A Conflict of Values:

When Perceived Compassion Decreases Trust

Matthew J. Lupoli[1], Min Zhang[2], Yidan Yin[2], & Christopher Oveis[2]

[1] Monash University

[2] University of California, San Diego

In press, Journal of Experimental Social Psychology

Abstract

Compassion benefits individuals, organizations, and society. As such, people may place greater trust in those who are perceived to be compassionate, believing that they will act with both benevolence and integrity. In some circumstances, however, acting with benevolence may seemingly require a sacrifice of integrity and vice versa. We propose that expectations of how compassionate people navigate these conflicts between benevolence and integrity influence trust of compassionate individuals. In five experiments (N = 1,744), we demonstrate that perceived compassion can either increase or decrease trust depending on the context. Specifically, perceived compassion decreased trust in individuals' integrity during benevolence-integrity conflicts (Experiments 2-5), but increased trust in their benevolence (Experiment 1) and integrity (Experiment 2) when these values were not in conflict. These effects were observed across several measures of trust, manipulations of perceived compassion, and experimental methods, including incentivized economic games (Experiments 1 and 2), realistic vignettes (Experiment 3), and incentivized organizational decisions (Experiments 4 and 5). Beliefs that compassionate individuals place a higher relative importance on benevolence versus integrity mediated the negative effect of perceived compassion on trust during benevolence-integrity conflicts (Experiment 5). Collectively, these results highlight a potential drawback of being seen as compassionate.

Keywords: compassion; emotion; trust; moral dilemmas; social perception

A Conflict of Values: When Perceived Compassion Decreases Trust

Imagine you will be giving a presentation at work that could determine your future at the company. You practice your presentation in front of a colleague who has offered to give you constructive feedback, which could be crucial in helping you to improve your talk. There is only one problem: Your colleague is widely known to be a compassionate person.

In this situation, would you trust your colleague to give you accurate feedback? On one hand, you might trust that your compassionate colleague would offer an honest opinion of your performance. Compassion is associated with benevolent intentions towards others (Goetz, Keltner, & Simon-Thomas, 2010), which research has shown to reliably predict trust (Colquitt, Scott, & LePine, 2007; Mayer, Davis, & Schoorman, 1995; Mayer & Davis, 1999).

However, the relationship between perceived compassion and trust may not be this straightforward. To return to the above example, choosing between providing positive but dishonest feedback or telling the hurtful truth is a common dilemma in both social and professional contexts. How are compassionate individuals expected to behave in situations where it seems difficult or impossible to act benevolently without sacrificing integrity and vice versa? Answering this question is critical, as conflicts between benevolence and integrity occur frequently (DePaulo, Kashy, Kirkendol, Wyer, & Epstein, 1996; Lupoli, Levine, & Greenberg, 2018; Moore, Mungia Gomez, & Levine, 2019) and trusting (or distrusting) others can have profound consequences (Dirks & Ferrin, 2001; Jachimowicz, Chafik, Munrat, Prabhu, & Weber, 2017; Yip & Schweitzer, 2015). In this work, we propose that although the perception of compassion in others can increase trust, when individuals are confronted with these benevolence-integrity conflicts, being seen as compassionate can *decrease* trust.

Compassion, Benevolence, and Integrity

To understand the relationship between perceived compassion and trust, one must first understand the nature of compassion. Compassion is defined as being emotionally motivated to alleviate the distress or suffering of others (Goetz et al., 2010). As such, compassion has been shown to drive prosocial behaviors such as generosity (Saslow et al., 2010), charitable giving (Small & Simonsohn, 2007), volunteerism (Omoto, Malsch, & Barraza, 2009), and helping (Batson & Shaw, 1991; Valdesolo & DeSteno, 2011). Collectively, this research highlights that compassion motivates people to act out of concern for the welfare of others.

Given that compassion increases prosocial intentions and behaviors, this emotion is clearly linked to benevolence, or the desire to do good for others (Mayer et al., 1995). A large body of research indicates that the perception of benevolence in others increases trust (for a review, see Colquitt et al., 2007). Thus, it seems likely that people would place trust in those known to be compassionate.

On the other hand, recent research indicates that compassionate individuals may not always be trustworthy. For instance, experiencing compassion can lead to the provision of unfair advantages for suffering individuals (Batson et al., 1995; Loewenstein & Small, 2007; Slovic, 2007; Slovic, Västfjäll, Erlandsson, & Gregory, 2017). Other work demonstrates that those who experience compassion use deception in order to help less fortunate others financially (Gino & Pierce, 2009). Additionally, compassion increases lying that is intended to protect individuals from emotional harm (Lupoli et al., 2017).

Together, this research indicates that compassionate individuals sometimes do not act with integrity. Integrity involves the adherence to moral principles such as fairness, honesty, and justice (Butler & Cantrell, 1984; Mayer et al., 1995). Like benevolence, the perception of integrity in others has also been shown to increase trust (Colquitt et al., 2007). Thus, to the extent

that people have an intuition about compassionate individuals' tendencies to sometimes forego integrity, it seems possible that perceived compassion could at times decrease trust.

How might these competing hypotheses regarding the relationship between perceived compassion and trust be reconciled? One important point to consider is that compassion-driven sacrifices of integrity are often made in service of benevolence: a suffering individual is given unfair advantages *to help that individual* (Batson et al., 1995); deception is used to *aid others financially* (Gino & Pierce, 2009); lies are told to *prevent emotional harm* (Lupoli et al., 2017). In all of these situations, compassion seems to push individuals towards benevolence when they are faced with an apparent tradeoff between behaving with integrity and acting with benevolence. We refer to these tradeoffs as benevolence-integrity conflicts and argue that they play a key role in the relationship between perceived compassion and trust.

Benevolence-Integrity Conflicts

Benevolence-integrity conflicts are situations in which behaving with integrity seemingly requires a sacrifice of benevolence and vice-versa. When confronted with these conflicts, individuals may feel that they are unable to uphold moral principles, such as being honest, while still caring for the welfare of others (Moore et al., 2019). Benevolence-integrity conflicts are both prevalent and important in real world contexts. For instance, people face these dilemmas when they decide whether to give hurtful but honest feedback to family or friends. Managers must sometimes choose whether to allocate employee rewards based on merit or to favor those in greater need. Professors confront these conflicts when they are asked to write recommendation letters for underqualified students. In all of these contexts, how these individuals decide to resolve the conflict could have critical effects both on the decision-maker and others. To revisit

the opening example, giving honest feedback may cause emotional pain to a colleague, but could also help to improve their performance and career in the long term.

It is important to note a few caveats concerning the nature of benevolence-integrity conflicts. First, the definition states that these situations *seemingly* require tradeoffs between benevolence and integrity. We are not claiming that individuals could never offer a prescription on how to best resolve these conflicts. Rather, we are claiming that these are contexts in which individuals might plausibly perceive a conflict in values, as it may not be obvious how to best navigate the conflict. Second, although these conflicts involve apparent tradeoffs between benevolence and integrity, individuals may at times produce creative solutions in an attempt to satisfy both values. For instance, a person might frame potentially hurtful feedback as suggestions for improvement and express confidence in the target's ability to reach their full potential (Levine, Roberts, & Cohen, 2019; Yeager et al., 2014). Lastly, although we dichotomize whether or not individuals are faced with benevolence-integrity conflicts for ease of experimentation, the extent to which individuals may perceive a conflict in values falls along a continuum. For example, an individual may experience a greater conflict between benevolence and integrity concerns when the target is and unable to take corrective action based on the feedback (e.g., criticizing a friend's outfit when s/he has no option to change; Levine, 2019), versus when the recipient can use the feedback to improve his/her situation.

We propose that benevolence-integrity conflicts play a critical part in the relationship between perceived compassion and trust. We next discuss the multi-faceted nature of trust to illuminate our proposition that perceived compassion can both increase and decrease trust, depending on whether a target faces a benevolence-integrity conflict the type of trust in question.

A Multi-Faceted View of Trust and the Role of Emotion

Consistent with prior research, we define trust as the willingness to be vulnerable to the actions of others based on positive expectations about their behavior (Rousseau, Sitkin, Burt, & Camerer, 1998). It is difficult to overstate the value of trust; trust is essential for the effective functioning of interpersonal relationships (Rempel, Holmes, & Zanna, 1985), organizations (De Jong, Dirks, & Gillespie, 2016; Dirks & Ferrin, 2001), and communities (Jachimowicz et al., 2017).

Given the multitude of situational factors (e.g., Kim, Ferrin, Cooper, & Dirks, 2004; Schweitzer, Ho & Zhang, 2018), characteristics of the truster (e.g., Buchan, Croson, & Solnick, 2008; Dunn & Schweitzer, 2005), and attributes of the trustee (e.g., Lount & Pettit, 2012) that contribute to trust provision, researchers have constructed theories to parsimoniously explain decisions to trust. One influential theory on trust that we draw upon in this research is the ability/benevolence/integrity model of trust (Mayer, Davis, & Schoorman, 1995). According to this theory, there are three trustee characteristics that foster trust: the trustee's perceived benevolence, and integrity, and ability. In this work, we focus specifically on trust decisions that relate to beliefs about a trustee's benevolence and integrity—that is, benevolence- and integrity-based trust.

Considering the importance of benevolence and integrity for establishing trust, scholars would benefit from understanding the cues that reliably signal these attributes. One such cue that may provide this information is the perception of emotion in others. According to emotions as social information theory, individuals' emotional expressions influence observers' affective responses, inferences, and behavior (Van Kleef, 2009; Van Kleef, de Dreu, & Manstead, 2010). In the current research, we draw on this theory by examining how perceived compassion influences both benevolence- and integrity-based trust. Importantly, we argue that one context in

which perceived compassion decreases integrity-based trust is when potential trustees are confronted with benevolence-integrity conflicts.

Hypotheses and Overview of Experiments

Considering the prevalence of benevolence-integrity conflicts in everyday life (DePaulo et al., 1996; Moore et al., 2019), and research demonstrating that compassion can reduce integrity when individuals encounter these conflicts (Batson et al., 1995, Gino & Pierce, 2009, Lupoli et al., 2017), it is likely that people have personal experience with compassionate others sacrificing integrity in the face of benevolence-integrity conflicts. We therefore hypothesized that when a target individual is confronted with a benevolence-integrity conflict, viewing that person as high in compassion will decrease integrity-based trust in that individual.

We also made several predictions about potential mechanisms underlying this effect. In general, trust decisions require the potential truster to make inferences about mental states (e.g., intentions and values) of the trustee. As such, it is possible that lay beliefs about the importance compassionate individuals place on different values would underlie this decrease in trust. Specifically, compassionate individuals may be believed to place a higher relative importance on benevolence (e.g., valuing harm mitigation) versus on integrity (e.g., valuing honesty). While possessing this hierarchy of values may normally not be seen as problematic, when a compassionate target is confronted with a benevolence-integrity conflict, it may signal a lower likelihood that the individual will act with integrity. Thus, we expected that when compassionate individuals face benevolence-integrity conflict (versus no conflict), beliefs that these individuals value benevolence over integrity would drive a decrease in integrity-based trust.

Another possibility is that compassion individuals are believed to be more concerned with impression management in general, and that this belief would underlie a decrease in trust

when compassionate individuals face benevolence-integrity conflicts (versus no conflict). There are two reasons for this. First, doing good things or others has reputational benefits and leads to other positive outcomes for the do-gooder (Flynn, Reagans, Amanatullah, & Ames, 2006; Hardy & Van Vugt, 2006). Because of this, people may believe that the benevolent actions of compassionate individuals are in part selfishly motivated (e.g., Cialdini et al., 1997). Second, research suggests that there are two subcomponents of agreeableness: (a) compassion and (b) politeness, which is associated with norm compliance (DeYoung, Quilty, & Peterson, 2007). Thus, it is possible that people conflate others' emotionally-driven concern for others' welfare (i.e., compassion) with their desires to meet society's expectations of politeness.

In addition to examining how benevolence-integrity conflicts influence the relationship between perceived compassion and trust, in this research we also explore how perceived compassion influences trust in the absence of such dilemmas. Given compassion's strong theoretical ties to benevolence, we hypothesized that perceived compassion would increase benevolence-based trust. But how would perceived compassion influence integrity-based trust when targets do not face a conflict of values? Past work has documented a halo effect, whereby social perceptions in one domain produce perceptions of the same valence in another domain (Nisbett & Wilson, 1977). Therefore, to the extent that compassionate individuals are known to prevent harm to others, they may also be thought to adhere to other moral principles, such as being honest and fair. As such, we hypothesized that when benevolence and integrity do not conflict, perceived compassion would increase integrity-based trust.

In five experiments, we examined how perceived compassion and benevolence-integrity conflicts influence trust. In Experiments 1 and 2, we used economic games to investigate the effect of perceived compassion on benevolence-based (Experiment 1) and integrity-based

(Experiment 2) trust. In Experiment 2, we also assessed the moderating role of benevolence-integrity conflicts. In Experiment 3, we tested how perceived compassion influences integrity-based trust in several realistic vignettes where targets faced benevolence-integrity conflicts. Experiment 4 further explored these effects using another operationalization of integrity-based trust: expected fairness. Lastly, in Experiment 5, we examined downstream consequences of these effects on hiring decisions and also tested potential mechanisms.

Experiment 1

In Experiment 1, we tested whether perceived compassion affects benevolence-based trust. Here, we manipulated perceived compassion by providing participants with personality test results of a partner. Then, participants played the trust game (adapted from Berg, Dickhaut, & McCabe, 1995) with their partner—a widely used and validated behavioral measure of benevolence-based trust (e.g., Levine & Schweitzer, 2015; Schweitzer, Hershey, & Bradlow, 2006).

Methods

Participants. Three hundred twenty-one participants were recruited on Amazon Mechanical Turk (MTurk) for this three-cell between-subjects design (high compassion/low compassion/control). Fifty-nine participants failed a comprehension check and were thus excluded (exclusions by condition for all studies are reported in the Supplemental Material). This left a final sample of 262 participants (M_{age} = 34.20, 43.13% female). For this study and all subsequent online studies, we planned to obtain 100 participants per between-subjects experimental condition a priori, which we expected would give us sufficient power to detect small-to-medium effect sizes. A sensitivity analysis (α = 0.05; ANOVA: fixed effects, omnibus, one-way) using G*Power indicated that the sample size of 262 gave us 80% power to detect an

effect size of $f = 0.19$. We report all measures, manipulations, and exclusions in these studies. No data was analyzed prior to completion of data collection.

Procedure. Participants were first informed that they would be paired with another MTurk worker (i.e., survey participant on MTurk). They learned that although they would not be interacting with this worker in real time, they would have the ability to influence the amount of a bonus payment that both themselves and this other worker could receive. In reality, there was no other worker, and all stimuli were pre-programmed in the experiment.

Next, we implemented the experimental manipulation of perceived compassion. Participants in the high compassion and low compassion conditions learned that their partner had completed a personality test in a previous survey to assess how compassionate, extraverted, and creative they are, and that they (participants) would view these results. Those in the high compassion condition then saw that the worker scored relatively high in compassion—a score of 90/100. Those in the low compassion condition observed that the worker was relatively low in compassion (9/100). Participants in both conditions saw that the worker had moderate levels of extraversion and creativity (55/100 in extraversion, 49/500 in creativity). We included these other personality dimensions in order to obscure the focal research question and thereby reduce potential demand effects. Participants also viewed commentary from the worker in response to their scores (see Figure 1). Those in the control condition learned of no such personality test and instead advanced to the next part of the study (the trust game, as described below). The control condition thus assessed benevolence-based trust of an anonymous individual without any prior knowledge of that individual.

A

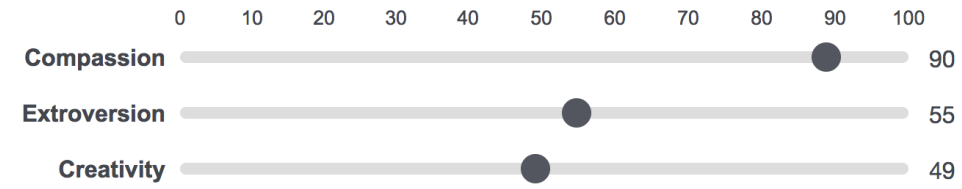

The score on each item represents how well that personality trait describes the test taker's personality. For example, a score of 80 on Extroversion means that person is more extroverted than 80% of the general population in the US.

Comment from the Worker:

"I think my scores are on point! I care a lot about people's feelings. I can't help but to sympathize with those who are going through hardship and misfortunes. I want to ease their suffering.

I think I'm relatively extroverted, but only when I'm in a small group. Creativity is important to me, so I always try to learn something new to give me new ideas."

B

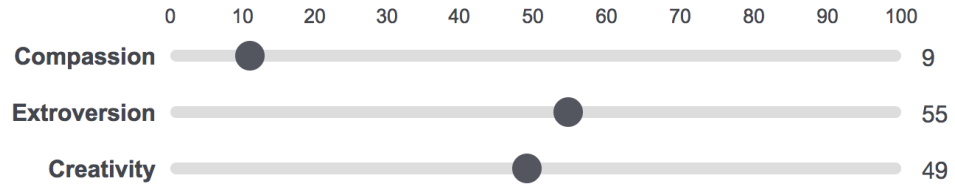

The score on each item represents how well that personality trait describes the test taker's personality. For example, a score of 80 on Extroversion means that person is more extroverted than 80% of the general population in the US.

Comment from the Worker:

"I think my scores are on point! I'm not particularly sensitive to people's feelings. Some of my friends had told me to sympathize with others more.

I think I'm relatively extroverted, but only when I'm in a small group. Creativity is important to me, so I always try to learn something new to give me new ideas."

Figure 1. In Experiment 1, participants in the high compassion (Panel A) and low compassion (Panel B) conditions viewed personality information about a worker with whom they would later play a trust game. Participants in the control condition received no information about a personality test.

Measure of benevolence-based trust. After the manipulation, we assessed participants' benevolence-based trust of the worker using a modified trust game (Berg et al., 1995). Participants were told that they were allocated a bonus of $10 and that they could choose to share some, none, or all of this money in $1 increments with the worker. Any amount they chose to send would be tripled, and the other worker would have the opportunity in a later survey to decide whether and how much to send back. This game measures benevolence-based trust; notably, there is no option to be honest or deceitful (which would implicate integrity-based trust). Rather, the decision to share money the game involves a willingness to make oneself vulnerable to another based on expectations of that person's benevolence (Levine & Schweitzer, 2015).

Participants indicated how much money ($0-$10) they chose to send to the worker. We informed them that 5% of participants would be randomly selected to actually receive the amount that the worker decided to send back. In actuality, those selected for the bonus received a randomly chosen amount from the possible amounts the worker could have sent back, depending on the amount participants sent originally.

Manipulation check. Lastly, we included a manipulation check of perceived compassion. We asked participants, "How compassionate is the Worker" (1 = *Not at all*, 7 = *Very much so*).[1] Demographics and attention/comprehension checks questions were also collected in all experiments but are not mentioned hereafter for brevity. Complete survey materials are posted on Open Science Framework (https://osf.io/ard8p/?view_only=d3120fc6824c4488a12fd94aa53e6de3) and further procedural details are provided in the Supplemental Material.

[1] In Experiments 1 and 2, we also asked about the perceived extroversion and creativity of the partner. Results with these items are available in the Supplemental Material.

Results and Discussion

Manipulation check. A one-way ANOVA revealed that the manipulation affected perceived compassion, $F(2, 259) = 193.00$, $p < .001$, $\eta_p^2 = .60$. Planned contrasts indicated that, as expected, those in the compassion condition ($M = 6.22$, $SD = 0.94$) viewed the worker as higher in compassion than those in the low compassion ($M = 2.56$, $SD = 1.47$; $p < .001$, $d = 3.00$) and control ($M = 4.44$, $SD = 1.22$; $p < .001$, $d = 1.63$) conditions. Those in the control condition also saw the worker as higher in compassion than those in the low compassion condition, $p < .001$, $d = 1.39$.

Benevolence-based trust. An ANOVA revealed a significant effect of condition, $F(2, 259) = 3.40$, $p = .035$, $\eta_p^2 = .03$. Planned contrasts indicated that participants shared more money with the high compassion worker ($M = 5.86$, $SD = 3.36$) than the low compassion worker ($M = 4.51$, $SD = 3.36$), $p = .010$, $d = 0.40$. The control worker ($M = 5.22$, $SD = 3.44$) was given an intermediate amount that did not differ significantly from that of either the high, $p = .207$, $d = 0.19$, or low compassion worker, $p = .174$, $d = 0.21$ (See Figure 2). Thus, the results of Experiment 1 showed that an individual high in compassion elicited more benevolence-based trust than an individual low in compassion.

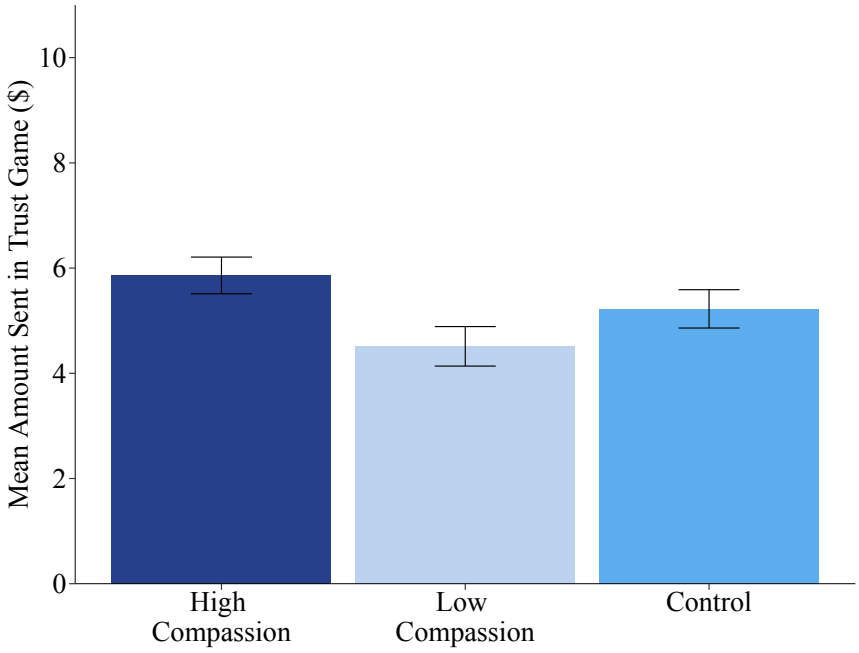

Figure 2. Effect of perceived compassion (high/low/control) on benevolence-based trust as measured by amount sent in the trust game, Experiment 1. Error bars represent standard errors.

Experiment 2

In Experiment 2, we examined how perceived compassion influences integrity-based trust. Here, we assessed participants' beliefs about whether an individual portrayed as high, moderate, or low in compassion acted honestly in an economic game. Additionally, in Experiment 2 we tested the moderating effect of benevolence-integrity conflicts, which were manipulated via payoffs in the game, as described below. We predicted that in the conflict condition, there would be a negative relationship between perceived compassion and integrity-based trust—i.e., that greater perceived compassion would lead to lower perceived honesty. In the no conflict condition, in contrast, we predicted perceived compassion would increase integrity-based trust.

Methods

Participants. Participants were 345 undergraduates at a Southwest United States university. We aimed to recruit as many participants as possible within the given laboratory time that was allotted for the study. Sensitivity analysis for a chi-squared goodness of fit test (α = 0.05) indicated that the sample size of 345 gave us 80% power to detect an effect size of w = 0.15. All participants were included in our analyses (M_{age} = 21.00, 55.36% female).

Procedure. Experiment 2 had a 3 (Between subjects: Perceived Compassion: high/moderate/low) x 2 (Within subjects: Benevolence-Integrity Conflict: conflict/no conflict) mixed design. In this version the sender-receiver game (adapted from Erat & Gneezy, 2012), participants (in the role of "Receiver") first viewed the information that an anonymous partner (in the role of "Sender") received in a previous study. Then, participants learned about the Sender's personality, which constituted the manipulation of compassion. After this, they indicated their beliefs about whether the Sender had acted honestly or dishonestly in the game and were incentivized for accuracy. As in Study 1, there was no Sender and all stimuli were pre-programmed.

Participants' task, they were told, would be to guess what decisions the Sender made in the game. We informed them that in order to make their guesses, they would first learn about the information that was ostensibly given to the Sender previously, which we then provided: (1) Both the Sender and Receiver would start out with a $1 bonus; (2) In each round of the game, a virtual coin would be flipped and the Sender would observe whether the coin landed on heads or tails; (3) The Sender would be asked to send a message to the Receiver: either "The coin landed on heads," or "The coin landed on tails"; (4) The Sender's message would determine how much additional bonus money both the Sender and Receiver could earn or lose for that round.

After reviewing the Sender's information, participants received the manipulation of perceived compassion, which was the ostensible personality test described in Experiment 1. The stimuli in the high and low compassion conditions were the same as those in Experiment 1. In the moderate condition, the Sender had a compassion score of 50/100 and creativity and extroversion scores that were identical to those in the high and low compassion conditions. The moderate condition included the same commentary from the Sender on his/her creativity and extroversion scores as did the other conditions, but no comments on the compassion scores were included.

Measure of integrity-based trust. Next, participants guessed the Sender's message in each of six rounds of the game. For each round, participants received the information that was available to the Sender, including the actual outcome of the coin flip and the amounts the Sender and the Receiver would earn depending on the Sender's message. Because participants learned the actual outcome of the coin flip, their guess about whether the Sender would say that the coin landed on heads or tails assessed their beliefs about the Sender's honesty. Thus, these guesses constituted our measure of integrity-based trust.

The manipulation of benevolence-integrity conflict was implemented via the structure of the potential payoffs in each round of the game, which are displayed in Table 1. In the no conflict rounds, the Receiver's payment remained unchanged regardless of whether the Sender was dishonest, but the Sender's payment would be increased by sending a dishonest message. Thus, the Sender's decision of whether to lie did not involve benevolence towards the Receiver, but did involve integrity. In the benevolence-integrity conflict rounds, the Sender received the same payment regardless of whether s/he was dishonest. However, a dishonest message in these rounds increased the Receiver's payment. Therefore, the Sender could not act both benevolently

(i.e., by increasing the Receiver's payoff) and with integrity.[2] Participants were told that for each correct guess of the Sender's message, they could earn an additional $0.50 bonus, with one bonus winner to be drawn from a lottery. Because there was no actual Sender in the game, we determined the bonus amount by randomly drawing from the range of potential winning amounts (maximum of $5.80).

Table 1. Potential payoffs in each round of the game in Experiment 2. In the no benevolence-integrity conflict condition, a dishonest message increased Senders' payoffs but had no effect on Receivers' payoffs. In the benevolence-integrity conflict condition, a dishonest message benefitted Receivers' payoffs but had no effect on Senders' payoffs.

Condition	Round	Message	Sender Payment ($)	Receiver Payment ($)
No Benevolence-Integrity Conflict	1	Honest	.5	.25
		Dishonest	1	.25
	2	Honest	.4	.2
		Dishonest	.8	.2
Benevolence-Integrity Conflict	3	Honest	.5	0
		Dishonest	.5	.25
	4	Honest	.4	0
		Dishonest	.4	.2
	5	Honest	.5	-.25
		Dishonest	.5	-0
	6	Honest	.4	-.2
		Dishonest	.4	-0

[2] Within the benevolence-integrity conflict rounds, we also included an exploratory manipulation of whether Sender dishonesty procured gains (2 rounds) or prevented losses (2 rounds) for the Receiver. Expected dishonesty did not differ significantly across gain/loss framing. Thus, we collapsed across the gain/loss framing into a single benevolence-integrity conflict condition. Full results are documented in the Supplemental Information.

Manipulation check. Following the measures of perceived dishonesty, we included a manipulation check that asked participants to indicate how compassionate they thought the Sender was (1 = *Not at all*, 7 = *Very much so*).

Results and Discussion

Manipulation check. A one-way ANOVA uncovered a significant effect of the perceived compassion manipulation on perceived compassion, $F(2, 340) = 341.53$, $p < .001$, $\eta_p^2 = .67$. Planned contrasts revealed that the Sender in the high compassion condition ($M = 6.23$, $SD = 0.88$) was perceived as more compassionate than that in the moderate ($M = 4.34$, $SD = 0.79$), $p < .001$, $d = 2.27$, and low compassion condition ($M = 2.75$, $SD = 1.28$), $p < .001$, $d = 3.17$, which also differed significantly from each other, $p < .001$, $d = 1.50$.

Integrity-based trust. We used a generalized linear mixed effect model (GLMM) to examine the effects of perceived compassion (between-subjects) and benevolence-integrity conflict (within-subjects) on the percentage of expect dishonest messages. Participants predicted that the Senders would send significantly more dishonest messages in the no conflict conditions ($M = 81.45\%$, $SD = 34.15\%$) than in the conflict conditions ($M = 67.03\%$, $SD = 34.86\%$), $X^2(1) = 77.71$, $p < .001$, $w = 0.47$. The main effect of perceived compassion was not significant, $X^2(1) = 4.23$, $p = .121$, $w = .01$. Importantly, the predicted interaction between perceived compassion and benevolence-integrity conflict was significant, $X^2(2) = 10.16$, $p = .006$, $w = 0.17$.

When there was a conflict between benevolence and integrity, there was a significant effect of perceived compassion, $X^2(2) = 12.24$, $p = .002$, $w = 0.19$. Consistent with our hypothesis, planned contrasts showed that Senders high in compassion ($M = 73.24\%$, $SD = 32.53\%$) were believed to send more dishonest messages than those low in compassion ($M = 63.48\%$, $SD = 35.71\%$), $p = .002$, and moderate compassion conditions ($M = 64.44\%$, $SD =$

35.68%), $p = .004$. The low and moderate compassion conditions did not differ in perceived dishonesty, $p = .700$.

There was also a significant effect of perceived compassion in the no conflict conditions, $X^2(2) = 29.35$, $p < .001$, $w = 0.29$. Confirming our hypothesis, planned contrasts indicated that Senders high in compassion ($M = 71.49\%$, $SD = 39.89\%$) were believed to send significantly fewer dishonest messages than those low in compassion ($M = 90.87\%$, $SD = 23.49\%$), $p < .001$. Senders of moderate compassion ($M = 81.90\%$, $SD = 34.51\ \%$) were predicted to send significantly more dishonest messages than those high in compassion, $p = .009$, and significantly less than those low in compassion, $p = .006$ (see Figure 3). A regression table is reported in the Supplemental Material.

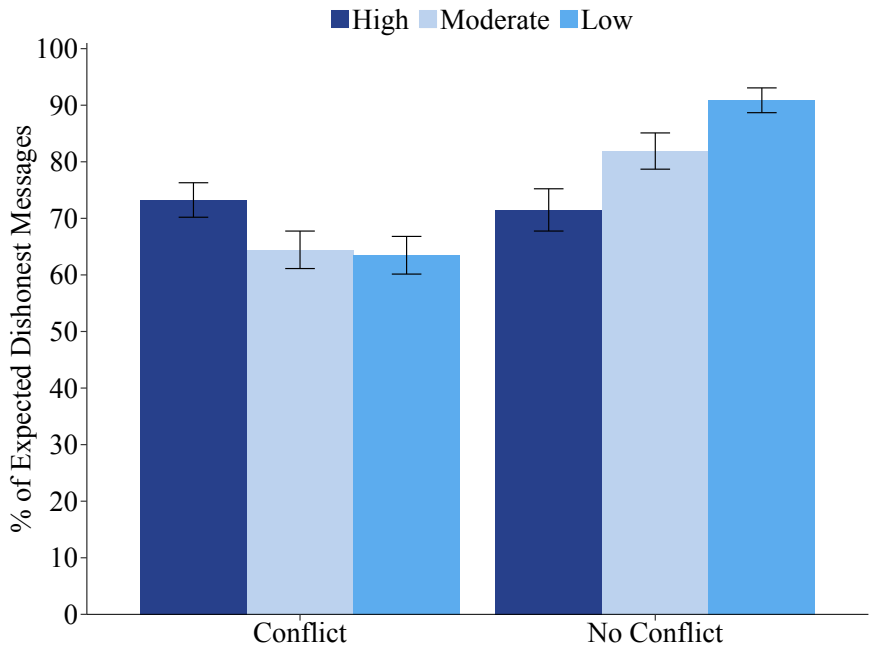

Figure 3. Effect of perceived compassion (high/moderate/low) and benevolence-integrity conflict (conflict/no conflict) on integrity-based trust as measured by expected dishonesty, Experiment 2. Error bars represent standard errors.

To sum up, Experiment 2 provided evidence that when individuals face a conflict between benevolence and integrity, perceived compassion reduces integrity-based trust of those individuals. In an economic game where being dishonest (i.e., violating integrity) resulted in economic gains for others (i.e., behaving benevolently), people expected an individual high in compassion to be less honest than a target low in compassion. In contrast, when dishonesty did not affect another's payment (i.e., when there was no benevolence-integrity conflict), compassionate individuals were thought to be more honest than those low in compassion.

Experiment 3

In Experiments 1 and 2, we used economic games to measure trust. While these games model decisions to trust in real world contexts (Murnighan & Wang, 2016), in Experiment 3, we sought to improve confidence in the generalizability of the effect. Thus, in this experiment, we explored the effect of perceived compassion on integrity-based trust in several benevolence-integrity conflicts depicted in realistic vignettes.

Methods

Participants. Three hundred sixty-nine participants were recruited on MTurk for this 2 (Perceived Compassion: compassionate/control) x 3 (Vignette: diagnosis/feedback/recommendation) mixed design. Three participants were excluded for failing an attention check, leaving a final sample of 366 (M_{age} = 37.90, 45.63% female). A sensitivity analysis (α = 0.05; ANOVA: repeated measures, between factors; correlation among repeated measures = 0.26) using G*Power indicated that the sample size of 366 gave us 80% power to detect an effect size of f = 0.14.

Procedure. Participants each read three hypothetical vignettes depicting a target who was presented with a realistic benevolence-integrity conflict. We manipulated perceived compassion

of the target between-subjects. Half of participants read vignettes in which the target was portrayed as compassionate; the other half received no information about the target's personality as a control.[3]

In the diagnosis vignette, participants read about a doctor, John, who faced a difficult decision about how to deliver a diagnosis of a terminal illness to a patient. In the compassionate condition, participants learned that "John is a doctor who is known to be a highly compassionate person. He is very sympathetic to his patients and cares a lot about easing their suffering." In the control condition, participants read only that "John is a doctor." The other vignettes depicted an employee charged with giving feedback to a poorly performing coworker (feedback vignette) and a professor who was asked to write a recommendation letter for a student (recommendation vignette). Complete vignettes are reprinted in the Appendix. All vignettes were displayed in randomized order.

Measure of integrity-based trust. We operationalized integrity-based trust as expected dishonesty, which we measured by assessing the extent to which the target was predicted to deviate from honest communication. For example, in the diagnosis vignette, we asked participants to predict how positive or negative John's diagnosis would be compared to what he truly believes (1 = *much more negatively*, 7 = *much more positively*). A score of 4 indicated an expected honest response—that "John will describe the diagnosis *exactly as positively or negatively* as what he truly believes"; scores above 4 indicated that John would be expected to give much more good news to the patient than is warranted. Participants responded on similar 7-

[3] We also manipulated the cost of dishonesty for the recipient of the lie (high/low cost) as an exploratory moderator. However, there were no significant effects or interactions with this manipulation. Thus, we collapsed across the two cost conditions. Full results are reported in the Supplemental Material.

point scales for the other vignettes to indicate expected dishonesty. For ease of interpretability, we subtracted participants' expected dishonesty scores by four. This way, when participants predicted that the target would be completely honest, the expected dishonesty score would be zero. When participants predicted the target would positively bias their communication (i.e., deliver overly good news), the expected dishonesty score would be positive.

Manipulation check. After indicating the perceived dishonesty of the target in each vignette, participants rated the degree to which "[Target] is a very compassionate person" (1= *Strongly disagree*, 7 = *Strongly agree*) as a manipulation check.

Results and Discussion

Manipulation check. A 2 (Perceived Compassion: compassionate/control) x 3 (Vignette: diagnosis/feedback/recommendation) mixed ANOVA on perceived compassion revealed a main effect of the perceived compassion manipulation, $F(1,364) = 11.54$, $p < .001$, $\eta_p^2 = .19$. The target in the compassionate condition ($M = 6.23$, $SD = 0.97$) was perceived as more compassionate that than in the control condition ($M = 5.33$, $SD = 0.92$), $d = 0.96$. There was no main effect of vignette, $F(2,728) = 0.96$, $p = .39$, $\eta_p^2 = .002$. There was a significant interaction between perceived compassion and vignette, $F(2,728) = 3.76$, $p = .024$, $\eta_p^2 = .03$, although the compassion manipulation was effective in all three vignettes, ps $< .001$, ds > 0.86 (full statistics are reported in the Supplemental Material for brevity).

Integrity-based trust. A 2x3 mixed ANOVA on expected dishonesty revealed a main effect of compassion, $F(1, 364) = 8.29$, $p = .004$, $\eta_p^2 = .01$. As expected, participants perceived the compassionate target as more dishonest ($M = 0.69$, $SD = 1.13$) than the control target ($M = 0.45$, $SD = 1.16$), $d = 0.21$. There was also a main effect of vignette, $F(2, 728) = 14.86$, $p < .001$, $\eta_p^2 = .02$ (full statistics reported in the Supplemental Material for brevity). There was no

compassion by vignette interaction, $F(2,728) = 0.39$, $p = .676$, $\eta_p^2 = .001$. These results are displayed in Figure 4.

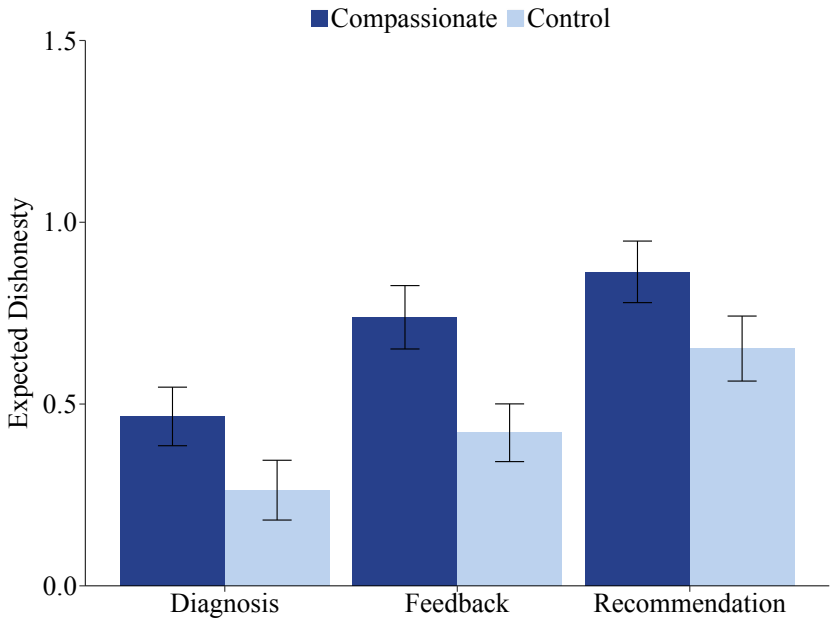

Figure 4. Effect of perceived compassion (compassionate/control) on integrity-based trust as measured by expected dishonesty across vignettes (diagnosis/feedback/recommendation), Experiment 3. Error bars represent standard errors.

Experiment 3 provided further evidence for the negative effect of perceived compassion on integrity-based trust. Although we caution that this effect was found in hypothetical vignettes, these vignettes portrayed common, realistic benevolence-integrity conflicts. When confronting these conflicts, compassionate individuals were thought to be more dishonest compared to when no information about whether target individuals were compassionate was provided.

Experiment 4

Experiment 4 aimed to extend the findings of Experiments 1-3 in several ways. In Experiments 2 and 3, we presented evidence that perceived compassion can reduce integrity-based trust in the form of anticipated dishonesty. However, integrity also encompasses principles

other than honesty, such as fairness (Mayer et al., 1995). Thus, in Experiment 4, we broadened our scope by testing whether perceived compassion decreases perceived fairness during benevolence-integrity conflicts. Additionally, in Experiments 1-3, we manipulated perceived compassion by giving participants direct insight into targets' trait compassion. In the real world, however, people often make inferences about others' traits based on observations of their behavior. Therefore, in Experiment 4, we manipulated perceived compassion in a target's self-reported account of a recent experience.

In this experiment, we simulated a workplace scenario involving employee evaluations of managers' fairness. Here, participants decided whether a compassionate manager or a neutral control manager would be more likely to fairly allocate a bonus from a zero-sum pool of money between two employees—the participant and a fellow MTurk worker—who had completed identical tasks. We manipulated benevolence-integrity conflict by varying whether the other MTurk worker was depicted as suffering. In this experimental context, the only way the suffering person could be helped is by the managers giving him/her a larger allocation of money. Thus, observing suffering should heighten the desire to act benevolently by allocating the sufferer more money, while fairness concerns would dictate an even allocation of money for equal work across employees. We hypothesized that being confronted with this benevolence-integrity conflict would decrease the perceived fairness of the compassionate manager relative to the control manager.

Participants. Participants were recruited on MTurk. We received 236 complete responses for this two-cell between-subjects design (benevolence-integrity conflict/no conflict). Thirteen participants were excluded for failing either of two attention checks, leaving a final sample of 223 (M_{age} = 35.32, 41.70% female). A sensitivity analysis (α = 0.05; chi-squared

goodness of fit test) using G*Power indicated that the sample size of 223 gave us 80% power to detect an effect size of $w = 0.19$.

Procedure. Participants first learned that they were assigned to a four-person group with three other MTurk workers. In this group, participants would play the role of Employee A, along with another employee (Employee B) and two managers (Managers A and B). In reality, there were no other workers.

Next, participants were told that "we are asking each group member to share with each other a recent meaningful experience in your life to help you get to know each other." They were also informed that the other three members of the group had already written about their recent experiences, which would be shared with participants. Participants then wrote about their recent meaningful experience with a reminder that this would be shared with their group members.

After this, participants read Employee B's recent experience, which constituted the between-subjects manipulation of benevolence-integrity conflict. Those in the conflict condition read a recent experience from Employee B that depicted him/her suffering after the death of a family member. This paragraph has been shown in previous work to elicit compassion (Lupoli et al., 2017). In contrast, those in the no conflict condition read about comparatively neutral passage about Employee B's recent trip to the zoo.

Then, participants read about both Managers A and B's recent experiences. Manager A (hereafter "the compassionate manager (A)") described the experience of being moved following witnessing homelessness. The experience of Manager B (hereafter "the control manager (B)") was comparatively neutral; this paragraph described the experience of trying yoga. All manipulations are reprinted in the survey materials on Open Science Framework.

Measures of integrity-based trust. Participants were then told that in a subsequent survey, both managers would view Employee A and B's recent experiences and then decide how to divide a $10 bonus between the two employees. We informed participants that since they and Employee B had completed identical tasks (i.e., a filler task prior to learning about the group task), "the fair allocation is to give each of you a $5 bonus." Participants then selected the manager they thought "is more likely to make a fair allocation of the $10 bonus between you and Employee B (that is, giving each of you $5)" and were incentivized for accuracy.[4] This choice constituted our dichotomous dependent measure of integrity-based trust.

Following this decision, we included a continuous measure of integrity-based trust. Participants responded to two questions assessing their best guess about how each manager would allocate the $10 between themselves and Employee B. These questions were 11-point continuous scales that specified each potential allocation between the two employees (i.e., 0 = *$0 to me, $10 to Employee B*; 1 = *$1 to me, $9 to Employee B*; etc.). A score of 5 on this scale indicated that an even allocation was expected (i.e., *$5 to me, $5 to Employee B*), while lower scores indicated the expectation that the suffering Employee B would be favored.

Manipulation checks. Participants then responded to manipulation check questions in which they evaluated their group members on several adjectives (1 = *Not at all*, 7 = *Very much*). For the manipulation check of benevolence-integrity conflict, participants rated Employee B on the items "in need," "suffering," and "vulnerable" (α = .88). They also indicated how compassionate they thought both the compassionate manager (A) (r = 0.80) and the control

[4] To incentivize accuracy, participants learned that one randomly selected participant would receive the average of the managers' allocations, along with an additional $1 bonus if the manager they selected made a fair (i.e. even split) allocation. A randomly selected participant was paid a bonus of $6 (one of the possible payment amounts that could result in the procedure).

manager (B) ($r = 0.69$) were using the items "compassionate" and "sympathetic." Additionally, participants indicated several other perceptions of their group members as robustness checks. For brevity, descriptions and analyses with these items are included in the Supplemental Material.

Results and Discussion

Manipulation check. A *t*-test indicated that those in the conflict condition ($M = 5.90$, $SD = 0.96$) reported greater levels of suffering in Employee B compared with those in the no conflict condition ($M = 2.98$, $SD = 1.66$), $t(221) = 15.85$, $p < .001$, $d = 2.12$. Furthermore, a paired *t*-test showed that participants rated the compassionate manager (A) ($M = 6.37$, $SD = 0.88$) as more compassionate than the control manager (B) ($M = 4.63$, $SD = 1.00$), $t(222) = 22.41$, $p < .001$, $d = 1.50$.

Integrity-based trust, dichotomous measure. In the no conflict condition, 83.62% of participants indicated that the compassionate manager (A) was more likely to make a fair allocation than the control manager (B). However, in the conflict condition (i.e., when Employee B was suffering), only 70.10% of participants reported that the compassionate manager (A) would make the fair allocation. A chi-square test revealed that this difference was significant, $X^2(1) = 5.03$, $w = 0.15$, $p = .025$. These results are displayed in Table 2.

Integrity-based trust, continuous measure. A 2 (between subjects: conflict/no conflict) x 2 (within-subjects: compassionate manager A/control manager B) mixed ANOVA revealed a main effect of condition, $F(1, 115) = 14.94$, $p < .001$, $\eta_p^2 = .12$. Participants in the conflict condition ($M = 4.15$, $SD = 1.82$) expected a less fair allocation than those in the no conflict condition ($M = 4.92$, $SD = 1.74$), $d = 0.43$. There was no effect of manager, $F(1, 115) = 0.24$, $p = .628$, $\eta_p^2 = .002$. However, there was a significant interaction, $F(1, 115) = 6.20$, $p = .014$, $\eta_p^2 = .05$. The difference in the compassionate manager A's expected allocation across conflict and no

conflict condition, $p < .001$, $d = 0.60$, was larger than the difference in control manager B's expected allocations across conflict and no conflict condition, $p = .049$, $d = 0.27$. Means and standard deviations are displayed in Table 2.

Table 2: Effect of benevolence-integrity conflict on integrity-based trust, Experiment 4.

	Dichotomous Measure		**Continuous Measure**	
	Conflict	**No Conflict**	**Conflict**	**No Conflict**
Compassionate Manager (A)	75 (70.10%)	97 (83.62%)	4.03 (1.95)	5.09 (1.59)
Control Manager (B)	32 (29.90%)	19 (16.38%)	4.28 (1.66)	4.75 (1.86)

Note. Dichotomous measure indicates the frequency and percentage (in parentheses) of participants believing that either manger (compassionate or control) is more likely to make a fair allocation. Continuous measure indicates the mean and standard deviation (in parentheses) of each manager's expected allocation to Employees A and B, where a mean of 5 indicates a fair allocation and < 5 indicates favoring the suffering Employee B.

In Experiment 4, we explored the relationship between perceived compassion and another instantiation of integrity-based trust: perceived fairness. We note that unlike the results of Study 2, the compassionate manager was trusted more than the control manager in both conflict and no conflict conditions. However, consistent with our expectations, perceived fairness of the compassionate manager significantly decreased when one employee was suffering—a situation that presented a benevolence-integrity conflict for the managers.

Experiment 5

In Experiment 5, we investigated potential downstream consequences of the effects of perceived compassion and benevolence-integrity conflicts on trust. Here, participants engaged in a mock hiring decision in which they were asked to evaluate the suitability of a job candidate. According to our theoretical framework, we would not expect compassionate individuals to be seen as less suitable for all positions. Rather, we hypothesized that compassionate people should

only be trusted less in positions that may require navigating benevolence-integrity conflicts. Here, we manipulated whether the job would entail benevolence-integrity conflicts by varying the extent to which the position involved giving feedback. We operationalized benevolence-integrity conflict in this way because feedback provision is common context in which such conflicts present themselves; a compassionate individual charged with giving feedback to others may be tempted to give overly-positive feedback so as to avoid causing emotional harm (Lupoli et al., 2017).

Additionally, in Experiment 5 we included a new control condition: perceived warmth. Most broadly, warmth involves the perception that others have prosocial intentions (Fiske, Cuddy, & Glick, 2007). Although the experience of compassion involves prosocial intentions, it also entails being emotionally driven to mollify the suffering of others—a feature not necessarily shared with warmth (Goodwin, Piazza, & Rozin, 2014). Thus, including warmth as a control allows us to determine if the effect is unique to perceived compassion rather than to other prosocial traits.

Furthermore, in Experiment 5 we assessed potential mechanisms. Given the results of Studies 1-4, it seems likely that perceivers make inferences about the mental processes of compassionate individuals that render these individuals less trustworthy during benevolence-integrity conflicts. Specifically, we expected that compassionate individuals would be seen to (a) place a higher importance on benevolence relative to integrity and (b) be more focused on creating a positive impression on others. However, we thought that either or both these perceptions would reduce trust only when compassionate individuals faced a benevolence-integrity conflict (versus no conflict).

Methods

Participants. We received 602 complete responses on MTurk for this 3 (Emotion: compassion/warmth/control) x 2 (Benevolence-Integrity Conflict: conflict/no conflict) between-subjects design. Fifty-four participants failed at least one of two attention checks throughout the study and were thus excluded. This left a final sample of 548 participants (M_{age} = 38.78, 41.97% female). A sensitivity analysis (α = 0.05; ANOVA: fixed effects, special, main effects and interactions) using G*Power indicated that this sample size gave us 80% power to detect an effect size of f = 0.13.

Procedure. Participants were first asked to imagine that they were a member of the hiring committee in an organization that is seeking to hire for the position of Web Manager. Their task, they learned, would be to evaluate the suitability of a candidate for the position. Participants read a job description for the advertised position, then viewed the resume of a candidate, along with that candidate's written response to a question that was asked during the application process (those in the neutral condition received no application question or written response). Finally, participants evaluated the suitability of the candidate and answered follow-up questions.

Those in the conflict condition read a description for a Web Manager position that discussed how the job would involve giving frequent feedback. It was made clear that the Web Managers would need to "communicate with programmers on a daily basis about how they can improve and that "it is essential that you give these programmers accurate and honest feedback." Those in the no conflict condition read a job description explicitly stating that although the job would involve working alongside other programmers, giving feedback would not be a part of the position (see survey materials on Open Science Framework for full descriptions).

After reading one of the two randomly assigned job descriptions, participants viewed the resume of the candidate. All participants saw the same resume, which was fictitious but realistic. All identifying information of the candidate was redacted so participants would not be influenced by the perceived gender, race, or socioeconomic status of the candidate.

Participants in the compassion and warmth conditions also saw the candidate's answer to an application question, which asked, "In 2-3 sentences, please tell us something about who you are as a person, above and beyond what is shown on your resume." Those in the neutral control condition saw only the candidate's resume. The candidate's response to this prompt constituted our emotion manipulation. Those in the compassion condition read the following response: "I consider myself to be a compassionate person. When I see or hear about people who are suffering, it evokes a powerful feeling of sympathy within me. One of my most sacred values is to prevent others from being harmed." Those in the warmth condition read: "I am a social person overall. I love working with and interacting with others, and generally feel at my best in the company of others. Being outgoing and friendly is an important part of who I am." Although scholars have debated the traits that constitute warmth (e.g., Kervyn, Fiske, & Yzerbyt, 2013), one influential model posits that there are two dimensions of warmth: a moral component, which includes traits such as kindness, lovingness, and trustworthiness, and a non-moral component, which includes being sociable, extroverted, and gregarious (Goodwin et al., 2014). To avoid potential overlap with either perceived compassion or trustworthiness, we manipulated the non-moral aspect of warmth in the perceived warmth condition. We note that this operationalization of warmth is a particularly well-suited control condition, as it seems plausible that an individual viewed as warm in this way may also be considered unlikely to give negative feedback.

Measure of candidate suitability. Next, participants evaluated the candidate by indicating the extent to which they thought the candidate "is suitable for the position"; "would likely have problems meeting the requirements of the position" (reverse scored); "would be likely to succeed at the position" (1 = *Strongly disagree*, 7 = *Strongly agree*; α = .78).

Mechanisms. After this, participants answered questions to measure potential mechanisms (1 = *Strongly disagree*, 7 = *Strongly agree*). The perceived importance the candidate placed on benevolence was measured with two items: "Preventing harm to others is important to Candidate 1"; "Not hurting others is important to Candidate 1" (r = .79). Importance of integrity was also measured with two items: "Integrity is important to candidate 1"; "Honesty is important to Candidate 1" (r = .74). We subtracted integrity scores from honesty scores to form a measure of the extent to which the candidate was believed to value benevolence over integrity. Additionally, we measured impression management concerns with two items: "Candidate 1 probably cares a lot about what others think about him/her"; "Candidate 1 is concerned with getting others to like him/her" (r = .65; adapted from Scheier & Carver, 1985).

Manipulation checks. For the manipulation checks, participants were then given the prompt, "To what extent is Candidate 1," and rated the candidate on a number of adjectives (1 = *Not at all*, 7 = *Very much*). The item "compassionate/sympathetic" constituted the manipulation check of perceived compassion. The perceived warmth manipulation check comprised the item "warm/sociable/extroverted." Additionally, the manipulation check of benevolence-integrity conflict included two items: "Performing the required duties for this position could potentially cause harm to others," and "Performing the required duties for this position could potentially make others feel badly" (1 = *Strongly disagree*, 7 = *Strongly agree; r* = .77). As in Study 4, we

also measured other perceptions of the candidate as robustness checks. Descriptions and analyses with these items are reported in the Supplemental Material for brevity.

Results and Discussion

Manipulation checks. Means and standard deviations for manipulation checks across emotion conditions are displayed in Table 3. A 3 (Emotion: compassion/warmth/control) x 2 (Benevolence-Integrity Conflict: conflict/no conflict) ANOVA on perceived compassion revealed a main effect of the emotion manipulation, $F(2, 542) = 123.58$, $p < .001$, $\eta_p^2 = .31$. A Tukey HSD test indicated that the candidate in the compassion condition was seen as more compassionate than the candidate in the warmth, $p < .001$, $d = 1.02$, and control condition, $p < .001$, $d = 1.65$, which also differed significantly from each other, $p < .001$, $d = 0.64$. Similarly, a 3x2 ANOVA indicated a main effect of the emotion manipulation on perceived warmth, $F(2, 542) = 85.27$, $p < .001$, $\eta_p^2 = .24$. The candidate in the warmth condition was believed to be more warm than that in the compassion, $p < .001$, $d = 0.47$, and control, $p < .001$, $d = 1.31$, which differed significantly from each other, $p < .001$, $d = 0.85$. There were no significant main effects of conflict or interactions for either perceived compassion or perceived warmth, ps > .200.

Furthermore, there was a main effect of conflict on the belief that the position could potentially involve hurting others, $F(1, 542) = 178.18$, $p < .001$, $\eta_p^2 = .24$. Participants in the conflict condition ($M = 3.86$, $SD = 1.77$) thought that the performing the duties of the position might involve hurting others to a greater extent than those in the no conflict condition ($M = 2.02$, $SD = 1.39$), $d = 1.15$. There was also an unpredicted main effect of emotion, $F(2, 542) = 5.18$, $p = .007$, $\eta_p^2 = .02$, as shown in Table 3 and reported further in the Supplemental Information. There was no significant interaction, $F(2, 542) = 1.74$, $p = .176$, $\eta_p^2 = .006$.

Table 3: Effect of emotion condition (compassion/warmth/control) on compassion, warmth, and benevolence-integrity conflict manipulation checks, Experiment 5.

Manipulation Check	Emotion Condition		
	Compassion	**Warmth**	**Control**
Compassion Check	6.56^{a} (0.77)	5.66^{b} (0.98)	4.98^{c} (1.10)
Warmth Check	5.80^{b} (0.98)	6.25^{a} (0.93)	4.90^{c} (1.12)
Benevolence-Integrity Conflict Check	3.19^{a} (2.04)	3.00^{ab} (1.84)	2.67^{b} (1.60)

Note. Numbers indicate means and standard deviation (in parentheses). Different letter superscripts within rows indicate significant differences ($p < .05$) across emotion condition (compassion/warmth/control) as determined by Tukey HSD tests.

Candidate suitability. A 3x2 ANOVA on candidate suitability indicated main effect of emotion, $F(2, 542) = 6.37$, $p = .002$, $= .02$, such that the warm ($M = 5.85$, $SD = 1.04$), $p = .009$, $d = 0.29$, and control candidate ($M = 5.85$, $SD = 0.97$), $p = .009$, $d = 0.30$, were seen as more suitable for the position than the compassionate candidate ($M = 5.50$, $SD = 1.33$), but did not differ from each other, $p = 1.00$. $d = .002$ (Tukey HSD tests). There was also a main effect of conflict, $F(1, 542) = 27.06$, $p < .001$, $\eta_p^2 = .05$, such that the candidate was rated as less suitable when there was a benevolence-integrity conflict ($M = 5.49$, $SD = 1.26$) versus no conflict ($M = 5.98$, $SD = 0.92$), $d = 0.45$. Most importantly, there was also a significant interaction, $F(2, 542) = 10.01$, $p < .001$, $\eta_p^2 = .04$. In the benevolence-integrity conflict condition, there was a significant effect of emotion, $F(2, 273) = 11.58$, $p < .001$, $\eta_p^2 = .08$. The compassionate candidate ($M = 5.01$, $SD = 1.46$) was viewed as less suitable for the job than both the warm candidate ($M = 5.68$, $SD = 1.06$), $p < .001$, $d = 0.53$, and the control candidate, ($M = 5.80$, $SD = 1.06$), $p < .001$, $d = 0.62$, who did not differ from each other, $p = .786$, $d = 0.11$ (Tukey HSD tests). However, in the no conflict condition, there was no effect of emotion, $F(2, 269) = 1.06$, $p = .348$, $\eta_p^2 = .008$. These results are depicted in Figure 5.

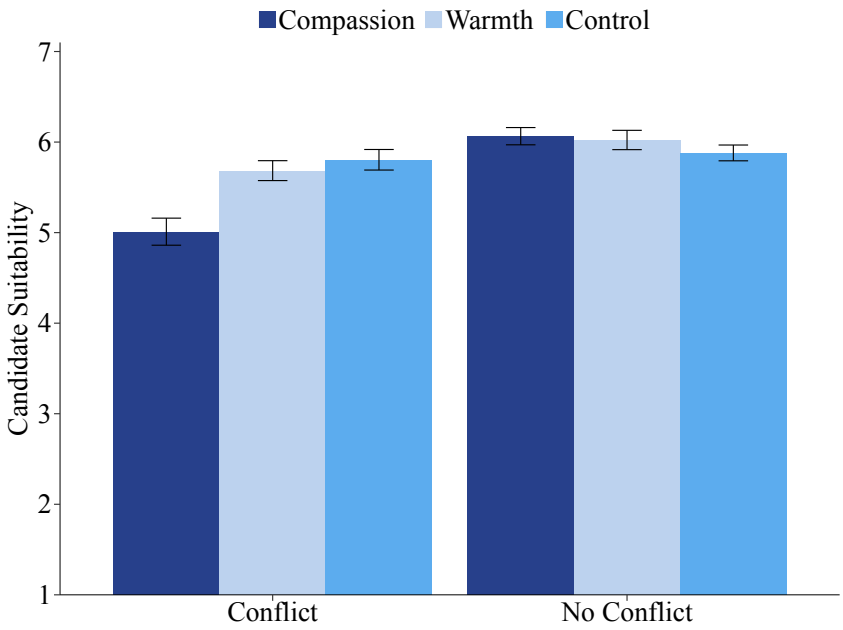

Figure 5. The effect of emotion (compassion/warmth/control) and benevolence-integrity conflict (conflict/no conflict) on candidate suitability ratings, Experiment 5.

Mechanisms. To assess mechanism, we ran a moderated mediation model with PROCESS in SPSS (Hayes, 2016; Model 14; moderator affects mediator to DV path). Emotion was entered as the independent variable, candidate suitability as the dependent variable, and benevolence-integrity conflict as the moderator. Perceived importance of benevolence relative to integrity and perceived importance of impression management were entered as mediators. Emotion was dummy coded (compassion = 1, warmth = 0, control = 0) to examine the effect of compassion relative to the mean of the warmth and control conditions.

This analysis revealed significant moderated mediation for the relative importance on benevolence versus integrity measure (95% CI [-.67, -.21]). When the job description did not involve giving feedback—that is, when there was no benevolence-integrity conflict—there was no significant indirect through the perceived relative importance of benevolence versus integrity (95% CI [-.03, .25]). However, when the job entailed giving feedback (benevolence-integrity

conflict condition), there was a significant indirect effect through this measure (95% CI [-.52, -.14]). The compassionate candidate was believed to place a higher importance on benevolence relative to integrity, $B = 1.12$, $p < .001$, and this belief in turn decreased ratings of the candidate's suitability in the conflict condition, $B = -0.29$, $p < .001$. There was no moderated mediation (95% CI [-.09, 02]) or any significant indirect effect for the perceived importance of impression management (conflict: 95% CI [-.08, 01]; no conflict: 95% CI [-.03, 03]).

The results of Experiment 5 highlight potential downstream consequences of the effect of perceived compassion on trust in an organizational context. A candidate high in compassion was seen as equally suitable for a job that did not involve giving feedback as a candidate high in warmth as well as a control candidate. However, when the job entailed potential conflicts between benevolence and integrity—that is, a requirement to give frequent feedback—the compassionate candidate was rated as less suitable than both warm and control candidates. We also obtained evidence for a mechanism in this study: The negative effect of perceived compassion on candidate suitability under a benevolence-integrity conflict (vs. no conflict) was driven in part by the belief that the compassionate candidate placed a higher importance on benevolence relative to integrity. Additionally, we provided evidence against an alternative explanation, which posited that this effect could be driven by beliefs that compassionate individuals are more concerned with impression management. In sum, these results cast light on lay beliefs about the values compassionate individuals use to navigate benevolence-integrity conflicts, as well as how these beliefs affect trust of those individuals.

General Discussion

Across five experiments implementing a variety of designs, we demonstrated that despite compassion's benefits, being perceived as compassionate can sometimes harm trust. While

perceived compassion did increase trust in others' benevolence (Experiment 1) and integrity (Experiment 2) when there was no conflict of values, compassionate individuals were believed to be less honest when they faced conflicts between benevolence and integrity (Experiments 2, 3, 5). Benevolence-integrity conflict also decreased the expected fairness compassionate individuals, though they were still viewed as more fair than control individuals (Experiment 4). We also obtained evidence for a mechanism behind the negative effect of perceived compassion on trust during benevolence-integrity conflicts: the belief that compassionate individuals place a higher value on benevolence relative to integrity (Experiment 5). Finally, we ruled out an alternative explanation: beliefs that compassionate individuals are more concerned with impression management did not drive this effect (Experiment 5).

Theoretical Implications

This work makes several theoretical contributions. First, this research extends literature highlighting potentially harmful effects of compassion. Past studies have identified negative effects that compassion can sometimes have on others, such as increasing dishonesty (Lupoli et al., 2017) and biasing people towards helping identifiable victims rather than victims in greater need (Loewenstein & Small, 2007; Slovic, 2007). In contrast, our research demonstrates for the first time an undesirable effect that being compassionate can have on the compassionate individuals themselves—that is, decreasing their trustworthiness in the eyes of others. Although we do not dispute the many prosocial effects of this emotion (e.g., Batson & Shaw, 1991; Dutton, Workman, & Hardin, 2014; Dutton, Worline, Frost, & Lilius, 2006), and note that we do not examine trust decisions following in-person interactions, this research helps to paint a more nuanced portrait of the role of compassion in social life.

This research also contributes to moral psychology via the theoretical development of benevolence-integrity conflicts. While some studies have examined decision making under benevolence-integrity conflicts (e.g., lying for the benefit of others; Levine et al., 2015; Lupoli et al., 2017, 2018), these studies have not provided a theoretical account of these conflicts across contexts. Researchers have recently begun to address these issues by introducing the construct of benevolence-integrity conflicts and offering theoretical perspectives on how individuals might navigate these dilemmas (Moore et al., 2019). Our research adds to and goes beyond this growing body of work by (a) experimentally testing the effects of benevolence-integrity conflicts on trust across multiple manipulations and experimental contexts; and (b) adopting an alternative perspective that explores not how individuals make decisions amidst these conflicts (e.g., Lupoli et al., 2017) or how people respond to these decisions (e.g., Lupoli et al., 2018), but rather lay theories about how and why others make these decisions. This is a useful perspective because, as these studies suggest, beliefs about how others navigate benevolence-integrity conflicts can affect important decisions, such as the choice of who to turn to for advice and feedback, as well as personnel selection in organizations.

Additionally, this research builds on emotions as social information theory (Van Kleef, 2009; Van Kleef et al., 2010) and compassion research more broadly. Although researchers have begun to explore the effects of perceived discrete emotions on observers (e.g., anger: Campellone & Kring, 2013; happiness: Centorrino, Djemai, Hopfensitz, Milinski, & Seabright, 2015), the effects of detecting compassion in others are not yet fully understood. Likewise, research on compassion has largely focused on the effects that experiencing compassion can have on the experiencer's decisions and on others (e.g., Condon & DeSteno, 2011; Batson & Shaw, 1991), while not addressing the causal effects of perceived compassion on observers.

These literature gaps are problematic because, as our findings indicate, the assumption that compassionate individuals are trusted more is not always accurate. Moreover, distrust of compassionate individuals could have unintended negative consequences. If compassionate people are trusted less for jobs that involve giving negative feedback (as Experiment 5 suggests), explicitly selecting against compassion for these positions could potentially result in perverse effects, such as a loss of organizational commitment stemming from employees seeing the organization as uncaring (Grant, Dutton, & Rosso, 2008). In general, more research is necessary to understand how perceived compassion influences social cognition and behavior.

Furthermore, the current studies broaden theory on trust. While most research drawing on the ability/benevolence/integrity model of trust (Mayer et al., 1995) looks at how situational and dispositional factors influence different types of trust (e.g., Kim et al., 2004; Levine & Schweitzer, 2015; Shao, 2019), none have examined when a unique trustee attribute may increase one type of trust but decrease another. By documenting contexts in which the extending of benevolence-and integrity-based trust does not co-occur, we highlight advantages of this multifaceted view of trust. In doing so, our findings also underscore benefits of this approach over self-report measures of overall trust (e.g., those that include statements such as, “I have complete trust in [target]”; Kirkpatrick & Locke, 1996), which do not capture the level of granularity associated with behavioral measures of benevolence- and integrity-based trust.

Practical Implications

These results also have important practical implications for compassionate individuals who wish to maintain and build trust. Given the importance of trust for relationships (Rempel et al., 1985), individuals would benefit from understanding the contexts in which their compassionate expression or reputation will lead people to question their integrity. In light of our

findings, compassionate individuals might consider strategies to mitigate or avoid skepticism about their integrity when they are confronted with benevolence-integrity conflicts. For example, those asked to give constructive feedback might demonstrate commitment to integrity simply by being honest and telling hard truths. They also might explain their reasoning behind favoring honesty despite its potential emotional costs, such as a concern for the long-term interests of the feedback recipient (e.g., performance improvement; Levine et al., 2019). This work opens new avenues for future research on how to help compassionate people—and individuals more broadly—navigate these difficult situations in way that preserves both their perceived benevolence and integrity.

In addition, these results have practical implications for organizations. Both compassion (e.g., Dutton et al., 2006, 2014) and trust (e.g., De Jong et al., 2016; Dirks & Ferrin, 2001) have been found to positively affect organizations, leading to a call for more compassion in organizations (Rynes, Bartunek, Dutton, & Margolis, 2012). However, because compassion can sometimes decrease trust, managers should develop strategies to maximize the benefits of compassion in the workplace while implementing policies and practices that buffer against the loss of trust. For example, managers might foster a culture of compassion by making accommodations for employees who face difficult situations at home, especially in light of the COVID-19 crisis (Gorlick, 2020). Similarly, organizations can promote compassion when colleagues make mistakes on tasks where they lack expertise. However, they may wish to de-emphasize or qualify the role of compassion when giving feedback or providing performance reviews. It will be a challenging but worthwhile task for researchers and managers to determine how to most effectively harness compassion in organizations.

Limitations and Future Directions

One limitation of this research is that we only examined trust between strangers (or rather fictitious individuals believed to be strangers). Future research should examine the association between perceived compassion and trust in closer relationships, such as those between friends or co-workers; it is possible that social distance, knowledge of past behavior, and/or relationship length (Levin, Whitener, & Cross, 2006) moderates the relationship between perceived compassion and trust, as more knowledge of a person introduces additional social perceptions that could interact with perceived compassion. Similarly, future research would benefit from testing the effects of perceived compassion on trust in dyadic interactions to better understand factors that moderate this effect in face-to-face communication.

Conclusion

It can be argued that compassion and trust are indispensable components of healthy relationships, organizations, and society. As such, researchers and individuals should strive to understand how both compassion and trust can be fostered. This research paints a nuanced portrait to further that understanding. While perceived compassion signals benevolent intentions, being seen as compassionate does not always inspire trust.

Open Practices

All data and survey materials can be found on Open Science Framework using the following link: https://osf.io/ard8p/?view_only=d3120fc6824c4488a12fd94aa53e6de3.

References

Batson, C. D., Klein, T. R., Highberger, L., & Shaw, L. L. (1995). Immorality from empathy-induced altruism: When compassion and justice conflict. *Journal of Personality and Social Psychology*, *68*(6), 1042-1054.

Batson, C. D., & Shaw, L. L. (1991). Evidence for altruism: Toward a pluralism of prosocial motives. *Psychological Inquiry*, *2*(2), 107-122.

Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, *10*(1), 122-142.

Butler Jr, J. K., & Cantrell, R. S. (1984). A behavioral decision theory approach to modeling dyadic trust in superiors and subordinates. *Psychological Reports*, *55*(1), 19-28.

Buchan, N. R., Croson, R. T., & Solnick, S. (2008). Trust and gender: An examination of behavior and beliefs in the Investment Game. *Journal of Economic Behavior & Organization*, *68*(3-4), 466-476.

Campellone, T. & Kring, A. M. (2013). Who do you trust? The impact of facial emotion and behavior on decision-making. *Cognition and Emotion*, 27, 603-620.

Centorrino, S., Djemai, E., Hopfensitz, A., Milinski, M., & Seabright, P. (2015). Honest signaling in trust interactions: Smiles rated as genuine induce trust and signal higher earning opportunities. *Evolution and Human Behavior*, *36*(1), 8-16.

Cialdini, R., Brown, S., Lewis, B., Luce, C., & Neuberg, S. (1997). Reinterpreting the Empathy–Altruism Relationship. *Journal of Personality and Social Psychology*, *73*(3), 481-494.

Colquitt, J. A., Scott, B. A., & LePine, J. A. (2007). Trust, trustworthiness, and trust propensity: a meta-analytic test of their unique relationships with risk taking and job performance. *Journal of Applied Psychology*, *92*(4), 909-927.

Condon, P., & DeSteno, D. (2011). Compassion for one reduces punishment for another. *Journal of Experimental Social Psychology*, *47*(3), 698-701.

De Jong, B. A., Dirks, K. T., & Gillespie, N. (2016). Trust and team performance: A meta-analysis of main effects, moderators, and covariates. *Journal of Applied Psychology*, *101*(8), 1134-1150.

DePaulo, B. M., Kashy, D. A., Kirkendol, S. E., Wyer, M. M., & Epstein, J. A. (1996). Lying in everyday life. *Journal of Personality and Social Psychology*, *70*(5), 979.

DeYoung, C. G., Quilty, L. C., & Peterson, J. B. (2007). Between facets and domains: 10 aspects of the Big Five. *Journal of Personality and Social Psychology*, *93*(5), 880-896.

Dirks, K. T., & Ferrin, D. L. (2001). The role of trust in organizational settings. *Organization Science*, *12*(4), 450-467.

Dunn, J. R., & Schweitzer, M. E. (2005). Feeling and believing: the influence of emotion on trust. *Journal of Personality and Social Psychology*, *88*(5), 736-748.

Dutton, J. E., Workman, K. M., & Hardin, A. E. (2014). Compassion at work. *Annual Review Organizational Psychology and Organizational Behavior, 1*(1), 277-304.

Dutton, J. E., Worline, M. C., Frost, P. J., & Lilius, J. (2006). Explaining compassion organizing. *Administrative Science Quarterly*, *51*(1), 59-96.

Erat, S., & Gneezy, U. (2012). White lies. *Management Science*, *58*(4), 723-733.

Fiske, S. T., Cuddy, A. J., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences*, *11*(2), 77-83.

Flynn, F. J., Reagans, R. E., Amanatullah, E. T., & Ames, D. R. (2006). Helping One's Way to the Top: Self-Monitors Achieve Status by Helping Others and Knowing Who Helps Whom. *Journal of Personality and Social Psychology*, *91*(6), 1123-1137.

Gino, F., & Pierce, L. (2009). Dishonesty in the name of equity. *Psychological Science*, *20*(9), 1153-1160.

Goetz, J. L., Keltner, D., & Simon-Thomas, E. (2010). Compassion: An evolutionary analysis and empirical review. *Psychological Bulletin*, *136*(3), 351-374.

Goodwin, G., Piazza, J., & Rozin, P. (2014). Moral Character Predominates in Person Perception and Evaluation. *Journal of Personality and Social Psychology*, *106*(1), 148-168.

Gorlick, A. (2020). The productivity pitfalls of working from home in the age of COVID-19. *Stanford News*. https://news.stanford.edu/2020/03/30/productivity-pitfalls-working-home-age-covid-19/.

Grant, A. M., Dutton, J. E., & Rosso, B. D. (2008). Giving commitment: Employee support programs and the prosocial sensemaking process. *Academy of Management Journal*, *51*(5), 898-918.

Hayes, A. F. (2016). The PROCESS macro for SPSS and SAS. http://www. processmacro. org/index.html.

Hardy, C. L., & Van Vugt, M. (2006). Nice guys finish first: The competitive altruism hypothesis. *Personality and Social Psychology Bulletin*, *32*(10), 1402-1413.

Jachimowicz, J. M., Chafik, S., Munrat, S., Prabhu, J. C., & Weber, E. U. (2017). Community trust reduces myopic decisions of low-income individuals. *Proceedings of the National Academy of Sciences*, *114*(21), 5401-5406.

Kervyn, N., Fiske, S. T., & Yzerbyt, V. Y. (2013). Integrating the stereotype content model (warmth and competence) and the Osgood semantic differential (evaluation, potency, and activity). *European Journal of Social Psychology*, *43*(7), 673-681

Kim, P. H., Ferrin, D. L., Cooper, C. D., & Dirks, K. T. (2004). Removing the shadow of suspicion: the effects of apology versus denial for repairing competence-versus integrity-based trust violations. *Journal of Applied Psychology*, *89*(1), 104-118.

Kirkpatrick, S. A., & Locke, E. A. (1996). Direct and indirect effects of three core charismatic leadership components on performance and attitudes. *Journal of Applied Psychology*, *81*(1), 36-51.

Levin, D. Z., Whitener, E. M., & Cross, R. (2006). Perceived Trustworthiness of Knowledge Sources: The Moderating Impact of Relationship Length. *Journal of Applied Psychology*, *91*(5), 1163-1171.

Levine, E. E. (2020). Community standards of deception. Working Paper.

Levine, E. E., Roberts, A. R., & Cohen, T. R. (2019). Difficult conversations: navigating the tension between honesty and benevolence. *Current Opinion in Psychology*, *31*, 38-43.

Levine, E. E., & Schweitzer, M. E. (2015). Prosocial lies: When deception breeds trust. *Organizational Behavior and Human Decision Processes*, *126*, 88-106.

Loewenstein, G., & Small, D. A. (2007). The scarecrow and the tin man: The vicissitudes of human sympathy and caring. *Review of General Psychology*,*11*(2), 112-126.

Lount Jr, R. B., & Pettit, N. C. (2012). The social context of trust: The role of status. *Organizational Behavior and Human Decision Processes*, *117*(1), 15-23.

Lupoli, M. J., Jampol, L., & Oveis, C. (2017). Lying because we care: Compassion increases prosocial lying. *Journal of Experimental Psychology: General*, *146*(7), 1026-1042.

Lupoli, M. J., Levine, E. E., & Greenberg, A. E. (2018). Paternalistic lies. *Organizational Behavior and Human Decision Processes*, *146*, 31-50.

Mayer, R. C., & Davis, J. H. (1999). The effect of the performance appraisal system on trust for management: A field quasi-experiment. *Journal of Applied Psychology*, *84*(1), 123-136.

Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, *20*(3), 709-734.

Moore, A. K., Munguia Gomez, D. M., & Levine, E. E. (2019). Everyday dilemmas: New directions on the judgment and resolution of benevolence–integrity dilemmas. *Social and Personality Psychology Compass*, e12472.

Murnighan, J. K., & Wang, L. (2016). The social world as an experimental game. *Organizational Behavior and Human Decision Processes, 136*(C), 89-94

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, *84*(3), 231-259.

Omoto A. M., Malsch, A. M., & Barraza J. A. (2009). Compassionate acts: Motivations for and correlates of volunteerism among older adults. In B. Fehr, S. Sprecher, & L.G. Underwood (Eds.), *The Science of Compassionate Love: Theory, Research, and Applications* (pp. 257–282). Malden, MA: Wiley-Blackwell.

Rempel, J. K., Holmes, J. G., & Zanna, M. P. (1985). Trust in Close Relationships. *Journal of Personality and Social Psychology*, *49*(1), 95-112.

Rousseau, D. M., Sitkin, S. B., Burt, R. S., & Camerer, C. (1998). Not so different after all: A cross-discipline view of trust. *Academy of Management Review*, *23*(3), 393-404.

Rynes, S. L., Bartunek, J. M., Dutton, J. E., & Margolis, J. D. (2012). Care and compassion through an organizational lens: Opening up new possibilities. *The Academy of Management Review*. *37*(4), 503-523.

Saslow, L. R., Willer, R., Feinberg, M., Piff, P. K., Clark, K., Keltner, D., & Saturn, S. R. (2013). My brother's keeper? Compassion predicts generosity more among less religious individuals. *Social Psychological and Personality Science*, *4*(1), 31-38.

Scheier, M. F., & Carver, C. S. (1985). The Self-Consciousness Scale: A Revised Version for Use with General Populations 1. *Journal of Applied Social Psychology*, *15*(8), 687-699.

Schweitzer, M. E., Hershey, J. C., & Bradlow, E. T. (2006). Promises and lies: Restoring violated trust. *Organizational Behavior and Human Decision Processes*, *101*(1), 1-19.

Schweitzer, M. E., Ho, T. H., & Zhang, X. (2016). How monitoring influences trust: A tale of two faces. *Management Science*, *64*(1), 253-270.

Shao, B. (2019). Moral anger as a dilemma? An investigation on how leader moral anger influences follower trust. *The Leadership Quarterly*, *30*(3), 365-382.

Slovic, P. (2007). "If I look at the mass I will never act": Psychic numbing and genocide. *Judgment and Decision Making*, *2*(2), 79-95.

Slovic, P., Västfjäll, D., Erlandsson, A., & Gregory, R. (2017). Iconic photographs and the ebb and flow of empathic response to humanitarian disasters. *Proceedings of the National Academy of Sciences*, *114*(4), 640-644.

Small, D. A., & Simonsohn, U. (2007). Friends of victims: Personal experience and prosocial behavior. *Journal of Consumer Research*, *35*(3), 532-542.

Van Kleef, G. A. (2009). How Emotions Regulate Social Life: The Emotions as Social Information (EASI) Model. *Current Directions in Psychological Science*, 184-188.

Van Kleef, G. A., de Dreu, C. K. W., & Manstead, A. S. R. (2010b). An interpersonal approach to emotion in social decision making: the emotions as social information model. *Advances in Experimental Social Psychology*, *42*, 45-96.

Valdesolo, P., & DeSteno, D. (2011). Synchrony and the Social Tuning of Compassion. *Emotion*, *11*(2), 262-266.

Yeager, D. S., Purdie-Vaughns, V., Garcia, J., Apfel, N., Brzustoski, P., Master, A., ... & Cohen, G. L. (2014). Breaking the cycle of mistrust: wise interventions to provide critical feedback across the racial divide. *Journal of Experimental Psychology: General*, *143*(2), 804-824.

Yip, J. A., & Schweitzer, M. E. (2015). Trust promotes unethical behavior: Excessive trust, opportunistic exploitation, and strategic exploitation. *Current Opinion in Psychology*, 6, 216-220.

Appendix

Study 3 Vignettes

Diagnosis Vignette	Feedback Vignette	Recommendation Vignette
John is a doctor [Compassion: who is known to be a highly compassionate person. He is very sympathetic to his patients and cares a lot about easing their suffering]. John just diagnosed his patient, Mary, with terminal cancer, meaning that the cancer is not curable and Mary will likely die. However, there is an experimental treatment she can get to reduce her pain in the future [High cost: but she will only be able to do so if she is fully informed of the diagnosis]. Mary does not need to be fully informed about the diagnosis to get the treatment. Mary can tell something is wrong and is very distressed. She asks John about her diagnosis.	Jane is an employee of a marketing company. [Compassion: Among her colleagues, Jane is known to be a highly compassionate person. She is very sympathetic to others' predicaments and is often concerned with alleviating others' suffering.] Jane has a colleague who will be delivering a presentation at work. [High cost: The presentation is a very important one. Upper management will be observing, and the colleague's performance on this presentation could determine future pay. It could also determine whether this colleague is allowed to continue working at the company or not.] The presentation is not a very important one. Upper management will not be observing, and the colleague's performance on this presentation will not affect future pay. It will have no bearing on whether this colleague is allowed to continue working at the company or not. This colleague gives a practice presentation in front of Jane a week before the actual presentation. The colleague was nervous about how the presentation went. At the end of the practice presentation, the colleague asked Jane how the presentation went	Taylor is a college professor. [Compassion: She is known to be a highly compassionate person. She cares about helping people who are suffering and when her students have difficulties, she is very sympathetic.] Adam, a student in Taylor's class, is applying for a graduate program. He asked Taylor for a recommendation letter. Taylor knew that her letter can have a strong influence on whether Adam gets admitted. Before she wrote the letter, Taylor thought about how attending this graduate program might influence Adam. The program Adam is interested in is known to be challenging and fast-paced. [High cost: The program is located all the way across the country, so Adam would have to move across the country to attend it. If he cannot follow the pace of the program, he would likely end up dropping out, and there would be no alternative programs for him in that area.] The program is located in the same city, so Adam does not need to move to attend it. If he cannot follow the pace of the program, he would likely end up dropping out, but there would be several alternative programs for him in that area.