Brief article

When it takes a bad person to do the right thing

Eric Luis Uhlmann [a],[*],[1], Luke (Lei) Zhu [b],[1], David Tannenbaum [c]

[a] HEC Paris School of Management, 1, Rue de la Libération, 78351 Jouy-en-Josas, France
[b] University of British Columbia, Sauder School of Business, 2053 Main Mall, V6T 1Z2, Vancouver, BC, Canada
[c] UCLA Anderson School of Management, 110 Westwood Plaza, 90095-1481, Los Angeles, CA, United States

ARTICLE INFO

ABSTRACT

Three studies demonstrate that morally praiseworthy behavior can signal negative information about an agent's character. In particular, consequentialist decisions such as sacrificing one life to save an even greater number of lives can lead to unfavorable character evaluations, even when they are viewed as the preferred course of action. In Study 1, throwing a dying man overboard to prevent a lifeboat from sinking was perceived as the morally correct course of action, but led to negative aspersions about the motivations and personal character of individuals who carried out such an act. In Studies 2 and 3, a hospital administrator who decided not to fund an expensive operation to save a child (instead buying needed hospital equipment) was seen as making a pragmatic and morally praiseworthy decision, but also as deficient in empathy and moral character.

1. Introduction

In his classic treatise on stigma, Goffman (1963) describes the case of James Berry, England's first professional hangman. Berry felt deeply stigmatized by his occupation: besides hiding his hangman's noose inside his clothing, he would often travel in public with his wife and son, believing that appearing as a "family man" would make him less conspicuous. In the first half of the 20th century, though, executing criminals received overwhelming public support in England as a means to deter crime (Bland, 1984). Why would Berry feel so stigmatized when he was engaging in behavior seen, in his cultural and historical context, as necessary for the greater good?

Goffman's example suggests that certain acts and organizational duties, even when viewed as necessary or even admirable, may nevertheless send negative moral signals about the person who carried them out. Perhaps some socially approved acts are assumed to require agents that are deficient in positive moral traits or are selected by persons

of inferior moral character. The present research accounts for this paradox by drawing on the distinction made in recent research between act-centered and person-centered moral judgments (Pizarro & Tannenbaum, 2011). When individuals make their moral judgments, they distinguish between the permissibility of the act in question and the moral character of the agent who carries it out. Acts are not just viewed as moral or immoral in-and-of-themselves: they also signal information about the presence (or absence) of moral virtues such as empathy and personal integrity. As a result, one can observe *act-person dissociations* in which an act is seen as relatively less immoral than a comparison act, yet signaling more negative information about the agent's character (Tannenbaum, Uhlmann, & Diermeier, 2011).

One of the primary factors distinguishing act-centered and person-centered judgments is their relative emphasis on pragmatic considerations such as aggregate material outcomes (Pizarro & Tannenbaum, 2011). Judgments of acts are particularly sensitive to consequentialist moral principles, where the outcomes of an act serves as the guiding basis for determining the rightness or wrongness of that activity (Royzman, Goodwin, & Leeman, 2011; Tannenbaum et al., 2011). Consequentialist morality often

* Corresponding author. Tel.: +33 (0)1 39 67 97 44.
   E-mail address: eric.luis.uhlmann@gmail.com (E.L. Uhlmann).
  [1] These authors contributed equally to this work.

focuses on maximizing positive material outcomes, traditionally along quantifiable metrics such as lives saved or monetary resources spent.

Nonconsequentialist ethics, by contrast, argue that consequences are not necessarily the primary determinant of moral worth. While nonconsequentialism can take many forms, a prominent school of thought focuses on human virtues—such as loyalty, fairness, interpersonal respect, and religious purity—that may not readily map onto outcomes (e.g., Anscombe, 1958; Aristotle, 4th Century B.C.E.). Suggestively, recent empirical studies have demonstrated that acts producing the worst material consequences are not necessarily viewed as most indicative of bad moral character (Tannenbaum et al., 2011). The present studies focus on prototypical moral dilemmas involving consequentialist vs. nonconsequentialist decisions, namely whether or not to actively harm one individual in order to prevent equal or greater harm from befalling a larger number of individuals.

Across three empirical studies, we provide evidence that consequentialist acts can lead to dissociated judgments of acts and persons. More specifically, the decision to sacrifice one life to save a greater number of lives, even when regarded as the morally correct course of action, can lead to negative inferences about the agent's moral character. We test this idea across two scenarios involving life-or-death decisions: the classic lifeboat dilemma and a modified version of the hospital administrator dilemma (Tetlock, Kristel, Elson, Green, & Lerner, 2000).

We also explore potential reasons for such act-person dissociations. The first such reason is that sacrificing one person to achieve positive aggregate outcomes can leave some ambiguity regarding the agent's motives (Bartels & Pizarro, 2011). A person may engage in consequentialist behavior because they are motivated by genuine (consequentialist) moral sentiments, or because of selfish motives such as a desire to share in the material benefits. In extreme cases, a sadist or psychopath may decide to inflict harm that later leads to positive material outcomes only because he or she actively enjoys inflicting pain on others. Thus, different dispositional profiles map onto consequentialist behavior in certain cases, rendering a consequentialist choice less informative of moral virtues. We test this *multiple-motives hypothesis* in Study 1. Of course, attributional ambiguity is not an inherent property of consequentialist decisions and can emerge in a variety of situations (for theoretical analyses of the factors that contribute to such ambiguity, see Kelley, 1967; Reeder & Brewer, 1979; Snyder, Kleck, Strenta, & Mentzer, 1979). One can readily imagine moral dilemmas in which a nonconsequentialist act is highly ambiguous as to the underlying motive (an issue we return to in the General Discussion). In addition, as shown in the present research, some consequentialist decisions (e.g., sacrificing a fellow passenger in the lifeboat dilemma) are more ambiguous than others (e.g., purchasing needed equipment rather than funding an expensive operation in the hospital administrator dilemma). Thus, ambiguity as to the underlying motive explains moral reactions to only a subset of consequentialist acts.

The second reason is that carrying out consequentialist dictates can require suppressing empathy for others. Consider an agent who, confronted with a moral dilemma, decides that killing one innocent person to save 10 others is morally justified. The agent is now faced with the grim task of actively taking an innocent life. People might assume that a truly virtuous person would be incapable of such an act, even if the situation demanded it. Such an inference would not be an unreasonable one, as empathy for others is especially likely to be elicited by a single identifiable victim (Kogut & Ritov, 2005; Small & Loewenstein, 2003) and high trait empathy has been linked to an aversion to consequentialist decisions (Choe & Min, 2011). This of course does not rule out the possibility that in other contexts, consequentialist decisions are driven by empathy for the many people who benefit when aggregate outcomes are maximized. However in the context of the moral scenarios we examine, in which one identifiable person must be sacrificed in order to save many faceless persons, the consequentialist decision may be considered necessary yet nevertheless promote the attribution that an agent is lacking in empathy and moral character. We test this *empathic suppression hypothesis* in all three studies.[2]

A third reason is that people value acts that take into account pragmatic constraints and goals, including a concern for maximizing aggregate outcomes. Thus, part of the appeal of sacrificing one life to save a greater number of lives is that doing so is both pragmatic and morally defensible based on consequentialist principles. An unwillingness to sacrifice one life regardless of the consequences may suggest high trait empathy and thus praiseworthy character, but at the same time fails to satisfy the standard of practicality to which acts are often held. We test this *moral pragmatism hypothesis* in Studies 1 and 3.

Finally, we examine the possibility that act-person dissociations do not extend to all act- or person-based judgments. Indeed, certain person judgments, in particular leadership evaluations, may require greater attention to acts. We hypothesized that leaders, whose acts have important practical implications at an aggregate level, are especially apt to be judged based on whether their decisions lead to positive aggregate outcomes. The accountability pressures of leadership positions demand decisions that confront the costs and benefits of different courses of action (Molinsky & Margolis, 2005)—and such pragmatic cost-benefit thinking is usually associated with consequentialist decision making (Baron, 1994; Sunstein, 2005). Under some circumstances people may prefer leaders willing to make the "hard choice" to sacrifice one life to save a greater number of lives. In such situations, people should favor pragmatic leaders who are comparatively low in positive moral traits such as empathy, and therefore more inclined to carry out consequentialist actions (Study 2).

---

[2] A pilot study provided some initial empirical support for the empathic suppression hypothesis. Although a medical research assistant who tested cancer treatments on animals was seen as engaging in more "morally praiseworthy" behavior than a pet store assistant ($p < .001$), she was at the same time rated as significantly lower in trait empathy ($p = .03$), $F_{intx}(1, 57) = 135.67$, $p < .001$.

2. Study 1

Our first study examined whether acts seen as morally preferable can nonetheless lead to negative character attributions in the context of the classic lifeboat dilemma. We hypothesized that throwing an injured man overboard to prevent a lifeboat from sinking would be perceived as the morally right thing to do, but that lifeboat passengers who undertook this course of action would be seen as worse people than if they refused to make such a sacrifice. To examine whether judgments of character were driven more by the consequentialist or nonconsequentialist decision, the study included a control condition in which participants were not provided with information about the passengers' final decision.

Study 1 also tested the multiple motive hypothesis. An agent may sacrifice one individual to save many on morally principled grounds, or for less principled reasons such as personal self-interest or even a sadistic desire to harm others (Bartels & Pizarro, 2011). We hypothesized that passengers who made the consequentialist decision would be viewed as potentially acting out of a desire to save themselves, and therefore less morally motivated than passengers who refused to sacrifice one of their own.

Although not the primary focus of Study 1, additional measures were also included to provide preliminary tests of the (complementary) empathic suppression and moral pragmatism hypotheses. We expected that inferences of empathy would show greater predictive power for judgments of moral character than for judgments of acts, helping explain the moral praise accorded to persons who refuse to actively harm an innocent person. On the other hand, we expected that inferences of pragmatic motives would show greater predictive power for judgments of acts than for judgments of character, helping explain why consequentialist acts are often valued.

2.1. Method

One hundred and thirty-seven adults (36% female; M_{age} = 30, range = 18–62) were recruited from Mechanical Turk (MTurk) in return for a small cash payment. The study employed a 2 (consequentialist vs. nonconsequentialist decision) × 2 (character vs. act evaluation) between-subjects design. Also included was a control condition in which subjects did not learn about the passengers' decision and evaluated them based on their moral character.

Participants were presented with a variation on the lifeboat dilemma, where shipwrecked passengers must decide whether to throw an injured man overboard to keep their lifeboat from sinking. In the *consequentialist* condition, the passengers decided to throw the injured man overboard and in the *nonconsequentialist* condition they decided not to do so. Participants were then randomly assigned to provide either act evaluations or character evaluations.

2.1.1. Moral judgments

In the act evaluation condition, participants were asked whether the passengers had made "the morally right decision" (1 = *extremely immoral decision*, 7 = *extremely moral*

decision), an "ethical decision" (1 = *extremely unethical decision*, 7 = *extremely ethical decision*), and "morally good decision" (1 = *extremely morally bad decision*, 7 = *extremely morally good decision*) (α = .93). In the character evaluation condition, they were asked whether the passengers had "good moral character" (1 = *extremely immoral people*, 7 = *extremely moral people*), were "ethical people" (1 = *extremely unethical people*, 7 = *extremely ethical people*), and were "morally good people" (1 = *extremely morally bad people*, 7 = *extremely morally good people*) (α = .91).

2.1.2. Trait ratings

All participants then evaluated on 7-point scales (1 = *not at all*, 7 = *extremely*) the target's empathy (*caring, warm, empathic*; α = .87), integrity (*principled, honorable, integrity*; α = .87), boldness (*bold, gutsy, fearless*; α = .78), and reasonableness (*reasonable, effective, realistic, rational*; α = .92). We chose to measure empathy, integrity, and boldness based on Walker and Hennig's (2004) prior work suggesting they represent the three core components of folk theories of moral character.

2.1.3. Perceived motives

Participants not in the control condition then responded to a set of items examining inferences about the underlying motives for the target's actions. They reported whether the passengers had acted out of "genuine moral concerns," "moral principle," and a "genuine moral stand," or based on "personal self-interest," "what was good for them personally," and "selfish reasons" (1 = *definitely not*, 7 = *definitely yes*). The self-interest items were reverse-coded and averaged into a composite along with the moral motives items (α = .89). Participants further completed items assessing whether the passengers acted based on "what was reasonable," and "what was rational to do," and took "a pragmatic stand" (1 = *definitely not*, 7 = *definitely yes*), which formed a reliable measure of pragmatic motives (α = .88). They were also asked whether the target's actions were "excessively moral" and "excessively pragmatic" (1 = *strongly disagree*, 7 = *strongly agree*).

2.1.4. Definition of moral decision

Finally, participants defined both a "good moral decision" and "ideal moral decision" as driven by moral considerations (1 = *least concerned with what is morally right*, 7 = *most concerned by what is morally right*) or pragmatic considerations (1 = *least concerned by what is practical*, 7 = *most concerned by what is practical*) and whether saving the most lives possible in the scenario was a "moral consideration" and "practical consideration" (1 = *not at all*, 7 = *extremely*).

2.2. Results and discussion

2.2.1. Moral judgments

The expected act-person dissociation emerged. There was a significant interaction between choosing the consequentialist option and whether the targets were evaluated based on their actions or their moral character, $F_{intx}(1,106)$ = 15.41, p < .001 (Fig. 1). Participants on average favored the consequentialist decision, indicating that

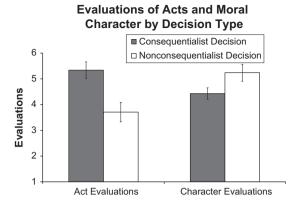## Evaluations of Acts and Moral Character by Decision Type

Fig. 1. Evaluations of acts and moral character by whether the target made a consequentialist or nonconsequentialist decision. Higher numbers reflect relatively more positive evaluations. Error bars represent standard errors.

throwing the injured man overboard was the morally right thing to do ($M = 5.35$, $SD = 1.68$ vs. $M = 3.71$, $SD = 1.84$), $t(49) = 3.32$, $p = .002$. But at the same time, they evaluated the moral character of passengers who took the consequentialist course of action more negatively than the moral character of passengers who refused to throw the man overboard ($M = 4.43$, $SD = 1.24$ vs. $M = 5.24$, $SD = 1.75$), $t(57) = 2.06$, $p = .04$.

2.2.2. Trait ratings

Consistent with the empathic suppression hypothesis, passengers who chose to throw the injured man overboard were viewed as reliably lower in empathy ($Ms = 4.13$ and 5.55, $SDs = 1.21$ and 1.35), $t(108) = 5.78$, $p < .001$. They were also viewed as lower in integrity ($Ms = 4.31$ and 5.31, $SDs = 1.32$ and 1.49), $t(108) = 3.75$, $p < .001$, but comparable in boldness to passengers in the nonconsequentialist condition ($Ms = 4.51$ and 4.74, $SDs = 1.53$ and 1.74), $t < 1$. At the same time, the consequentialist passengers were also viewed as more reasonable people ($Ms = 5.51$ and 3.18, $SDs = 1.40$ and 1.78), $t(108) = 7.68$, $p < .001$.

2.2.3. Perceived motives

Of particular interest, empirical support emerged for the multiple motives hypothesis. Participants were less likely to attribute morally principled reasons to passengers who threw the injured man overboard ($M = 3.29$, $SD = 1.24$) than if they refused to throw the man overboard ($M = 5.67$, $SD = 1.27$), $t(108) = 9.94$, $p < .001$. Conversely, the consequentialist decision was significantly more likely to be seen as driven by pragmatic motives than the nonconsequentialist decision ($Ms = 5.36$ and 2.62, $SDs = 1.37$ and 1.50), $t(108) = 10.00$, $p < .001$. The nonconsequentialist decision ($M = 5.27$, $SD = 1.67$) was also more likely to be viewed as excessively moral than throwing the man overboard ($M = 3.81$, $SD = 1.72$), $t(106) = 4.49$, $p < .001$. Conversely, throwing him overboard was significantly more likely to be rated as excessively pragmatic ($Ms = 4.84$ and 3.65, $SDs = 1.59$ and 1.95), $t(108) = 3.53$, $p = .001$.

2.2.4. Definition of moral decision

Paired-sample *t*-tests indicated that participants were more likely to rate a "good moral decision" as based on moral considerations than on pragmatic considerations ($Ms = 5.64$ and 4.29, $SDs = 1.41$ and 1.97), $t(108) = 5.26$, $p < .001$. This was also true for an "ideal moral decision" ($Ms = 5.20$ and 4.64, $SDs = 1.73$ and 1.84), $t(109) = 2.17$, $p = .03$. Saving the most lives possible was significantly less likely to be rated a moral issue than a practical issue ($Ms = 5.56$ and 5.90, $SDs = 1.49$ and 1.37), $t(108) = 2.02$, $p < .05$, although the high means (both above 5.5 on a scale of 1–7) indicate it was clearly considered both.

2.2.5. Tests of underlying processes

We further examined the extent to which ratings of empathy, moral motives, and pragmatic motives drove moral judgments of acts and persons. Consistent with the empathic suppression hypothesis, ratings of empathy positively predicted person judgments, $r(58) = .72$, $p < .001$, but not act judgments, $r(50) = .12$, $p = .39$. Consistent with the mixed motives present in Study 1's lifeboat dilemma, moral motives likewise predicted person judgments, $r(58) = .45$, $p < .001$, but not act judgments, $r(50) = -.09$, $p = .55$. Supporting the moral pragmatism hypothesis, pragmatic motives showed the opposite pattern, significantly predicting act judgments, $r(50) = .76$, $p < .001$, but not person judgments, $r(58) = -.01$, $p = .97$.

Moderated mediation analyses (Preacher, Rucker, & Hayes, 2007) with 10,000 re-samplings revealed that empathy mediated person judgments, 95% CI [$-2.01, -.76$], but not act judgments, 95% CI [$-1.29, .05$]. Moral motives likewise drove person judgments, 95% CI [$-2.33, -1.01$], but not act judgments, 95% CI [$-1.52, 0.71$]. Conversely, perceived pragmatic motives drove act judgments, 95% CI [$1.92, 3.62$], but not assessments of moral character, 95% CI [$-.06, 1.23$].

2.2.6. Comparisons with the control condition

A final set of analyses examined whether assessments of the passengers' moral character were driven more by the consequentialist or nonconsequentialist decision. Judgments of moral character were significantly less positive in the consequentialist decision condition than in the control condition ($Ms = 4.43$ and 5.10, $SDs = 1.24$ and 1.07), $t(56) = 2.18$, $p = .03$. In contrast, character evaluations were quite similar in the control condition and nonconsequentialist decision condition ($Ms = 5.10$ and 5.24, $SDs = 1.07$ and 1.75), $t < 1$. This suggests that judgments of moral character were driven more by the decision to throw the injured man overboard than the decision not to do so.

In sum, participants favored a consequentialist act that maximized positive outcomes, but made negative character attributions about agents who carried out that same act. The results also suggest that moral character evaluations were influenced by an asymmetry in the putative motivations for each decision. When the passengers refused to throw the injured man overboard, they were seen as clearly acting on moral principle. But when passengers chose to sacrifice the injured man, it seemed possible they were only acting out of self-interest rather than for morally principled reasons. Providing initial support for the

empathic suppression and moral pragmatism hypotheses, inferences of empathy exhibited greater predictive power for judgments of moral character than for judgments of acts, whereas pragmatic motives showed greater predictive power for judgments of acts than of character.

3. Study 2

Our second study examined the idea that ambiguous motives are not a necessary condition for morally praiseworthy acts to result in negative character attributions. Even when an agent is seen to sacrifice one life to save many others based on genuine moral principle, doing so may still suggest a suppressed sense of empathy. Study 2 therefore assessed reactions to a hospital administrator who must decide whether to spend a large sum of money to save one little boy or on new hospital equipment that would save a greater number of lives in the long run. In this context, personal self-interest and sadistic tendencies are unlikely reasons for pursuing a consequentialist course of action. We expected that participants would endorse the consequentialist decision, but view an administrator who made that same decision as comparatively lacking in empathy and moral character.

Of additional interest was the possibility that leadership evaluations of the hospital administrator would mirror judgments of his acts rather than his character. Because leadership decisions have important practical implications at an aggregate level, people should value leaders able to suppress their empathy and make decisions that maximize collective outcomes (Molinsky & Margolis, 2005).

3.1. Method

Ninety-two adults (57% female; M_{age} = 34, range = 18–68) were recruited from MTurk. The study employed a 2 (consequentialist vs. nonconsequentialist decision) × 3 (act vs. moral character vs. leadership evaluation) between-subjects design.

All participants read about a hospital administrator in Deeg, India, who had to decide whether to spend $2 million on a life-saving operation for a little boy, or use the money to purchase better equipment that would save 500 future lives. The scenario was similar to one used by Tetlock et al. (2000), but differed in that it specified a very large number of future lives saved. In order to test the hypothesis that people can make negative attributions about consequentialist agents even when they prefer consequentialist behavior, one must examine contexts in which consequentialist acts are generally preferred. However, in the original Tetlock et al. scenario, participants strongly prefer the nonconsequentialist act (i.e., saving the little boy). Accordingly, we increased the number of future lives saved to 500 to make the consequentialist decision more attractive. We further moved the location to a poor Indian community to make it believable that $2 million in hospital equipment could save so many lives. In the *nonconsequentialist* condition the hospital administrator decided to save the boy, and in the *consequentialist* condition the administrator decided to purchase the new equipment.

3.1.1. Moral and leadership evaluations

After reading about the decision, participants were randomly assigned between-subjects to evaluate either whether he did "the morally right thing" (1 = *extremely bad act*, 7 = *extremely good act*), had "good moral character" (1 = *extremely bad person*, 7 = *extremely good person*) or was "a good leader" (1 = *extremely bad leader*, 7 = *extremely good leader*).

3.1.2. Removed from job

All participants then indicated whether they felt the hospital administrator should be "removed from his job" and "allowed to continue in his role as the director of the hospital" (1 = *definitely not*, 7 = *definitely yes*). The latter item was reverse coded and averaged with the "removed" item (r = .91).

3.1.3. Perceived motives

Moral motives were then assessed as in Study 1 (α = .81).

3.1.4. Trait ratings

Finally, perceived empathy, integrity, and boldness were measured using the same items as in Study 1 (αs were .90, .90, and .82, respectively).

3.2. Results and discussion

3.2.1. Perceived motives

We first tested the multiple motive hypothesis by examining whether the target that made a consequentialist decision was less likely to be viewed as acting out of moral principal. Unlike Study 1, no reliable difference in motive attributions was found. The hospital administrator who chose to buy new equipment and the administrator who chose to save the little boy were both seen as driven by moral motives (Ms = 5.14 and 5.34, SDs = 1.22 and 1.09), t < 1.

3.2.2. Moral and leadership evaluations

Regardless, an act-person dissociation emerged. There was a significant interaction between whether the hospital administrator chose the consequentialist or nonconsequentialist option and whether he was evaluated based on his actions, moral character, or leadership, $F_{intx}(2, 86)$ = 8.74, $p < .001$ (Fig. 2). The hospital administrator who decided to buy new equipment was perceived as having made a more moral decision than the administrator who saved little Ravi (Ms = 5.38 and 4.05, SDs = 1.12 and 1.63), t(32) = 2.60, p = .01. But at the same time, the administrator who made the consequentialist decision was viewed as a worse person than the administrator who chose to save the little boy (Ms = 5.00 and 6.00, SDs = 1.37 and 1.08), t(27) = 2.15, p = .04.

As expected, leadership evaluations mirrored act judgments rather than assessments of moral character. Even though participants viewed the consequentialist administrator as a less moral person, they viewed him as a better leader than the administrator who saved little Ravi (Ms = 5.69 and 3.69, SDs = 1.30 and 1.84), t(27) = 3.41, p = .002.

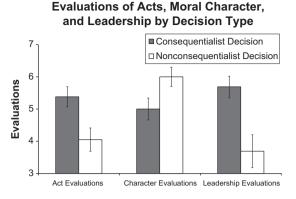## Evaluations of Acts, Moral Character, and Leadership by Decision Type

Fig. 2. Evaluations of acts, moral character, and leadership by whether the target made a consequentialist or nonconsequentialist decision. Higher numbers reflect relatively more positive evaluations. Error bars represent standard errors.

3.2.3. Removed from job

Consistent with the results for leadership evaluations, participants were significantly less likely to feel that the hospital administrator who chose to buy new equipment should be removed from his job ($M = 2.29$, $SD = 1.38$) than the administrator who saved the little boy ($M = 3.45$, $SD = 1.83$), $t(90) = 3.41$, $p = .001$.

3.2.4. Trait ratings

The administrator who chose to buy new hospital equipment was viewed as relatively low in empathy ($Ms = 4.79$ and 6.18, $SDs = 1.19$ and 1.12), $t(90) = 5.81$, $p < .001$, but not integrity ($Ms = 5.14$ and 4.99, $SDs = 1.23$ and 1.54), $t < 1$. Unexpectedly, the administrator who made the consequentialist decision was also perceived as less bold ($Ms = 4.89$ and 5.52, $SDs = 1.34$ and 1.33), $t(90) = 2.28$, $p = .03$. Although speculative, participants may have found it bold to spend so much of the organization's resources to save one patient, potentially risking the wrath of other stakeholders. However, perceived boldness did not significantly predict evaluations of moral character, $r(28) = .10$, $p = .60$, acts, $r(33) = -.07$, $p = .69$, or leadership, $r(28) = .32$, $p = .09$.

3.2.5. Tests of underlying processes

As hypothesized, perceived empathy positively predicted evaluations of moral character, $r(28) = .76$, $p < .001$, but not act judgments, $r(33) = -.13$, $p = .48$, or leadership judgments, $r(28) = .24$, $p = .22$. Mediation analyses (Preacher et al., 2007) with 10,000 re-samplings found that empathy significantly mediated person judgments, 95% CI $[-2.32, -0.82]$, and leadership judgments, 95% CI $[-1.20, -0.16]$, but not act judgments, 95% CI $[-1.14, 0.23]$.

In sum, even in the absence of ambiguous motives, the expected act-person dissociation emerged. Participants viewed a hospital administrator who chose to buy needed equipment rather than save a little boy as having acted on principle and as making the morally right decision. But at the same time, they also saw him as lacking in empathy and as a morally worse person. Further consistent with the empathic suppression hypothesis, inferences of empathy predicted judgments of moral character, but not acts. Despite the fact he was seen as less empathic and a worse person, the hospital administrator who decided not to save the little boy was evaluated more positively as a leader. This suggests that laypeople sometimes prefer leaders who are capable of suppressing feelings of empathy and making pragmatic decisions that maximize positive collective outcomes.

4. Study 3

Our final study addressed some important questions left unanswered by the results from Study 2. One particularly pressing goal was to formally test the moral pragmatism hypothesis in the context of the hospital administrator dilemma, as Study 2 was run chronologically prior to Study 1 and did not include measures of pragmatism. An additional aim was to establish whether judgments of character in this scenario are driven more by the decision to save the little boy or to buy new hospital equipment. The study therefore included a control condition in which participants were not told which course of action the hospital administrator decided to take. Finally, in order to determine whether choosing not to save the little boy truly made the hospital administrator seem like more of a bad person (rather than just less of a good person), participants evaluated the "badness" and "goodness" of his moral character on separate scales. We hypothesized that the hospital administrator who chose the consequentialist course of action would be significantly more likely to be seen as a bad person.

4.1. Method

Seventy-seven adults (32% female; $M_{age} = 29$, range = 18–60) were recruited from MTurk. Participants were randomly assigned to a 2 (consequentialist vs. nonconsequentialist decision) × 2 (act vs. character evaluation) between-subjects design. Also included was a control condition in which participants did not receive information about the hospital administrator's decision.

4.1.1. Moral judgments

In the act evaluation condition, participants indicated on separate unipolar scales whether the administrator made a "morally bad decision" (1 = not at all morally bad decision, 7 = extremely morally bad decision) and "morally good decision" (1 = not at all morally good decision, 7 = extremely morally good decision), and further evaluated the administrator's decision on a bipolar scale ranging from 1 (extremely immoral act) to 7 (extremely moral act). In the person evaluation condition, participants were asked to indicate on separate unipolar scales whether the administrator was a "morally bad person" (1 = not at all morally bad person, 7 = extremely morally bad person) and "morally good person" (1 = not at all morally good person, 7 = extremely morally good person), as well as to rate the administrator on a bipolar scale ranging from 1 (extremely immoral person) to 7 (extremely moral person).

4.1.2. Trait ratings

They further rated the empathy (α = .93), integrity (α = .85), boldness (α = .84), and reasonableness (α = .93) of the administrator.

4.1.3. Perceived motives

Participants who were not in the control condition then completed our measures of moral motives (α = .82) and pragmatic motives (α = .94). As in Study 1, these participants were further asked whether the target's actions were excessively moral and excessively pragmatic.

4.1.4. Definition of moral decision

Finally, participants who were not in the control condition defined a good moral decision and ideal moral decision as driven by moral vs. pragmatic considerations, and rated whether saving the most lives possible in the scenario was a moral vs. practical consideration.

4.2. Results and discussion

4.2.1. Moral judgments

An act-person dissociation again emerged. A 2 (consequentialist vs. nonconsequentialist decision) × 2 (act vs. character evaluation) ANOVA revealed a significant interaction between the target's decision and whether he was evaluated based on his actions or moral character. This interaction emerged for the bipolar moral-immoral item, $F_{intx}(1,53)$ = 13.61, p = .001, the unipolar "morally bad" item, $F_{intx}(1,53)$ = 7.58, p = .008, and the unipolar "morally good" item $F_{intx}(1,53)$ = 7.52, p = .008.

Examination of the bipolar moral-immoral item revealed a pattern of simple effects identical to that observed in Study 2. The hospital administrator who decided to buy new equipment was perceived as having made a more moral decision than the administrator who chose to save the little boy (Ms = 5.94 and 4.70, SDs = 1.00 and 1.64), $t(26)$ = 2.51, p = .02. But the same time, he was also seen as having worse moral character (Ms = 4.21 and 5.53, SDs = 1.42 and 1.19), $t(27)$ = 2.72, p = .01.

Analyses of the unipolar items indicated that deciding to buy the equipment was viewed as a significantly more "morally good" decision than saving the little boy (Ms = 5.78 and 4.40, SDs = 1.00 and 1.71), $t(26)$ = 2.70, p = .01, but not a significantly less "morally bad" decision (Ms = 2.17 and 2.90, SDs = 1.38 and 1.37), $t(26)$ = 1.35, p = .19. Conversely, the hospital administrator who decided to buy the equipment was seen as a significantly more "morally bad" person (Ms = 3.50 and 2.07, SDs = 1.83 and 1.16), $t(27)$ = 2.54, p = .02, but not a significantly less "morally good" person (Ms = 4.93 and 5.40, SDs = 1.14 and 1.24), $t(27)$ = 1.06, p = .30. The latter pattern of results indicates that relative to the nonconsequentialist decision, the consequentialist decision did make the hospital administrator seem like more of a bad person (rather than just less of a good person). However, it is worth noting that the effect of choosing the consequentialist course of action on being perceived as a morally bad person was not significantly greater than for being seen as a morally good person (p = .17). Thus, it cannot be concluded that our effects on

character evaluations were driven solely by perceived badness.

4.2.2. Trait ratings

The administrator who chose to buy the equipment was viewed as lower in empathy (Ms = 5.13 and 5.84, SDs = 1.12 and 1.30), $t(55)$ = 2.23, p = .03, comparable in integrity (Ms = 5.33 and 5.09, SDs = 1.24 and 1.27), t < 1, and lower in boldness (Ms = 4.83 and 5.56, SDs = 1.41 and 1.22), $t(54)$ = 2.05, p < .05, but a more reasonable person (Ms = 5.76 and 4.32, SDs = 1.12 and 1.80), $t(55)$ = 3.69, p < .001, than the administrator who chose to save the little boy.

4.2.3. Perceived motives

The consequentialist and nonconsequentialist decision were equally likely to be viewed as driven by moral motives (Ms = 5.14 and 5.32, SDs = 1.11 and 1.29), t < 1, but the consequentialist decision was significantly more likely to seen as driven by pragmatic motives (Ms = 5.68 and 3.72, SDs = 1.12 and 1.82), $t(55)$ = 5.01, p < .001. Moreover, saving little Ravi (M = 4.36, SD = 1.35) was no more likely to be perceived as excessively moral than buying the equipment (M = 4.94, SD = 1.46), $t(55)$ = 1.53, p = .13. However, buying the equipment was significantly more likely to be rated as excessively pragmatic (Ms = 4.75 and 3.48, SDs = 1.57 and 1.53), $t(55)$ = 3.07, p = .003.

4.2.4. Definition of moral decision

Paired-sample t-tests indicated that participants were more likely to define a "good moral decision" as driven by moral considerations than pragmatic considerations (Ms = 5.62 and 4.82, SDs = 1.48 and 1.78), $t(54)$ = 3.17, p = .003, which was likewise true for an "ideal moral decision" (Ms = 5.55 and 4.87, SDs = 1.32 and 1.63), $t(54)$ = 2.51, p = .02. Saving the most lives possible in the scenario was less likely to be deemed a moral consideration than a practical consideration (Ms = 5.58 and 5.88, SDs = 1.36 and 1.21), $t(56)$ = 2.04, p < .05, although the relatively high means suggest it was considered both.

4.2.5. Tests of underlying processes

For ease of analysis and presentation, we computed a three-item measure of moral evaluations averaging the bipolar immoral-moral item, the unipolar morally good item, and (after reverse scoring it) the unipolar morally bad item (α = .82). As expected, empathy positively predicted person judgments, $r(28)$ = .73, p < .001, but not act judgments, $r(27)$ = .05, p = .80. Conversely, pragmatic motives predicted favorable act judgments, $r(27)$ = .67, p < .001, but not evaluations of moral character, $r(28)$ = .19, p = .33. Using empathy as the mediator, mediation analyses with 10,000 re-samplings revealed that empathy drove person judgments, 95% CI [0.05, 1.12], but not act judgments, 95% CI [−.23, 0.49]. Further, using pragmatic motives as the mediator, mediation analyses showed that pragmatic motives drove act judgments, 95% CI [−2.03, −.65], but not person judgments, 95% CI [−1.23, 0.11].

4.2.6. Comparisons with the control condition

A final set of analyses suggest that character evaluations were driven more by the decision to buy the equipment rather than the decision to save little Ravi. Judgments of moral character were significantly less positive in the consequentialist decision condition than in the control condition on the bipolar moral-immoral person item ($Ms = 4.21$ and 5.75, $SDs = 1.42$ and 1.07), $t(32) = 3.59$, $p = .001$, and unipolar "morally good person" item ($Ms = 4.93$ and 5.90, $SDs = 1.14$ and 1.12), $t(32) = 2.47$, $p = .02$, and significantly more negative on the unipolar "morally bad person" item ($Ms = 3.50$ and 1.80, $SDs = 1.83$ and 1.06), $t(32) = 3.43$, $p = .002$. In contrast, character evaluations were comparable between the control condition and the nonconsequentialist decision condition on the bipolar immoral-moral person item ($Ms = 5.75$ and 5.53, $SDs = 1.07$ and 1.19), $t < 1$, unipolar "morally good person" item ($Ms = 5.90$ and 5.40, $SDs = 1.12$ and 1.24), $t(33) = 1.25$, $p = .22$, and the unipolar "morally bad person" item ($Ms = 1.80$ and 2.07, $SDs = 1.06$ and 1.16), $t < 1$.

In sum, choosing to buy new hospital equipment rather than save one little boy was seen to satisfy pragmatic goals and as the best moral decision under the circumstances, yet led to negative empathy and moral character ratings. Further consistent with the empathic suppression and moral pragmatism hypotheses, inferences of empathy predicted judgments of moral character (but not acts), whereas pragmatic motives predicted judgments of acts (but not character). Finally, comparisons with a control condition suggest that character judgments were driven more by the consequentialist decision to buy the equipment than the nonconsequentialist decision to save the child.

5. General discussion

Even when people endorse consequentialist acts, they can nonetheless make negative attributions about the moral character of agents who carry out those very same acts. In Study 1, throwing a dying man overboard to prevent a lifeboat from sinking was viewed as the morally correct course of action, but led to negative aspersions about the motivations and personal character of individuals who carried out such an act. In Studies 2 and 3, a hospital administrator who chose to buy needed equipment rather than fund an expensive operation to save one little boy was seen as having made a pragmatic and morally praiseworthy decision, but also as deficient in empathy and moral character.

These studies further explored some of the mechanisms that can give rise to such act-person dissociations. One reason is that consequentialist acts can present a multiple-motives problem: an agent may inflict harm to achieve good material outcomes because they genuinely endorse consequentialist moral principles, or because they wish to cause harm for less noble reasons (Study 1). Another reason is that upholding consequentialist moral principles by sacrificing one person to save many can require suppressing the empathy naturally elicited by identifiable victims (Kogut & Ritov, 2005; Small & Loewenstein, 2003).

Even in the absence of mixed motives, empathic suppression led to negative character attributions based on consequentialist acts (Studies 2 and 3). Finally, consequentialist acts are often morally appealing because they satisfy pragmatic goals, and pragmatic concerns are weighted more heavily in judgments of acts than of persons (Studies 1 and 3).

Interestingly, our results for leadership evaluations suggest that not all person-based judgments are one and the same. Even though he was viewed as lacking in empathy and moral character, a hospital administrator who chose to buy new equipment rather than save one little boy was seen as the better leader (Study 2). Although this suggests people may prefer leaders who have a somewhat suppressed sense of empathy (and are thus able to make tough decisions that maximize aggregate outcomes), it is important not to take this argument too far. Anecdotally, at least, voters seem to care about whether politicians express empathic outrage under the right circumstances. It seems likely that people want a leader who has a fully intact emotional-response profile but is also able to suppress such emotions in order to make pragmatic and effective decisions.

Future research should examine moral judgments of acts and persons in the context of a larger variety of moral dilemmas. The present studies suggest that sacrificing one person to save many, even when seen as the morally appropriate course of action, leads to negative aspersions about the agent's moral character in two reasonably representative situations: the classic lifeboat dilemma and a modified version of the hospital administrator dilemma (Tetlock et al., 2000). Replicating this effect across a greater number of scenarios is not just an important avenue for future research: it is a necessary condition for drawing firmer conclusions both about the generality of act-person dissociations and also the potential underlying mechanisms we have explored in this paper.

In particular, multiple motives concerns are almost certainly not unique to consequentialist outcomes, but are instead situation dependent. One can easily imagine cases in which it is unclear whether an agent acted based on a nonconsequentialist principle (e.g., truthfully telling a competitor that her performance was lackluster out of a genuine principle against lying) or out of selfish motives (e.g., a desire to rattle her confidence). It also seems unlikely that all consequentialist acts are thought to require a lack of empathy in order to carry them out. Indeed, some consequentialist decisions may be driven by an empathic concern for others' welfare (e.g., for the many individuals who are helped when aggregate outcomes are maximized).

At the same time, it remains possible that many moral tradeoffs — even those that do not pit consequentialist against nonconsequentialist values — can give rise to dissociated moral judgments of acts and the agents who carry them out. To the extent social perceivers are more likely to emphasize pragmatic considerations with regards to the moral status of acts as opposed to persons, they may value "balanced" decisions that satisfy multiple moral values at once, yet respect persons who remain loyal to a single ethical principle.

Just as the present studies add to our understanding of nonconsequentialist decision making, they also contribute to the emerging literature on person-centered moral judgments. Unlike any previous research, the present research documents act-person dissociations in the context of morally praiseworthy acts. Participants endorsed consequentialist courses of action as the morally right thing to do, but also made negative attributions about agents who carried out such acts. The present research also suggests that leadership evaluations are unlike other person judgments in that they require more attention to acts and their aggregate consequences. As a result, pragmatic and effective leadership can be associated with certain negative moral traits, such as a suppressed sense of empathy. Taken together, these findings further underscore a fundamental insight about the nature of human moral judgment: evaluations of acts and the people who carry them out can be very different.

References

Anscombe, G. E. M. (1958). Modern moral philosophy. *Philosophy, 33*, 1–19.

Aristotle (4th Century, B.C.E./1998). *The Nicomachean ethics*. Oxford: Oxford University Press.

Baron, J. (1994). Nonconsequentialist decisions. *Behavioral and Brain Sciences, 17*, 1–10.

Bartels, D., & Pizarro, D. A. (2011). The mismeasure of morals: Antisocial personality traits predict utilitarian responses to moral dilemmas. *Cognition, 121*, 154–161.

Bland, J. (1984). *The common hangman: English and Scottish hangmen before the abolition of public executions*. Hornchurch, UK: I. Henry Publications.

Choe, S. Y., & Min, K.-H. (2011). Who makes utilitarian judgments? The influences of emotions on utilitarian judgments. *Judgment and Decision Making, 6*, 580–592.

Goffman, E. (1963). *Stigma: Notes on the management of spoiled identity*. Upper Saddle River, NJ: Prentice-Hall.

Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska symposium on motivation* (pp. 192–238). Lincoln: University of Nebraska Press.

Kogut, T., & Ritov, I. (2005). The "identified victim" effect: An identified group, or just a single individual? *Journal of Behavioral Decision Making, 18*, 157–167.

Molinsky, A. L., & Margolis, J. D. (2005). Necessary evils and interpersonal sensitivity in organizations. *Academy of Management Review, 30*, 245–268.

Pizarro, D. A., & Tannenbaum, D. (2011). Bringing character back: How the motivation to evaluate character influences judgments of moral blame. In P. Shaver & M. Mikulincer (Eds.), *The social psychology of morality: Exploring the causes of good and evil* (pp. 91–108). Washington, DC: American Psychological Association.

Preacher, K. J., Rucker, D. D., & Hayes, A. F. (2007). Addressing moderated mediation hypotheses: Theory, methods, and prescriptions. *Multivariate Behavioral Research, 42*, 185–227.

Reeder, G. D., & Brewer, M. B. (1979). A schematic model of dispositional attribution in interpersonal perception. *Psychological Review, 86*, 61–79.

Royzman, E. B., Goodwin, G. P., & Leeman, R. F. (2011). When sentimental rules collide: Norms with feelings in the dilemmatic context. *Cognition, 121*, 101–114.

Small, D. A., & Loewenstein, G. (2003). Helping a victim or helping the victim: Altruism and identifiability. *Journal of Risk and Uncertainty, 26*, 5–16.

Snyder, M. L., Kleck, R. E., Strenta, A., & Mentzer, S. J. (1979). Avoidance of the handicapped: An attributional ambiguity analysis. *Journal of Personality and Social Psychology, 37*, 2297–2306.

Sunstein, C. (2005). Moral heuristics. *Behavioral and Brain Sciences, 28*, 531–542.

Tannenbaum, D., Uhlmann, E. L., & Diermeier, D. (2011). Moral signals, public outrage, and immaterial harms. *Journal of Experimental Social Psychology, 47*, 1249–1254.

Tetlock, P. E., Kristel, O., Elson, B., Green, M., & Lerner, J. (2000). The psychology of the unthinkable: Taboo trade-offs, forbidden base rates, and heretical counterfactuals. *Journal of Personality and Social Psychology, 78*, 853–870.

Walker, L. J., & Hennig, K. H. (2004). Differing conceptions of moral exemplarity: Just, brave, and caring. *Journal of Personality and Social Psychology, 86*, 629–647.