other options) in the public compared to the private context *because* they wanted to signal positive moral character. We used PROCESS Model 4 (10,000 iterations; Hayes, 2017) to test this with context (private versus public) as the predictor variable and the desire to signal positive moral character as the mediator. However, it is possible that their behavior informed what they believed

⁷ We conducted meta-analyses of the two effects reported in this paragraph (including Studies 2 and 3 and the supplemental study) using the methodology of McShane and Böckenholt (2017) modified as per McShane and Böckenholt (2020) to account for our trichotomous dependent measure. A larger proportion of participants chose the lower-power beneficiary in public versus private (increase of 21.2%, $SE = 6.2\%, N = 941, z = 3.41, p = .001$). Furthermore, the *change* in proportion of those choosing the lower-power beneficiary was larger than the *change* in proportion of those choosing the higher-power beneficiary (difference between these changes in proportion = 21.3%, $SE = 4.4\%, N = 941, z = 4.84, p < .001$).

⁸ The analyses reported in this paragraph were altered from the pre-registration plan because they provide a more precise test of our predictions. The pre-registered analyses also yield statistically significant results. Multinomial logistic regression results: coefficient for give to the assistant (vs. keep), $B = -1.70$ ($SE B = 0.37$), $p < .001$; coefficient for give to the VP (vs. keep), $B = -1.49$ ($SE B = 0.46$), $p = .001$.

their own motive to be. We think this is less likely and less intuitive, but it remains a plausible alternative because our mediator and outcome variable were measured in the same study.

We recoded the choice variable into 0 = “keep” or “give to VP” and 1 = “give to Assistant”. Supporting our prediction, the indirect effect was significant, $b = 1.18$, $SE = 0.31$, 95% CI [0.71, 1.91]. Because the ANOVA on desire to avoid signaling a negative motive was affected by context, we also ran an exploratory analysis to test whether context affected bonus choice through a desire to avoid a negative reputation. The result was not significant, $b = -0.11$, $SE = 0.08$, 95% CI [-0.29, 0.01].

Study 2 yielded results consistent with our predictions: significantly more participants chose to offer their bonus to the lower-power beneficiary in the public condition compared to the private condition. This choice was driven by an instrumental desire to signal positive moral character traits to an uninvolved third party. Participants also chose the higher-power beneficiary more often in the public condition, but the increase for lower-power beneficiaries was marginally larger. This suggests that prosocial actors are particularly prone to choose lower-power beneficiaries to strategically bolster their own reputation. These results also support the notion that adding social visibility (i.e., the public condition) affects resource redistribution in social hierarchies. Although both high- and low-power individuals were more likely to receive monetary favors in public, lower-power individuals benefitted marginally more.

One interesting result from Study 2 was the relatively low proportion of participants who chose to give the bonus to their supervisor (an average of 15.3% across conditions). Given existing research that highlights the general tendency to ingratiate higher-power counterparts (e.g., Jones, 1964), we were surprised that this number wasn't higher. This may be because offering money seemed awkward or even inappropriate in the context of ingratiation. At the same time, if the tendency to ingratiate was low among participants, this may have artificially inflated the propensity to instead instrumentally choose the lower-power beneficiary to signal positive moral character in the public condition. In Study 3, therefore, we focus on a different way of offering kindness to another person.

8. Study 3

We have proposed that individuals hold a lay theory that selecting a lower-power beneficiary will be particularly effective in signaling positive moral character to others, and thus are more likely to choose lower-power beneficiaries when reputation concerns are heightened. Studies 1 and 2 offered support for some aspects of this theory, but did not directly test whether participants hold our predicted lay theory, nor did Studies 1 and 2 demonstrate the lay theory's importance in driving beneficiary choice. Therefore, our first goal in Study 3 was to directly test whether participants believe that choosing a lower-power beneficiary is more effective than choosing a higher-power beneficiary for fostering a positive reputation. We predicted that this belief would be stronger in public conditions, where third parties will know about one's choice of beneficiary, and therefore can serve to build one's reputation.

In addition, we sought to enrich our understanding of other motives that may operate in power hierarchies. To this end, we also tested lay theories linking beneficiary choice (higher- versus lower-power) with a desire to ingratiate and a desire to altruistically help others. In doing so, we aimed to paint a fuller and more complete picture of the motives operating within power hierarchies, and presumably driving beneficiary choice. Indeed, after reporting these lay theories, participants then chose a beneficiary from the power hierarchy in which we placed them for the experiment.

Another objective of Study 3 was to conceptually replicate the effects of context on beneficiary choice with a new form of prosocial behavior – namely, volunteering time rather than money. We assumed that in work groups, people are incentivized to build a moral reputation because

colleagues have the ability to offer their assistance (e.g., time, expertise) to those with positive reputations. That is, we did not believe an explicit financial incentive was necessary to elicit the hypothesized results.

We also sought to conceptually replicate our effects in a more ecologically valid setting. At the time we ran this study, work-from-home laws were being lifted in the United Kingdom, and workers were being encouraged to return to the office. We took advantage of this real-world situation and designed a study in which we asked participants to imagine returning to face-to-face work in the office. Furthermore, we asked participants to imagine that, being relatively new to the company, they had not met their colleagues in person before. In sum, we created a procedure to mimic what many workers were experiencing in real life at that time.

A final goal of Study 3 was to identify an experimental paradigm in which more participants chose to ingratiate their boss in the private condition. We did this to rule out the possibility that the results of Study 2 could be attributed to a general reluctance to choose higher-power beneficiaries. Offering money to one's boss might have felt awkward and unusual, so we changed the operationalization of prosocial behavior in Study 3. Specifically, we gave them the option of choosing to donate time to either their boss or subordinate (or to choose to not donate their time).

Importantly, even if the proportion of individuals choosing the higher-power beneficiary is higher in the private condition compared to what it was in Study 2 (i.e., higher than 10.2%), our predictions remain the same. We predict that the proportion of individuals choosing the lower-power beneficiary will increase in public compared to private, and this increase will be larger than the increase (or decrease) in the proportion of individuals choosing the higher-power beneficiary.

This study was pre-registered.⁹

8.1. Method

8.1.1. Participants

We recruited 500 participants with the aim of having >462 valid participants in order to achieve 80% power to detect the mediated effect, assuming small effect sizes for both the “a” and “b” paths (Fritz & MacKinnon, 2007). Four-hundred seventy-nine working adults (44% men, 55% women, 1% Other) were recruited through prolific.co.uk. Using a sensitivity power analysis in G*Power, we found that with 479 participants, the smallest effect size we could detect at 80% power ($\alpha = 0.05$) would be $f = 0.13$. Since our study materials ask participants to choose an activity for the upcoming weekend, we recruited participants on a Thursday (all of them completed the study on that day) to increase the realism. All reported living in the United Kingdom, which was important because our materials contained UK-specific information. Also, we selected working adults because we wanted participants who could more easily imagine real-world organizational experiences similar to the experience described in our materials. Their mean age was 38.73 years ($SD = 12.19$). All participants were paid £0.70.

8.1.2. Procedure

Participants were asked to imagine they work on the Project Management team in a medium-sized construction firm that builds multi-family residences. They read that they joined the team four months earlier, but they had been working at home because of COVID policies. Then, they read that all employees had been asked to return to the office full-time a few weeks earlier, and that they were just getting to know their co-workers. Participants learned that they held a middle-management role on the Blue Project Management team, with a higher-power boss and a lower-power assistant.

Power was manipulated by informing participants that their boss assessed their performance each year, and that this determines their

⁹ https://aspredicted.org/S8V_5SS

wage increase and promotion potential. The same dynamic was true of the relationship between the participant and their assistant.

Next, participants learned about two other teams in the Project Management group (the Red team and the Yellow team), each having the same three-person structure as the Blue team. They read that all the teams work in the same area of the building, either in closed-door offices (*private* condition) or in an open-plan structure (*public* condition). All participants then read that they interacted with all members of the Project Management team on a regular basis and therefore wanted to make a good impression on both their boss and the Project Management team more generally.

In the next portion of the study, participants were told that the company encourages giving back to the community and that their boss and assistant would each be volunteering at different homeless shelters that weekend. They were told that each of these individuals would appreciate having someone else join them, since they care a lot about their chosen charity.

In the *private* condition, participants were told that, if they chose to volunteer, they would need to discuss logistics with the person, but could do it over email so no one else would know they were volunteering. In the *public* condition, participants were told that, if they chose to volunteer, they would need to discuss logistics and, given the open plan in the office, everyone else on the Project Management team would know that they chose to volunteer with that person.

Finally, participants completed the dependent measures, were debriefed and were then paid.

8.1.3. Measures

8.1.3.1. Lay theories. Participants were asked, if they chose to volunteer with their *Assistant*, to what extent would: (1) they build a better reputation in the group (“...others on the Project Management team think better of you” and “...you build a good reputation in the group”; $r(479) = 0.75, p < .001$), (2) they ingratiate their *Lead* (“...your *Lead* think better of you;” 1 item), (3) they ingratiate their *Assistant* (“...your *Assistant* think better of you;” 1 item), and (4) they feel that they had done a good deed (“...feel that you had helped someone” and “feel that you had done a good deed”; $r(479) = 0.73, p < .001$).

Participants answered the same questions, but in the context of choosing to volunteer with their *Lead* (build reputation ($r(479) = 0.71, p < .001$), ingratiate *Lead*, ingratiate *Assistant*, make them feel they had done a good deed ($r(479) = 0.76, p < .001$)).

All items were reported on 5-point scales. The order of the target was randomized: Half of the participants responded to the questions about choosing their *Assistant* first and then questions about choosing their *Lead* second, and half of the participants saw the reverse order.

8.1.3.2. Beneficiary choice. Participants answered how likely they are to volunteer with their *Lead*, and how likely they are to volunteer with their *Assistant* (5-point scales). They then answered a forced choice item about what they would do (three options: not volunteer, volunteer with *Lead*, volunteer with *Assistant*).

8.1.3.3. Manipulation checks. To check for the effectiveness of our context manipulation, participants reported how many people would know about their actions if they chose to volunteer (1 = *Only the person I volunteer with*, 5 = *Almost everyone on the team*). We also asked participants how much power they had in the organization compared to the *Lead*, and compared to their *Assistant* (1 = *I have less power than my Lead/Assistant*, 3 = *I have the same amount of power as my Lead/Assistant*, 5 = *I have more power than my Lead/Assistant*).

8.2. Results

8.2.1. Manipulation checks

A one-way ANOVA with context (*private* versus *public*) predicting the context manipulation check revealed a strong and significant effect, $F(1, 477) = 6022.85, p < .001, \eta_p^2 = 0.93$. Participants believed that more people would know about their actions in the *public* condition ($M = 4.83, SD = 0.55$) compared to the *private* condition ($M = 1.11, SD = 0.50$).

A repeated measures ANOVA comparing the two power manipulation check items revealed a strong and significant effect, $F(1, 469) = 3161.93, p < .001, \eta_p^2 = 0.87$. As expected, when answering about their *Assistant* ($M = 4.46, SD = 0.61$), participants reported having relatively more power than when answering about their *Lead* ($M = 1.50, SD = 0.67$).

8.2.2. Lay theories

We predicted that participants would believe that choosing the lower-power beneficiary would be more effective in fostering a group-based reputation in the *public* versus the *private* condition. Supporting our prediction, a one-way ANOVA revealed a significant effect, $F(1, 477) = 49.11, p < .001, \eta_p^2 = 0.09$ ($M_{private} = 3.22, SD = 1.14; M_{public} = 3.85, SD = 0.73$).

Furthermore, we predicted that participants would believe the reputational benefits of choosing a lower-power beneficiary in *public* versus *private* were greater than the reputational benefits of choosing a higher-power beneficiary. A mixed between- and within-participant ANOVA, with context as the predictor variable and belief that choosing the lower-power versus choosing the higher-power beneficiary would build a better reputation as the within-participant outcome variable, revealed a significant interaction, $F(1, 477) = 17.36, p < .001, \eta_p^2 = 0.04$. As predicted, the reputation boost from choosing a lower-power beneficiary in *public* compared to *private* ($M_{private} = 3.22, SD = 1.14; M_{public} = 3.85, SD = 0.73; F(1, 477) = 49.11, p < .001, \eta_p^2 = 0.09$) was greater than the boost from choosing a higher-power beneficiary ($M_{private} = 3.18, SD = 1.10; M_{public} = 3.48, SD = 0.88; F(1, 477) = 10.32, p = .001, \eta_p^2 = 0.02$). This analysis also revealed a significant main effect of context, $F(1, 477) = 31.80, p < .001, \eta_p^2 = 0.06$, such that participants believed that they would more effectively build their reputation in the *public* context versus the *private* one. Finally, there was a significant main effect of beneficiary choice (the within-participant variable), $F(1, 477) = 27.05, p < .001, \eta_p^2 = 0.05$, with participants believing that choosing the lower-power beneficiary would lead to a better reputation than choosing a higher-power beneficiary.

8.2.3. Beneficiary choice

8.2.3.1. Continuous measures. We predicted that participants would report being more likely to choose the lower-power beneficiary in *public* compared to *private*. A one-way ANOVA with context predicting reported likelihood of choosing the lower-power beneficiary revealed a significant effect, $F(1, 477) = 4.38, p = .037, \eta_p^2 = 0.01$, with the means in the direction predicted ($M_{private} = 3.55, SD = 1.05; M_{public} = 3.74, SD = 0.93$).

Furthermore, we predicted that the increased likelihood of choosing the lower-power beneficiary in the *public* versus the *private* condition would be larger than the change in likelihood of choosing the higher-power beneficiary. A mixed within- (likelihood of choosing higher-power versus lower-power beneficiary) and between-participants (context: *private* versus *public*) ANOVA revealed a significant interaction, $F(1, 477) = 6.23, p = .013, \eta_p^2 = 0.01$. The results suggest that, as predicted, the increased likelihood of choosing the lower-power beneficiary in the *public* condition compared to the *private* condition

Table 2
Count data and percentages by context condition (Study 3).

Volunteering Choice:	Don't volunteer	Volunteer with HP beneficiary	Volunteer with LP beneficiary	Total
Condition:				
Private Context	24 (9.3%)	110 (42.6%)	124 (48.1%)	258
Public Context	23 (10.4%)	71 (32.1%)	127 (57.5%)	221

(HP = Higher-power; LP = Lower-power)

($M_{\text{private}} = 3.55$, $SD = 1.05$; $M_{\text{public}} = 3.74$, $SD = 0.93$; $F(1, 477) = 4.38$, $p = .037$, $\eta_p^2 = 0.01$) was greater than the increased (in this case, decreased) likelihood of choosing the higher-power beneficiary in public versus private ($M_{\text{private}} = 3.50$, $SD = 1.08$; $M_{\text{public}} = 3.35$, $SD = 1.02$; $F(1, 477) = 2.48$, $p = .116$, $\eta_p^2 = 0.01$).

A significant main effect of beneficiary also emerged, $F(1, 477) = 10.07$, $p = .002$, $\eta_p^2 = 0.02$ in that participants were more likely to say they would choose to volunteer with their assistant ($M_{\text{lower-power}} = 3.63$, $SD = 1.00$) compared with their Lead ($M_{\text{higher-power}} = 3.43$, $SD = 1.05$).

8.2.3.2. Categorical measure. A chi-square analysis on the 2 (context: private, public) x 3 (choice: not volunteer, higher-power beneficiary, lower-power beneficiary) contingency table revealed a marginal effect, $\chi^2(2, N = 479) = 5.64$, $p = .060$ (Pearson test). Participants were more likely to choose the lower-power beneficiary in the public condition (57.5%) compared to the private condition (48.1%). In contrast, participants were less likely to choose the higher-power beneficiary in the public (32.1%) versus private condition (42.6%). See Table 2.

We predicted that the proportion of individuals choosing a lower-power beneficiary would be larger in the public condition (57.5%) compared to the private condition (48.1%). The difference was statistically significant, $z = 2.07$, $p = .039$. We also predicted that the increase in the proportion of participants choosing the lower-power beneficiary in public versus private would be larger than the increase (or decrease) in the proportions choosing the higher-power beneficiary. 9.4% more participants chose the lower-power beneficiary in public compared to private; 10.5% fewer participants chose the higher-power beneficiary in public compared to private. The difference was statistically significant, $z = 2.34$, $p = .019$. Taken together, this again supports the prediction that, as reputational concerns become more salient, lower-power beneficiaries are a more attractive choice than higher-power beneficiaries are.^{10,11}

8.2.4. Indirect effects

We predicted that participants would be more likely to choose the lower-power beneficiary in the public compared to the private condition because they were more likely to believe doing so would benefit their reputation. We ran PROCESS Model 4 with 10,000 bootstrap samples to test this. For both models, we inserted lay theories as the mediator and choice as the outcome variable. However, it is possible that people's

¹⁰ We conducted meta-analyses of the two effects reported in this paragraph (including Studies 2 and 3 and the supplemental study), using the methodology of McShane and Böckenholt (2017) modified as per McShane and Böckenholt (2020) to account for our trichotomous dependent measure. A larger proportion of participants chose the lower-power beneficiary in public versus private (increase of 21.2%, $SE = 6.2\%$, $N = 941$, $z = 3.41$, $p = .001$). Furthermore, the change in proportion of those choosing the lower-power beneficiary was larger than the change in proportion of those choosing the higher-power beneficiary (difference between these changes in proportion = 21.3%, $SE = 4.4\%$, $N = 941$, $z = 4.84$, $p < .001$).

¹¹ The analyses reported in this paragraph were altered from the pre-registration plan because they provide a more precise test of our predictions. The pre-registered analyses also yield statistically significant results. Test for the equality of two proportions results: $\chi^2(2, N = 479) = 7214.99$, $p < .001$, 95% CI [0.195, 0.204].

choice preferences affect their lay theories, rather than people's lay theories affecting their choice preferences. We think this is less likely and less intuitive, but it remains a plausible alternative because our mediator and outcome variables were measured in the same study.

In our first model, context was inserted as the independent variable, the belief that choosing the lower-power beneficiary would build a better reputation was entered as the mediator, and likelihood of choosing the lower-power beneficiary (continuous variable) was the outcome variable. The results revealed a significant indirect effect, 0.22, $SE = 0.04$, 95% CI [0.15, 0.30], supporting our predictions. We also tested a similar model, but this time with participants' choice of beneficiary (lower-power versus not lower-power) as a dichotomous outcome variable. Again, the result revealed a significant indirect effect, 0.22, $SE = 0.07$, 95% CI [0.10, 0.36].

8.2.5. Exploratory analyses

The core predictions of this paper focus on choosing beneficiaries within a power hierarchy to boost one's reputation. However, as mentioned in the introduction, a series of additional motives may operate and drive participants' choice of beneficiary in hierarchical contexts. For example, participants who chose to volunteer with the higher-power beneficiary may have been trying to ingratiate this person, so that they could gain privileged access to the resources that their higher-power colleague controls (e.g., promotions, raises). Alternatively, participants may have chosen the lower-power beneficiary because they truly wanted to help the person that had the least access to desired resources. To further investigate these motives, we ran a series of exploratory analyses investigating lay theories related to pleasing one's boss and helping others.

The first set of analyses focused on lay theories related to ingratiating a higher-power person in the hierarchy, and choice of beneficiary. In the first analysis, we tested which choice of beneficiary (higher-power versus lower-power) participants believed would make their supervisor think better of them (i.e., which would be more effective in ingratiating this person). We ran a repeated measures ANOVA with two levels: "choosing my assistant will make my Lead think better of me" versus "choosing my Lead will make my Lead think better of me". The results revealed a significant effect, $F(1, 478) = 267.23$, $p < .001$, $\eta_p^2 = 0.36$. Participants believed that their Lead would think better of them if they chose this person (i.e., the Lead themselves; $M = 4.11$, $SD = 0.80$) over the choosing the lower-power assistant ($M = 3.16$, $SD = 1.13$). We next tested if participants' lay theory that choosing their Lead would make their Lead think better of them influenced their choice of beneficiary. We regressed likelihood of choosing the higher-power beneficiary on this lay theory. A regression revealed a significant effect, $\beta = 0.33$, $t(477) = 7.60$, $p < .001$. In sum, consistent with existing research, participants believed that the best way to ingratiate one's boss would be to offer them kind acts (rather than offering them to someone else). Furthermore, this belief drove participants' tendency to choose this person as a beneficiary.

Our second set of analyses investigated lay theories related to altruism and choice of beneficiary. A repeated-measures ANOVA testing whether participants believed choosing their Lead versus their Assistant would make them feel they had done a good deed revealed a significant effect, $F(1, 478) = 42.27$, $p < .001$, $\eta_p^2 = 0.08$. Participants believed that choosing the lower-power beneficiary was more likely to make them feel

they had done a good deed ($M = 4.32$, $SD = 0.68$) than choosing the higher-power beneficiary ($M = 4.14$, $SD = 0.85$). Finally, we tested whether participants were more likely to report choosing the lower-power beneficiary if they felt that doing so would make them feel they had done a good deed. A regression revealed a significant positive effect, $\beta = 0.27$, $t(477) = 6.06$, $p < .001$.

8.3. Discussion

The results of Study 3 offer several important contributions. First, we empirically demonstrate the role of lay theories in driving participants' choice of the lower-power beneficiary. Specifically, participants were more likely to choose the lower-power beneficiary in public (compared to private) because they believed this choice would be more effective in bolstering their reputation in the group. Second, this study also provides evidence for other motives operating in the context of choosing a beneficiary for a kind act offered within a power hierarchy. Namely, we show evidence for ingratiation and altruism motives: participants believed that choosing to volunteer with their boss would be more likely to make the boss see them more positively. It also made participants more likely to select the boss as a beneficiary. In addition, participants believed that choosing the lower-power beneficiary was most effective in doing a good deed – and this belief drove their choice of the lower-power beneficiary. However, above and beyond these dynamics, we were able to demonstrate empirically that self-interested reputation motives drove the choice of the lower-power beneficiary, especially in contexts where reputation motives are salient (i.e., public contexts).

An additional contribution of Study 3 was to conceptually replicate the results of our prior study in a context that involved volunteering one's time, rather than donating money. We effectively created a realistic paradigm in which large portions of participants selected each of the choice options: refraining from offering generosity, choosing the lower-power beneficiary and choosing the higher-power beneficiary. Therefore, unlike Study 2, which may have made choosing the higher-power beneficiary unattractive, Study 3 seemed to allow for participants to pursue a variety of motives. Furthermore, like Study 2, the results here revealed that social visibility prompted a redistribution of resources in the power hierarchy. In Study 3, the pattern was more extreme however: participants in the public condition were less likely to choose the higher-power beneficiary (compared to in private) and more likely to choose the lower-power beneficiary. Thus, resources were taken away from the top of the hierarchy and given to those at the bottom.

9. General discussion

Existing research convincingly demonstrates that in social hierarchies, people are assumed to offer kind acts to higher-power beneficiaries for self-interested reasons, whereas they are assumed to offer kind acts to lower-power beneficiaries for altruistic reasons. In the current research, however, we propose that under some conditions, particularly when one's reputation is at stake, people choose lower-power beneficiaries for strategic, self-interested reasons (i.e., to boost one's standing within the group). In Study 1, we found that people believe observers will evaluate their kind acts more positively when the beneficiary has less versus more power than they do. In Study 2, we demonstrated the effects of this lay theory on beneficiary selection. Using real money that participants could keep or give to a higher- or lower-power beneficiary, we found that the increased tendency to give the money away in public versus private contexts was more focused on the lower-power beneficiary than the higher-power one. Furthermore, the increased tendency to give to lower-power beneficiaries in the public condition was explained by the participant's desire to signal positive moral character to others. In Study 3, we conceptually replicated the results of Study 2 using a different act of kindness – offering one's time instead of money – and examined the roles of different lay theories in beneficiary selection, as a more direct test of why people choose lower-

power beneficiaries in the public than private condition. Specifically, we found that people were more likely to select lower-power beneficiaries in public because they believed doing so would bolster their reputation.

9.1. Theoretical and practical contributions

Our findings offer noteworthy contributions to several areas. Perhaps most centrally, our research introduces a counterintuitive way in which people strategically use social hierarchies to signal desired qualities to observers. In contrast to the widespread belief that people who choose to offer kindness to high-power individuals are motivated by self-interest, and those who instead offer kindness to lower-power individuals are altruistically motivated, we reveal that people can also choose lower-power beneficiaries for self-interested reasons. Specifically, we propose that people consider the effect that their choice of beneficiary may have on their moral reputation, and then choose the beneficiary that best signals their own moral character. Therefore, when reputation concerns are more salient (e.g., in public versus private contexts), people are more likely to choose lower-power beneficiaries because they believe doing so sends a better signal of their own moral reputation. In introducing this perspective, we demonstrate that outside observers – despite overestimating others' self-interested motives much of the time (Miller & Ratner, 1998) – have a blind spot to a certain form of self-interestedness. That is, observers may assume that giving “downward” is altruistic when, oftentimes, doing so is actually self-interested.

Indeed, a second contribution of this work is to fundamentally reframe our understanding of the benefits offered by power hierarchies. As has been well documented, the resources that are unequally distributed within power hierarchies (e.g., money, network access) undoubtedly motivate strategic ingratiation to be directed at higher-power individuals. Here, we reveal a second important incentive that power hierarchies offer: the ability to build a positive reputation. By highlighting this second, significant currency that operates in power hierarchies, we can better understand what motivates prosocial giving. In fact, these findings suggest an inherent struggle that (strategic) prosocial actors face when choosing a beneficiary within a hierarchy: should they choose a beneficiary with more power, who can offer resources in return, or should they choose a lower-power beneficiary, in the hopes of bolstering their reputation within the group? Our findings suggest that people may navigate this challenge at least partially by attending to the context (i.e., whether their actions are public or private).

Third, our findings reveal the importance of social characteristics – namely, the beneficiary's position in a hierarchy – when selecting beneficiaries of kind acts. Social psychological research has demonstrated that power can play an important role in the interpersonal perception process, affecting tendencies to stereotype (Fiske, 1993), objectify (Gruenfeld, Inesi, Magee, & Galinsky, 2008), and sexually harass (Williams, Gruenfeld, & Guillory, 2017) others, among other outcomes. Here, we show for the first time that people hold lay theories about how potential beneficiaries' power (or lack thereof) may impact judgments of their actions, and that these lay theories can in turn affect their choice of beneficiary.

Research on the role of power in beneficiary selection is particularly important for understanding how inequalities are maintained or dismantled in society. For example, women typically have lower power and fewer resources than men in social groups, and in many cultural contexts, certain ethnic and racial groups have more power and resources than others. Therefore, the fact that prosocial actors use social characteristics (i.e., power) to choose beneficiaries of their kindness could either magnify or diminish existing inequalities. The results presented here suggest that this dynamic may sometimes diminish inequalities, in that prosocial actors are increasingly likely to select lower-power beneficiaries when their actions are observed by others. In fact, it is because lower-power individuals have fewer resources to offer in return that they make better beneficiaries, from a signaling perspective. Many societies have established policies and mechanisms aimed at helping those with less power, including mentoring programs, homeless

shelters, and welfare systems more generally. Here, we argue that social groups have a built-in, reputation-based process that serves this very goal.

This work also provides a significant contribution to research on interpersonal cooperation. Existing literature suggests that motives to bolster one's reputation drive the choice of *whether or not* to engage in prosocial behavior. In these paradigms, participants are typically presented with one potential beneficiary, about which they know nothing (Jordan et al., 2016; Wedekind & Milinski, 2000). In the real world, however, we often have many possible beneficiaries, and have some knowledge about them. In such contexts, how do prosocial actors choose beneficiaries? Our research extends existing work to propose that actors will be more likely to choose beneficiaries that they believe will better communicate their own moral character. We investigated this proposition in the context of power hierarchies, because existing research has shown that observers judge kind acts more positively when they are offered to lower- versus higher-power beneficiaries (Inesi et al., 2021). We found that prosocial actors increasingly choose lower-power beneficiaries when their actions are seen by third parties because they think it offers a better signal of their own moral character.

Finally, our findings contribute to our understanding of how people use hierarchies to signal personal characteristics. According to the compensatory theory of status seeking, people who are motivated to seek status prefer proximity to products that are associated with power and status. For example, people who desire status tend to purchase loud luxury goods from prominent brands to signal to third parties that they themselves are wealthy, and not poor (Han et al., 2010; Rucker & Galinsky, 2008). Here, we show the opposite: that individuals prefer third parties to hear about their own past associations with a lower-power rather than a higher-power individual. Taken together, these findings suggest that people leverage hierarchies in various ways to communicate unobservable characteristics to others. When offering kindness, individuals want to signal their altruism and beneficence, and therefore choose lower-power beneficiaries. However, when engaging in actions like socializing, individuals may instead choose higher-power others as a means of signaling to third parties that they are part of the high-status ingroup. This would be an interesting avenue for future research to investigate.

9.2. Limitations and future research directions

Although our findings offer important contributions, there are limitations to the generalizability that deserve mention. First, none of our studies included real-time third-party observers. In Study 2, we deceived participants into believing that the other individuals in the hierarchy and the third-party observers were real, when in fact they were not. In Studies 1 and 3, we asked working adults to imagine a common situation in which they wanted to create a good impression on their colleagues. Therefore, although we attempted in various ways to replicate the experience of choosing a beneficiary in the context of real observers, it would be important for future work to ensure the same effects generalize to real-world situations. Another limitation to the generalizability of our findings is that all studies were run on online samples. Although we wanted to run experiments in the laboratory, this was significantly impeded by the ongoing global pandemic. Thus, it is important for future research to replicate these results in face-to-face interactions.

In addition, Studies 2 and 3 only considered situations in which participants had two potential beneficiaries with different levels of power. We did not test whether the relative power of a beneficiary would affect the decision to give if only a single potential beneficiary were present. For example, imagine that an employee is leaving the office for the evening with a group of colleagues. She sees her boss working late, and she thinks it would be kind to offer to stay and help. The boss is the only potential beneficiary. How likely is the employee to offer her boss help, rather than go out with her colleagues? And more importantly for the current research, would the propensity to offer help be different if the person

working late was a subordinate instead of a boss? We predict that it would - however, this is an empirical question meriting future attention. More generally, this topic would benefit from a greater understanding of relevant moderators. In the current set of studies, we tested the proposed psychological mechanism through mediation. However, it would be useful to identify and test theoretically driven moderators as a way of providing additional support for our proposed mechanism. For example, it would be interesting to test whether the expectation of future interactions with the audience in the public condition would alter our findings. Because we have proposed that this effect is driven by a desire to gain social rewards by fostering a reputation for moral character, it seems likely that the effect would be weaker when there is no expectation of future contact with the observers.

Another area for future research would be to clarify when people are more versus less likely to choose the higher-power beneficiary in public. In Study 2, participants were more likely to choose higher-power beneficiaries when in public (albeit less than lower-power ones). In Study 3, social visibility made participants people less likely to choose higher-power beneficiary. One factor driving this difference may be the perceived value of ingratiation. Ingratiation seemed more attractive in Study 3 (based on the number of people who chose it in private). Therefore, participants may have a lay theory that - were people to know about it - the self-interested motive would be obvious and would lead to negative reputational effects. Another possible factor is the type of observer. In Study 2, the observer was a person who was uninvolved in the hierarchical work group. In Study 3, the "observer" was the other members of their work group. Participants may have believed that members of the work group would judge ingratiation more negatively than an uninvolved third party would. These possibilities are ripe areas for future research.

A related opportunity for future research is to understand the downstream consequences of this redistribution of resources. Studies 2 and 3 both demonstrated that social visibility leads to a shift in where acts of kindness are directed in the hierarchy. In both cases, lower-power individuals were more likely to be chosen as beneficiaries by prosocial actors, suggesting that this dynamic benefits lower-power individuals. At the same time, the ultimate consequence of this choice may be that observers reward the actor for his or her kind acts (Inesi et al., 2021). Therefore, it may be that, whereas lower-power groups may initially benefit from the actor's desire to signal moral character, in the long run higher-power actors may benefit even more due to reputational gains. The long-term resource redistribution implications of kind acts offered within hierarchies are beyond the scope of the current research. However, the current research does have important implications for the maintenance of inequity, and thus deserves attention.

10. Conclusion

Cooperation is a critical component of successful social groups. Although prior work has demonstrated that people's acts of kindness to lower-power beneficiaries are often assumed to stem from altruistic motives, we show that reputation-based motives may also be at play. Our studies demonstrate that people increasingly select lower-power beneficiaries when their actions are public (versus private) because they believe doing so sends a positive signal of their own moral character. Ironically, the result of this self-interested behavior is - at least in the short term - a redistribution of resources from individuals with higher power to individuals with lower power.

OPEN PRACTICES STATEMENT: Studies 1, 2 and 3 were all formally pre-registered (Study 1: https://aspredicted.org/NVH_RRX; Study 2: https://aspredicted.org/UAS_RQA; Study 3: https://aspredicted.org/S8V_5SS). Our Research Ethics Committee did not give permission to post data online, but the corresponding author will share data upon request. Verbatim materials are available on OSF (Supplemental Materials A): https://osf.io/mzkdx/?view_only=9564f554ef2a48ce903b4e35ae27f823

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of Competing Interest

None.

Data availability

Data will be made available on request.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jesp.2022.104441>.

References

- Ariely, D., Bracha, A., & Meier, S. (2009). Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *American Economic Review*, 99(1), 544–555. <https://doi.org/10.1257/aer.99.1.544>
- Batson, C. D. (1998). Altruism and prosocial behavior. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *Vol. 2. The handbook of social psychology* (pp. 282–316). Boston: McGraw-Hill.
- Bénabou, R., & Tirole, J. (2006). Incentives and prosocial behavior. *American Economic Review*, 96(5), 1652–1678. <https://doi.org/10.1257/aer.96.5.1652>
- Bereczkei, T., Birkas, B., & Kerekes, Z. (2007). Public charity offer as a proximate factor of evolved reputation-building strategy: An experimental analysis of a real-life situation. *Evolution and Human Behavior*, 28(4), 277–284.
- Berman, J. Z., Levine, E. E., Barasch, A., & Small, D. A. (2015). The braggart's dilemma: On the social rewards and penalties of advertising prosocial behavior. *Journal of Marketing Research*, 52(1), 90–104. <https://doi.org/10.1509/jmr.14.0002>
- Berman, J. Z., & Silver, I. (2022). Prosocial behavior and reputation: When does doing good lead to looking good? *Current Opinion in Psychology*, 43, 102–107.
- Bohra, K. A., & Pandey, J. (1984). Ingratiation toward strangers, friends, and bosses. *The Journal of Social Psychology*, 122(2), 217–222.
- Carlson, R. W., & Zaki, J. (2018). Good deeds gone bad: Lay theories of altruism and selfishness. *Journal of Experimental Social Psychology*, 75, 36–40. <https://doi.org/10.1016/j.jesp.2017.11.005>
- Cottrell, C. A., Neuberg, S. L., & Li, N. P. (2007). What do people desire in others? A sociofunctional perspective on the importance of different valued characteristics. *Journal of Personality and Social Psychology*, 92(2), 208–231. <https://doi.org/10.1037/0022-3514.92.2.208>
- Cowgill, C. M., Rios, K., & Simpson, A. (2017). Generous heathens? Reputational concerns and atheists' behavior toward Christians in economic games. *Journal of Experimental Social Psychology*, 73, 169–179.
- Critcher, C. R., Inbar, Y., & Pizarro, D. A. (2013). How quick decisions illuminate moral character. *Social Psychological and Personality Science*, 4(3), 308–315. <https://doi.org/10.1177/1948550612457688>
- De Freitas, J., DeScioli, P., Thomas, K. A., & Pinker, S. (2019). Maimonides' ladder: States of mutual knowledge and the perception of charity. *Journal of Experimental Psychology: General*, 148(1), 158–173. <https://doi.org/10.1037/xge0000507>
- Derfler-Rozin, R., Baker, B., & Gino, F. (2018). Compromised ethics in hiring processes? How referrers' power affects employees' reactions to referral practices. *Academy of Management Journal*, 61(2), 615–636. <https://doi.org/10.5465/amj.2015.1355>
- Drory, A., & Zaidman, N. (2007). Impression management behavior: Effects of the organizational system. *Journal of Managerial Psychology*, 22(3), 290–308.
- Emerson, R. M. (1962). Power-dependence relationships. *American Sociological Review*, 27(1), 31–41. <https://doi.org/10.2307/2089716>
- Emler, N. (1990). A social psychology of reputation. *European Review of Social Psychology*, 1(1), 171–193.
- Faul, F., Erdfelder, E., Buchner, A., & Lang, A. G. (2009). Statistical power analyses using G*power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, 41, 1149–1160.
- Feinberg, M., Willer, R., & Schultz, M. (2014). Gossip and ostracism promote cooperation in groups. *Psychological Science*, 25(3), 656–664.
- Fiske, S. T. (1993). Controlling other people: The impact of power on stereotyping. *American Psychologist*, 48(6), 621.
- French, J. R. P., & Raven, B. (1959). The basis of social power. In D. Cartwright (Ed.), *Studies in social power* (pp. 529–569). Ann Arbor, MI: University of Michigan Press.
- Friesen, J. P., Kay, A. C., Eibach, R. P., & Galinsky, A. D. (2014). Seeking structure in social organization: Compensatory control and the psychological advantages of hierarchy. *Journal of Personality and Social Psychology*, 106(4), 590.
- Fritz, M. S., & MacKinnon, D. P. (2007). Required sample size to detect the mediated effect. *Psychological Science*, 18(3), 233–239.
- Griskevicius, V., Tybur, J. M., & Van den Bergh, B. (2010). Going green to be seen: Status, reputation, and conspicuous conservation. *Journal of Personality and Social Psychology*, 98(3), 392–404. <https://doi.org/10.1037/a0017346>
- Gruenfeld, D. H., Inesi, M. E., Magee, J. C., & Galinsky, A. D. (2008). Power and the objectification of social targets. *Journal of Personality and Social Psychology*, 95(1), 111.
- Gruenfeld, D. H., & Tiedens, L. Z. (2010). Organizational preferences and their consequences. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.), *The handbook of social psychology* (pp. 1252–1287). John Wiley & Sons, Inc. <https://doi.org/10.1002/9780470561119.socpsy002033>
- Han, Y. J., Nunes, J. C., & Drèze, X. (2010). Signaling status with luxury goods: The role of brand prominence. *Journal of Marketing*, 74(4), 15–30. <https://doi.org/10.1509/jmkg.74.4.015>
- Hardy, C. L., & Van Vugt, M. (2006). Nice guys finish first: The competitive altruism hypothesis. *Personality and Social Psychology Bulletin*, 32(10), 1402–1413.
- Hayes, A. F. (2017). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach* (2nd ed.). New York, NY: Guilford Publications.
- Inesi, M. E., Adams, G. S., & Gupta, A. (2021). When it pays to be kind: The allocation of indirect reciprocity within power hierarchies. *Organizational Behavior and Human Decision Processes*, 165, 115–126. <https://doi.org/10.1016/j.obhdp.2021.04.005>
- Inesi, M. E., Gruenfeld, D. H., & Galinsky, A. D. (2012). How power corrupts relationships: Cynical attributions for others' generous acts. *Journal of Experimental Social Psychology*, 48(4), 795–803. <https://doi.org/10.1016/j.jesp.2012.01.008>
- Jones, E. E. (1964). *Ingratiation: A social psychological analysis*. New York, NY: Appleton-Century-Crofts.
- Jordan, J. J., Hoffman, M., Nowak, M. A., & Rand, D. G. (2016). Uncalculating cooperation is used to signal trustworthiness. *Proceedings of the National Academy of Sciences*, 113(31), 8658–8663. <https://doi.org/10.1073/pnas.1601280113>
- Kay, A. C., Wheeler, C., Bargh, J. A., & Ross, L. (2004). Material priming: The influence of mundane physical objects on situational construal of competitive behavioral choice. *Organizational Behavior and Human Decision Processes*, 95(1), 83–96. <https://doi.org/10.1016/j.obhdp.2004.06.003>
- Keltner, D., Gruenfeld, D. H., & Anderson, C. (2003). Power, approach, and inhibition. *Psychological Review*, 110(2), 265.
- Kipnis, D., Schmidt, S. M., & Wilkinson, J. (1980). Intraorganizational influence tactics: Explorations in getting one's way. *Journal of Applied Psychology*, 65(4), 440.
- Kunstman, J. W., Fitzpatrick, C. B., & Smith, P. K. (2018). Poisoned praise: Discounted praise backfires and undermines subordinate impressions in the minds of the powerful. *Social Psychological and Personality Science*, 9(4), 470–480. doi: 10.1177/21948550617712028.
- Magee, J. C., & Galinsky, A. D. (2008). Social hierarchy: The self-reinforcing nature of power and status. *Academy of Management Annals*, 2(1), 351–398.
- McShane, B. B., & Böckenholt, U. (2017). Single paper Meta-analysis: Benefits for study summary, theory-testing, and replicability. *Journal of Consumer Research*, 43(6), 1048–1063.
- McShane, B. B., & Böckenholt, U. (2020). Enriching Meta-analytic models of summary data: A thought experiment and case study. *Advances in Methods and Practices in Psychological Science*, 3(1), 81–93.
- Milinski, M., Semmann, D., & Krambeck, H. J. (2002). Reputation helps solve the 'tragedy of the commons'. *Nature*, 415, 424–426. <https://doi.org/10.1038/415424a>
- Miller, D. T., & Ratner, R. K. (1998). The disparity between the actual and assumed power of self-interest. *Journal of Personality and Social Psychology*, 74(1), 53.
- Nowak, M. A., & Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, 437, 1291–1298. <https://doi.org/10.1038/nature04131>
- Rucker, D. D., & Galinsky, A. D. (2008). Desire to acquire: Powerlessness and compensatory consumption. *Journal of Consumer Research*, 35(2), 257–267. <https://doi.org/10.1086/588569>
- Simpson, B., & Willer, R. (2008). Altruism and indirect reciprocity: The interaction of person and situation in prosocial behavior. *Social Psychology Quarterly*, 71(1), 37–52. <https://doi.org/10.1177/019027250807100106>
- Thomas, K. A., & Clifford, S. (2017). Validity and mechanical Turk: An assessment of exclusion methods and interactive experiments. *Computers in Human Behavior*, 77, 184–197. <https://doi.org/10.1016/j.chb.2017.08.038>
- Tiedens, L. Z., & Fragale, A. R. (2003). Power moves: Complementarity in dominant and submissive nonverbal behavior. *Journal of Personality and Social Psychology*, 84(3), 558.
- Vonk, R. (1998). The slime effect: Suspicion and dislike of likeable behavior toward superiors. *Journal of Personality and Social Psychology*, 74(4), 849–864. <https://doi.org/10.1037/0022-3514.74.4.849>
- Wedekind, C., & Milinski, M. (2000). Cooperation through image scoring in humans. *Science*, 288, 850–852. <https://doi.org/10.1126/science.288.5467.850>
- Williams, M. J., Gruenfeld, D. H., & Guillory, L. E. (2017). Sexual aggression when power is new: Effects of acute high power on chronically low-power individuals. *Journal of Personality and Social Psychology*, 112(2), 201.
- Wu, J., Balliet, D., & Van Lange, P. A. M. (2016). Reputation, gossip and human cooperation. *Social and Personality Psychology Compass*, 10(6), 350–364. <https://doi.org/10.1111/spc3.12255>
- Zahavi, A., & Zahavi, A. (1997). *The handicap principle: A missing piece of Darwin's puzzle*. New York, NY: Oxford University Press.
- Zitek, E. M., & Phillips, L. T. (2020). Ease and control: The cognitive benefits of hierarchy. *Current Opinion in Psychology*, 33, 131–135.
- Zitek, E. M., & Tiedens, L. Z. (2012). The fluency of social hierarchy: The ease with which hierarchical relationships are seen, remembered, learned, and liked. *Journal of Personality and Social Psychology*, 102(1), 98–115. <https://doi.org/10.1037/a0025345>