



Changing impressions in competence-oriented domains: The primacy of morality endures[☆]

Andrew Luttrell^{a,*}, Simona Sacchi^b, Marco Brambilla^b

^a Ball State University, USA

^b University of Milano-Bicocca, Italy

ARTICLE INFO

Editor: Joris Lammers

Keywords:

Moral character
Impression updating
Moral primacy
Social evaluation
Organizational behavior

ABSTRACT

The Moral Primacy Model proposes that throughout the multiple stages of developing impressions of others, information about the target's morality is more influential than information about their competence or sociability. Would morality continue to exert outsized influence on impressions in the context of a decision for which people view competence as the most important attribute? In three experiments, we used an impression updating paradigm to test how much information about a target's morality versus competence changed perceivers' impressions of a job candidate. Despite several pilot studies in which people said they would prioritize competence over morality when deciding to hire a potential employee, results of the main studies reveal that impressions changed more when people received new information about a target's immorality than about his incompetence. This moral primacy effect held both for global impressions and willingness to hire the target, but direct effects on evaluations of the target as an employee did not consistently emerge. When the new information about the target was positive, we did not reliably observe a moral primacy effect. These findings provide important insight on the generalizability of moral primacy in impression updating.

1. Introduction

Our social landscape is enormously varied. It includes people who are virtuous and people who are unethical, people who are skilled and people who are inept, people who are friendly and people who are distant. So, who do we like and who do we dislike?

Emerging research has emphasized the importance of moral character in social evaluations; above all else, we like moral people more than immoral people (Brambilla, Sacchi, Rusconi, & Goodwin, 2021; Goodwin, 2015). But does this extend to every context and type of judgment? Indeed, despite evidence that people treat morality as essential to their identities (Strohming & Nichols, 2014), other studies show that they often fail to prioritize moral traits when it comes to their own self-improvement (Sun & Goodwin, 2020). Perhaps there are occasions when people similarly fail to prioritize morality in their judgments of others as well. We examine the reliability of the moral primacy effect in a context in which people expect competence to be the primary dimension of judgment: updating one's evaluations of a prospective

employee.

1.1. The structure of social evaluation

Several prominent theories of person perception have proposed that social evaluations can be reduced to two underlying dimensions (see Abele, Ellemers, Fiske, Koch, & Yzerbyt, 2021; Fiske, 2018; Koch, Yzerbyt, Abele, Ellemers, & Fiske, 2021). First, evaluations are informed by a *warmth* dimension that tracks a person's positive or negative intentions and their social skills. Traits that correspond to this dimension include friendliness, sincerity, and trustworthiness. Second, evaluations are also informed by a *competence* dimension that tracks a person's ability to enact their intentions and their general assertiveness. Traits that correspond to this dimension include skill, confidence, and capability.¹

More recently, however, researchers have suggested that “warmth” (as it has been conceptualized in the past) is actually comprised of two distinct dimensions: *sociability* and *morality* (see Brambilla & Leach,

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

* Corresponding author at: Department of Psychological Science, Ball State University, 2000 W University Ave, Muncie, IN 47306, USA.

E-mail address: alluttrell@bsu.edu (A. Luttrell).

¹ Other two-dimension models of social evaluation offer somewhat different accounts of these dimensions, using labels such as “horizontal” versus “vertical” traits or “agency” versus “communion.” See Abele et al. (2021) for an attempt to theoretically integrate across such models.

2014; Goodwin, 2015). Whereas sociability refers to perceived cooperation and social connectedness (e.g., “friendliness,” “likability”), morality refers to perceptions of a target's intentions to pursue virtuous ends (e.g., “honesty,” “trustworthiness”). Although these two dimensions have often been subsumed under the umbrella of “warmth,” they are empirically distinct (Brambilla, Sacchi, Rusconi, Cherubini, & Yzerbyt, 2012; Goodwin, Piazza, & Rozin, 2014; Leach, Ellemers, & Barreto, 2007).

1.2. The primacy of morality in impression development

All three dimensions—morality, sociability, and competence—account for variance in social evaluations, but is one more important than the others? Although some two-dimensional models had proposed a primacy of warmth, such claims received either little empirical attention or came with theoretical caveats (Abele et al., 2021). However, once warmth is further distinguished as sociability and morality, claims about dimensional primacy become much stronger.

Scads of studies have now documented the primacy of moral information in impression development. The Moral Primacy Model (MPM) in particular holds that impressions unfold over multiple stages—gathering information, forming a first impression, and revising the impression—and moral information is theorized to be dominant at each stage (Brambilla et al., 2021). For example, consider how people form an initial impression of someone based on minimal information. Not only do people openly report that moral traits are more important than other traits when evaluating another person, global impressions of a target are influenced by moral traits more than other traits (Goodwin et al., 2014; Wojciszke, Bazinska, & Jaworski, 1998). Similarly, moral traits dominate in shaping impressions of social groups as well (Brambilla et al., 2012; Leach et al., 2007).

But we do not usually learn a bit about a person, then form an impression and maintain that impression indefinitely. A more typical social experience involves learning several things about a person over time, which would prompt a perceiver to update their previous impression in light of new information. Contrary to some expectations that first impressions are hard to change, the right information can indeed change even automatic impressions (Brannon & Gawronski, 2017; Ferguson, Mann, Cone, & Shen, 2019), and these revised impressions persist over time (Cone, Flaharty, & Ferguson, 2021). In a recent set of studies, Brambilla, Carraro, Castelli, and Sacchi (2019) tested whether information about a person's morality, sociability, or competence most strongly compels people to update their first impressions. People formed initial impressions of a target after learning about something he had done and then reported their evaluations of him again after learning about another thing he had done. Across studies, people's impressions of a target changed more when the new information they learned was about his morality than when it was about his competence or sociability.

The particular potency of moral information seems tied to its diagnosticity. A person's moral or immoral actions are taken as an especially clear reflection of who the person is. Indeed, just as diagnosticity is a key determinant of implicit impression change (Cone & Ferguson, 2015; Skowronski & Carlston, 1987), Brambilla et al. (2019) also found that the moral primacy effect in impression updating was driven by people's perception that morality-relevant information was especially informative of the target's intentions. Although these studies did not find a difference between positive and negative information on degree of impression updating (see also Brannon, Sacchi, & Gawronski, 2017), others have argued that *immorality* may particularly influence people's impressions because a person acting immorally is more diagnostic than acting morally (e.g., Reeder & Brewer, 1979).

1.3. Is information about morality always dominant?

The research on impression formation has repeatedly supported the

primacy of moral traits, but might there be situations where people prioritize another dimension before morality? The prior research on various two-dimensional models of impressions suggests this is plausible. Although warmth-related traits had seemed dominant, we now know that such dominance can vanish under certain conditions (see Abele et al., 2021). For example, “warmth” traits influence evaluations less than competence traits when people evaluate relationally close targets or evaluate themselves (Abele & Wojciszke, 2007; Wojciszke & Abele, 2008) or when perceivers have higher power (Cislak & Cichocka, 2018).

Nevertheless, although the general dimension of warmth may sometimes take a backseat, the specific dimension of moral character may not. Whereas competence and sociability perceptions can vary in their influence, Landy, Piazza, and Goodwin (2016) found that moral targets were always preferred over immoral targets. But even this effect can depend on other factors. Melnikoff and Bailey (2018) proposed that salient goals could lead to a preference for *immoral* targets. In one study, for example, they instructed participants that they would play a game in which their job was to choose a spy in a battle of counterintelligence. Sometimes they had to choose a spy that would work for their side; sometimes they chose a spy for their opponent. In a choice between a candidate who was described as relatively moral and one described as relatively immoral, people opted for the more moral spy when choosing for the other side and opted for the more immoral spy when choosing their own. Because an immoral spy would help the player win in this game by being adept at lying and deception, that spy was preferred on both explicit and implicit measures of overall liking. These studies provide powerful evidence that the effects of moral character on impression formation can be moderated.

Whereas Melnikoff and Bailey's research focused on salient goals that make immorality desirable, other situations might also undermine moral dominance by prioritizing *other* goals. In the present research, we examined contexts in which *competence* would be expected to dominate. That is, people may sometimes be oriented to select the most competent person at any expense, which would prioritize competence information in impression formation. Specifically, we looked to a hiring context. A desire to maximize productivity and profits would presumably render “competence” a primary dimension of importance, and prior writing in organizational psychology suggests as much. Indeed, some studies have shown stronger effects of perceived competence (versus perceived warmth) in hiring judgments (Cuddy, Fiske, & Glick, 2004; Krings, Sczesny, & Kluge, 2011) and interview outcomes (Amaral, Powell, & Ho, 2019; Barrick, Swider, & Stewart, 2010; Graves & Powell, 1988; Howard & Ferris, 1996), particularly in the absence of motivations to discount otherwise qualified candidates (Phelan, Moss-Racusin, & Rudman, 2008).

Because hiring decisions prototypically emphasize a desire for competent candidates, we were able to test two competing hypotheses about impression updating. On one hand, if the prior evidence for moral primacy arose because those studies did not make any specific dimension salient in the judgment context, we would expect this hiring context to provoke a *competence primacy* effect. That is, if the relative priorities of morality and competence are malleable, then we should not expect a moral primacy effect in every judgment context.

Abele et al.'s (2021) theoretical account of when a given dimension is prioritized in social evaluation also predicts that moral primacy can be moderated. They identify two mechanisms by which a dimension's priority could change. First, the “subjective weight” of a trait dimension can vary; various aspects of the perceiver and the judgment can shape the perceived importance assigned to that type of information. Second, the “pragmatic diagnosticity” of a trait dimension can also vary; different types of information may appear to be differently reliable and valid as a basis for judgments in one context versus another.

Both mechanisms suggest that information about a target's morality may be less dominant in a hiring context. Indeed, our own pilot testing (which we will review shortly) shows that people view competence-

related information as more important in this environment, suggesting greater subjective weight on this dimension over others. Further, [Brambilla, Rusconi, Sacchi and Cherubini \(2011, Study 1\)](#) showed the malleability of pragmatic diagnosticity. They asked people to evaluate the relevance of various traits to accomplishing different goals. These goals included forming an overall impression of a person as well as rendering judgments specific to competence, morality, and sociability. A main effect across goal conditions showed that information about a person's morality dominated in terms of its perceived relevance. However, results also supported an interaction such that people were sensitive to the information that was most goal-relevant. That is, they indicated that information about a person's competence was more relevant than information about his morality when they were given a competence-relevant goal (i.e., deciding to employ someone for a research program).

On the other hand, if morality is truly primary in interpersonal impressions, we would expect it to still exert disproportionate influence on judgments even in competence-oriented contexts. Indeed, in at least some ways, people continue to prioritize morality over competence when making first impressions in organizational settings. For instance, when a group of employees evaluated a prospective boss, information about the potential manager's morality had a larger influence on the employees' impressions than information about his competence ([Pagliaro, Brambilla, Sacchi, D'Angelo, & Ellemers, 2013](#)). Similarly, people's interest in working for an organization were more strongly influenced by information about that organization's morality than information about its competence ([van Prooijen & Ellemers, 2015](#)). Although these findings highlight how moral primacy can occur within an organization context, they focus on the climate preferred by employees. It remains to be seen whether competence continues to be treated with secondary importance in contexts that more directly highlight the value of recruiting competent employees.

2. The present research

We conducted three experiments to test our competing hypotheses about whether morality-related information would continue to exert more influence than competence-related information in a judgment context for which people perceive competence concerns to dominate (i.e., evaluating a job candidate). Specifically, we examined these hypotheses in an impression updating paradigm. That is, although the impression formation work has suggested that moral primacy can extend to organizational settings ([Pagliaro et al., 2013](#); [van Prooijen & Ellemers, 2015](#)), perhaps once an initial impression is established, people attend more critically to goal-relevant information. We also saw impression updating as a particularly important phase of impression formation, which has been understudied ([Brambilla et al., 2021](#)). In organizational contexts especially, people are likely to constantly be updating their impressions of others, given the cooperative nature of the workplace and the importance of thoroughly considering a candidate before hiring them.

In all studies, participants formed an initial impression of a hypothetical job candidate based on limited information. These participants then received new information either about the target's moral character or about his competence. We tested whether information about a person's morality still has a disproportionate influence on how much perceivers update their evaluations even in the context of judging the target's suitability as an employee.

Participants in Experiment 1 formed initial impressions of a job candidate based on information about his sociability and then received new information either about his morality or his competence. Experiment 2 replicated the design of Experiment 1 and considered whether the target's status as a supervisor or subordinate would moderate people's use of morality-related versus competence-related information in revising their impressions. In a final pre-registered study, Experiment 3 tested the relative effects of morality-related and competence-related

information on first impressions of a job candidate and on changes to those evaluations. This study also directly manipulated whether the target was applying to a non-profit versus for-profit organization to further highlight goals to recruit competent employees in the latter condition. We conclude with an internal meta-analysis to assess the consistency of our findings across studies and also across the kinds of judgments participants made (i.e., global evaluations versus evaluations specific to the target's employment suitability). Across all studies, we deliberately chose stimuli that would reliably convey information relevant to either morality or competence that was also equivalent in evaluative extremity.

In sum, our work contributes to the literature on the key role of morality in social perception in several ways. First, the present work aims to extend prior research evidence testing whether morality has a leading role in impression updating and whether the social context in which the evaluation is made influences the dimension relevance. Indeed, a long tradition of research has shown that the same information might be treated differently at different stages of impression development. In fact, the raw information available about a person is key to forming an initial impression; however, at the impression updating stage, additional motivational factors and more elaborate cognitive processes can intervene ([Brambilla et al., 2021](#); [Trobe & Liberman, 1996](#)). Thus, showing that morality dominates over sociability and competence in shaping first impressions does not warrant the conclusion that morality plays a leading overall role in the development of impressions over time. Thus, our work complements prior work showing the importance of morality in shaping first impressions by testing its relevance when people are asked to revise their first evaluations.

Further, our analysis is useful for better capturing how morality works in organizational contexts. Indeed, we are not aware of any experimental work that has tested how morality influences impression updating in organizational contexts. Rather, the prior studies on organizational morality have only considered first impressions. Such studies have also mainly considered the climate preferred by employees, leaving unexplored the role of morality when recruiting individual employees. These prior studies have also never tested whether important factors related to the organization context (i.e., the target's potential status and the profit versus nonprofit status of the organization) would moderate the primacy of morality. Our studies aim to fill these gaps.

See the online supplement and the project's page on the Open Science Framework (OSF) for all materials across studies, including full stimuli and question wording. We report all measures, manipulations, and exclusions included in the reported experiments. All sample sizes were determined prior to data analysis.

3. Pilot studies

We conducted a few brief pilot surveys to assess lay beliefs about which characteristics matter most when making hiring decisions (see the online supplement for full details). Each survey listed "three general qualities" of a potential job applicant: "morality (e.g., honesty, trustworthiness)," "competence (e.g., intelligence, capability)," and "sociability (e.g., friendliness, cooperativeness)." The options appeared in a random order, and participants were asked: "which of these do you think is the most important quality a company should look for in a potential employee?" We fielded this online survey in three samples corresponding to the samples we would use in our main experiments: respondents in the United States ($N = 250$), respondents in Italy ($N = 250$), and a sample without geographic restrictions ($N = 500$). In all three samples, "competence" was chosen most often (59% - 71%), followed by "morality" (20% - 31%). As anticipated, "sociability" was rarely chosen (7-9%). These pilot studies show that in the context of evaluating job applicants, people indeed tend to think that competence is of primary importance.

4. Experiment 1

The goal of Experiment 1 was to compare information about a person's morality versus competence when people update their impressions of a job candidate. Specifically, all participants formed an initial impression of a target based on information about his sociability, a dimension that people often see as less relevant to evaluating people in a hiring context. Next, participants were given new information about that person's morality or competence. We assessed whether one type of information or the other led to bigger changes in impressions of the target, evaluations of him as a potential employee, and probability of hiring him.

If morally relevant information is still dominant in this hiring context, we would expect a main effect of the type of information presented following first impressions whereby receiving new information about the candidate's morality would change evaluations of him more than receiving new information about his competence. Indeed, moral primacy would hold regardless of whether the new information was positive or negative; however, previous research has shown that moral information can be particularly influential when it is negative (Reeder & Brewer, 1979; Skowronski & Carlston, 1987; for a review, Rusconi, Sacchi, Brambilla, Capellini, & Cherubini, 2020). Thus, it was also possible that we would find an interaction whereby moral primacy is even more pronounced when the new information is negative.

4.1. Method

4.1.1. Participants

To recruit participants with some familiarity with workplace practices, an open invitation to participate in a study on social judgment in the workplace was disseminated on LinkedIn. Responses were closed eight weeks after the link was posted. In total, 134 employees working in different types of organizations participated in the survey ($M_{\text{age}} = 42.60$; $SD = 12.67$; 47.8% female). A sensitivity analysis shows that this sample provides 80% power to detect medium-sized ($d = 0.49$) overall differences between morality-related and competence-related information.² The survey was presented in Italian.

4.1.2. Procedure

Participants were asked to imagine being a member of a company HR department who must select a new employee. They were asked to evaluate a candidate named Fabio. The first information they received about him (Time 1; T1) was in the form of a recommendation letter ostensibly written by his previous employer, which spoke either favorably or unfavorably of his sociability. Participants indicated how much they thought the letter discussed the candidate's sociability, morality, and competence, reported their overall first impressions of Fabio, evaluated him specifically as a potential employee, and indicated how likely they would be to hire him.

After completing this first section of the questionnaire, participants were given additional information about Fabio (Time 2; T2) that varied in its valence (negative vs. positive) and type (morality vs. competence). After receiving this additional information, participants again reported their general impressions of the candidate, their evaluations of him as an

employee, and their probability of hiring him, using the same measures as before.

In sum, the experiment employed a 2 (T1 Information Valence: Positive vs. Negative) \times 2 (T2 Information Valence: Positive vs. Negative) \times 2 (T2 Information Type: Morality vs. Competence) design with all the factors varying between participants.

4.1.3. Independent variables

The fabricated recommendation letter was presented to instill relatively positive or negative first impressions of the target based primarily on his sociability. Participants saw a scanned image of a letter, much of which had been redacted to focus attention on the information about Fabio's sociability while maintaining the realism of this task. In the positive condition, the readable portion of the letter described Fabio as a person able to work in a team and likely to develop a friendship with his colleagues. In the negative condition, he was described as a person who prefers individual work and is likely to build only formal relationships in the workplace.

After forming their initial impressions based on the target's sociability, participants saw additional information about the target in the form of a list of qualities or behaviors that were true of him. For example, one of the negative moral qualities was that Fabio had lied to his boss whereas one of the negative competence qualities was that he rarely completes a job independently. One of the positive moral qualities was that he takes responsibility for errors, whereas a positive competence quality was that he is able to complete assigned work independently. A pre-test confirmed that these pieces of information were balanced for favorability, ruling out confounds pertaining to the extremity of the morality-relevant versus competence-relevant information.³

4.1.4. Dependent measures

Perceived Type of T1 Information. After reading the recommendation letter, participants were asked how much the information they read was about Fabio's (im)morality, (in)competence, and (un)sociability, regardless of whether it was positive or negative. Responses were provided on three 7-points scales anchored at "not at all" and "very much."

General Impressions. After reading both the initial recommendation letter and the follow-up information, participants responded to six questions about their overall impressions of Fabio (see Brambilla, Sacchi, Pagliaro, & Ellemers, 2013). These included items such as "I'd like to meet Fabio" and "If I met Fabio, I would try to keep him at a distance" (reversed). Responses were provided on a 7-point scale anchored at "not at all" and "very much." These items showed good internal reliability at both T1 ($\alpha = 0.77$) and T2 ($\alpha = 0.86$) and were thus averaged at each time point to form indices of overall impressions.

Evaluations as an Employee. To measure people's evaluations of Fabio more directly in the context of a hiring decision, participants were told to imagine they had to decide whether or not to hire Fabio and respond to five additional questions. These included statements such as "Fabio would be an excellent resource for the company" and "workers like Fabio create more problems than opportunities" (reversed).

² Although we ultimately analyze the data with factorial ANOVAs, we opt for the simplicity of presenting power based on simple group differences (Cohen's d). If we instead conduct a sensitivity analysis for the main effect of information type in the context of the other independent variables, we similarly find that this sample provides 80% power to detect a medium-sized effect ($\eta_p^2 = 0.06$). Further, we note that across studies, we test and interpret higher-order interactions between independent variables, which are not accounted for in our focal sensitivity analyses; however, across studies and tested models, sample sizes provide 80% power to detect interactions greater than $\eta_p^2 = 0.059$ in Experiment 1, $\eta_p^2 = 0.013$ in Experiment 2, and $\eta_p^2 = 0.014$ in Experiment 3.

³ We conducted a pilot study to select a set of morality and competence behaviors balanced for valence. Fifty-one students were approached at random in university libraries. Participants were exposed to three negative and three positive behaviors pertaining to morality or competence. Participants were asked to rate the valence of such behaviors on a scale ranging from 1 (very negative) to 5 (very positive). To verify that morality and competence-related behaviors were properly balanced, we conducted a 2 (dimension: morality vs competence) \times 2 (valence: positive vs. negative) ANOVA on the perceived valence. As intended, the ANOVA showed a significant effect of valence, $F(1, 48) = 483.19, p < .001, \eta_p^2 = 0.91$: positive behaviors were perceived as more positive ($M = 4.43, SD = 0.52$) than negative behaviors ($M = 1.85, SD = 0.49$). Importantly, there was no interaction between dimension and valence, $F(1, 48) = 2.09, p = .15, \eta_p^2 = 0.04$.

Responses were again provided on a scale from 1 (not at all) to 7 (very much) and showed good internal reliability at both T1 ($\alpha = 0.83$) and T2 ($\alpha = 0.89$), so they were averaged to form composite indices at each time point.

Hiring Probability. Finally, at each time point, participants were asked how likely they would be to hire Fabio for a job at their company. Responses were given on a scale from 0% to 100% with response options in increments of 10%. Although responses were given on an 11-point survey scale, we transformed responses to conform to the 0–100 scale so summary statistics align with the units of the scale labels.

4.2. Results

4.2.1. Statistical approach

For this and the two studies that follow, unless otherwise noted, we conducted between-subjects ANOVAs (Type 3) to test for main effects and interactions, and we decomposed significant interactions using regression models to center a given variable on the condition(s) of interest, effects coding the other independent variables. Across studies, we adopt the same effects coding scheme for valence (−1: negative, +1: positive) and information type (−1: competence, +1: morality). Full analysis scripts are available on the project's OSF page.

4.2.2. Preliminary analyses on first impression

To verify that the first description was perceived as related to sociability, a within-participants ANOVA was conducted, comparing ratings of the information's relevance to sociability, morality, and competence. Indeed, the omnibus ANOVA was significant, $F(2, 264) = 208.32, p < .001, \eta_p^2 = 0.61$, and specifically, Tukey post-hoc comparisons show that the information in the recommendation letter was perceived as more related to sociability ($M = 5.50, SD = 1.50$) than to morality ($M = 2.50, SD = 1.89$), $p < .001$, and competence ($M = 2.36, SD = 1.58$), $p < .001$.

Then, a series of *t*-tests compared general impressions, evaluations as an employee, and hiring probability between positive and negative T1 conditions. Indeed, general impressions towards the candidate were more favorable in the positive condition ($M = 5.09, SD = 0.90$) than in the negative condition ($M = 4.51, SD = 0.95$), $t(132) = 3.63, p < .001, d = 0.63, 95\% \text{ CI}: [0.28, 0.98]$. The same pattern arose on evaluation as an employee ($M_{pos} = 5.17, SD_{pos} = 1.00; M_{neg} = 3.90, SD_{neg} = 1.03$), $t(132) = 7.28, p < .001, d = 1.26, 95\% \text{ CI}: [0.89, 1.64]$, and hiring probability ($M_{pos} = 60.55, SD_{pos} = 21.53; M_{neg} = 44.92, SD_{neg} = 21.03$), $t(132) = 4.23, p < .001, d = 0.73, 95\% \text{ CI}: [0.38, 1.09]$.

4.2.3. Impression updating

To examine our primary hypotheses, we computed indices of impression updating by subtracting T1 responses from T2 responses for each of our three dependent variables. Following prior research on impression updating and to more directly assess effects on magnitude of change, these scores were considered in their absolute values. Thus, higher scores on an index of updating reflects greater change after being exposed to the new information.

We conducted 2 (T1 Valence) \times 2 (T2 Valence) \times 2 (T2 Type) between-participants ANOVAs on each updating index. We report all statistically significant effects.

First, for changes in general impressions, the ANOVA revealed a main effect of T2 information type $F(1, 126) = 5.52, p = .02, \eta_p^2 = 0.04$. Consistent with the moral primacy hypothesis, general impressions of the candidate changed more when the second set of information was related to morality ($M = 1.13, SD = 1.07$) rather than to competence ($M = 0.72, SD = 0.69$). The analysis also yielded a main effect of T2 information valence, $F(1, 126) = 63.18, p < .001, \eta_p^2 = 0.33$. In line with the literature on the negativity bias (Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001; Skowronski & Carlston, 1987), the negative information ($M = 1.35, SD = 0.97$) changed impressions more than the positive information ($M = 0.38, SD = 0.34$). Moreover, there was a

significant interaction between T2 information type and valence, $F(1, 126) = 8.49, p = .004, \eta_p^2 = 0.06$ (see Fig. 1). When T2 information was positive, the impact of morality- and competence-related information on impression change was analogous, $B = -0.04, t(126) = -0.38, p = .70, 95\% \text{ CI}: [-0.22, 0.15]$. In contrast, it was when T2 information was negative that morality-related information led to greater change than competence-related information, $B = 0.33, t(126) = 3.91, p < .001, 95\% \text{ CI}: [0.16, 0.49]$. The ANOVA did not yield other significant effects, $ps > 0.09$.

The ANOVAs on changes in evaluations of the candidate as an employee and hiring probability did not show analogous main effects of T2 information type or an interaction effect between T2 type and valence, $ps > 0.50$. Both models, however, showed the same main effects of T2 valence as reported in the previous analysis, $ps < 0.001$.

4.3. Discussion

Taken together, these findings show that even in the context of considering someone for a job, which people typically say is a context where competence should be prioritized, information about a person's morality still had a bigger effect on changes to general impressions than information about competence. The results, however, did not support direct effects of information type on judgments that were more closely related to the hiring context itself, namely evaluations of the target as a potential employee or perceived likelihood that he would be hired.

Finally, in line with prior studies showing a robust negativity bias in the moral domain (see Rusconi et al., 2020), we only found evidence of moral dominance in impression updating when the new information was negative. That is, new information suggesting that the target is *immoral* carried greater weight than new information suggesting he was *incompetent*. When the new information was positive, however, we found no evidence for one type of information being more influential than the other on impression updating.

5. Experiment 2

The results of Experiment 1 generally supported the primacy of moral information in impression updating even in a context prototypically defined by competence concerns. Our findings came with two caveats, however: the effect of information type only emerged for negative information and only occurred directly on general impressions but not evaluations specific to the hiring decision. The relatively small sample size in Experiment 1 may have provided insufficient power to detect small effects for positive information or on domain-specific judgments that are nevertheless reliable. Therefore, Experiment 2 was conducted to assess the replicability of findings with a larger sample.

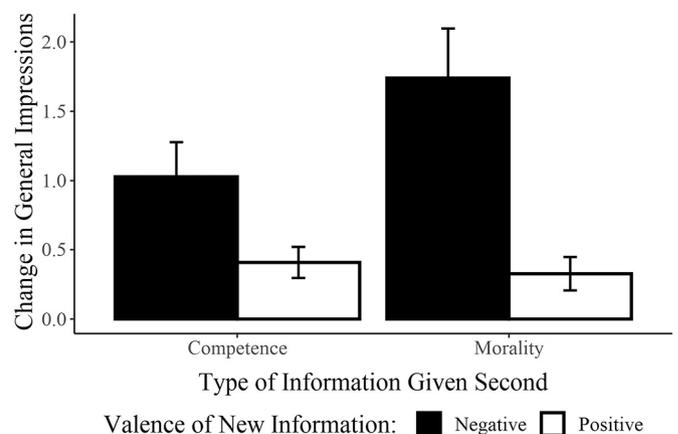


Fig. 1. An interaction between the valence and type of information at T2 on changes in people's general impressions of the job candidate (Study 1).

Experiment 2 also aimed to extend the results of the previous study by testing the effect of the target's social status. That is, participants again formed a first impression of an imaginary future colleague based on information pertaining to his sociability, but this target was presented as either striving to be the participant's subordinate or supervisor. Prior research has identified a relationship between status and perceived competence (e.g., Caprariello, Cuddy, & Fiske, 2009; Fiske, Cuddy, Glick, & Xu, 2002), so we tested the possibility that people would weight competence-related information more heavily in their impressions of higher (vs. lower) status targets. Nevertheless, other evidence highlights that status more directly tracks agency and not competence per se (Carrier, Louvet, Chauvin, & Rohmer, 2014), so we included this manipulation primarily for exploratory purposes.

The remainder of the study was otherwise nearly identical to Experiment 1.

5.1. Method

5.1.1. Participants

Participants were recruited as in Experiment 1. We advertised the study online and closed responses after three months. In total, 630 employees ($M_{\text{age}} = 30.43$; $SD = 8.37$; 42.7% female) participated in the survey. A sensitivity analysis shows that this sample size provided 80% power to detect rather small overall differences between receiving new morality-related versus competence-related information ($d = 0.22$). Given that we observed a moral primacy effect only for new negative information, the half of this study's sample ($N = 315$) that would see negative T2 information still provides 80% power to detect small-medium effects of information type ($d = 0.32$).

To assess how well participants could realistically engage with the task, we also asked about their employment. Most of them (84.9%) were employed in a private company, and the other participants worked in different types of organizations (e.g., public company [6.4%]; non-governmental organization [2.2%]; third sector [5.1%]). The survey was presented in Italian.

5.1.2. Procedure

The experimental procedure was nearly identical to the procedure that was employed in Experiment 1. The key difference was that we additionally manipulated the job candidate's status at T1. Thus, the experiment employed a 2 (Target Status: Lower vs. Higher) \times 2 (T1 Information Valence: Positive vs. Negative) \times 2 (T2 Information Valence: Positive vs. Negative) \times 2 (T2 Information Type: Morality vs. Competence) design with all the factors varying between participants.⁴

5.1.3. Independent variables

As before, participants were asked to evaluate a new colleague, Fabio, likely to be hired in their company, but in the *lower-status* condition, this new colleague was described as striving for a position subordinate to the participant's own hypothetical position (i.e., the participant would be Fabio's supervisor), and in the *higher-status* condition, he was striving for a position higher than the participant's (i.e., Fabio would be the participant's supervisor). Then participants saw the same fake recommendation letter displaying either positive or negative information related to his sociability (T1) and later saw additional information that varied in valence and relevance to morality versus competence (T2). These materials were the same as Experiment 1.

5.1.4. Dependent measures

Perceived Type of T1 Information. Perceptions of the initial

⁴ Some participants in this study did not provide responses for at least one item that measured our dependent variables. We assume these few cases are random and for each analysis we present the results for which the relevant data are available.

information was measured as in Experiment 1.

General Impressions. Overall impressions of the candidate were measured with the same items as in Experiment 1; internal reliability was again good at both time points ($\alpha > 0.81$).

Evaluations as an Employee. To assess evaluations of the target as a potential employee more fully, we expanded this to a 12-item scale. New items included, for example, "I believe that work performance could improve thanks to the collaboration with Fabio." Internal reliability was good at both T1 ($\alpha = 0.86$) and T2 ($\alpha = 0.93$).

Hiring Probability. Finally, rather than ask people to imagine that they oversaw hiring decisions and to estimate the probability that they would hire Fabio (as we did in Experiment 1), we instead asked participants to estimate the probability that Fabio would be hired if he was being considered at the company the participants themselves worked for. Response scales were the same as in Experiment 1.

5.2. Results

5.2.1. Preliminary analyses on first impression

As in Experiment 1, results of a within-participants ANOVA showed that perceptions of the initial information differed in the qualities it communicated, $F(2, 1236) = 1195.12$, $p < .001$, $\eta_p^2 = 0.66$. Tukey post-hoc comparisons show that the information was perceived as more related to sociability ($M = 5.79$, $SD = 1.45$) than to morality ($M = 2.36$, $SD = 1.68$), $p < .001$, and competence ($M = 2.38$, $SD = 1.60$), $p < .001$.

Then, a series of 2 (Status) \times 2 (T1 Valence) between-participants ANOVAs were conducted on T1 ratings. First, results for general impressions of the target showed that they were more favorable when the initial information was positive ($M = 5.40$, $SD = 0.81$) than when it was negative ($M = 4.81$, $SD = 0.98$), $F(1, 626) = 66.49$, $p < .001$, $\eta_p^2 = 0.10$. The analysis yielded neither a main effect of status, $F(1, 626) = 0.49$, $p = .48$, $\eta_p^2 < 0.001$, nor an interaction effect, $F(1, 626) = 2.53$, $p = .11$, $\eta_p^2 = 0.004$.

Second, the ANOVA on evaluations of the target as an employee showed a main effect of T1 information valence ($M_{\text{pos}} = 5.15$, $SD_{\text{pos}} = 0.80$; $M_{\text{neg}} = 4.38$, $SD_{\text{neg}} = 0.83$), $F(1, 626) = 139.80$, $p < .001$, $\eta_p^2 = 0.18$, and a main effect of status ($M_{\text{high}} = 4.67$, $SD_{\text{high}} = 0.92$; $M_{\text{low}} = 4.84$, $SD_{\text{low}} = 0.88$), $F(1, 626) = 5.17$, $p = .02$, $\eta_p^2 = 0.01$. The T1 valence \times status interaction was not significant, $F(1, 626) = 1.68$, $p = .20$, $\eta_p^2 = 0.003$.

Finally, the ANOVA on hiring probability also showed a main effect of valence: participants thought it was more likely that Fabio would be hired when they saw positive ($M = 44.72$, $SD = 12.31$) versus negative ($M = 31.65$, $SD = 12.31$) information at T1, $F(1, 619) = 204.87$, $p < .001$, $\eta_p^2 = 0.25$. There was also a significant valence \times status interaction, $F(1, 626) = 8.19$, $p = .004$, $\eta_p^2 = 0.01$, such that the effect of valence was greater for the higher status target, $B = 7.82$, $t(619) = 12.20$, $p < .001$, 95% CI: [6.56, 9.08], compared to the lower status target, $B = 5.22$, $t(619) = 8.06$, $p < .001$, 95% CI: [3.95, 6.49]. The main effect of status was not significant, $F(1, 626) = 0.04$, $p = .84$, $\eta_p^2 < 0.001$.

5.2.2. Impression updating

Updating indices were computed as in Experiment 1. A 2 (Status) \times 2 (T1 Valence) \times 2 (T2 Valence) \times 2 (T2 Type) between-participants ANOVA was conducted on the indices of changes in general impressions, employee evaluations, and hiring probability.

Since all the analyses showed neither a significant main effect of status nor two-, three-, or four-way interaction effects involving this variable ($ps > 0.05$), data have been collapsed over this factor. Therefore, the following results refer to the 2 (T1 Valence) \times 2 (T2 Valence) \times 2 (T2 Type) ANOVAs.

On the updating index of general impressions, in line with Experiment 1, the ANOVA revealed a main effect of the type of information presented at T2, $F(1, 618) = 7.17$, $p = .008$, $\eta_p^2 = 0.01$. Impressions of the target changed more when the new information was related to morality ($M = 0.98$, $SD = 0.97$) than when it was related to competence

($M = 0.80$, $SD = 0.73$). The analysis also yielded a main effect of T2 valence, $F(1, 618) = 153.89$, $p < .001$, $\eta_p^2 = 0.20$, such that impressions changed more when the new information was negative ($M = 1.27$, $SD = 0.97$) than when it was positive ($M = 0.52$, $SD = 0.55$). Moreover, the T2 information valence \times type interaction was significant, $F(1, 618) = 17.31$, $p < .001$, $\eta_p^2 = 0.03$ (see Fig. 2A). When information at T2 was positive, impression updating did not depend on whether that information was related to morality or competence, $B = -0.04$, $t(618) = -1.05$, $p = .29$, 95% CI: $[-0.13, 0.04]$. In contrast, when the information at T2 was negative, impressions changed more when that information was about the target's moral character ($M = 1.48$, $SD = 1.04$) than when it was about his competence ($M = 1.05$, $SD = 0.80$), $B = 0.21$, $t(618) = 4.83$, $p < .001$, 95% CI: $[0.12, 0.29]$.

Unlike Experiment 1, the ANOVA also supported an unexpected interaction between T1 and T2 information valence, $F(1, 618) = 14.00$, $p < .001$, $\eta_p^2 = 0.02$. Positive information at T2 had a similar impact on impressions regardless of whether the prior impression was based on a positive ($M = 0.46$, $SD = 0.46$) or negative description ($M = 0.57$, $SD = 0.62$), $B = -0.06$, $t(618) = -1.31$, $p = .19$, 95% CI: $[-0.14, 0.03]$. In contrast, additional negative information had a greater impact when the prior impression was based on a positive ($M = 1.44$, $SD = 1.01$) rather than a negative description ($M = 1.09$, $SD = 0.85$), $B = 0.17$, $t(618) = 3.97$, $p < .001$, 95% CI $[0.09, 0.25]$. The result suggests that the greater influence of negative information is particularly strong when it is challenging a prior impression rather than supporting it, in some ways consistent with Rothbart and Park's (1986) finding that "favorable traits are hard to acquire but easy to lose" (p. 137). Nevertheless, T1 valence did not moderate the overall effect of T2 information type ($p = .29$) nor the T2 valence \times T2 type interaction ($p = .18$).

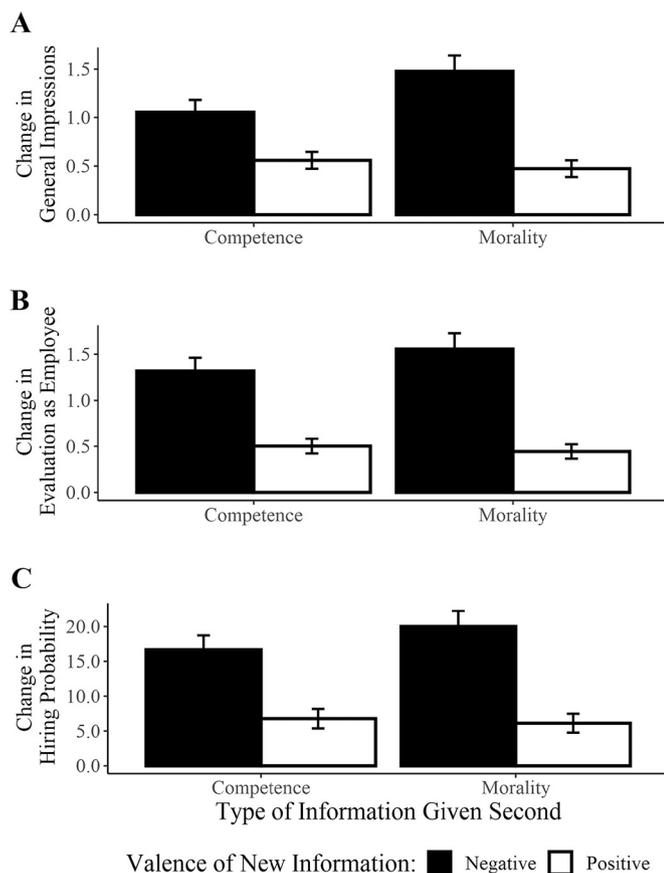


Fig. 2. Interactions between the valence and type of information at T2 on (A) changes in people's general impressions of the job candidate, (B) changes in people's evaluation of the target as an employee, and (C) changes in people's judgments of the candidate's likelihood of being hired (Study 2).

In this experiment, the same T2 valence \times type interaction emerged for both changes to evaluations of the target as a potential employee, $F(1, 619) = 5.46$, $p = .02$, $\eta_p^2 = 0.01$ (Fig. 2B), and changes to perceived hiring probability, $F(1, 609) = 4.42$, $p = .04$, $\eta_p^2 = 0.007$ (Fig. 2C). The T1 valence \times T2 valence interaction also emerged on both of the other indices of impression updating ($ps < 0.001$), but once again T1 valence did not moderate any effects of T2 information type ($ps > 0.62$).

5.3. Discussion

The results of Experiment 2 largely reinforce those of Experiment 1. Once again, we found that information about a person's moral character has a bigger influence on changes to one's overall impression of that person than information about the person's competence, particularly when that information is negative. Notably, this study also found evidence that negative, morally relevant information is also particularly influential for changes in judgments more directly tied to the hiring decision context, namely evaluations of the target as a potential employee and the perceived likelihood that the person would be hired. Finally, results of this study did not provide evidence that these effects depended on the relative status that the job candidate is striving to achieve.

6. Experiment 3

Thus far, two studies have supported the moral primacy hypothesis in impression updating even in a context where a desire for competence would be particularly salient. Although the effects have been especially robust for overall impressions of a person, the results of Experiment 2 also showed effects on judgments more directly relevant to the hiring context.

In addition to further examining these effects in another national context (the United States), Experiment 3 had two primary aims. First, we manipulated the display of morality-related versus competence-related information at T1 and T2. Therefore, we were able to test for moral primacy in this context both at the stage of forming first impressions and during impression updating. Support for moral primacy in first impressions would take the form of an interaction between the valence and type of information presented at T1; impressions should correspond to the valence of information more when that information is about morality versus competence. By contrast, support for moral primacy in impression updating would take the form of a main effect of T2 information type on absolute change in impressions; impressions should change more when the new information is about morality versus competence. However, as we have seen in the previous two studies, these dynamics may be particularly true for negative T2 information.

Second, we manipulated the type of organization the target was apparently applying to. Prior research has shown that people often stereotype for-profit companies as higher in competence and non-profit companies as higher in warmth (Aaker, Vohs, & Mogilner, 2010). Therefore, when the organization's non-profit versus for-profit status is clarified, we might expect to see people prioritize morality-related information when judging an applicant's suitability for a non-profit organization but instead prioritize competence-related information when the organization is clearly focused on profit.

This study was pre-registered on OSF (<https://osf.io/pe7rc>).

6.1. Method

6.1.1. Participants

We set out to recruit 600 participants from the United States using the online survey platform, *Mechanical Turk* (via *CloudResearch*). Participants were paid \$0.50 to complete a brief survey. In the end, $N = 598$

provided complete data ($M_{age} = 36.30$, $SD = 10.86$; 41.3% female). Of these, 41 people failed a simple attention check in which they were asked: "Using the box below, please type the first letter of this sentence."⁵ Despite setting CloudResearch to block duplicate IP addresses from participating, an additional three cases were excluded for having an IP address that matched that of an earlier participant. Following exclusions, this study's sample size was $N = 554$. Using the same sensitivity analysis as the previous studies, this sample provided 80% power to detect a relatively small moral primacy effect across the sample ($d = 0.24$) and within the half of the sample who would receive negative information at T2 ($d = 0.34$). The survey was presented in English.

6.1.2. Procedure

Participants were told that they would participate in a study to examine first impressions based on minimal information about a person. They were asked to imagine that they were a recruiter for either a non-profit or for-profit organization, and they were considering a particular candidate for a job opening. To reinforce the manipulation, participants read that the "primary goal of your organization is to maximize [community impact/annual profits]."

Next, they saw a photo of a White adult man facing forward and read the following instructions: "You learn about this individual by interacting with him directly, talking to his references, and observing him at recruitment events. Through these channels, you learned about his past behavior. To see something you learned about this individual, please continue to the next screen."

On the following screen, the survey displayed the man's headshot again and randomly presented a negative or positive behavior that conveyed the applicant's morality or competence. To focus attention on the information, the survey did not allow participants to advance until four seconds had passed. Based only on this limited information, participants rated their impressions of the applicant and their willingness to hire him.

Participants were then invited to continue to the next screen to "see something else you learned about this individual." The behavior they saw on the following page was opposite in valence and type to the behavior they saw previously. For example, if the participant initially saw something negative about the person's morality, they would then see something positive about his competence. Once again, the option to advance to the next screen did not appear until four seconds had passed. After reading about this second behavior, participants again rated their impressions and indicated their willingness to hire the applicant using the same scales as before. All participants then responded to an attention check, reported on several demographic variables, and concluded the study.

In sum, the experiment employed a 2 (Organization: Non-Profit vs. For-Profit) \times 2 (T1 Information Valence: Positive vs. Negative) \times 2 (T1 Information Type: Morality vs. Competence) design with all the factors varying between participants. Note that we will also present analyses in which T2 valence and T2 type are the independent variables; however, these are just the exact opposites of their T1 counterparts.

6.1.3. Stimuli

The first piece of information participants learned about the applicant (T1) was drawn randomly from 16 statements that varied in valence and relevance to morality versus competence. For example, one of the negative moral behaviors was that this person had lied to his parents whereas a negative competence behavior was that he had flunked an

⁵ We initially pre-registered a different attention check that we ultimately did not include in the survey. We decided that the attention check we included in the pre-registration was rather complex for such a brief survey and may serve to over-exclude participants who were otherwise attending to the brief stimuli of interest. Instead, we opted for this simpler attention check, and we enforced a clear criterion for exclusion.

aptitude test at work. A positive moral behavior was that he had given money to a charity, and a positive competence behavior was that he graduated at the top of his class in college. These statements were selected based on pilot testing which included an expanded set of behaviors. We deliberately chose stimuli that reliably conveyed information uniquely about a person's competence or morality and did so in a way that implied similarly extreme positivity and similarly extreme negativity.⁶

The second piece of information participants saw (T2) was drawn randomly from the four behaviors that were opposite the valence and type of information they saw initially.

6.1.4. Measures

General Impressions. After the first and second pieces of information, participants reported their "overall feelings about the job candidate" on three semantic differential items anchored at "dislike," "unfavorable," and "negative" on the low end (-4) and "like," "favorable," and "positive" at the high end ($+4$), respectively. The three items showed good internal reliability at both T1 ($\alpha = 0.98$) and T2 ($\alpha = 0.97$) and were thus averaged at each time point to form indices of overall impressions.

Hiring Probability. Following each impression measure were three items assessing participants' assessment of how likely the applicant was to be hired. Specifically, they indicated how likely it was that they "would interview the candidate," they "would personally hire the candidate," and "the candidate would be hired for the job." Responses were provided on five-point scales anchored at "not at all likely" and "extremely likely." The three items showed good internal reliability at both T1 ($\alpha = 0.94$) and T2 ($\alpha = 0.93$) and were thus averaged at each time point to form indices of hiring probability.

6.2. Results

6.2.1. Deviations from pre-registered analysis plan

To match the analysis method for first impressions, our pre-registered analysis plan for impression updating focused on directional impression change (i.e., the raw difference between T1 and T2 responses) as the outcome variable. We opted instead to focus on absolute

⁶ We collected additional pilot data to ensure that the stimuli varied in their perceived valence and relevance to morality and competence. Indeed they did. However, we found a significant behavior valence \times type interaction on perceived valence. To probe this interaction, we focused on the *extremity* of each behavior's perceived valence to account for the possibility that more extreme information is the proximal predictor of impression updating. The analysis yielded a main effect of behavior valence: positive behaviors were judged to be more extreme ($M = 2.08$, $SD = 0.69$) than negative behaviors ($M = 1.45$, $SD = 0.73$), $F(1, 99) = 181.53$, $p < .001$, which works against the negativity bias we observe across studies. However, we also found a behavior valence \times type interaction on perceived valence extremity, $F(1, 99) = 14.34$, $p < .001$. The interaction is driven by the fact that negative moral behaviors seemed more extreme ($M = 1.76$, $SD = 0.61$) than negative competence behaviors ($M = 1.14$, $SD = 0.71$). Although this perhaps suggests that the observed dominance of negative moral information in this study can be reduced to the greater extremity that happens to characterize the behaviors we selected, two additional analyses provide evidence against this. First, controlling for the average extremity of each behavior, as assessed in this pilot study, does not affect the statistical significance of the impression updating effects in Experiment 3. In addition, inspection of the pilot data shows that it is just two negative moral behaviors that are more extreme than the negative competence behaviors. Omitting participants who those particular behaviors at T2 also does not affect the statistical significance of impression updating effects in Experiment 3. (See online supplement for complete analyses). Together, these pilot data and additional analyses provide further evidence that the dominance of moral information is about the qualitative distinction in the type of behavior and not a quantitative distinction in the extremity of stimuli chosen for the study.

change, as in the previous two studies, because (1) it better situates our findings in the established literature on moral primacy in impression updating, (2) it allowed us to continue testing for valence asymmetries, and (3) it provided higher power by situating the hypothesized effect as a main effect rather than an interaction. However, following the pre-registered analysis plan leads to the same conclusions. See the online supplement for these analyses.

6.2.2. First impressions

To assess whether morality or competence information carried greater weight when forming first impressions and whether this was moderated by the decision context, we conducted two 2 (Organization Type) \times 2 (T1 Valence) \times 2 (T1 Type) between-subjects ANOVAs on Time 1 impressions and hiring probability.

First, results for overall impressions supported the expected information valence \times type interaction, $F(1, 546) = 21.60, p < .001, \eta_p^2 = 0.04$. Although positive information about competence led to more favorable impressions ($M = 2.04, SD = 1.43$) than negative information about competence ($M = -0.13, SD = 1.78$), $B = 1.08, t(546) = 11.12, p < .001, 95\% \text{ CI}: [0.89, 1.27]$, the effect of valence was more pronounced when the information communicated the target's moral character ($M_{\text{neg}} = -1.21, SD_{\text{neg}} = 1.96; M_{\text{pos}} = 2.23, SD_{\text{pos}} = 1.23$), $B = 1.72, t(546) = 17.90, p < .001, 95\% \text{ CI}: [1.53, 1.91]$. The interaction was not significantly moderated by the type of organization, although a marginal three-way interaction emerged such that the valence \times type effect was somewhat stronger in the for-profit condition, $F(1, 546) = 3.04, p = .08, \eta_p^2 = 0.01$; however, the valence \times type interaction is significant in each organization condition, $ps < 0.05$. There was a main effect of organization, however, such that people's impressions of the candidate tended to be more favorable when he was applying for a job at a for-profit organization ($M = 0.97, SD = 2.21$) than at a non-profit organization ($M = 0.47, SD = 2.14$), $F(1, 546) = 13.86, p < .001, \eta_p^2 = 0.03$.

We next repeated the above analyses to test the effects on perceived hiring probability following the first piece of information. Similar to the effects for impressions, people expressed that the candidate was more likely to be hired when he was applying for a job at a for-profit organization ($M = 2.96, SD = 1.15$) than at a non-profit organization ($M = 2.71, SD = 1.09$), $F(1, 546) = 11.43, p < .001, \eta_p^2 = 0.02$; he also seemed more likely to be hired when the information about him was positive ($M = 3.53, SD = 0.75$) versus negative ($M = 2.13, SD = 0.99$), $F(1, 546) = 359.35, p < .001, \eta_p^2 = 0.40$; and hiring probability was marginally higher when the information was about competence ($M = 2.93, SD = 1.10$) than morality ($M = 2.74, SD = 1.14$), $F(1, 546) = 3.60, p = .06, \eta_p^2 = 0.01$. No interaction effects were significant in this model, $ps > 0.15$.

6.2.3. Impression updating

We next examined the effects on absolute changes in general impressions following the second piece of information. In line with the moral primacy hypothesis, impressions changed more when the new information was about the target's morality ($M = 2.16, SD = 1.94$) than when it was about his competence ($M = 1.57, SD = 1.51$), $F(1, 546) = 15.98, p < .001, \eta_p^2 = 0.03$. Additionally, as in the previous studies, impressions changed more when the new information was negative ($M = 2.16, SD = 1.94$) than when it was positive ($M = 1.57, SD = 1.50$), $F(1, 546) = 16.31, p < .001, \eta_p^2 = 0.03$. These effects were qualified, however, by an information valence \times type interaction, $F(1, 546) = 13.75, p < .001, \eta_p^2 = 0.03$ (Fig. 3A). When T2 information was positive, the impact of morality- and competence-related information on impression change was analogous, $B = 0.02, t(546) = 0.20, p = .84, 95\% \text{ CI}: [-0.18, 0.22]$. In contrast, it was when T2 information was negative that morality-related information led to greater change ($M = 2.70, SD = 2.18$) than competence-related information ($M = 1.59, SD = 1.47$), $B = 0.55, t(546) = 5.45, p < .001, 95\% \text{ CI}: [0.35, 0.75]$. This interaction was not further qualified by the type of organization, $F(1, 546) = 1.07, p = .30, \eta_p^2 = 0.002$.

The same ANOVA was conducted on hiring probability. Overall,

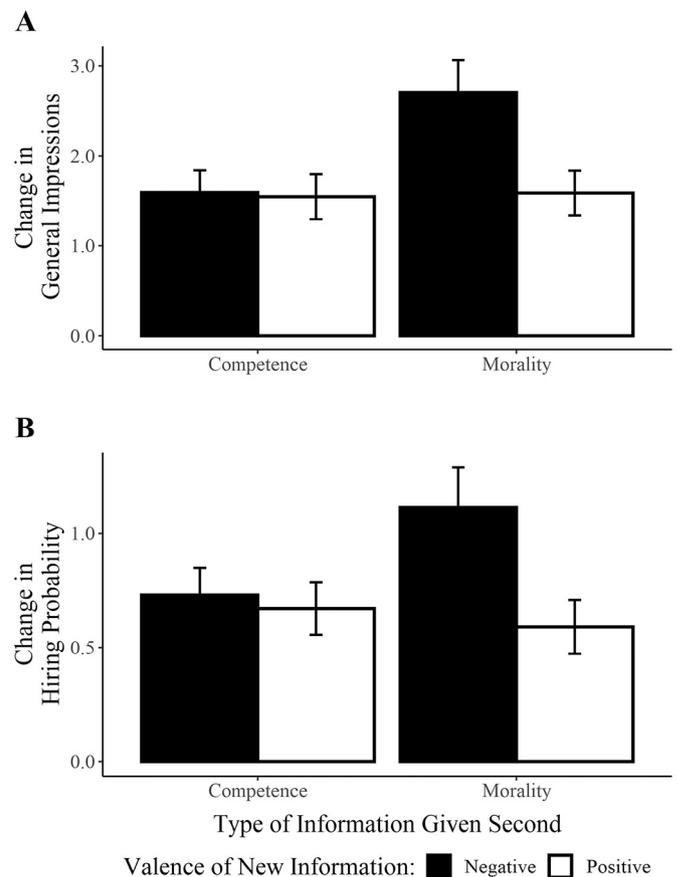


Fig. 3. Interactions between the valence and type of information at T2 on (A) changes in people's general impressions of the job candidate and (B) changes in people's judgments of the candidate's likelihood of being hired (Study 3).

judgments changed more when the new information was about the target's morality ($M = 0.86, SD = 0.94$) than when it was about his competence ($M = 0.70, SD = 0.71$), $F(1, 546) = 4.77, p = .03, \eta_p^2 = 0.01$. Additionally, judgments of hiring probability changed more when the new information was negative ($M = 0.93, SD = 0.93$) than when it was positive ($M = 0.63, SD = 0.70$), $F(1, 546) = 17.78, p < .001, \eta_p^2 = 0.03$. These effects were qualified, however, by an information valence \times type interaction, $F(1, 546) = 11.21, p < .001, \eta_p^2 = 0.02$ (Fig. 3B). When T2 information was positive, the impact of morality- and competence-related information on impression change was analogous, $B = -0.04, t(546) = -0.82, p = .41, 95\% \text{ CI}: [-0.14, 0.06]$. In contrast, it was when T2 information was negative that morality-related information led to greater change ($M = 1.11, SD = 1.07$) than competence-related information ($M = 0.73, SD = 0.71$), $B = 0.19, t(546) = 3.91, p < .001, 95\% \text{ CI}: [0.10, 0.29]$.

6.3. Discussion

In sum, the results of Experiment 3 corroborated findings in the previous two studies, namely that negative information about a target's moral character had a bigger effect on impression updating than negative information about their competence. In addition, this study replicated this effect in another national context and found little evidence that it was moderated by the organization's non-profit versus for-profit nature. Further, we found that even in a context particularly likely to make competence salient—hiring someone to work at a for-profit company (with explicit instructions to hire someone who will maximize annual profits)—information about the candidate's moral character had a larger effect on both first and revised impressions than

information about that person's competence.

Finally, after conducting Experiment 3, we also collected additional data to test the possibility that our moral primacy effects are being driven more simply by the perceived frequency of the stimuli in our design. Although we argue that the perceived diagnosticity of behavioral information is tied to qualitative differences between morality and competence information, others have argued that diagnosticity results from quantitative differences in the frequency of such behaviors in the world (e.g., Mende-Siedlecki, Baron, & Todorov, 2013). That is, perhaps negative moral behaviors are especially influential because they seem relatively uncommon and thus are more informative. Therefore, we displayed all sixteen behaviors to a new group of 100 participants and asked them to evaluate how common each one was using a 4-item measure from prior research (Brambilla et al., 2019). However, we found a main effect of behavior type such that the eight moral behaviors (regardless of valence) were seen as more common than the eight competence behaviors, on average. This pattern suggests that the moral primacy effect is not driven by a perception that the negative moral behaviors are relatively rare (see also Brambilla et al., 2019). We report this additional pilot study in full in the online supplement.

7. Internal meta-analysis

Finally, we conducted an internal meta-analysis of our findings across these three studies. These were the only studies conducted in this line of research, and a meta-analysis will help estimate the robustness of our effects across the available data. In addition, because we presented indications of statistical power for main effects of behavior type despite testing and interpreting interactions, meta-analysis allows us to maximize power and estimate all effects across studies. However, we also wanted to use this analysis to test whether the primacy of morally relevant information differs for general, context-free impressions versus judgments more directly tied to the competence-oriented domain of hiring. That is, previous work in this area has emphasized that the

primacy of moral information holds for overall impressions of a person (Brambilla et al., 2021; Landy, Piazza, & Goodwin, 2018). Indeed, this may well hold even when the context might prioritize other dimensions like competence; however, in line with the potential goal-dependence of moral dominance (e.g., Melnikoff & Bailey, 2018), moral information may play a less dominant role when it comes to judgments that are directly relevant to a competence-oriented domain. Therefore, we used multilevel meta-analysis to compare the size of the moral primacy effect on general impressions versus hiring-specific judgments across studies, accounting for shared variance in effect sizes that come from the same samples.

Because we tended to find support for moral primacy only when T2 information was negative, we focus on effect sizes for receiving new information about the target's *immorality* versus *incompetence* on all dependent measures in each study. See the online supplement for a full meta-analysis across both T2 valence conditions, which supports the valence asymmetry across effect size estimates. Because all experiments included additional between-participants manipulations (i.e., T1 valence, status, and organization type), we computed separate effect sizes for each condition as each corresponds to a unique sample of participants. This allows us to include multiple independent samples in the analysis while avoiding averaging across experimental manipulations. We then conducted a multilevel meta-analysis using the *metafor* package for R (Viechtbauer, 2010), entering sample as a random effect.

Fig. 4 presents effect sizes for each sample and the meta-analytic effect across samples for each dependent measure. Across samples and dependent measures, an overall meta-analysis shows a reliable moral primacy effect, $\bar{d} = 0.37, t = 7.47, p < .001, 95\% \text{ CI}: [0.27, 0.47]$. In a subsequent model, we tested dependent variable as a moderator of the effect size. The results provide marginal support for such moderation, $F(2, 19) = 3.39, p = .05$. The effect of morality-related versus competence-related information at T2 on changes in general impressions was significantly larger than the effect on changes in evaluations of the target as an employee, $b = -0.32, t = -2.42, p = .03$, and marginally

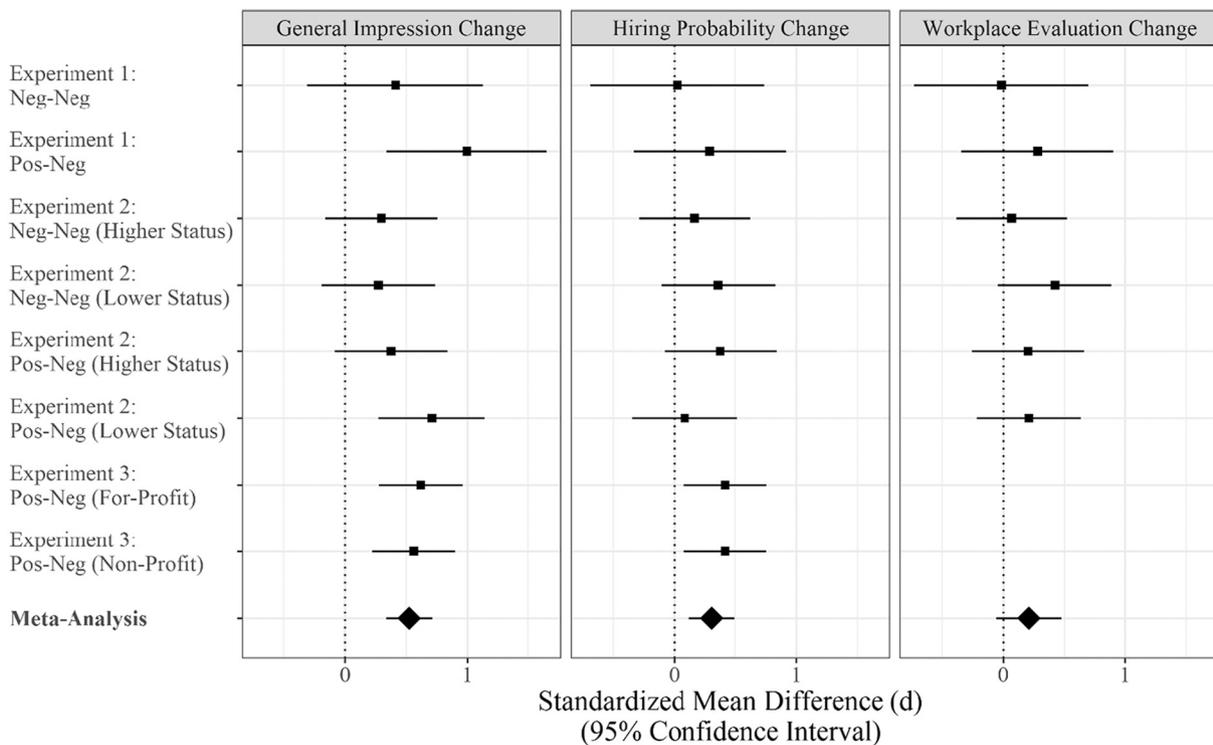


Fig. 4. Internal meta-analysis across studies. This meta-analysis includes effects of morality-related versus competence-related negative information presented at T2. More positive effects indicate a stronger effect of morality-related (vs. competence-related) information on each dependent variable. Labels for each sample on the y-axis present the valence presented at T1 and the valence presented at T2, respectively.

larger than the effect on changes in hiring probability, $b = -0.22$, $t = -1.92$, $p = .07$.

Indeed, in separate analyses on each dependent variable, the data support a reliable, moderately sized moral primacy effect on general impression change, $\bar{d} = 0.52$, $t = 6.57$, $p = .003$, 95% CI: [0.34, 0.71], a smaller but still reliable effect on changes in hiring probability, $\bar{d} = 0.31$, $t = 3.91$, $p = .006$, 95% CI: [0.12, 0.50], but a nonsignificant effect on changes in evaluations of the target as an employee, $\bar{d} = 0.21$, $t = 2.01$, $p = .10$, 95% CI: [-0.06, 0.47]. None of these separate analyses find evidence that the effect size varies across samples (i.e., across additional manipulations), $ps > 0.58$.

8. General discussion

Across three studies, we found that even though people often identify competence as the primary trait of interest when it comes to choosing who to hire as an employee (see pilot studies), information about a job candidate's (im)morality consistently exerted more influence on impression updating than information about his (in)competence. This moral primacy effect was apparent both for global evaluations of the target as well as more domain-specific judgments like probability of hiring the person. These effects were also unmoderated by the target's potential status or by the organization's non-profit versus for-profit status (and corresponding instructions to prioritize community impact versus profits). The moral primacy effects we observed for negative information, however, did not extend to effects of positive information, consistent with prior research on negativity biases in impression development (e.g., Skowronski & Carlston, 1987; see also Riva, Brambilla, & Vaes, 2016).

Notably, the results of an internal meta-analysis show that the primacy of negative moral information emerged most strongly on global impressions. These measures did not specifically invoke considerations directly relevant to the hiring decision itself, which may explain why morality continued to play a dominant role. Nevertheless, a reliable moral primacy effect also emerged for people's judgments of the candidate's hireability, which is presumably more directly tied to the hiring context that should highlight concerns for competence. Although we also observed moral primacy for revisions to people's evaluations of the target as a potential employee in Experiment 2, this effect was not significant in Experiment 1 nor was it reliable in the meta-analysis.

In some ways, the patterns we observed across studies is consistent with the hypothesis that moral primacy would be attenuated in a judgment context that makes competence an important dimension of evaluation. First, we found evidence for moral primacy more strongly on changes in global impressions than on changes to domain-specific evaluations. It seems the dominance of moral information can be attenuated when the dependent measure renders another dimension (i.e., competence) especially relevant. In fact, when we examined effects on *first* impressions in Experiment 3, we found evidence for moral primacy on overall impressions but not on initial judgments of hiring probability. Second, contrary to prior research on impression updating (Brambilla et al., 2019), we failed to find a moral primacy effect when new information was *positive*. In other words, in this hiring context, information about a candidate's positive moral character did not carry more influence than information about his competence. Such a valence asymmetry is evident in some stages of impression development. For instance, negative morality tends to exert a greater influence in shaping first impressions and behavioral responses (Brambilla, Sacchi, Menegatti, & Moscatelli, 2016; Menegatti, Moscatelli, Brambilla, & Sacchi, 2020; Skowronski & Carlston, 1987; see also Brambilla & Leach, 2014). The evidence of valence asymmetries in impression updating is more mixed (Brambilla et al., 2019; Reeder & Coovert, 1986) and it would be important to systematically investigate the factors that might promote or suppress the negativity effect on morality when studying impression updating (e.g., the context of evaluation).

Notably, prior research on the confirmability of trait concepts has documented a *negativity* bias for morality information but also a *positivity* bias for competence information (Mende-Siedlecki et al., 2013; Skowronski & Carlston, 1987; Trafimow & Trafimow, 1999). Honest individuals are expected to behave almost exclusively honestly, while dishonest individuals could engage in both honest and dishonest behaviors (see Trafimow & Trafimow, 1999). As a result, negative (vs. positive) moral traits are more diagnostic of an individuals' underlying moral character. By contrast, incompetent individuals are expected to behave almost exclusively inefficiently, while competent individuals could engage in both competent and incompetent behaviors (i.e., even a genius can fail sometimes). Yet, although the negativity bias for moral information in impression formation is well documented, prior research has failed to offer compelling evidence that competence carries more weight than incompetence in predicting impressions (for reviews, Brambilla et al., 2021; Rusconi et al., 2020). Indeed, a growing body of work fails to find such an effect when global impressions are taken into account, perhaps reflecting that competence remains a dimension with relatively little impact on overall impressions (for a review, Brambilla et al., 2021).

Our conclusions about the boundaries of moral primacy are necessarily tentative, however. First, we caution against over-interpreting the differences we identify between dependent measures. Although these measures are conceptually distinct in meaningful ways, they are also empirically distinct in wording and response scales. Effects may be weaker on one dependent variable because that variable was measured less reliably and not because the actual effect on the conceptual variable is smaller. Second, null effects for the positive T2 information do not necessarily mean that morality and competence information were treated equally, only that we cannot conclude that one had a stronger effect than the other.

Putting these relative differences and valence asymmetries aside, perhaps the most striking conclusion we can draw from these studies is that people revise their impressions of a job candidate more after seeing an indication that he is *immoral* than after seeing an indication that he is incompetent. This occurs both for global impressions and for judgments of his hireability. It even occurs when participants were specifically instructed to maximize profits for a for-profit company (Experiment 3). These findings stand as a potent demonstration that people can continue to weight moral information more heavily even for decisions that they might otherwise view as chiefly about finding the most competent person.

8.1. Limitations and future directions

Although this research has benefited from the use of samples that varied in national origin, employment status, and recruitment method, we acknowledge that further diversity in participants would be an asset as it would in much psychological research. Further, however, we also acknowledge that across studies, the social target was always male and either implicitly or explicitly from the dominant ethnic group in the sample's society. Recently, moral psychologists have argued for the importance of deliberately considering social identity in research on moral judgment (Harris, Pärnamets, Brady, Robertson, & Van Bavel, 2021; Hester & Gray, 2020). Perhaps the persistent dominance of moral character in employment settings depends on the target's identity; future work could probe this question more directly. However, prior work has demonstrated moral primacy both for ingroup (Leach et al., 2007) and outgroup evaluations (Brambilla et al., 2012; Brambilla, Hewstone, & Colucci, 2013; Brambilla, Sacchi, et al., 2013). In fact, moral character may even matter more when evaluating marginalized targets in employment contexts. For example, whereas competence is indeed valued for male managerial applicants, hiring decisions become more dependent on perceived sociability when the applicant is a competent woman (Phelan et al., 2008). These shifting standards suggest that moral character may be even more important for women or other marginalized

social groups despite the presumption that competence should be the more strongly valued dimension.

We also acknowledge that our studies drew upon a limited sample of stimuli. Although we developed different stimuli for different studies and randomly drew on a bank of possible stimuli in Experiment 3, it is possible that our findings depend on idiosyncratic features of our stimuli rather than the dimensions of morality and competence per se. Nevertheless, we conducted several analyses to rule out the influence of potentially confounding variables (i.e., extremity and perceived behavior frequency), but perhaps the stimuli in our studies conveyed an aspect of morality and competence that does not generalize to all possible behaviors along those dimensions. We invite further work that samples even more behavioral stimuli to test the bounds of moral primacy.

These studies also relied on hypothetical hiring scenarios; perhaps people are more willing to abandon morality judgments in favor of competence when an actual hiring decision is on the line. Indeed, some evidence suggests that people abide by moral norms in hypothetical scenarios but are more likely to demonstrate self-interest at the expense of such moral norms when the decisions are real and consequential (FeldmanHall et al., 2012). Other work, however, has found behavioral consequences of moral primacy (Brambilla et al., 2016; Menegatti et al., 2020), extending that work beyond self-report outcome measures, but generalizing to actual hiring decisions remains an open question. Nevertheless, it remains striking that people both express a priority for competence in one hypothetical survey but still succumb to the impact of negative moral information in another.

In addition to exploring beyond hypothetical scenarios, researchers might also gain from examining impression updating in other competence-oriented contexts. Judgments of political candidates, for instance, are often informed more by their associations with competence than warmth (Ksiazkiewicz, Vitriol, & Farhart, 2018; Vitriol, Ksiazkiewicz, & Farhart, 2018). Perhaps in this or other domains, the dominant impact of morality judgments may fade as competence judgments take center stage.

8.2. Conclusion

In sum, we tested the boundaries of moral primacy and found that even in a context where other dimensions could dominate, information about a job candidate's immorality continued to have disproportionate influence on general impressions of him and evaluations of his suitability as an employee. However, our findings further show that the relative effect of negative moral versus competence information on domain-specific judgments tended to be smaller than effects on general impressions. In addition, unlike prior research on impression updating (Brambilla et al., 2019), we observed no evidence for moral primacy in this context when the new information was positive (although this pattern may be indicative of a more general valence asymmetry in the effects of morally relevant information). Together, these findings provide an important extension of the Moral Primacy Model but also provide useful insight on the generalizability of the effect.

Open practices

Data and analysis scripts for all reported results are available on this project's page on the Open Science Framework: <https://osf.io/3he4a/>. Also included is the article's online supplement, which contains all materials for reproducing the reported methodology. The pre-registration for Experiment 3 is available at: <https://osf.io/pe7rc>.

Funding

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

Acknowledgements

The authors thank Skylar Brannon for her integral involvement in the conception and design of Experiment 3.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jesp.2021.104246>.

References

- Aaker, J., Vohs, K. D., & Mogilner, C. (2010). Nonprofits are seen as warm and for-profits as competent: Firm stereotypes matter. *Journal of Consumer Research*, 37(2), 224–237. <https://doi.org/10.1086/651566>
- Abele, A. E., Ellemers, N., Fiske, S. T., Koch, A., & Yzerbyt, V. (2021). Navigating the social world: Toward an integrated framework for evaluating self, individuals, and groups. *Psychological Review*, 128(2), 290–314. <https://doi.org/10.1037/rev0000262>
- Abele, A. E., & Wojciszke, B. (2007). Agency and communion from the perspective of self versus others. *Journal of Personality and Social Psychology*, 93(5), 751–763. <https://doi.org/10.1037/0022-3514.93.5.751>
- Amaral, A. A., Powell, D. M., & Ho, J. L. (2019). Why does impression management positively influence interview ratings? The mediating role of competence and warmth. *International Journal of Selection and Assessment*, 27(4), 315–327. <https://doi.org/10.1111/ijasa.12260>
- Barrick, M. R., Swider, B. W., & Stewart, G. L. (2010). Initial evaluations in the interview: Relationships with subsequent interviewer evaluations and employment offers. *Journal of Applied Psychology*, 95(6), 1163–1172. <https://doi.org/10.1037/a0019918>
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, 5(4), 323–370. <https://doi.org/10.1037/1089-2680.5.4.323>
- Brambilla, M., Carraro, L., Castelli, L., & Sacchi, S. (2019). Changing impressions: Moral character dominates impression updating. *Journal of Experimental Social Psychology*, 82, 64–73. <https://doi.org/10.1016/j.jesp.2019.01.003>
- Brambilla, M., Hewstone, M., & Colucci, F. P. (2013). Enhancing moral virtues: Increased perceived outgroup morality as a mediator of intergroup contact effects. *Group Processes & Intergroup Relations*, 16(5), 648–657. <https://doi.org/10.1177/1368430212471737>
- Brambilla, M., & Leach, C. W. (2014). On the importance of being moral: The distinctive role of morality in social judgment. *Social Cognition*, 32(4), 397–408. <https://doi.org/10.1521/soco.2014.32.4.397>
- Brambilla, M., Rusconi, P., Sacchi, S., & Cherubini, P. (2011). Looking for honesty: The primary role of morality (vs. sociability and competence) in information gathering. *European Journal of Social Psychology*, 41, 135–143. <https://doi.org/10.1002/ejsp.744>
- Brambilla, M., Sacchi, S., Menegatti, M., & Moscatelli, S. (2016). Honesty and dishonesty don't move together: Trait content information influences behavioral synchrony. *Journal of Nonverbal Behavior*, 40(3), 171–186. <https://doi.org/10.1007/s10919-016-0229-9>
- Brambilla, M., Sacchi, S., Pagliaro, S., & Ellemers, N. (2013). Morality and intergroup relations: Threats to safety and group image predict the desire to interact with outgroup and ingroup members. *Journal of Experimental Social Psychology*, 49(5), 811–821. <https://doi.org/10.1016/j.jesp.2013.04.005>
- Brambilla, M., Sacchi, S., Rusconi, P., Cherubini, P., & Yzerbyt, V. Y. (2012). You want to give a good impression? Be honest! Moral traits dominate group impression formation. *British Journal of Social Psychology*, 51(1), 149–166. <https://doi.org/10.1111/j.2044-8309.2010.02011.x>
- Brambilla, M., Sacchi, S., Rusconi, P., & Goodwin, G. (2021). The primacy of morality in impression development: Theory, research, and future directions. *Advances in Experimental Social Psychology*, 64, 187–262. <https://doi.org/10.1016/bs.aesp.2021.03.001>
- Brannon, S. M., & Gawronski, B. (2017). A second chance for first impressions? Exploring the context-(in)dependent updating of implicit evaluations. *Social Psychological and Personality Science*, 8(3), 275–283. <https://doi.org/10.1177/1948550616673875>
- Brannon, S. M., Sacchi, D. L. M., & Gawronski, B. (2017). (In)consistency in the eye of the beholder: The roles of warmth, competence, and valence in lay perceptions of inconsistency. *Journal of Experimental Social Psychology*, 70, 80–94. <https://doi.org/10.1016/j.jesp.2016.12.011>
- Caprariello, P. A., Cuddy, A. J. C., & Fiske, S. T. (2009). Social structure shapes cultural stereotypes and emotions: A causal test of the stereotype content model. *Group Processes & Intergroup Relations*, 12(2), 147–155. <https://doi.org/10.1177/1368430208101053>
- Carrier, A., Louvet, E., Chauvin, B., & Rohmer, O. (2014). The primacy of agency over competence in status perception. *Social Psychology*, 45(5), 347–356. <https://doi.org/10.1027/1864-9335/a000176>
- Cislak, A., & Cichocka, A. (2018). Power, self-focus, and the big two. In *Agency and communion in social psychology* (p. 12). Routledge. <https://doi.org/10.4324/9780203703663-12>
- Cone, J., & Ferguson, M. J. (2015). He did what? The role of diagnosticity in revising implicit evaluations. *Journal of Personality and Social Psychology*, 108(1), 37–57. <https://doi.org/10.1037/pspa0000014>

- Cone, J., Flaharty, K., & Ferguson, M. J. (2021). The long-term effects of new evidence on implicit impressions of other people. *Psychological Science*, 32(2), 173–188. <https://doi.org/10.1177/0956797620963559>
- Cuddy, A. J. C., Fiske, S. T., & Glick, P. (2004). When professionals become mothers, warmth doesn't cut the ice. *Journal of Social Issues*, 60(4), 701–718. <https://doi.org/10.1111/j.0022-4537.2004.00381.x>
- FeldmanHall, O., Mobbs, D., Evans, D., Hiscox, L., Navrady, L., & Dalgleish, T. (2012). What we say and what we do: The relationship between real and hypothetical moral choices. *Cognition*, 123(3), 434–441. <https://doi.org/10.1016/j.cognition.2012.02.001>
- Ferguson, M. J., Mann, T. C., Cone, J., & Shen, X. (2019). When and how implicit first impressions can be updated. *Current Directions in Psychological Science*, 28(4), 331–336. <https://doi.org/10.1177/0963721419835206>
- Fiske, S. T. (2018). Stereotype content: Warmth and competence endure. *Current Directions in Psychological Science*, 27(2), 67–73. <https://doi.org/10.1177/0963721417738825>
- Fiske, S. T., Cuddy, A. J. C., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology*, 82(6), 878–902. <https://doi.org/10.1037/0022-3514.82.6.878>
- Goodwin, G. P. (2015). Moral character in person perception. *Current Directions in Psychological Science*, 24(1), 38–44. <https://doi.org/10.1177/0963721414550709>
- Goodwin, G. P., Piazza, J., & Rozin, P. (2014). Moral character predominates in person perception and evaluation. *Journal of Personality and Social Psychology*, 106(1), 148–168. <https://doi.org/10.1037/a0034726>
- Graves, L. M., & Powell, G. N. (1988). An investigation of sex discrimination in recruiters' evaluations of actual applicants. *Journal of Applied Psychology*, 73(1), 20–29. <https://doi.org/10.1037/0021-9010.73.1.20>
- Harris, E. A., Parnaments, P., Brady, W. J., Robertson, C., & Van Bavel, J. J. (2021). The role of generalizability in moral and political psychology. *OSF Preprints*. <https://doi.org/10.31219/osf.io/nfjct>
- Hester, N., & Gray, K. (2020). The moral psychology of raceless, genderless strangers. *Perspectives on Psychological Science*, 15(2), 216–230. <https://doi.org/10.1177/1745691619885840>
- Howard, J. L., & Ferris, G. R. (1996). The employment interview context: Social and situational influences on interviewer decisions. *Journal of Applied Social Psychology*, 26(2), 112–136. <https://doi.org/10.1111/j.1559-1816.1996.tb01841.x>
- Koch, A., Yzerbyt, V., Abele, A., Ellemers, N., & Fiske, S. T. (2021). Social evaluation: Comparing models across interpersonal, intragroup, intergroup, several-group, and many-group contexts. *Advances in Experimental Social Psychology*, 63, 1–68. <https://doi.org/10.1016/bs.aesp.2020.11.001>
- Krings, F., Szesny, S., & Kluge, A. (2011). Stereotypical inferences as mediators of age discrimination: The role of competence and warmth. *British Journal of Management*, 22(2), 187–201. <https://doi.org/10.1111/j.1467-8551.2010.00721.x>
- Ksiazkiewicz, A., Vitriol, J., & Farhart, C. (2018). Implicit candidate-trait associations in political campaigns. *Political Psychology*, 39(1), 177–195. <https://doi.org/10.1111/pops.12398>
- Landy, J. F., Piazza, J., & Goodwin, G. P. (2016). When it's bad to be friendly and smart: The desirability of sociability and competence depends on morality. *Personality and Social Psychology Bulletin*, 42(9), 1272–1290. <https://doi.org/10.1177/0146167216655984>
- Landy, J. F., Piazza, J., & Goodwin, G. P. (2018). Morality traits still dominate in forming impressions of others. *Proceedings of the National Academy of Sciences of the United States of America*, 115(25), Article E5636. <https://doi.org/10.1073/pnas.1807096115>
- Leach, C. W., Ellemers, N., & Barreto, M. (2007). Group virtue: The importance of morality (vs. competence and sociability) in the positive evaluation of in-groups. *Journal of Personality and Social Psychology*, 93(2), 234–249. <https://doi.org/10.1037/0022-3514.93.2.234>
- Melnikoff, D. E., & Bailey, A. H. (2018). Preferences for moral vs. immoral traits in others are conditional. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, 115(4), E592–E600. <https://doi.org/10.1073/pnas.1714945115>
- Mende-Siedlecki, P., Baron, S. G., & Todorov, A. (2013). Diagnostic value underlies asymmetric updating of impressions in the morality and ability domains. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 33(50), 19406–19415. <https://doi.org/10.1523/JNEUROSCI.2334-13.2013>
- Menegatti, M., Moscatelli, S., Brambilla, M., & Sacchi, S. (2020). The honest mirror: Morality as a moderator of spontaneous behavioral mimicry. *European Journal of Social Psychology*, 50(7), 1394–1405. <https://doi.org/10.1002/ejsp.2670>
- Pagliaro, S., Brambilla, M., Sacchi, S., D'Angelo, M., & Ellemers, N. (2013). Initial impressions determine behaviours: Morality predicts the willingness to help newcomers. *Journal of Business Ethics*, 117(1), 37–44. <https://doi.org/10.1007/s10551-012-1508-y>
- Phelan, J. E., Moss-Racusin, C. A., & Rudman, L. A. (2008). Competent yet out in the cold: Shifting criteria for hiring reflect backlash toward agentic women. *Psychology of Women Quarterly*, 32(4), 406–413. <https://doi.org/10.1111/j.1471-6402.2008.00454.x>
- van Prooijen, A.-M., & Ellemers, N. (2015). Does it pay to be moral? How indicators of morality and competence enhance organizational and work team attractiveness: Does it pay to be moral? *British Journal of Management*, 26(2), 225–236. <https://doi.org/10.1111/1467-8551.12055>
- Reeder, G. D., & Brewer, M. B. (1979). A schematic model of dispositional attribution in interpersonal perception. *Psychological Review*, 86(1), 61–79. <https://doi.org/10.1037/0033-295X.86.1.61>
- Reeder, G. D., & Coovert, M. D. (1986). Revising an impression of morality. *Social Cognition*, 4(1), 1–17. <https://doi.org/10.1521/soco.1986.4.1.1>
- Riva, P., Brambilla, M., & Vaes, J. (2016). Bad guys suffer less (social pain): Moral status influences judgements of others' social suffering. *British Journal of Social Psychology*, 55(1), 88–108. <https://doi.org/10.1111/bjso.12114>
- Rothbart, M., & Park, B. (1986). On the confirmability and disconfirmability of trait concepts. *Journal of Personality and Social Psychology*, 50(1), 131–142. <https://doi.org/10.1037/0022-3514.50.1.131>
- Rusconi, P., Sacchi, S., Brambilla, M., Capellini, R., & Cherubini, P. (2020). Being honest and acting consistently: Boundary conditions of the negativity effect in the attribution of morality. *Social Cognition*, 38(2), 146–178. <https://doi.org/10.1521/soco.2020.38.2.146>
- Skowronski, J. J., & Carlston, D. E. (1987). Social judgment and social memory: The role of cue diagnosticity in negativity, positivity, and extremity biases. *Journal of Personality and Social Psychology*, 52(4), 689–699. <https://doi.org/10.1037/0022-3514.52.4.689>
- Strohinger, N., & Nichols, S. (2014). The essential moral self. *Cognition*, 131(1), 159–171. <https://doi.org/10.1016/j.cognition.2013.12.005>
- Sun, J., & Goodwin, G. P. (2020). Do people want to be more moral? *Psychological Science*, 31(3), 243–257. <https://doi.org/10.1177/0956797619893078>
- Trafimow, D., & Trafimow, S. (1999). Mapping perfect and imperfect duties onto hierarchically and partially restrictive trait dimensions. *Personality and Social Psychology Bulletin*, 25(6), 686–695. <https://doi.org/10.1177/0146167299025006004>
- Trope, Y., & Liberman, A. (1996). Social hypothesis testing: Cognitive and motivational mechanisms. In *Social psychology: Handbook of basic principles* (pp. 239–270). The Guilford Press.
- Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, 36(1), 1–48. <https://doi.org/10.18637/jss.v036.i03>
- Vitriol, J. A., Ksiazkiewicz, A., & Farhart, C. E. (2018). Implicit candidate traits in the 2016 U.S. presidential election: Replicating a dual-process model of candidate evaluations. *Electoral Studies*, 54, 261–268. <https://doi.org/10.1016/j.electstud.2018.04.009>
- Wojciszke, B., & Abele, A. E. (2008). The primacy of communion over agency and its reversals in evaluations. *European Journal of Social Psychology*, 38(7), 1139–1147. <https://doi.org/10.1002/ejsp.549>
- Wojciszke, B., Bazinska, R., & Jaworski, M. (1998). On the dominance of moral categories in impression formation. *Personality and Social Psychology Bulletin*, 24(12), 1251–1263. <https://doi.org/10.1177/01461672982412001>