Journal of Personality and Social Psychology
2017, Vol. 113, No. 5, 730–752

0022-3514/17/$12.00 http://dx.doi.org/10.1037/pspi0000103

Hypocritical Flip-Flop, or Courageous Evolution? When Leaders Change Their Moral Minds

Tamar A. Kreps
University of Utah

Kristin Laurin
University of British Columbia

Anna C. Merritt
Stanford University

How do audiences react to leaders who change their opinion after taking moral stances? We propose that people believe moral stances are stronger commitments, compared with pragmatic stances; we therefore explore whether and when audiences believe those commitments can be broken. We find that audiences believe moral commitments should *not* be broken, and thus that they deride as hypocritical leaders who claim a moral commitment and later change their views. Moreover, they view them as less effective and less worthy of support. Although participants found a moral mind changer especially hypocritical when they disagreed with the new view, the effect persisted even among participants who fully endorsed the new view. We draw these conclusions from analyses and meta-analyses of 15 studies (total $N = 5{,}552$), using recent statistical advances to verify the robustness of our findings. In several of our studies, we also test for various possible moderators of these effects; overall we find only 1 promising finding: some evidence that 2 specific justifications for moral mind changes—citing a personally transformative experience, or blaming external circumstances rather than acknowledging opinion change—help moral leaders appear more courageous, but no less hypocritical. Together, our findings demonstrate a lay belief that moral views *should* be stable over time; they also suggest a downside for leaders in using moral framings.

Keywords: hypocrisy, leadership, morality, opinion change, social perception

Supplemental materials: http://dx.doi.org/10.1037/pspi0000103.supp

Taking a moral stance can make leaders more effective. A leader who gives moral arguments for her positions, for instance by saying she supports the social safety net because it advances human dignity, rather than because it makes good economic sense, appears to have better moral character (Van Zant & Moore, 2015) and to be more authentic and committed to the cause (Kreps & Monin, 2014), and she has a better chance of persuading, and inspiring commitment in, her followers (Brown et al., 2005; Conger & Kanungo, 1998; Van Zant & Moore, 2015). Moral stances can also reassure followers by reinforcing their worldviews and imbuing issues with a sense of meaning (Greenberg et al., 1990; Heine et al., 2006).

Leaders may therefore choose to take moral stances, believing that this will improve audiences' perceptions. But all people, even leaders, change their minds sometimes. Consider the many political leaders, from Hillary Clinton (James, 2014) to George H.W. Bush (Levine, 2015), who have recently changed their minds about same-sex marriage. How might audiences react when a leader changes her mind, and does it matter whether the leader initially took a moral stance? Will they deride a leader who changes her moral mind as a flip-flopper with no consistent core (how can we trust her integrity if she changed her mind about an opinion she claimed was based on her fundamental moral values)? Or will they in fact approve of leaders who are willing to change their moral minds, and see them as demonstrating virtues such as open-mindedness and courage (should we not admire the open-mindedness it takes to realize that one might have been wrong on something so personal as a moral issue, and the courage it takes to admit it)? Or does it depend on which position the leader changed from and to, with people viewing their leaders through partisan lenses, deriding those who abandoned a moral and "right" position but lauding those who abandoned a moral and "wrong" position?

The aim of the current research is to test these possibilities—and, in so doing, to shed light on how people interpret moral stances and what they expect from their leaders.

Moral Views Likely Convey Commitment

Leaders who take moral stances probably do so for many reasons (Skitka, Bauman, & Sargis, 2005; Graham, Haidt, & Nosek, 2009; Kreps & Monin, 2011)—including strategic reasons, as doing so can make them appear more effective and likable (Brown et al., 2005; Conger & Kanungo, 1998; Greenberg et al.,

This article was published Online First June 8, 2017.

Tamar A. Kreps, Department of Management, University of Utah; Kristin Laurin, Department of Psychology, University of British Columbia; Anna C. Merritt, Department of Psychology, Stanford University.

This work was supported through Stanford Graduate School of Business faculty research funding. The funding source was in no way involved in designing or conducting the research. We thank Ellen Reinhart and Lauren Agnew for their assistance with data collection, and the members of the Monin Morality Lab for their comments on a draft of this article. We also thank Jeremy Biesanz for generously and patiently sharing his statistical expertise; Elizabeth Dunn, Elizabeth Tenney, and Jessica Tracy for helping us make the manuscript as readable as possible.

Correspondence concerning this article should be addressed to Tamar A. Kreps, Department of Management, University of Utah, 1655 East Campus Center Drive, Salt Lake City, UT 84112. E-mail: tamar.kreps@utah.edu

1990; Heine et al., 2006; Van Zant & Moore, 2015). That said, when a leader takes a moral stance her intended message is that she truly has a strong moral conviction about the issue (Skitka, Washburn, & Carsel, 2015). Lay people may in turn be aware that individuals with strong moral convictions are committed to maintaining their positions even when doing so is difficult or unpopular (Aramovich et al., 2012; Cole Wright, Cullum, & Schwab, 2008; Mullen & Nadler, 2008; Skitka & Morgan, 2014; Skitka & Mullen, 2002; Skitka et al., 2005). If so, then audiences who hear a leader take a moral stance may take it as a commitment that the stance will endure into the future. Consistent with this idea, the single paper to date that directly examined how lay people perceive moral stances finds that people judge speakers who take a moral stance to be more committed to their positions (Kreps & Monin, 2014). Thus, people may believe moral stances reveal stronger tendencies to maintain the same view in the future than do pragmatic ones.

How Will Audiences Respond to Moral Mind Changes?

So, people tend to stick to views that are morally based, and we have proposed that third-party observers may be aware of this pattern. But do they also think that people *ought* to be more committed to their morally based views? We argue that this question has at least three possible answers.

The Hypocrisy Hypothesis

If people take moralizing an attitude to be a commitment to maintaining that attitude forever, much as a marriage is a commitment to maintaining a current partnership forever, one possibility is that people believe that moralizing an attitude *should* constrain the moralizer's future behavior. Just as people may believe that a marital commitment is sacred and should never be broken, people may believe that once a leader takes a moral position, he is making an unbreakable promise to maintain that position forever. If this is the case, then people likely see going back on a moral principle as not only surprising, but morally deceitful. Accordingly, we propose that moral mind changes may generate perceptions of *hypocrisy*.

Hypocrisy is "saying one thing but doing another" (Barden, Rucker, & Petty, 2005, p. 1464), or "feigning to be what one is not or to believe what one does not" (Merriam-Webster, 2014). If audiences believe that claiming to espouse a moral principle should convey an especially enduring future commitment, then later abandoning that commitment meets both these definitions of hypocrisy: Either the initial claim (or the commitment to endure that it implied) was false, or the new claim must be. If so, audiences may suspect that one of these claims was made for self-serving reasons. By contrast, an initially pragmatic "flip-flopper," who made no pledge to endure, need not have feigned anything. Thus, we propose:

The hypocrisy hypothesis. Audiences view initially moral opinion change as revealing more hypocrisy than initially pragmatic opinion change.

Previous research on hypocrisy has focused on people who do not "practice what they preach" (Stone & Fernandez, 2008), and on people who claim a position that they have belied with their previous behavior (cf. Effron & Miller, 2015). Our hypocrisy hypothesis enriches these conceptions of hypocrisy in two ways. Building on the first conception, we propose that the discrepancy need not be between what one *practices* and what one *preaches*. Instead, it may be that holding one position and then later holding the opposite position can be enough to warrant the hypocrisy label, depending on the basis for one's original position. Similarly, researchers using the second conception have studied how contradictory *behavior* in one's past can make a claim appear hypocritical. Here, we explore whether, when a leader changes her mind, the basis for her earlier position—as moral or pragmatic—determines whether that holding that position was enough of a contradictory prior "behavior" to make a later claim seem hypocritical.

The Courageous Evolution Hypothesis

A different possibility is that, although people believe that (also much like marriage) moralizing an attitude *does* constrain the future actions of the moralizer, they do not necessarily believe that it *should*. In other words, just as a marriage may sometimes prevent people from ending a relationship that is damaging to one or both parties, people may believe that the commitment of moralizing prevents leaders from changing their views even under circumstances when they really should. Thus, maybe "marrying" a particular position is seen not as a commitment that *should* be maintained, the way marrying a particular human being often is, but rather as an obstacle to updating one's views and opinions when required. Such a belief would be reasonable given evidence on commitment and consistency in general (Cialdini, Cacioppo, Bassett, & Miller, 1978; Cialdini & Trost, 1998): Once people have committed themselves to a particular position, particularly in public (Cialdini & Trost, 1998), they are unlikely to move away from it, even when they should, for fear of seeming inconsistent. Therefore, if moral claims imply extra commitment, moral claimers may indeed by extra unwilling to update their positions.

If people recognize moral stances as constraining leaders' future positions without believing that this fact is particularly desirable, then they may view speakers who deviate from moral stances as superior in some ways—specifically, as more courageous and more flexible—than those who deviate from less sticky pragmatic stances. Put differently, if audience members recognize all of the forces working against changes from initially moral stances, and yet value and appreciate leaders' capacity to evolve, they may especially admire leaders who can overcome the odds and acknowledge that they have changed.

Consistent with this reasoning, anecdotally, moral mind changers do sometimes encounter praise instead of criticism. Many commentators hailed President Obama's courage in announcing that he had come to support same-sex marriage—calling it, for example, "the most courageous thing he [had] done since he entered the White House" (Kaiser, 2012)—after he had earlier opposed same-sex marriage on religious grounds, calling marriage a "sacred union" (Miller, 2015). Similarly, some of the most respected policy changes in U.S. history occurred because of leaders' opinion changes from previously moral stances—for example, Lincoln's change from a principled defender of states' rights to the president who abolished slavery across the United States (Fredrickson, 2008). Surely this evolution, rather than making us see Lincoln as a hypocrite, has contributed to our positive view of him (Grant, 2015).

Despite the appeal of the hypocrisy hypothesis, therefore, it also seems conceivable that, at least under certain conditions, audience members give extra credit to initially moral mind changers. Van Zant & Moore (2015, Study 4) found, consistent with this intuition (although they explained their findings in a different way), that leaders who took a moral stance were later evaluated more positively even after they announced they would not follow through on that stance. Thus, if people believe that initial moral stances *do* indicate a tendency to maintain one's position unwaveringly, without believing that they *should* do so, they may view moral mind changers as more courageous and more flexible than pragmatic mind changers.

The courageous evolution hypothesis. Audiences view initially moral opinion change as revealing more courage and/or flexibility than initially pragmatic opinion change.

The Partisanship Hypothesis

A third possibility is that particular circumstances moderate these hypotheses. Indeed, there are some circumstances under which most people likely agree that a marital commitment should be broken (such as when one or both partners are abusive), and others under which they likely agree it should not (such as when one partner is experiencing fleeting nostalgia over unsown wild oats). Likewise, it seems plausible that there are circumstances under which audiences believe more strongly that moral stances *should* persist into the future and others under which they believe more strongly that they should *not*.

In particular, it seems plausible that people's beliefs follow their partisan preferences. All of the anecdotes we listed above as supporting the courageous evolution hypothesis involve leaders changing from a view many people see as "incorrect" to one they see as "correct" (cf. Pew Research Center, 2015, on same-sex marriage). To return to Lincoln, perhaps our very reverence for his abolitionist policy is what saves him from appearing hypocritical, and allows us to see him, instead, as a courageous and open-minded hero. Similarly, opponents of same-sex marriage probably did not find Obama's change "courageous." Thus, it may be that audiences' tendency to attribute hypocrisy or to grant extra credit to moral mind changers depends on their relationship to the leader's initial and new positions—whether, from their perspective, leaders are changing from the "incorrect" view to the "correct" one, or vice versa.

This possibility resonates with research showing that human beings are notoriously partisan animals (Pronin, Puccio, & Ross, 2002; Ross & Ward, 1996). For example, people idealize those they like more than those they dislike (Murray, Holmes, & Griffin, 1996), they believe their own group's offenses are less severe than another group's (Abrams, de Moura, & Travaglino, 2013; Hastorf & Cantril, 1954), and they believe unjust preferential treatment is fair so long as it helps people they empathize with (Blader & Rothman, 2014). Moral preferences in particular can color people's interpretations of others' actions: Those who approve of another person's actions believe those actions are motivated by the "true self," but those who disapprove of the same actions see the other person as biased by external forces (Newman, Bloom, & Knobe, 2014, see also Christy et al., 2017). In our context, this may mean that when a leader changes his moral mind to a view that audiences endorse, they are more likely to view that as an authentic expression of truth, or that when a leader changes his moral mind to a view that audiences find abhorrent, they are more likely to assume he is lying under external pressure.

The partisanship hypothesis. Perceptions of hypocrisy and/or courageous evolution depend on the direction of leaders' change relative to audiences' views. As audiences disagree more with a leader's ultimate view, they become more likely to view initially moral opinion change as revealing hypocrisy; and/or as audiences agree more with a leader's ultimate view, they become more likely to view initially moral opinion change as revealing courage, relative to pragmatic opinion change.

Ultimately, the hypotheses we offer here relate to different assumptions about the norm of continuing to support moralized positions: Is it merely a descriptive norm about the way things are, or is there a corresponding prescriptive, or injunctive, norm about the way things ought to be (e.g., Cialdini, Reno, & Kallgren, 1990)? The hypocrisy hypothesis presumes that the norm is prescriptive: That people believe moralizers *should* stick to their beliefs over time, no matter what. The courageous evolution hypothesis presumes that the norm is only descriptive: That people merely perceive that moralizers *do* tend to stick to their beliefs even in the face of adversity, without endorsing this tendency as right and good. And the partisanship hypothesis posits that people's interpretations of norms as descriptive or prescriptive are flexible and responsive to psychological motivations (cf. Kay et al., 2009). Thus, they may interpret a change from moral "wrong" to "right" as a brave, authentic violation of an unnecessarily restrictive descriptive norm, while at the same time interpreting a change from moral "right" to "wrong" as a hypocritical violation of an important prescriptive norm.

Downstream Consequences for Leaders

Our final prediction is that there will be further downstream consequences for leaders of changing their moral minds. More specifically, we propose that, if the hypocrisy hypothesis is correct, moral mind changers will seem less effective as leaders and will lose audience support. Conversely, we propose that if the courageous evolution hypothesis is correct, moral mind changers will seem more effective as leaders and will gain audience support. Finally, we also propose that if the partisanship hypothesis is correct, perceptions of moral mind changers' effectiveness, and support for moral mind changers, will depend on audiences' partisan biases.

We predict that these downstream consequences will follow directly from the character perceptions we laid out in the previous section. No literature has directly focused on how perceptions of hypocrisy and courage play into leadership. However, many theoretical perspectives on leadership suggest that perceived hypocrisy may especially harm, and perceived courage and flexibility may especially benefit, leaders. First, hypocrisy means that a leader will be hard to predict, and therefore hard to coordinate with and ineffective at serving the important function of helping to solve coordination problems (e.g., Van Vugt, Hogan, & Kaiser, 2008). Hypocrisy also means that a leader is unreliable and therefore unable to competently guide a group, and unworthy of cognitive trust (McAllister, 1995); it also means that a leader is dishonest, perhaps willing to misrepresent herself for personal gain, and hence unworthy of affective trust (McAllister, 1995).

Second, a leader who seems courageous for updating her views and humble for admitting she was wrong may seem especially effective at serving the group, and may ultimately reap the rewards of competitive altruism (Hardy & Van Vugt, 2006).

Taken together, these literatures suggest that, if moral mind changers seem hypocritical, they will therefore also seem less effective as leaders and will lose audience support; in contrast, if moral mind changers seem courageous and flexible, they will therefore also seem more effective as leaders and will gain audience support. They also suggest that if people's perceptions of the character of a moral mind-changer—his hypocrisy and his courage—depend on partisanship, then so will their perceptions of his effectiveness and their support: If initially moral mind changers seem more hypocritical only to audience members who disagree with their later views, then they will therefore also seem less effective and receive less support from these audience members specifically; if initially moral mind changers seem more courageous only to audience members who agree with their later views, then they will therefore also seem more effective and receive more support from these audience members specifically.

Downstream effects. When audiences see initially moral mind changers as more hypocritical, they will therefore also see them as less effective and support them less; when instead they see them as more courageous, they will therefore also see them as more effective and support them more.

Overview of Studies

In the process of refining our methods, we conducted a set of 15 distinct studies, each testing these alternative hypotheses. Rather than selectively present a subset of our studies, we report the results of all 15 studies here (Figure 1). In other words, we present every data point we have obtained (excluding two studies whose manipulation checks failed and one study which used different dependent measures, see Footnote 4—although including these studies yields very similar results). This means that our meta-analyses across these 15 studies provides the most accurate estimate of the true size of these effects that we could produce (e.g., Rothstein et al., 2006). We also hope that fully describing all of the exploratory tests we have attempted will help stimulate further research on this topic.

Because each of our studies uses a very similar design when it comes to testing our primary hypotheses, and because there are a large number of them, we do not describe each study individually.

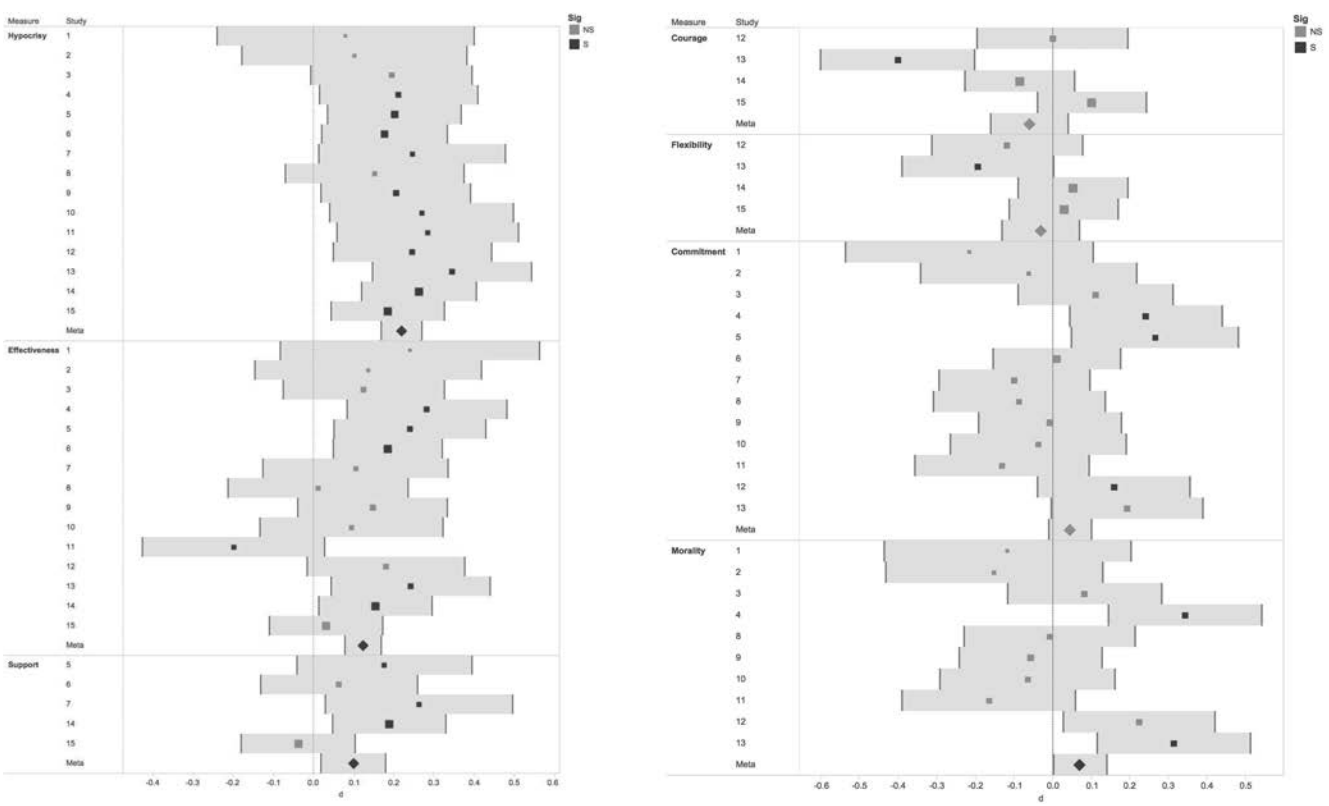

Figure 1. *d* effect sizes for main effects of initial framing on key dependent measures, all studies and meta-analysis. Gray shaded areas show 95% confidence intervals of the *d* based on the variance calculated according to Johnson and Eagly's (2000) method. Significance tests for individual studies are based on our original regression analyses. For the individual studies, the size of the square corresponds to the weight of the estimate in the meta-analysis. "Support" refers to intentions to support the candidate and to vote for the candidate. We reverse scored effectiveness, support, commitment and morality such that higher coefficients indicate more negative views of the initially moral leader.

Instead, we first describe the methods and results of an Example Study—specifically, Study 14 from our larger set—to orient the reader to our general approach. We then provide a summary description of our entire set of studies, numbered in chronological order, noting where the methods deviated from the Example Study (these differences between studies are summarized in Table 1). Finally, we provide a summary of the results from the individual studies, and then metaanalyze and discuss the overall results. Full details, materials, and statistical tests for all studies are in supplemental online material (SOM).

We began, however, with a pilot study to evaluate the descriptive and prescriptive norms people hold regarding moral stances, described in detail in the SOM. We asked participants to think about positions held for moral reasons and about positions held for pragmatic reasons, and to rate them each separately on items designed to capture both descriptive and prescriptive norm perceptions. On the one hand, we found some evidence for the prescriptive norm implied by the hypocrisy hypothesis: Participants reported that people *should* stick to their initial position more when that position was moral rather than pragmatic. On the other hand, we found some evidence for the descriptive norm implied by the courageous evolution hypothesis: Participants reported others *did* stick to initially moral positions even more strongly than they reported that they *should*. These findings confirm our reasoning about the norms that people might have regarding moral mind changes, but they do not clearly favor one hypothesis over the other; we therefore proceeded to directly test these hypotheses in the subsequent 15 studies.

In each of these studies, participants learned about political or business leaders who changed their opinion on an issue. We manipulated whether the leaders were changing from an initially moral or pragmatic stance, to test the hypocrisy and courageous evolution hypotheses. We also measured participants' agreement with the leader's final position, to test the partisanship hypothesis. Our primary dependent measures were participants' perceptions of the leader (e.g., hypocrisy, effectiveness). We kept these elements of our study designs constant, but changed several other elements to demonstrate robustness. For example, across studies we used different contexts (e.g., marriage equality, immigration reform, the death penalty, environmental initiatives), and different designs (between- vs. within-participant s); we also systematically manipulated and measured several different exploratory variables (e.g., moral relativism, leader gender), to test potential moderators of our primary effects.

Table 1
Summary of Methods, Designs, and Measures, All Studies

Study	Issue	Context	Mediator measures	Downstream measures[a]	N[b]	Position manip?	Manipulated moderators	Measured moderators	Notes
1	Environmental impact	Organization	Hypocrisy	E, M, C	151	N	None	Issue moralization	
2	Environmental impact	Organization	Hypocrisy	E, M, C	198	N	None	Issue moralization	
3	Ritual drug use	Politics (island)	Hypocrisy	E, M, C	387	Y	Later moral/prag framing	Issue moralization	
4	Same-sex marriage	Politics (US)	Hypocrisy	E, M, C	398	Y	Later moral/prag framing	Issue moralization	
5	Ritual drug use	Politics (island)	Hypocrisy, surprise	E, S, C	167	Y	Later moral/prag framing	Issue moralization	Primary manipulation within subjects
6	Same-sex marriage	Politics (US)	Hypocrisy, surprise	E, S, C	206	Y	Later moral/prag framing	Issue moralization	Primary manipulation within subjects
7	Same-sex marriage	Politics (US)	Hypocrisy	E, S, ad effectiveness	297	N	None	Issue moralization	Stimuli conveyed via realistic attack ads
8	Environmental impact	Org/Politics (US)	Hypocrisy	E, M, C	301	N	Reliance on popular support	None	
9	Environmental impact	Organization	Hypocrisy	E, M, C	450	N	Reliance on popular support	None	
10	Same-sex marriage	Politics (US)	Hypocrisy, surprise	E, M, C	303	Y	Tying later view to same value	Issue moralization	
11	Same-sex marriage	Politics (US)	Hypocrisy, surprise	E, M, C	306	Y	Tying later view to same value	Issue moralization	
12	Immigration reform	Politics (US)	Hypocrisy, courage, flex.	E, M, C	404	Y	Later moral/prag framing	Issue moralization	
13	Sexualized ads	Organization	Hypocrisy, courage, flex.	E, M, C	401	Y	Later moral/prag framing	None	
14	Death penalty/ SSM	Politics (US)	Hypocrisy, courage, flex.	E, S	777	Y	Later moral/prag framing; leader gender	Moral relativism	Used validated E and S measures; controlled for issue
15	Death penalty	Politics (US)	Hypocrisy, courage, flex.	E, S	806	Y	Later framing: moral/ prag/external constraints/ transformational story	Issue moralization	Pre-registered; used validated E and S measures

[a] E = leader effectiveness, S = support for leader, M = morality, C = commitment to issue. [b] Total sample size after omitting participants who failed any attention checks.

Example Study (Study 14)

We begin by describing an Example Study, which we chose from within our larger set because its design allows us to clearly illustrate the core elements of our design that were common to all studies, while its results reasonably approximate those emerging from our cross-study analyses. This Example Study is Study 14—that is, the 14th study that we ran.

In this Example Study, participants read about a political leader's views on a political issue—either the death penalty or same-sex marriage. We used two different political issues to increase the robustness and generalizability of our findings. Participants read an earlier view, where the leader expressed one position (e.g., pro death penalty), and a later view, where the leader expressed the opposite view (e.g., anti death penalty) while acknowledging the departure from the earlier view. Our primary manipulation was the framing of the earlier view: Some participants read that the leader framed the initial view using moral arguments (e.g., citing concerns for justice and violent criminals getting their just deserts), whereas others read that the leader framed the initial view using pragmatic arguments (e.g., citing concerns for the expense of paying for violent criminals' life in prison). In each case, we also manipulated what side of the issue the leader switched from and to, and measured participants' own views. These latter procedures allowed us to index participants' agreement with the leader's ultimate view, and thus test the hypocrisy hypothesis.

General Note on Methodology

In our Example Study, and also in all of our studies, we determined sample size in advance based on the study design and our prior experience with moralization manipulations (including, as our research proceeded, the previous studies in this project). After two initial studies, we aimed for at least $n = 75$ per cell to test our core initial framing $\times$ agreement with the leader's final position design—that is, a total sample size of at least 300 in between-subjects designs. We sometimes aimed for larger sample sizes when we included exploratory manipulations we thought might moderate these effects. We stopped collecting data once we had recruited the intended number of participants, and we did not look at the data before data collection was complete. All participants were Mechanical Turk workers based in the United States, who had completed at most 100 HITs and had at least a 70% approval rating; each participant completed only one of our studies. We excluded only participants who failed attention checks (as summarized below and described in each study in SOM), and we made these exclusions in every study that included such checks, before performing analyses. We did not force participants to answer all questions, so small discrepancies in the degrees of freedom we report are due to missing data on individual questions. We report all conditions and all measures for all of our studies (although exploratory measures and manipulations are described and analyzed in greater detail in the Supplement). Data for all studies are available at https://osf.io/st72f/. We preregistered Study 15—the last one we conducted—on the Open Science Framework (see: https://osf.io/rp8ua/?view_only=fa565703babd4cca9c40da57469f672f).

Example Study Method

Participants and design. Eight hundred seventeen Mechanical Turk workers participated for payment. We excluded data from 40 participants who did not respond correctly to an attention check ("Please leave this item blank; do not select anything on this line," see Oppenheimer et al., 2009). Thus, the final sample size was 777 participants (458 female, 314 male, 2 "agender," 1 "genderqueer," 1 "Q," 1 "none").

Our main manipulation was the leader's framing (moral vs. pragmatic) of the initial position. We also manipulated the leader's position (pro- to anti-same-sex-marriage or death penalty, vs. anti- to pro-same-sex-marriage or death penalty), the latter of which we used along with a measure of participants' own attitudes toward the relevant issue to compute a continuous index of their match with the leader's final position. In this study we also included two exploratory manipulations and one exploratory measure, described in detail below, which we intended to test as potential moderators.

Procedure. Participants first indicated their own position on the issue the target in their study would be discussing (the death penalty or same-sex marriage), from 1, *Strongly Oppose* to 7, *Strongly Support*.

Participants next read initial comments made by a Midwestern congressperson arguing for or against either the death penalty or same-sex marriage, on the basis of either moral arguments (citing justice concerns in the case of the death penalty, or equality/respect for tradition in the case of same-sex marriage) or pragmatic arguments (citing what would be best for the economy). For example, participants reading the comments of a politician supporting the death penalty for moral reasons read the following statement:

> I support the death penalty, and it's a moral issue for me. It's a matter of justice, and a question of people who have committed an unforgivable wrong getting what they deserve. From a purely ethical perspective, we need to consider the death penalty as a legitimate piece of our justice system.

As another example, participants reading the comments of a politician opposed to same-sex marriage for pragmatic reasons read the following statement:

> I oppose gay marriage, but it's not a moral issue for me. It's a matter of not having to invest in the cost of changing government systems that are already in place and working fine. I support the gay community in many ways, but I do not think we should expand the definition of marriage. From a purely economic perspective, same-sex couples should not be able to marry.

As one of our exploratory manipulations, we varied the gender of the leader (by giving the leader's name as either Carl or Kathryn). We were interested in gender because we wondered whether people might hold different expectations for male versus female leaders.

Participants then completed baseline measures of their perceptions of the leader. Although our interest was in how the leader's initial framing would influence audience perceptions *after* the leader changed his or her mind, we worried that the initial moral or pragmatic framings could influence audience perceptions even before the mind change occurred, which would contaminate our postmind change measures. We therefore included these baseline measures to serve as covariates; as noted below, however, results were similar regardless of whether or not we included them.

Specifically, participants completed measures of two outcome measures: leader effectiveness (e.g., "Based on what you know about X, how effective do you think [he/she] is as a leader?"; adapted from Rice, Instone, & Adams, 1984; from 1, *Not at all effective*, to 5, *Extremely effective*) and political support for the leader (e.g., "Given everything you know about X, if you were in his/her congressional district, how likely would you be to vote for X in an election?," adapted from e.g., Paunesku, Akhtar, & Tormala, 2013; from 1, *Very unlikely*, to 7, *Very likely*). Next, participants completed measures of our three character perceptions: Hypocrisy, courage, and flexibility. Items were intermixed, and the entire set was preceded by the text "What does X's statement about [issue] say about the kind of person [he/she] is? To what extent does it reveal each of the following in [his/her] character. . . ." Sample items were ". . . hypocrisy?" for hypocrisy; ". . . bravery" for courage, and ". . . flexibility" for flexibility." Participants rated these items on a scale ranging from 1, *Not at all*, to 5, *Very much*.

Next, participants learned that the leader had spoken again, later, about the same issue, and read the leader's new comments. The new comments always took the opposing position to the initial one, and always acknowledged the previous position as well as the mind change. As an exploratory manipulation, we varied the framing of this new position: Half of the participants read a new position framed morally, while the other half read a new position framed pragmatically. For example, participants who read that the leader now *opposed* the death penalty for *moral* reasons read the following statement:

> I know that at an earlier time I supported the death penalty, but I've given it some more thought, and my views have changed. I oppose the death penalty now, and I'll tell you why. It's still a moral issue for me, and a question of justice. I've realized, though, that we can never be 100% certain that the convicted party is guilty, and truly defending justice means never taking the risk of killing an innocent victim. Putting all economic concerns aside, you have to acknowledge that the death penalty is not a good thing.

As another example, participants who read that the leader now *supported* same-sex marriage for *moral* reasons read the following statement:

> I know that at an earlier time I opposed gay marriage, but I've given it some more thought, and my views have changed. I support gay marriage now, and I'll tell you why. It's become a moral issue for me. Allowing gay people to marry is ultimately the best way to give them rights they deserve and achieve equality. Putting all economic concerns aside, you have to acknowledge that same-sex marriage is a good thing."

We included this second exploratory manipulation as a potential moderator: We aimed to see whether one or the other of these later framings might mitigate the hypocrisy or courage that participants perceived in leaders who deviated from an initially moral stance. For instance, perhaps changing *to* a pragmatic stance would exaggerate the perceived hypocrisy of a moral mind changer, because it would make it seem as if he had abandoned moral values entirely.[1]

Finally, participants completed the same measures of outcome variables (leader effectiveness and support), and the same measures of character perceptions (hypocrisy, courage, and flexibility). At the end of the survey, they reported demographic information, and completed an exploratory measure of moral relativism (Forsyth, 1980; e.g., "what is ethical varies from one situation and society to another"). We included this measure also as a potential moderator; wondering, for instance, whether individuals who view morality as relative would be more tolerant of a moral mind change.

Example Study Results

Analytic strategy. We used R (R Core Team, 2014; packages from Bates, Maechler, Bolker, & Walker, 2014; Kuznetsova, Brockhoff, & Christensen, 2016; Revelle, 2014) to analyze results, predicting each of our dependent variables (as measured after the mind change) using initial framing, participants' agreement with the leader's final position, and their interaction. We calculated the participants' agreement, which we refer to as "position match," based on the leaders' final position together with participants' self-reported attitudes. Specifically, this variable ranged from -3 (the participant's attitude was at the scale endpoint on the *opposite* side from the leader's later view) through 0 (the participant's attitude was at the midpoint) to 3 (the participant's attitude was at the scale endpoint on the *same* side as the leader's later view). Thus, higher values reflect greater agreement with the leader's final opinion and greater disagreement with the initial opinion.

To examine the mediating role of perceptions of hypocrisy, courage, and flexibility in explaining effects on our downstream outcome measures, we then used our basic models to test indirect effects, using bootstrap mediation analyses with 1000 iterations. We report results for all key statistics emerging from these analyses for this study and all individual studies in Table 2.

[1] This later framing manipulation clearly bears some resemblance to our primary variable of interest, which is whether the leader framed his or her *initial* view as moral or pragmatic. As a result, it may seem as though the arguments we made earlier regarding initial framing could apply to later framing as well: If a leader comes to adopt a new moral view he did not claim to hold before, audiences may suspect hypocrisy if they believe that moral view *should* be lifelong and unchanging, or they may look approvingly at him for being courageous and open-minded enough to have overcome his prior insensitivity to the moral value in question. However, our marriage analogy helps illustrate why we see this possibility as related but quite definitely separate. We view a moral stance, like getting married, as a prospective commitment, or a pledge of continued future support, not a retroactive commitment, or an implication that one's past behavior always honored the commitment. It may seem like a hypocritical violation of a commitment, or like a courageous decision to end an unhealthy state of affairs, if a person decides she no longer wants to be married *after* having committed to it. But it can be neither hypocritical nor courageous to have not wanted to be in the marriage at some point *before* having taken the marital vows. Likewise, we have argued that a moral stance, *once taken*, can be perceived as either requiring (hypocrisy hypothesis) or making it hard to avoid (courageous evolution hypothesis) sticking with that stance into the future. Whether or not a moral stance carries with it the implication that one has *always* held that stance, even in the past, is an interesting, but different, question. For this reason, we focus our paper on how audiences perceive mind changes depending on whether the *initial* stance was moral or pragmatic. Our findings thus complement prior literature, which demonstrates mostly benefits of moral stances (e.g., Brown et al., 2005; Conger & Kanungo, 1998; Kreps & Monin, 2014), by exploring the dangers that may lurk in the future for leaders who take these stances, if they later change their minds.

Table 2
Summary of Individual Study Results, All Studies

DV	Study	*M*(*SD*) by initial framing condition		*df* for coeff.[a]	Test of hypocrisy hypothesis: Effect of initial framing *b* [CI_{95}]	Indirect effect through hypocrisy [CI_{95}]	Test of partisanship hypothesis: Interac. with position match *b* [CI_{95}]	Zone of significance[b] *position match scores*
		Moral	Pragmatic					
Hypocrisy	1	5.25 (1.63)	5.13 (1.36)	146	.12 [−.37, .62]	N/A	.10 [−.30, .50]	N/A
	2	5.28 (1.62)	5.13 (1.49)	192	.16 [.28, .59]		−.24 [−.61, .13]	N/A
	3	5.08 (1.70)	4.82 (1.63)	381	.32† [−.01, .64]		−.13 [−.29, .03]	N/A
	4	5.58 (1.55)	5.23 (1.62)	394	**.32* [.02, .62]**		−.07 [−.19, .06]	N/A
	5	4.70 (1.34)	4.43 (1.44)	162	**.29* [.07, .52]**		.10† [−.02, .21]	−.71 to +3
	6	5.04 (1.50)	4.76 (1.45)	198	**.28* [.05, .51]**		**−.10* [−.19, −.004]**	−3 to +.61
	7	4.78 (1.09)	4.57 (.89)	285	**.31* [.02, .60]**		−.01 [−.13, .11]	N/A
	8	5.11 (1.44)	4.88 (1.45)	308	.22 [−.10, .54]		.002 [−.30, .31]	N/A
	9	4.78 (1.32)	4.58 (1.44)	445	**.28* [.03, .53]**		.15 [−.09, .39]	N/A
	10	4.58 (1.64)	4.17 (1.79)	295	**.47* [.07, .86]**		.07 [−.15, .28]	N/A
	11	4.38 (1.52)	3.95 (1.71)	300	**.46* [.10, .82]**		−.09 [−.29, .10]	N/A
	12	3.06 (1.29)	2.76 (1.17)	397	**.29* [.06, .53]**		−.04 [−.16, .08]	N/A
	13	2.98 (1.20)	2.59 (1.17)	393	**.41*** [.18, .64]**		−.10 [−.24, .03]	N/A
	14	3.24 (1.34)	2.99 (1.26)	763	**.32* [.15, .49]**		−.07† [−.15, .005]	−3 to +1.54
	15	2.51 (1.18)	2.29 (1.09)	768	**.21* [.05, .40]**		.04 [−.12, .05]	N/A
Effectiveness	1	3.08 (1.15)	3.36 (1.25)	144	−.30 [−.70, .10]	[−.30, .17]	.05 [−.27, .38]	N/A
	2	3.22 (1.27)	3.29 (1.29)	189	−.18 [−.54, .18]	[−.21, .08]	−.08 [−.39, .23]	N/A
	3	3.90 (1.41)	3.95 (1.28)	376	−.16 [−.42, .09]	[−.26, .01]†	−.09 [−.22, .04]	N/A
	4	3.04 (1.39)	3.23 (1.43)	388	**−.40** [−.66, −.14]**	[−.28, .01]†	.03 [−.08, .14]	N/A
	5	3.36 (1.33)	3.74 (1.40)	163	**−.39** [−.63, −.14]**	**[−.30, −.04]***	−.07 [−.18, .05]	N/A
	6	2.91 (1.40)	3.18 (1.34)	200	**−.27** [−.47, −.07]**	**[−.33, −.03]***	.04 [−.04, .12]	N/A
	7	3.42 (1.08)	3.52 (.88)	287	−.10 [−.33, .12]	**[−.22, −.01]***	−.02 [−.11, .07]	N/A
	8	3.49 (1.25)	3.41 (1.12)	302	−.02 [−.29, .25]	[−.25, .01]†	.23 [−.03, .49]	N/A
	9	3.48 (1.18)	3.54 (1.13)	441	−.17 [−.38, .04]	[−.18, .01]†	**−.26* [−.46, −.06]**	−1.84 to +3
	10	3.35 (1.41)	3.34 (1.36)	293	−.13 [−.41, .14]	**[−.42, −.03]***	.07 [−.07, .22]	N/A
	11	3.72 (1.39)	3.21 (1.39)	298	**.28* [.02, .54]**	**[−.34, −.05]***	**.24** [.09, .38]**	Moral > prag +1.51 to +3
	12	3.74 (1.46)	3.95 (1.33)	394	−.24† [−.50, .01]	**[−.31, −.03]***	**.27* [.01, .52]**	−3 to −.14
	13	4.23 (1.41)	4.56 (1.21)	393	**−.33* [−.60, −.06]**	**[−.33, −.08]****	**.18** [.02, .33]**	−3 to +.46
	14	2.24 (1.05)	2.35 (2.42)	769	**−.16* [−.29, −.02]**	**[−.25, −.08]****	.05 [−.12, .02]	N/A
	15	2.62 (.93)	2.66 (.89)	769	−.03 [−.15, .10]	**[−.13, −.02]***	.05 [−.01, .12]	N/A
Support	5	T1 = 35.3%; T2 = 19.8%	T1 = 49.1%; T2 = 32.3%	N/A	**−.63* [−1.18, −.10]**	**[−.06, −.01]***	−.12 [−.39, .14]	N/A
	6	T1 = 42.7%; T2 = 17.6%	T1 = 17.4%; T2 = 18.0%	N/A	−.23 [−.82, .36]	[−.03, .003]	.06 [−.18, .30]	N/A
	7	2.92 (1.56)	3.32 (1.23)	282	**−.35* [−.65, −.04]**	[−.31, .03]	.09 [−.04, .22]	N/A
	14	2.32 (1.53)	2.52 (1.52)	769	**−.29** [−.49, −.09]**	**[−.34, −.10]****	.07 [−.02, .16]	N/A
	15	2.98 (1.61)	2.94 (1.55)	762	.06 [−.17, .28]	**[−.23, −.04]****	−.01 [−.13, .10]	N/A
Courage	12	2.54 (1.15)	2.57 (1.12)	394	−.02 [−.23, .20]	N/A	−.02 [−.14, .09]	N/A
	13	2.29 (1.00)	2.60 (1.13)	391	**−.43*** [−.64, −.23]**		.05 [−.06, .17]	N/A
	14	3.19 (1.21)	3.30 (1.10)	763	−.13 [−.28, .03]		−.03 [−.10, .04]	N/A
	15	3.48 (1.08)	3.38 (1.10)	766	.11 [−.04, .26]		−.001 [−.08, .08]	N/A
Flexib.	12	3.08 (1.15)	3.21 (1.02)	396	−.12 [−.33, .08]		.02 [−.09, .12]	N/A
	13	3.35 (.98)	3.68 (.94)	394	**−.19* [−.38, −.004]**		.04 [−.07, .14]	N/A
	14	3.23 (1.17)	3.20 (1.12)	760	.06 [−.10, .21]		−.05 [−.12, .02]	N/A
	15	3.54 (.95)	3.53 (.98)	767	.03 [−.11, .17]		.01 [−.06, .08]	N/A
Commitment	1	2.70 (1.36)	2.31 (1.19)	143	.23 [−.25, .70]	[−.22, .12]	.08 [−.23, .40]	N/A
	2	2.83 (1.41)	2.31 (1.29)	190	.03 [−.43, .49]	[−.17, .06]	.02 [−.02, .07]	N/A
	3	3.41 (1.48)	3.45 (1.42)	377	−.16 [−.45, .13]	[−.24, .01]†	−.04 [−.19, .10]	N/A
	4	2.51 (1.44)	2.56 (1.54)	391	**−.45* [−.81, −.09]**	[−.23, .01]†	−.002 [−.12, .12]	N/A
	5	3.11 (1.32)	3.47 (1.43)	163	**−.38*** [−.60, −.16]**	**[−.27, −.03]***	−.05 [−.16, .05]	N/A
	6	2.67 (1.43)	2.67 (1.39)	199	.02 [−.18, .21]	**[−.25, −.04]****	.01 [−.07, .08]	N/A
	7	2.41 (1.48)	2.35 (1.36)	287	.14 [−.18, .45]	**[−.23, −.01]***	−.04 [−.17, .11]	N/A
	8	2.93 (1.39)	2.81 (1.27)	306	.15 [−.17, .47]	[−.23, .02]†	.0004 [−.28, .28]	N/A
	9	2.88 (1.20)	2.66 (1.30)	442	.01 [−.24, .26]	[−.18, .004]†	**−.22* [−.43, −.002]**	Nowhere significant
	10	2.80 (1.36)	2.63 (1.24)	294	−.06 [−.39, .28]	**[−.24, −.02]***	.05 [−.20, .11]	N/A
	11	3.34 (1.47)	2.61 (1.36)	298	.21 [−.12, .53]	**[−.33, −.05]***	.03 [−.13, .19]	N/A
	12	3.10 (1.73)	3.36 (1.68)	392	**−.32* [−.58, −.07]**	**[−.30, −.03]***	.17 [−.08, .42]	N/A
	13	3.59 (1.65)	3.87 (1.55)	395	−.31† [−.63, .01]	**[−.39, −.10]****	.16† [−.02, .34]	−3 to +.13

(*table continues*)

Table 2 (*continued*)

DV	Study	*M*(*SD*) by initial framing condition: Moral	*M*(*SD*) by initial framing condition: Pragmatic	*df* for coeff.[a]	Test of hypocrisy hypothesis: Effect of initial framing *b* [CI_{95}]	Indirect effect through hypocrisy [CI_{95}]	Test of partisanship hypothesis: Interac. with position match *b* [CI_{95}]	Zone of significance[b] *position match scores*
Morality	1	3.14 (1.11)	3.12 (1.10)	143	−.13 [−.54, .27]	[−.27, .14]	.02 [−.28, .32]	N/A
	2	3.19 (1.23)	3.06 (1.13)	190	.20 [−.18, .59]	[−.21, .08]	.01 [−.03, .05]	N/A
	3	3.61 (1.58)	3.62 (1.31)	377	−.14 [−.43, .15]	[−.28, .01]†	−.04 [−.18, .09]	N/A
	4	2.91 (1.36)	3.06 (1.42)	388	**−.51*** [−.79, −.24]**	[−.26, .03]	.06 [−.05, .16]	N/A
	8	3.47 (1.17)	3.37 (1.13)	307	−.01 [−.30, .28]	[−.22, .03]	−.03 [−.27, .21]	N/A
	9	3.51 (1.08)	3.30 (1.05)	443	−.07 [−.28, .15]	[−.15, .005]†	−.12 [−.31, .06]	N/A
	10	3.22 (1.37)	3.01 (1.23)	289	−.08 [−.35, .20]	**[−.36, −.03]***	.10 [−.04, .25]	N/A
	11	3.63 (1.32)	3.02 (1.29)	289	.22† [−.03, .47]	**[−.33, −.05]****	**.19** [.05, .32]**	Moral > prag +1.74 to +3
	12	3.58 (1.39)	3.76 (1.33)	394	−.26 [−.57, .06]	**[−.32, −.02]***	.27† [−.04, .59]	−3 to −.59
	13	3.80 (1.40)	4.13 (1.23)	392	**−.42** [−.67, −.16]**	**[−.45, −.13]****	**.22** [.07, .36]**	−3 to +.81
AE	7	4.23 (1.11)	3.78 (1.18)	286	**.41** [.15, .67]**	**[.03, .29]***	.02 [−.08, .13]	N/A
SfO	7	3.54 (1.31)	3.52 (1.23)	287	−.03 [−.32, .26]	[−.03, .25]†	−.02 [−.14, .11]	N/A

Note. For ease of interpretation, we have **bolded** significant results. AE: Ad effectiveness (Study 7 attack ads); SfO: Support for Opponent (Study 7 attack ads). [a] For Studies 5 and 6, because we use mixed models with degrees of freedom determined according to the Satterthwaite approximation, different coefficients within the same model can have different degrees of freedom, and degrees of freedom can be non-integers. In this table, we present the degrees of freedom for the main effect of initial framing, rounded to the nearest integer. Exact degrees of freedom for all coefficients in these models are in the Analyses SOM. Also for Studies 3a/b, there are no degrees of freedom for the Voting measure because we used a binomial mixed model (to predict this binary measure), and hence have *z* scores instead of *t* scores. [b] By "Zone of significance," we mean the results of floodlight analyses conducted when we found significant interactions. Each floodlight analysis identifies the range of values of position match at which there is a significant effect consistent with hypocrisy hypothesis.
† $.05 < p < .10$. * $.01 < p < .05$. ** $.001 < p < .01$. *** $p < .001$.

Because, in this study, we had taken baseline measures of the dependent measures after the initial statement, we controlled for these baseline measures in our models (though results were similar if we did not control for baseline scores). Moreover, given that, in this study, we examined two separate issues, we also controlled for a main effect of issue (though results were similar if we instead performed all of the same analyses without controlling for issue and simply collapsing across them, or including interactions with issue in addition to the main effects).

To test the effects of exploratory moderators (in this study, leader gender, participant moral relativism, and later moral vs. pragmatic framing), we created additional models that each added one of these variables—with continuous moderators centered—and all possible interactions. However, because across studies results for these exploratory moderators were generally null, we do not report these here in our Example Study results, but return to them in our summary of Additional Results below.

Strong support for the hypocrisy hypothesis. According to the hypocrisy hypothesis, we would expect main effects of initial framing such that leaders who framed their initial views morally would seem more hypocritical, less effective and less worthy of support. As predicted, the initially moral leader seemed more hypocritical, $b = .32$, $CI_{95} = [.15, .49]$, $p < .001$, less effective, $b = -.16$, $CI_{95} = [-.29, -.02]$, $p = .027$, and less worthy of support, $b = -.29$, $CI_{95} = [-.49, -.09]$, $p = .005$, compared with the initially pragmatic leader (see SOM for full regression tables). Furthermore, tests of indirect effects showed that hypocrisy accounted for the effects of initial framing on leader effectiveness, $CI_{95} = [-.25, -.08]$, and support, $CI_{95} = [-.34, -.10]$.

No support for the courageous evolution hypothesis. According to the courageous evolution hypothesis, we would expect main effects of initial framing such that leaders who framed their initial views morally would seem more courageous and flexible, more effective and more worthy of support. We found no effects of initial framing on perceptions of either courage, $b = -.13$, $CI_{95} = [-.28, .03]$, $p = .105$ or flexibility, $b = .06$, $CI_{95} = [-.10, .21]$, $p = .471$, and, as reported in the previous paragraph, effects on the downstream measures were opposite to the ones predicted by the courageous evolution hypothesis.

Mixed support for the partisanship hypothesis. According to the partisanship hypothesis, we would expect initial framing to interact with participants' match with leaders' views, such that any hypocrisy effects would be stronger for participants who disagreed with leaders' later views, and any courageous evolution effects would be stronger for participants who agreed with leaders' later views. We found a marginally significant interaction consistent with this hypothesis in predicting hypocrisy, $b = -.07$, $CI_{95} = [-.15, .005]$, $p = .065$ (but not effectiveness, $b = .05$, $CI_{95} = [-.12, .02]$, $p = .514$, or support, $b = .07$, $CI_{95} = [-.02, .16]$, $p = .137$), such that the hypocrisy effect got stronger as participants disagreed more with the leader's new view. More specifically, floodlight analyses showed that participants found the moral mind changer significantly more hypocritical than the pragmatic mind changer so long as participants' agreement with the leader's final position was below 1.54 (on our scale ranging from −3 to 3), at which point $b = .21$, $t(763) = 1.96$, $p = .050$. The effect never reversed: Even participants who were in complete agreement with leaders' new views (i.e., where position match = 3) found initially moral leaders (nonsignificantly) more hypocritical, simple effect $b = .10$ [−.18, .39], $t(763) = .71$, $p = .479$.

Example Study Discussion

This Example Study illustrates our general methodology and analytic approach. Its results also approximate our findings across studies as summarized below: We found small but significant effects consistent with the hypocrisy hypothesis. In other words, participants viewed initially moral mind changers as more hypocritical, and therefore as less effective and less worthy of support, compared with initially pragmatic mind changers. We also found weaker evidence that these effects might be attenuated among individuals who strongly agree with the leader's ultimate view. With these example methods and results in mind, we turn now to summarizing our entire set of 15 studies.

Summary of Methods Across Remaining 14 Studies

In this section, we summarize our methods across each of our 14 additional studies. In each case, our methods were very similar to those described in the Example Study. Therefore, we summarily describe the general procedure while noting points of difference between studies. First, however, we provide a narrative summary of the studies, in the order in which we conducted them, to provide the reader with a sense of how the research program progressed.

The first two studies we ran (Studies 1 and 2) used a business context, with a managerial leader changing his mind about whether or not to adopt a new environmentally friendly supply system. The effects we found did not reach significance, though they were in the same direction as and of a size comparable to those from the Example Study. In our next two studies, we switched over to a political context, with a leader in a small traditional society changing his mind about drug use (Study 3), or a Western political leader changing his mind about same-sex marriage (Study 4). We also, for the first time, had leaders explain their new positions, manipulating whether these explanations were moral or pragmatic. We decided a priori that if we found no significant results there, we would abandon the project. However, our results were either partially (Study 3) or fully (Study 4) significant and in support of the hypocrisy hypothesis, so in our next two studies (Studies 5 and 6) we replicated similar designs using a within-subjects manipulation, where participants read about two political leaders rather than just one. Growing confident in our findings, which continued to generally support the hypocrisy hypothesis, we designed our next study (Study 7) to test our hypotheses using realistic and ecologically valid, but experimenter-designed, attack ads, accusing a political candidate of changing his mind after first adopting a moral or pragmatic position.

At that point, given that we obtained relatively (though not perfectly) consistent results in Studies 3 through 7, we wondered whether our first two studies had shown null results because they used a business context. We reasoned that perhaps moral mind changes were seen as hypocritical especially when the mind changer was clearly reliant on public support, and could be suspected of lying to appease supporters. Therefore, in our next two studies, we returned to the environmental issue from Studies 1 and 2, manipulating whether the leader was a manager or a senator (Study 8), or whether the leader was a manager who had a high or low need for public support (Study 9). We found no evidence that these need for public support manipulations changed our findings, and the effects in these studies were again in the direction predicted by the hypocrisy hypothesis (though the effects in Study 8 were not fully significant). We next wondered whether we had been tipping the scales in favor of the hypocrisy hypothesis by having the moral mind changer appear to abandon not only his initial position, but the very basis for that position. In our next two studies (Studies 10 and 11), we therefore included additional conditions where the leaders explained their new positions using the same underlying value—for instance, the leader might have started out opposed to same-sex marriage out of a respect for American traditions, but then later decided to support same-sex marriage because he realized the more important American tradition to respect was equal rights for all. Both studies produced partial support for the hypocrisy hypothesis, and there was no evidence that the effects looked different in our new justification conditions, compared with our standard conditions.

In our final three studies (in addition to our Example Study, Study 14), we included measures of not only hypocrisy, but also courage and flexibility, in each case finding at least partial support for the hypocrisy hypothesis and none for the courageous evolution hypothesis. We conducted these additional studies, and measured courage and flexibility alongside hypocrisy, to eliminate any demand characteristics created by asking participants only about hypocrisy, and not about the potential benefits of moral mind-changes. In these studies, as in our most conservative conditions from Studies 10 and 11, we had the leader justify his new position using the same underlying principle where possible (i.e., when he changed from a moral position to a moral position, or from a pragmatic position to a pragmatic position). In the first (Study 12), we used a political context and the issue of immigration reform. In the next (Study 13), we returned to a business context, this time having a manager change his mind from one ad campaign to the next about whether to use images with vaguely sexual undertones. (In our Example Study, Study 14, recall that we used two different issues, the death penalty and same-sex marriage.)

Finally, in our last study (Study 15), we again used the issue of the death penalty, and also included two additional conditions where the leader gave different justifications for his new position: He either described a transformative experience that spurred his mind change, or blamed his change on external forces beyond his control. We explain these framings in greater detail below; we included them because existing theorizing led us to wonder whether they might help mitigate some of the damage experienced by moral mind changers. Throughout our studies, we also included additional potential moderators of our findings (e.g., participant gender, participants' moralization of the issue at hand, participants' moral relativism vs. absolutism).

Participants and Design Across Studies

A total of 5,678 participants participated in our 15 studies. We recruited all participants from Amazon's Mechanical Turk. In most studies, as in our Example Study, we included an attention check (usually it was the one described above, see Oppenheimer et al., 2009; in two studies, we used checks that tested attention to specific content in the stimuli; see SOM). In each case, we excluded data from participants who responded incorrectly (126 across our studies, or 2.2% of our total sample), prior to performing any analyses, leaving a total sample of 5,552. Designs and sample sizes for all studies are summarized in Table 1, and described in full in SOM; studies are numbered in the chronological order in which we ran them.

As in our Example Study, our primary manipulation was always the leader's initial framing. In most studies (13 of the 15), this manipulation was between subjects, as described in the Example Study: Each participant learned about a single leader who took either a moral or a pragmatic initial stance. In one of these 13 studies (Study

7), our manipulation took the form of realistic-seeming political ads attacking a particular candidate for changing his mind after having taken an initially moral or an initially pragmatic stance on the issue of same-sex marriage (see SOM for details). In two additional studies (Studies 5 and 6), we instead manipulated initial framing within subjects: Each participant read about two leaders, one taking an initially moral stance and one taking an initially pragmatic stance.

In 10 of 15 studies, as described in the Example Study, we also manipulated which side of the issue the leader switched from and to (e.g., whether he switched from supporting to opposing the death penalty, or vice versa), to ensure a wide range in terms of participants' agreement with the leader's final position. As noted in our summary above, in several studies we included additional manipulations as tests of prospective moderators; we describe those below in the Additional Procedures section.

Procedure Across Studies

As described in the Example Study, participants always began by reading (or viewing, in the case of our attack ad study) a leader's initial position on a particular issue, which we manipulated to be framed morally or pragmatically. All our framing manipulations are presented in the SOM; however, below we provide additional examples of some of our morally framed positions:

> From a moral perspective, it's the right thing to do. It doesn't matter whether you think this improves the company's bottom line; what matters most is our moral duty to do what's best for the environment. (Study 1)

> I oppose an expanded path to citizenship, and it's a moral issue for me. It's a matter of fairness and treating all people equally, and not giving unfair advantages to those who have broken the law. From a purely ethical perspective, we cannot create more paths to citizenship for undocumented residents. (Study 12)

> I do NOT think that we should go with the woman-in-a-sports-bra option for the sweatpants ad. We want to support women's rights as a moral principle, so I do not think we should do anything that contributes in any way to their objectification. (Study 13)

And of some of our pragmatically framed positions:

> From a practical perspective, it's the right thing to do. It doesn't matter whether this matters on a moral level; what matters most is doing what's best for the company's bottom line. (Study 1)

> I oppose an expanded path to citizenship, and it's an economic issue for me. Allowing more people to become citizens and paying for the costs of administration and benefits for them would simply be too costly for our economy to withstand. From a purely practical perspective, we cannot create more paths to citizenship for undocumented residents. (Study 12)

> I do NOT think that we should go with the woman-in-a-sports-bra option for the sweatpants ad. I've done the math, and it turns out that this kind of sexualized ad generates more negative than positive buzz about our products. It simply makes more business sense to avoid that. (Study 13)

In all 15 studies, participants then completed manipulation checks in which they rated the leader's message on its pragmatism and morality (these checks indicated that all manipulations were successful: Participants perceived the moral framing as relatively moral and the pragmatic framing as relatively pragmatic; see Analyses in SOM).

In 13 of 15 studies, as described in the Example Study, participants then completed baseline measures of their perceptions of leaders, described below in the Measures section. As noted above, we included these measures so that we would be able to account for any influences of the moral framing on participants' perceptions that had occurred *before* the mind change; however results were similar regardless of whether or not we controlled for these baseline measures.

Next, participants read (or viewed, in the case of our attack ad study) a second statement in which the leader expressed a change of opinion and took the opposite stance from his or her initial statement. Our later positions are also available in the SOM; however see additional examples here:

> I know I've been saying I support the new system, but I've been giving it some more thought, and *I actually do not think we should adopt this.* (Study 1)

> I know that at an earlier time I supported an expanded path to citizenship, but I've given it some more thought, and my views have changed. I oppose an expanded path to citizenship, and I'll tell you why. It's become [still] a moral issue for me, and a question of fairness and treating all people equally. I've realized [, though,] that equality means not giving unfair advantages to those who have broken the law. Putting all economic concerns aside, you have to acknowledge that a more accessible path to citizenship is not a good thing. (Study 12)

> I actually think we should go with the version of the sneaker campaign where the models are wearing FEWER clothes. I know that I objected to the sexualized image in the sweatpants campaign, but I've changed my mind. I've realized that, to put it bluntly, sex sells [It may create negative buzz, but it turns out it also increases sales by a pretty important factor]. To protect our bottom line, we have to get on board with this trend. (Study 13)

In five studies (Studies 1, 2, 7, 8 and 9), the leader provided no explanation for his mind change (see example of Study 1 above). In two studies (Studies 10 and 11), the leader always justified his later statement using the same framing (moral or pragmatic) as his initial statement. In the remaining eight studies, we manipulated whether this later position was framed morally or pragmatically, as described in the Example Study. In three of the studies where leaders did explain their mind change, they explained how their new position reflected the same underlying principle as their initial position (as in Study 12 above).

After reading this later statement, participants completed our key measures of character perceptions (hypocrisy, and in some studies courage and flexibility) and outcomes (perceptions of leaders' effectiveness, and in some studies their support for the leaders, as well as additional exploratory measures described in the Additional Procedures section below).[2]

Participants always indicated their own position on the issue the target in their study would be discussing, usually from 1, *Strongly*

[2] In some studies, participants then completed exploratory questions meant to assess their perceptions of the leader's true views following his or her mind change, for example, "If you had to guess, what do you think is [leader's] TRUE, personal opinion on [issue]?" We describe detailed results for these measures in Analyses SOM.

Oppose, to 7, *Strongly Support* as in the Example Study. In the Example Study, this measure came before the rest of the materials; in all other studies it came after the experimental stimuli and dependent measures.[3] Participants completed demographic measures at the end of each study.

Main Dependent Measures

We provide sample items for each of our measures here, and Table O in Studies SOM shows the specific wording for all items used in all of our studies.

Hypocrisy. We measured perceived hypocrisy in each study using one or two items (e.g., "[Leader] is a hypocrite," "[Leader] lacks integrity"). In Studies 1–11, the hypocrisy measure appeared intermixed with the downstream variable measures; in Studies 12–15, it appeared intermixed with courage and flexibility, after the downstream variable measures.

Courage and flexibility. We measured perceived courage (e.g., "To what extent does [leader's statement] reveal courage to admit mistakes?") and flexibility (e.g., "To what extent does [leader's statement] reveal flexibility") in Studies 12–15, using two items each.

Effectiveness. In all studies, we measured the leader's perceived effectiveness as a leader. The wording of this measure differed across our studies. In Studies 1–13, we used our own original measure (e.g., "[Leader] seems like a good leader"), and in Studies 14 and 15, we used the established measure described in the Example Study (adapted from Rice, Instone, & Adams, 1984; e.g., "How effective do you think [leader] is . . . as a leader?").

Support. In five of our studies, we measured participants' intentions to support the leader politically. In Studies 14 and 15, we adapted the established measure described in the Example Study (e.g., Paunesku, Akhtar, & Tormala, 2013; e.g., "If you were in [leader's] district, how likely would you be to . . . vote for [leader] in an election?"). In Study 7, we used our own original continuous single-item measure ("How would this ad affect your likelihood of voting for [candidate]?"). In Studies 5 and 6, which used within-subjects designs, we asked participants to indicate whether they would vote for the initially moral candidate, the initially pragmatic candidate, or neither, and used their intentions to either vote or not vote for each candidate as a binary measure of support.

Additional Procedures (Varying Across Studies)

In addition to this basic procedure and design, which were similar across all studies and allowed us to test our key hypotheses, in several studies we also included additional measures or manipulations, most of which are described in our narrative summary above. Results for these exploratory tests are summarized in the Additional Results section; as noted above, we find little evidence for any moderators of our primary findings, or for effects on exploratory measures.

Additional manipulations. Our narrative summary, along with Table 1, describes the additional manipulations we used to test exploratory moderators (see Table 1). To restate, in eight of 15 studies, we manipulated whether the leader's *later* view, post-opinion-change, was moral or pragmatic. In one study (the Example Study), we manipulated leaders' gender, to test whether people responded differently to moral mind changes in men versus women. In two studies, we manipulated the leader's reliance on popular support, to test whether hypocrisy effects would be most pronounced when leaders had a motivation to base their positions on public opinion. In two studies, we manipulated whether the leader explicitly based his ultimate opinion on the same principle as his initial opinion, thereby trying to exhibit some degree of consistency, to test whether this consistency would attenuate the hypocrisy effects. Finally, in one study, we included two additional conditions, designed in part to address an inconsistency with a recent study (Van Zant & Moore, 2015); we describe this study in greater detail below.

Additional moderating measures. In several studies we included additional measures to test for interactions with our independent variables of interest (see Table 1). In 10 of 15 studies, we measured participants' own moralization of the issue at hand (using two items from Skitka & Morgan, 2014: e.g., "To what extent is your position on the death penalty . . . a reflection of your core moral beliefs and convictions?") to test whether those who themselves moralized the issue at hand would be especially upset by moral mind changes, as might be predicted by the literature on taboo tradeoffs (Tetlock et al., 2000). In one study (the Example Study, i.e., Study 14), we measured participants' general degree of moral relativism versus absolutism, to test whether people would be more likely to view moral mind changes as courageous if they believed that multiple different moral stances could simultaneously be valid.

Additional dependent measures. Our main focus was on how leaders' mind changes affected audiences' perceptions of leaders' hypocrisy, courage, flexibility and effectiveness, and their willingness to support the leaders. However, in several studies we also measured three other perceptions of leaders as exploratory measures (see Table 1).

Commitment. In 12 studies, we measured perceptions of leaders' commitment, usually using two items (e.g., "Aaron Watson is committed to his view on same-sex marriage"; "Aaron Watson cares deeply about same-sex marriage"). Prior literature has suggested that leaders who use moral framings appear more committed to the issue (Kreps & Monin, 2014), so we wondered whether changing from a moral view would eliminate this effect, or perhaps even reverse it.

Morality. In 10 studies, we measured perceptions of leaders' morality, usually using three items (e.g., "Aaron Watson seems like a moral person"). Prior literature has suggested that leaders who use moral framings appear to have higher moral character (Kreps & Monin, 2014; Van Zant & Moore, 2015), so we wondered whether changing from a moral view would eliminate this effect, too, or perhaps even reverse it.

Surprise. In four studies, we also included a measure of participants' surprise at the mind change (two items, e.g., "Aaron Watson's second statement surprised me"), to test whether mere surprise was a key mediator of the effects on downstream percep-

[3] In studies in which we measured participants' attitudes following the experimental manipulations, we tested whether attitudes were affected by the manipulations. Generally speaking they were not, all $ps > .2$, with one exception: in Study 13, attitudes were marginally affected by the leader's position such that participants agreed more with the later position than the initial one, $M(SD) = 3.83(1.74)$ versus $3.50(1.69)$ $t(399) = 1.95$, $p = .052$.

tions, and whether, as we predicted, our theoretically derived character perceptions would explain these downstream perceptions above and beyond mere surprise.[4]

Results Concerning Our Key Hypotheses

In describing these results, we first summarize results as we obtained them, analyzing each study individually. We then present a series of analyses across the entire set of studies, which can provide a more specific estimate of all our effect sizes and answer questions regarding the consistency of that effect size across studies. This latter set of analyses provides a more definitive answer to the question of which of the effects from our individual studies are robust.

Individual Study Results

Analytic strategy. The strategy we followed when initially analyzing data from each of our individual studies was the strategy described above in our Example Study. Briefly, we used R (R Core Team, 2014) to compute regression analyses predicting each dependent measure using initial framing, participants' agreement with the leader's final position (calculated through a comparison of participants' own position with the leader's final position), and their interaction, and to test indirect effects using bootstrap mediation analyses with 1000 iterations. In Studies 5 and 6, in which our primary manipulation was within subjects, we used mixed models instead of linear regressions, and quasi-Bayesian Monte Carlo simulations instead of bootstraps. For our tests of exploratory moderators (results reported below in Additional Results), we computed additional regression analyses, each time including one exploratory moderator model at a time as an additional independent variable, and all possible interactions.

Results. All results are presented in detail in the SOM; we have summarized key results—main effects of initial framing and its interaction with position match, and tests of indirect effects—in Table 2. We generally found support for the hypocrisy hypothesis. More specifically, in 11 of 15 studies, we found that initially moral leaders were perceived as significantly more hypocritical (all 15 of the effects were in the predicted direction). In 5 of 15 studies, we found that initially moral leaders were perceived as less effective (14 of the 15 effects were in the predicted direction. And in 3 of 5 studies, we found that initially moral leaders received less support (4 of 5 effects were in the predicted direction). Regarding the courageous evolution hypothesis, one single time we found that the initially moral leader was perceived as *more* effective, and we never found that the initially moral leader was perceived as more courageous or more flexible.

Additionally, we found occasional interactions (significant in 4 of 15 studies) supporting the partisanship hypothesis, such that participants who disagreed with leaders' later views were especially likely to view the initially moral leader in more negative terms. However, in 2 of 15 studies we found interactions in the reverse direction such that it was participants who *agreed* with the leaders' later views who viewed the initially moral leader more negatively. Thus, our individual studies provide support, albeit inconsistent, for the hypocrisy hypothesis, virtually no support for the courageous evolution hypothesis, and mixed evidence regarding the partisanship hypothesis: The most common pattern we found was that people viewed moral mind changers as more hypocritical, less effective, and less worthy of support, although, again, these differences did not always reach significance.

Cross-Study Results

Analytic strategy. Having conducted these tests across our 15 studies, and having found results that differed somewhat from study to study, we were left with the question of how to evaluate the totality of these results. On the one hand, we observed effects that were nearly always (33 of 35 times) in the direction predicted by the hypocrisy hypothesis. On the other hand, we also observed multiple nonsignificant results, and results from our individual tests of the partisanship hypothesis were especially hard to interpret. What do our null results mean for the overall reliability, consistency and replicability of our findings, and for the likelihood that the hypocrisy hypothesis and/or the partisanship hypothesis is true?

As researchers call for increasingly open and transparent reporting of results to mitigate publication bias (e.g., Schimmack, 2012), several methodology experts have noted that, given the relatively low statistical power of typical psychology studies, one would expect to see some studies fail to find significant results simply because of Type II error (i.e., false negative findings). Lakens and Etz (2017), in particular, asserted in a recent paper that "mixed results are not only likely to be observed in lines of research, but when observed, mixed results often provide evidence for the alternative hypothesis, given reasonable levels of statistical power and an adequately controlled low Type 1 error rate" (p. 2). But how can we evaluate whether our particular pattern of mixed results does in fact provide evidence for any hypotheses, or whether it instead reflects the vagaries of chance?

[4] In addition to the 15 studies we describe here, we ran three additional experiments with large methodological limitations or differences from the current set of studies, and which we therefore omit from this paper because they do not provide adequate tests of the same hypotheses. If we *do* include them, the meta-analytic results remain the same, both in terms of their approximate size, their significance and their lack of heterogeneity. Nevertheless, we believe their results should not be included in an estimation of the true effect size, and we explain our reasoning here. In the very first study we ran, the manipulation check suggested that participants in the initially pragmatic condition still saw the leader as being overall more moral than pragmatic ($M = 2.89$ on a scale where 1 = *moral* and 5 = *pragmatic*). This result suggested our manipulation was not clear, and we improved it for subsequent studies. Similarly, in the first study we attempted using video attack ads (like Study 7, and run just prior to that study), 41.5% of participants (spread evenly in the two initial framing conditions) responded incorrectly to a binary manipulation check asking whether the candidate's initial view was moral or pragmatic. This result again suggested our manipulation was not clear—participant s' accuracy on this item was not far off chance levels, suggesting that even many of those who *did* respond correctly may have been guessing. We edited the videos (e.g., we slowed them down and added emphasis in key places) before using them again in Study 7. Finally, in one additional study, we used a primary dependent measure that was too different to permit comparison to our other studies; in particular we measured broad competence (e.g., how intelligent is X) which is not clearly linked to hypocrisy in the same way as leader effectiveness should be, and cannot be properly compared with our traditional measure of effectiveness as a leader specifically; we did not measure leader support.

To answer this question, we conducted two sets of cross-study analyses. First, we conducted a series of meta-analyses across our entire set of 15 studies, to evaluate whether any of the key effects predicted by our hypotheses emerged significantly across our samples. Second, we evaluated the consistency of our results across studies. The claim that, say, the hypocrisy hypothesis is likely to be true would be supported if (a) the meta-analytic tests of the effects it predicts are significant, and (b) the individual estimates of these effects in our individual studies emerge consistently with statistically comparable size. Such a significant and consistent pattern would suggest that the effects in each of our studies are individual estimates of the same true effects, whose sizes are best estimated by the meta-analyses. To formally test the consistency, we used a method recently developed by Lakens and Etz (2017), which calculates the relative likelihood of obtaining results if a given hypothesis is true versus under the null hypothesis, and we also calculated the heterogeneity of the effect sizes across studies in the form of Cochran's Q.

Finally, to visually demonstrate the lack of publication bias in our 15 reported studies, we also present funnel plots (e.g., Johnson & Eagly, 2000) for all of our main effect meta-analyses in the Analyses SOM.

Meta-analytic tests of effect size. First, we conducted meta-analytic tests for each of our key dependent measures, following Johnson and Eagly's (2000) methods for computing, weighting, and estimating overall effect sizes (pooling the data and using a multilevel modeling approach can introduce problems that the weighted mean approach avoids; Bravata & Olkin, 2001; nevertheless we also analyzed the data in this way and found statistically similar results). In particular, to test both the main effect of initial framing (for the hypocrisy and courageous evolution hypotheses) and its interaction with position match, we used their recommended approach based on F values from ANOVA. Following this approach, we ran ANOVAs for each of our individual studies, in each case including initial framing condition as our independent variable, position match as a covariate interacting with this independent variable (i.e., as a second, continuous, independent variable).[5] We recoded the direction of effects on the downstream measures such that positive values correspond to the direction predicted by the hypocrisy hypothesis (i.e., positive values reflect *less* effectiveness and support for the initially moral mind changer). To meta-analyze indirect effects, we calculated standardized indirect effects with bootstrapped confidence intervals (to evade problems of skew and kurtosis associated with standard confidence intervals; cf. Preacher & Hayes, 2004). We then treated these as rs and meta-analyzed them using Johnson and Eagly's (2000) method. All statistics used in these meta-analyses are in Tables 3 and 4.

Consistent support for the hypocrisy hypothesis. All three of our key dependent measures showed the pattern of effects predicted by this hypothesis, such that people believed an initially moral, compared with an initially pragmatic, mind changer was more hypocritical ($d = .22$ [.17, .27], $p < .001$), less effective ($d = .14$ [.08, .19], $p < .001$), and less worthy of support ($d = .11$, [.03, .19], $p < .010$). Furthermore, the meta-analysis of indirect effects showed significant indirect effects of initial framing via hypocrisy on both effectiveness ($\beta = .22$ [.17, .27], $p < .001$) and support ($\beta = .21$ [.13, .28], $p < .001$).

No support for the courageous evolution hypothesis. The courageous evolution hypothesis was contradicted by the effects on effectiveness and support in the direction consistent with the hypocrisy hypothesis, but we note that there was also no meta-analytic effect of initial framing on perceptions of courage ($d = -.06$ [$-.16$, .04]) or flexibility ($d = -.03$ [$-.13$, .07]). Thus we found no evidence that people view moral mind changers as more courageous or more flexible than their pragmatic counterparts.

Mixed support for the partisanship hypothesis, with the hypocrisy hypothesis holding robustly. The partisanship hypothesis was not supported when we examined the downstream measures: There was no interaction on effectiveness ($d = .04$ [$-.01$, .09]) or support ($d = .03$ [$-.01$, .07]). However, we did find a significant interaction when considering perceived hypocrisy ($d = -.05$ [$-.10, -.003$], $p < .050$), in the direction predicted by the partisanship hypothesis. Because of this partisanship effect, we thought it worthwhile to examine simple effects, to test whether any degree of partisanship would prevent people from viewing moral mind changers as more hypocritical. We conducted parallel ANOVAs to determine the simple effects of initial framing when position match was at its highest (3; where participants wholeheartedly agreed with the leader's new position), thus offering the most conservative test of the hypocrisy hypothesis, and then meta-analyzed these simple effects. Even in this most conservative of tests, focusing on participants who very strongly agreed with the leader's later view (at later agreement = 3), the moral leaders continued to appear more hypocritical, $d = .09$ [.04, .14], $p < .001$.

Thus, even participants who enthusiastically agreed with a leader's new position found that leader more hypocritical if he had first taken a moral stance. The effect of initial position on perceived hypocrisy, although it was smaller among participants who agreed with the new position, emerged at all levels of (dis)agreement. Moreover, people viewed moral mind changers as equally ineffective and unworthy of support regardless of their agreement with these leaders' ultimate positions. We therefore conclude that although we found one small hint of evidence consistent with the partisanship hypothesis, the hypocrisy hypothesis is what best describes people's responses to moral mind changers.

Tests for consistency and robustness. Next, we conducted two separate tests to examine the consistency of these effects. As noted above, if the effect sizes turned out to be consistent, we would not be concerned that the varying levels of significance across studies indicated true variation in the effects we

[5] In these analyses, as in the key regression analyses we report for each study, we controlled for baseline dependent measure scores in those studies where we measured them. We then followed Johnson & Eagly's (2000, p. 509) recommendation for reconstituting the error term to make these studies statistically comparable to those without a baseline measure. To check for robustness, we additionally performed two other sets of meta-analyses: one where we excluded the baseline measure covariates (also using ANOVA), and another, for the main effect tests only, based on the raw means and standard deviations in the initially moral and initially pragmatic groups. Across our key variables, results are the same in all three methods of meta-analysis. There was one instance where these methods yielded different results when it came to our exploratory variables; we note that instance in the text.

Table 3
Statistics for Meta-Analyses of Key Dependent Measures

DV	Study	N_{moral}	N_{prag}	F main	d main[a]	F interaction	d interaction[a]	Weight[b]
Hypocrisy	1	76	74	.25	.08	.24	.08	37.46
	2	98	98	.52	.10	1.70	−.18	48.94
	3	194	191	3.69	.20	2.66	−.17	95.79
	4	200	198	4.51	.21	1.15	−.11	98.94
	5	162	162	6.70	.20	2.67	.13	82.33
	6	197	197	6.22	.18	3.68	−.14	101.34
	7	142	147	4.41	.25	.02	−.02	71.68
	8	151	161	1.84	.15	.00	.00	77.69
	9	221	228	4.77	.21	1.55	.12	111.63
	10	144	155	5.46	.27	.36	.07	73.98
	11	157	147	6.18	.28	.90	−.11	75.16
	12	200	201	6.14	.25	.40	−.06	99.49
	13	203	195	11.92	.35	2.28	−.15	98.00
	14	381	388	13.38	.26	2.77	−.12	190.58
	15	386	386	6.67	.19	.74	−.06	192.17
		Meta-analytic estimates:			.22***		−.06*	
Effectiveness	1	75	74	2.17	.24	.10	−.05	36.98
	2	97	97	.92	.14	.24	−.07	48.39
	3	193	188	1.52	.13	1.84	−.14	95.05
	4	198	195	7.90	.28	.29	.05	97.27
	5	165	165	9.66	.24	1.14	−.08	82.90
	6	199	199	6.91	.19	.91	.07	101.81
	7	143	148	.82	.11	.17	−.05	72.63
	8	147	160	.01	.01	2.92	.19	76.61
	9	220	226	2.45	.15	6.11	−.23	111.18
	10	143	155	.69	.10	.71	−.10	74.29
	11	155	147	2.97	−.20	7.05	.30	75.08
	12	200	199	3.28	.18	3.91	.20	99.34
	13	201	195	5.81	.24	7.56	.27	98.26
	14	384	391	4.71	.16	.36	.04	193.15
	15	386	387	.20	.03	2.47	.11	193.22
		Meta-analytic estimates:			.14***		.04	
Support	5	165	165	5.23	.18	2.86	−.13	83.10
	6	202	202	.85	.07	.11	.02	103.00
	7	139	147	4.95	.26	1.79	.16	70.83
	14	384	391	6.93	.19	1.86	.10	192.87
	15	381	385	.26	−.04	.05	−.02	191.46
		Meta-analytic estimates:			.11**		.03	
Courage	12	200	201	.00	.00	.15	−.04	100.25
	13	202	194	16.01	−.40	.76	.09	97.01
	14	379	390	1.40	−.09	.94	−.07	192.04
	15	386	384	2.00	.10	.00	.00	192.25
		Meta-analytic estimates:			−.06		−.01	
Flexibility	12	200	201	1.39	−.12	.09	.03	100.08
	13	204	195	3.74	−.19	.42	.07	99.24
	14	377	389	.52	.05	2.40	−.11	191.39
	15	384	387	.16	.03	.05	.02	192.73
		Meta-analytic estimates:			−.03		−.05	

[a] We coded all variables such that positive *d* main effect values would be in the direction consistent with the hypocrisy hypothesis; we coded courage and flexibility such that positive main effects would be in the direction consistent with the courageous evolution hypothesis (because these measures were not used to test the hypocrisy hypothesis). [b] In estimating both the main effect and the interaction, we calculated weights for each study using Johnson and Eagly's (2000) formula, based on the sample size and the main effect *d*.

uncovered. Instead, consistent, nonheterogenous effect sizes would mean that these varying levels of significance are merely indicative of what we would expect given a small but robust effect size in underpowered studies (despite our efforts to power them adequately by averaging nearly 200 participants per cell). We focused our tests on the consistency of the effects predicted by the hypocrisy hypothesis, since these are the ones that our meta-analyses suggest are plausibly real.

Lakens and Etz likelihood ratio. Our first test applies a brand new method developed by Lakens and Etz (2017). Their technique offers a way of calculating—given a pattern of significant and nonsignificant results, a Type I error rate, and an

Table 4
Statistics for Meta-Analyses of Additional Dependent Measures

DV	Study	N_{moral}	N_{prag}	F main	d main[a]	F interaction	d interaction[a]	Weight[b]
Commitment	1	76	74	1.75	−.22	.00	.00	37.28
	2	98	98	.19	−.06	.65	.11	48.98
	3	194	188	1.18	.11	.39	−.06	95.33
	4	200	196	5.81	.24	.00	.00	98.27
	5	163	163	11.62	.27	.90	−.07	82.27
	6	201	201	.04	.01	.04	.01	102.75
	7	143	148	.72	−.10	.30	−.06	72.64
	8	150	161	.59	−.09	.00	.00	77.58
	9	221	226	.00	−.01	3.90	−.19	111.74
	10	144	155	.10	−.04	.32	−.07	74.64
	11	156	147	1.31	−.13	.11	.04	75.52
	12	200	199	2.51	.16	2.94	.17	99.44
	13	204	196	3.75	.19	2.94	.17	99.50
		Meta-analytic estimates:			.05		.001	
Morality	1	76	74	.51	−.12	.00	−.01	37.43
	2	98	98	1.13	−.15	.80	−.13	48.86
	3	195	191	.66	.08	.37	−.06	96.41
	4	200	198	11.76	.34	.81	.09	98.05
	8	151	161	.00	−.01	.05	−.02	77.92
	9	221	228	.35	−.06	1.68	−.12	112.18
	10	144	155	.31	−.06	1.59	.15	74.61
	11	157	147	2.08	−.17	5.14	.26	75.66
	12	200	201	5.01	.22	1.92	.14	99.63
	13	202	195	9.81	.31	8.29	.29	98.01
		Meta-analytic estimates:			.07*		.07	

[a] We coded commitment and morality such that positive d main effect values would reflect less commitment and less morality in the initially moral condition. [b] In estimating both the main effect and the interaction, we calculated weights for each study using Johnson and Eagly's (2000) formula, based on the sample size and the main effect d.

average statistical power—a likelihood ratio comparing the likelihood of the observed pattern emerging under the alternative hypothesis (in this case, under the hypocrisy hypothesis) to the likelihood of the observed pattern emerging under the null hypothesis. In other words, this likelihood ratio captures how much more likely the observed pattern is if the researchers' hypothesis is true than if it is false. Using this method, a likelihood ratio between 8 and 32 provides moderate evidence for an effect, and one over 32 provides strong evidence.

To calculate our likelihood ratio, we first had to calculate our post hoc statistical power in each study, using the effect sizes from our meta-analysis. These effects, although overall significant in our cross-study analyses, were smaller than we anticipated at the outset of our research—especially the effects on our downstream measures of effectiveness and support. Therefore, despite our efforts to recruit reasonable-sized samples, our average post hoc power turned out to be quite low: it was 56% for the main effect on hypocrisy, 27% for effectiveness, and 23% for support (note that, given the observed effect size for support, even a sample of 400 per cell provided only around 33% power). The resulting likelihood ratios supported the notion that these patterns of mixed significant and nonsignificant results were perfectly consistent with the hypocrisy hypothesis being true. For the effect on hypocrisy, which was significant in 11 of 15 studies, the likelihood ratio was over 16 billion (strong evidence for our hypocrisy hypothesis); for effectiveness, significant in 5 of 15, it was 329.58 (also strong evidence); for support, significant in 3 of 5, it was 63.95 (strong evidence yet again). Therefore, alongside the robustly significant meta-analytic tests, these likelihood ratios demonstrated that the null findings in some of the individual studies should not be taken as evidence of the hypocrisy hypothesis being untrue or unreliable, but rather that they reflect the (ultimately low) power of our studies, in spite of their large samples.

Tests of Cochran's Q. Another way of addressing the consistency of our findings is to test whether there is significant heterogeneity in the effect sizes across studies. For each meta-analysis, we computed Cochran's Q statistic to test for heterogeneity (e.g., Hedges, 1981). This test was not significant for hypocrisy, $Q(14) = 4.91$, $p = .987$; effectiveness, $Q(14) = 17.40$, $p = .234$; or support, $Q(4) = 7.54$, $p = .110$, nor for the indirect effect meta-analysis for effectiveness, $Q(14) = 8.36$, $p = .869$, or support, $Q(4) = 2.89$, $p = .576$. There also was no significant heterogeneity in the partisanship interaction on hypocrisy, $Q(14) = 13.08$, $p = .520$. Thus, there was no evidence that the size of these effects differed across our studies, and we concluded that any small variation was due to unsystematic sampling error. These tests further support the interpretation that, although not every individual study showed significant results consistent with the hypocrisy hypothesis, the apparent inconsistency was due to relatively low statistical power, rather than to hidden moderators, which would cause a greater diversity of effect sizes. We can therefore have confidence in the robustness of the hypocrisy hypothesis in the contexts captured across our studies.

Additional Results

Having thoroughly tested our key predictions, we now offer for the interested reader a summary of our exploratory findings. Full details are available in the SOM; overall, we found little evidence for any consistency in these exploratory findings, with one possible exception described at the end of this section.

Additional downstream measures and mediators: Commitment, morality, and surprise. Despite the fact that initially moral leaders appeared more hypocritical and less effective and worthy of support, these effects did not extend to participants' perceptions of their ultimate commitment to the issue ($d = .05$ [−.01, .11]). Initially moral leaders did ultimately appear less moral, as predicted by the hypocrisy hypothesis, $d = .07$ [.003, .14], $p < .050$. However, as we describe in Footnote 5, we ran two slightly different versions of the same meta-analyses to test for robustness, and these secondary analyses did not replicate this significant effect on morality. Thus we treat this effect with great caution. There also were no partisanship interactions on either commitment ($d = .001$ [−.06, .06]) or morality ($d = .07$ [−.01, .14]).

Perhaps there were no reliable total effects on commitment or morality because, given prior literature (e.g., Kreps & Monin, 2014), these two perceptions are likely to be strongly *positively* affected by a leader's initial moral stance. Thus, the hypocrisy of the mind change may have been enough of a blemish to eliminate these differences (as suggested also by the robust indirect effects) but not to reliably reverse them. Clearly, however, perceptions of hypocrisy do more strongly affect other, more global perceptions that many leaders care about, as evidenced in the significant total effects on effectiveness and support.

As for surprise, we did find that, overall in the four studies in which we measured this variable, initial moral framings increased surprise, $d = .16$ [.05, .27], $p < .010$. However, analyses including both hypocrisy and surprise supported the role of hypocrisy as the superior mediator. There were no significant meta-analytic indirect effects via surprise when controlling for hypocrisy, but there were still in some cases significant indirect effects via hypocrisy when controlling for surprise; in addition, in the regression analyses, hypocrisy remained a strong and consistent predictor when controlling for surprise, but the converse was not true (see Analyses SOM). Therefore, although it is true that moral mind changers are more surprising than pragmatic mind changes, we found that perceived hypocrisy in particular did a better job of explaining the effects on downstream measures, above and beyond mere surprise.

Exploratory moderators.

Generally null results across studies. We found surprisingly little evidence that any of our exploratory manipulated or measured moderators interacted with our effects of interest. Although we acknowledge that we did not test most of these individual moderators as many times as we tested most of our primary hypotheses, and therefore these null results could be masking real but small moderating effects, we meta-analyzed these exploratory moderators as well to try to increase our power (summary in Table 5; full results available in Analyses SOM). Moreover, at least in the case of the exploratory moderators that we manipulated, we have successful manipulation checks (see SOM), indicating the failure of these manipulations to moderate our findings cannot be due to a methodological failure of the manipulation to capture the conceptual variable in question. Instead, the generally null results of these analyses, in conjunction with the lack of heterogeneity we reported above for our primary hypothesis tests, suggest that the conceptual variables we considered were likely ineffective in attenuating the hypocrisy effects. To summarize:

We found no evidence that individuals who endorsed moral relativism—that is, who acknowledged that different moral views

Table 5
Statistics for Meta-Analyses of Indirect Effects via Hypocrisy

DV	Study	N_{moral}	N_{prag}	β	d	Weight
Effectiveness	1	75	74	.03	.07	37.48
	2	97	97	.05	.10	49.18
	3	193	188	.09	.18	95.87
	4	198	195	.08	.16	98.94
	5	165	165	.12	.24	82.92
	6	199	199	.13	.27	101.32
	7	143	148	.08	.16	73.98
	8	147	160	.10	.19	77.30
	9	220	226	.07	.15	111.67
	10	143	155	.15	.31	74.45
	11	155	147	.14	.28	75.72
	12	200	199	.12	.23	100.31
	13	201	195	.17	.34	98.29
	14	384	391	.16	.32	191.55
	15	386	387	.08	.16	200.85
	Meta-analytic estimate:				$.22^{***}$	
Support	5	165	165		.15	83.26
	6	202	202	.08	.12	102.80
	7	139	147	.06	.19	72.60
	14	384	391	.09	.30	192.03
	15	381	385	.15	.18	198.91
	Meta-analytic estimate:				$.21^{***}$	
Commitment	1	76	74	.02	.05	37.74
	2	98	98	.03	.07	49.47
	3	194	188	.08	.16	95.70
	4	200	196	.07	.14	98.76
	5	163	163	.10	.20	82.57
	6	201	201	.09	.19	102.30
	7	143	148	.04	.09	74.14
	8	150	161	.15	.30	76.80
	9	221	226	.07	.14	111.95
	10	144	155	.09	.18	75.30
	11	156	147	.12	.25	75.86
	12	200	199	.09	.19	100.55
	13	204	196	.16	.32	98.70
	Meta-analytic estimate:				$.18^{***}$	
Morality	1	76	74	.03	.06	37.47
	2	98	98	.05	.10	49.18
	3	195	191	.09	.18	96.11
	4	200	198	.08	.16	98.42
	8	151	161	.10	.19	77.56
	9	221	228	.07	.14	112.20
	10	144	155	.15	.30	74.06
	11	157	147	.14	.28	75.19
	12	200	201	.12	.25	99.49
	13	202	195	.22	.45	97.03
	Meta-analytic estimates:				$.19^{***}$	

[a] We coded the variables such that positive effect estimates would be in the direction consistent with the hypocrisy hypothesis. [b] We calculated weights for each study using Johnson and Eagly's (2000) formula, based on the sample size and the *d*.

could be equally valid—saw the initially moral leader as less hypocritical, effective, or worthy of support. We also found no consistent evidence that individuals who themselves moralized the issue at hand were any harsher on the initially moral leader. We also found no evidence that framing the later view morally or pragmatically mitigated the hypocrisy effects.

We found no evidence that leaders who were more reliant on popular support, and hence might be perceived as more motivated to lie to please potential supporters, were more vulnerable to the hypocrisy hypothesis.

We found no evidence that leaders who explicitly tied their later positions to the same values as in their initial moral positions were spared the negative judgment of audiences. There were no moderating effects, individually or meta-analytically, in the two studies that tested this idea directly; we also replicated the hypocrisy effects both in studies where moral-moral leaders used the same underlying value, and in studies where they did not.

Finally, we did not find consistent or meta-analytically significant interactions between initial moral framing and participant gender. We did find limited evidence that the leader's gender might matter, with the nature of this effect appearing to depend on the issue being discussed; we leave it to future research to clarify the nuances of that relationship.

One promising finding. All in all, these results paint a glum picture for initially moral leaders: When leaders have taken a moral position, there appears to be very little they can do to avoid being perceived as hypocritical should they later change their minds. The one glimmer of hope we found came in our final study (Study 15), where we included two additional later framing conditions, testing possible justifications a moral mind changer might use to mitigate the damage. In one such condition, leaders tied their (moral or pragmatic) mind change to a narrative about a personally transformative experience. For example, in one condition the leader said:

> I know that at an earlier time I supported the death penalty, but I've recently had a life-changing experience that's made me rethink things. I spent some time with a death row inmate, and saw what a truly unjust system we have. This man has been on death row for two decades. In that time he's earned a Master's degree, started a youth program to help curb school bullying, and given input on six new Congressional initiatives addressing youth crime. And we're going to kill this man? After this experience, I believe that the death penalty is not a good thing.

We thought providing such a narrative might help, because more transparent leaders are more trusted (Norman, Avolio, & Luthans, 2010); because people (at least in the United States) find personal narratives to be the most acceptable justifications (Miller & Ratner, 1998; Wuthnow, 1993); and because audiences respond positively to opinion changes when they can understand why the shift has occurred (Reich & Tormala, 2013; see also Effron & Miller, 2015).

In the second new condition in Study 15, leaders tied their apparent (moral or pragmatic) flip-flop to external circumstances, denying that their true attitude had changed; essentially, they claimed that no hypocrisy had taken place, but that they had simply been prevented from acting through no fault of their own. For example, in one condition the leader said:

> I know that at an earlier time I opposed the death penalty, and I still hold the same opinion. However, we're not going to be able to remove the death penalty from our state's justice system right now. Unfortunately, due to circumstances beyond my control, my colleagues in the legislature have refused to put this issue on our agenda for the current calendar year. As a result of this situation, this year I will have to focus on achieving our goals on the other agenda items, and I will return to the death penalty at a later time.

We included this condition as a bridge to prior literature on leaders' commitment to a course of action (e.g., Knight, 1984; Medcof & Evans, 1986; Staw & Ross, 1980). Recently, Van Zant and Moore (2015, Experiment 4) found—in stark contrast to our studies in which leaders changed their minds—that initially moral leaders who changed course due to "unanticipated economic difficulties" retained their greater initial audience support compared with initially pragmatic leaders. We wondered, therefore, whether explaining a flip-flop as merely a matter of externally constrained behavior would lead to more positive outcomes than acknowledging a change of opinion.

Detailed analyses are available in the SOM. To summarize: Although attributing a moral mind change to a transformative experience or to external constraints did not seem to make leaders seem any less hypocritical, we did find some mixed evidence that in both of these conditions, *in addition* to seeming more hypocritical, the moral mind changers *also* seemed more courageous, which indirect effects showed was associated with seeming more effective and worthy of support.

However, before recommending these two approaches to leaders preparing to change from an initially moral stance, we raise three caveats. First, these courageous evolution effects emerged alongside the hypocrisy effects, and typically did not produce significant total effects on the downstream measures (in other words, initially moral leaders still looked no better off). Second, attributing a mind change to external constraints was generally detrimental to all leaders, whether their initial position was moral or pragmatic, even if it was slightly less bad for initially moral leaders; therefore, this approach seems unwise. (In contrast, attributing a mind change to a transformative experience generally improved perceptions of all leaders, and initially moral leaders especially.) And third, these results are based on a single study, and emerged with imperfect consistency; therefore they should be viewed with much less confidence than the key findings we have emphasized thus far. Nevertheless, these findings offer the only hope we were able to uncover that people might sometimes be willing to view a moral mind change in a positive light, and we invite future research to explore them further.

General Discussion

When judging a leader who has changed her mind after taking a moral stance on an issue, audiences could potentially see the change as a sign of inspiring strength and flexibility, or of damning hypocrisy—or it could depend on whether the leader moved from "right" to "wrong" or vice versa. Our aim in this paper was to test these three hypotheses. Across our 15 studies, we found the strongest support for the hypocrisy hypothesis: Leaders who changed their moral minds were seen as more hypocritical, and not as any more courageous or flexible, than those whose initial view was amoral. They were also seen as less effective and less worthy of support, and indirect effects suggested that these effects were due to the effects on hypocrisy. Overall, these findings indicate that people believe not only that moral stances *do* endure over time, but that they *should*. The fact that we obtained these results when analyzing across our entire set of existing data (i.e., with an empty file drawer; see Footnote 4), and that they held up to several different tests of statistical robustness, makes them all the more likely to be meaningful and replicable psychological findings.

In contrast to this robust hypocrisy effect, we found few hints of support for the countervailing courageous evolution effect. In one outlier study (Study 11), initially moral leaders actually seemed more effective overall than initially pragmatic leaders (despite also seeming more hypocritical). Perhaps more interestingly, in our final study (Study 15), initially moral leaders who justified their later actions based on either inescapable external constraints or tales of personal transformation made up for their greater hypocrisy by also seeming more courageous, thus ending up *equally* effective and worthy of support. Thus, although leaders may, with effort and skill, cast their evolution in a positive enough light to overcome some of the cost of their hypocrisy, we found far more pervasive evidence to the contrary: That no matter what a leader does, and no matter who her audience is, changing her moral mind will cost her.

As indicated by our meta-analyses, the best estimates we have concerning the ultimate effects of the hypocrisy hypothesis suggest that they range from *d* of .11 to *d* of .22—from just over a tenth to nearly a quarter of a standard deviation. Is there a reason we should care about these robust, but undeniably small, effects? In fact, there are both theoretical and practical reasons why these effects are important. On the theoretical side, ours is among the only work demonstrating downsides for leaders in using moral framings. Amid a sea of research showing that moral leaders are seen as more authentic and virtuous (Kreps & Monin, 2014; Van Zant & Moore, 2015), inspiring (Brown et al., 2005; Conger & Kanungo, 1998), and reassuring (Greenberg et al., 1990; Heine et al., 2006), our studies show that initially moral leaders who change their minds actually end up worse off than if they had never taken the stance to begin with. Despite their small size, these effects add important nuance to the literature on moral leadership, and may help explain why some real-life leaders who value flexibility may be reluctant to take moral stances (Bird & Waters, 1989). On the practical side, with political candidates competing at the center of the electorate and many consequential races won by margins as tiny as a fraction of a percent of voters (for instance, multiple states in the 2016 U.S. presidential race)—the psychological processes demonstrated in our studies could potentially have large real-world effects.

The main findings in this paper make four contributions to the literatures on social psychology, leadership, and organizational behavior. First, as we have just noted, prior literature (e.g., Brown et al., 2005; Conger & Kanungo, 1998; Van Zant & Moore, 2015) has mostly focused on how leaders are better and more effective when they take moral stances; we demonstrate a boundary to this overall effect and thus a caveat to the overall conclusion. Second, we contribute to the literature on hypocrisy (e.g., Effron & Miller, 2015) by demonstrating how perceived hypocrisy can be a relevant dimension when leaders simply change their positions, depending on how their initial positions were framed—not just when they specifically guide others to behave in ways they themselves are failing to behave. Third, we contribute to the literature on lay conceptions of morality by providing the first demonstration that lay people believe attitudes based on moral values *should* endure over time. And, finally, we contribute to the leadership literature (e.g., Van Vugt et al., 2008; McAllister, 1995), which has not yet considered hypocrisy as an important dimension of leader perceptions, by showing that leaders' statements can affect their perceived hypocrisy, and that these perceptions of hypocrisy can further affect how effective leaders seem and how much potential followers would intend to support them. In the remainder of this General Discussion, we discuss further findings from our collection of studies, and explore more deeply the implications of our findings and the questions they raise surrounding lay perceptions of morality.

Exploring Possible Moderators

Across our studies, we tested a number of variables as potential moderators for our effects. Our strongest prediction was the partisanship hypothesis: that participants would view as most hypocritical leaders who changed their moral minds in the "wrong" direction, and as most courageous leaders who changed their moral minds in the "right" direction. Our data as a whole provide some support for this prediction with respect to perceptions of hypocrisy, but not downstream perceptions of effectiveness or support. We were surprised that, even when it emerged, this partisanship effect was dwarfed by the overall main effect: Participants continued to see initially moral leaders as more hypocritical than initially pragmatic ones, even when those leaders switched over to participants' own side.

In addition to this confirmatory partisanship prediction, we also tested exploratory predictions about other potential moderators: Giving a moral versus a pragmatic justification for the *later* view, the leader's gender, participants' beliefs about moral relativism versus absolutism, and their moral conviction for the issue. In short, we did not find strong evidence that any of these variables were reliable moderators of our effects (though, again, we conducted each of these tests across only a subset of our broader set of studies). Of particular note, the analyses involving later framing and participants' moral conviction help address the relationship between our work and prior work on sacred values and taboo trade-offs. If our hypocrisy effect had merely been due to taboo trade-off outrage, then we would have expected the effect to hold more strongly—or perhaps only—for audience members who moralized the issue in question, and we would have expected that giving a moral justification for the later view would eliminate the effect (bringing the mind change into the realm of tragic, not taboo, trade-offs). The fact that we found neither of these interactions, therefore, suggests that the sacred values literature cannot fully explain our findings.

Does the Hypocrisy Hypothesis Hold for Changing to a Moral Framing?

In the overview of studies, we explained why we believed the hypocrisy hypothesis would apply only to changing *from* a moral framing. In particular, we do not assume that audiences expect lifetime consistency in moral views; we propose that they view taking a moral stance, like a marriage vow, as a commitment that extends into the future but not into the past. Because of this logic, we treated later moral framing merely as an exploratory moderator of our effects of primary interest, which concern initial framing. However, a separate, tangentially relevant, and certainly interesting question is whether later moral framing exerts any *main* effects on perceptions of hypocrisy and leader support. In other words, if a leader has changed his mind—regardless of how he framed his initial stance—does taking a moral stance after the change help or hurt? The answer seems to be that, if anything, it helps: When we found effects of later framing, they indicated that moral framings made flip-flopping leaders (regardless of their initial framing) seem if anything *less* hypocritical (in 3 of 8 studies; meta-analytic $d = .16$ [.10, .22], $p < .001$), and better overall (in 9 of 18 tests predicting our downstream measures; effectiveness: meta-analytic $d = .16$ [.10, .22], $p < .001$; support: meta-analytic $d = .24$ [.16, .33], $p < .001$). Why is this pattern so different from the robust effect of initial moral framings? Much prior work suggests that moral stances help leaders in general, outside the context of a mind change, so perhaps later moral framings help simply because such framings are generally appealing. It would be interesting to explore whether this was the only reason for these later framing effects, or whether using moral framings helps specifically following a mind change—perhaps because it addresses audiences concerns about inconsistency by promising an especially consistent self from now on. (Of course, we would expect that leaders who changed their minds a second time would again be worse off if they changed from a moral stance.)

Lay Views About Moral Attitudes and Mind Changes

Our findings speak to people's everyday understandings of moral claims. In particular, they highlight two aspects of these lay beliefs.

Moral positions *should* endure into the future. We noted that our three predictions ultimately tested whether people believe that others who take moral stances *should* subsequently stick to their views, or merely that they *do* stick to their views. As we reviewed, prior research has demonstrated that people who hold a moral view tend to act consistently with that view (Aramovich et al., 2012; Luttrell, Petty, Briñol, & Wagner, 2016; Skitka et al., 2005); if people are aware of this pattern, they should believe that, descriptively, others who moralize are likely to act consistently (see Kreps & Monin, 2014). But do they go further and apply the corresponding prescriptive norm? Our findings suggest that they do: Audiences not only expect, but *demand* greater consistency from speakers who moralize, and derogate them if they fail to exhibit it. We not only document the descriptive expectation, but also are the first to show that observers hold the corresponding prescriptive norm.

Wherein lies the hypocrisy? This finding raises the question: If leaders violate this prescription by breaking from their commitment to moral stances, audience members clearly believe that some misrepresentation has occurred—but where? Do they doubt the first statement, the second, or both? In two of our studies, we asked participants to guess leaders' true views, and to report their confidence in those guesses. Although length considerations precluded detailed reports of these findings in the main text (they are in the Analyses SOM), we found that participants were generally more inclined to doubt the later view and stick to the initial one when the initial stance was morally based. On the other hand, they also saw initially moral leaders as having changed more. There was no difference in how confident participants ultimately felt about the true views of initially moral versus initially pragmatic leaders. Thus, it appears that participants are not attributing greater hypocrisy because they are especially skeptical of the initial view and see the first moral statement as mere pandering; nor do they assume that initially moral mind changers have no core self whatsoever. Instead, they appear to continue viewing the moral stance as more authentic (cf. Kreps & Monin, 2014; Van Zant & Moore, 2015), and perceive hypocrisy in the leader's dramatic switch away from this view.

We were particularly surprised to find no evidence for the perception that moral mind changers have no true core, because, in real U.S. election contexts, some leaders have been widely thought to have lost substantial support by appearing to be chronic "flip-floppers" (e.g., Capehart, 2012 on Mitt Romney). Perhaps only a leader who changed multiple times would appear to lack any core whatsoever, and a single mind change is not enough. Or perhaps, even though the change from an initially moral opinion is especially hypocritical, the personal tone of the original stance nonetheless carries over and mitigates the appearance of lacking a self. It would be interesting to examine why some politicians *do* ultimately earn reputations as chronic flip-floppers, and how their use of moral framings might contribute.

One possibility we have not yet considered empirically is that people view moral mind changers as hypocrites not because they doubt either of their *positions*, but because they doubt the *justifications* for those positions. In other words, they may not doubt that a leader who changed her moral mind previously supported the death penalty, or that she now is opposed to it, but they may doubt that her reason for initially supporting it were in fact moral: They may assume that moral commitments are so strong they are in fact impossible to break, and so therefore if she appeared to break one, it must not have been real to begin with. Our data cannot speak to this possibility, but we consider it a worthy topic for future research.

When *Is* Changing One's Moral Mind Acceptable?

Returning to our primary finding, though, what strikes us most about it is how hard it was to eliminate. Audience members continued to perceive moral mind changes as hypocritical even when we designed justifications for moral mind changes that we thought, based on prior research, audience members would view positively (although they did also appear to acknowledge leaders' courage). Returning to the marriage analogy we used in the introduction, it seems that there *must* be predictable instances when people view moral mind changes in a more unequivocally positive light. Even if, as a rule, people who abandon their marriage commitments appear hypocritical, there are certainly times—for instance, in the case of an abusive relationship—when these individuals receive praise, at least from some quarters. By analogy, it seems there must be circumstances under which a person who changes his moral mind is viewed positively, at least by some. Van Zant and Moore's (2015) finding that people continued to praise initially moral course-changers as having superior character also provides evidence that this can sometimes occur.

Knowing that these positive responses might be more elusive than the negative ones, we designed our studies specifically to make the moral mind changes seem as praise-worthy as possible. Our moral mind changers typically acknowledged their previous position, rather than not mentioning it in apparent hope that audiences would not remember. We also frequently had the mind changers provide thoughtful explanations for why they changed their minds—in many cases, we even had the moral mind-changers tie their initial and later views to the same moral value (e.g., opposing and then later supporting sexually explicit advertising based on different interpretations of feminism), so their mind change did not mean they had abandoned an underlying moral value. Even so, participants still reacted more negatively these moral reversals than to pragmatic ones. And we found no evidence that people were more forgiving of moral mind changes when they themselves did not find the moral value in question compelling, or when they believed that morality was relative and that different values can be equally valid.

Nevertheless, we state again: It must be true that *some* people *sometimes* recognize courage and open-mindedness in a moral mind change. Future research might examine moderating effects of other features of the leader (e.g., has he previously shown himself to be an exceptional moral hero?), the audience (e.g., are they intuitive thinkers or analytic thinkers?), the issue (e.g., do people view some values as more legitimately changeable than others?), or the mind change (e.g., did the leader apologize for her earlier vehemence?). One particularly promising idea might be to consider the timing of the mind-change announcement. Audience members may believe that a sincere opinion change takes time—perhaps longer if the change involves reevaluating one's previous moral values (cf. Tetlock et al., 2000, on the preference for lengthy deliberation in matters of tragic tradeoffs). Thus, they may be suspicious of relatively quick changes, believing that a sincere mind change about a weighty issue should take longer than it did for the quickly evolving leader. Thus, perhaps simply lengthening the time frame would make a moral mind change seem more acceptable. If so, clever leaders who find themselves experiencing quick but true moral mind changes (which do happen; cf. Broockman & Kalla, 2016) may do well to keep their new opinions to themselves for a time. Future research could test this possibility. That said, many of the leaders in our studies did take up to two years between their contradictory statements, so this would either need to test even longer timeframes, or explore different aspects of time. For instance, in addition to examining the role of the time between the initial and later statements, research could examine the role of how long the change itself takes—did the leader suddenly declare she had changed her mind, or did she signal over the course of weeks or even months that she was reconsidering her position?

Limitations and Future Considerations

One notable limitation of our studies is that all of them were conducted online. Our measures were all self-reported perceptions and intentions, and the format in which participants learned about flip-flops may have been different from how people learn about leaders' mind changes in other contexts (although we suspect that many people do learn about leaders' stated opinions via the Internet). However, we did our best to make the messages realistic, using, for example, the types of messages found on real politicians' websites, and political advertisement videos resembling real ads. Given that we wanted to test complex hypotheses with adequate power, we decided to use online participant sources that would allow us to recruit relatively large samples (e.g., around 800 participants for Studies 14 and 15); future research may explore similar questions using field methods.

A second notable feature is that we focused exclusively on how flip-flops affect *leaders* (not individuals in other contexts). We chose to explore these hypotheses using leaders because of the specific additional consequences that hypocrisy might have for people in leadership roles. For example, being hypocritical is likely to cost leaders their followers' trust (McAllister, 1995), damage their moral leadership influence (Brown et al., 2005), and impair their ability to solve coordination problems (Van Vugt et al., 2008). By virtue of their role, leaders are therefore likely to suffer interpersonal damage from moral mind changes beyond what people in nonleadership roles would experience. Leaders' moral mind changes are also more likely to be noticeable to others, because leaders are in a position to take public stances on organizational and government policy, and to justify these stances to large and diverse audiences. For all of these reasons, we believe that the effects we documented are especially important for leaders to understand. That said, the theoretical reasoning underlying our predictions for perceptions of hypocrisy, courage and flexibility—based on people's believe in a prescriptive norm of moral consistency, not just a descriptive norm—could have applied to regular people just as well as to leaders. Future research may examine how this same psychological mechanism guides people's responses to moral mind changers who are not leaders.

Conclusion

Leaders who take a moral stance may appear more authentic and inspiring, but they also face a risk: We find that, if they later change their minds, audiences are likely to see such leaders as especially hypocritical, and hence unlikable and unworthy of support. Far from being appreciated for their maturity and courage in acknowledging the evolution in their views, such leaders face backlash across a variety of contexts and circumstances—although, in some cases, not among audience members who believe they have switched from "wrong" to "right." In other words, people not only expect that others are likely to stick to their moral views; they see this commitment as a moral obligation and turn against those who break it. In this sense, moral talk is not cheap—those who deviate from their initial moral views pay a price by appearing ultimately more hypocritical, less effective, and less worthy of support—and leaders may do well to avoid it in the absence of true, enduring conviction.

References

Abrams, D., Randsley de Moura, G., & Travaglino, G. A. (2013). A double standard when group members behave badly: Transgression credit to ingroup leaders. *Journal of Personality and Social Psychology, 105,* 799–815. http://dx.doi.org/10.1037/a0033600

Aramovich, N. P., Lytle, B. L., & Skitka, L. J. (2012). Opposing torture: Moral conviction and resistance to majority influence. *Social Influence, 7,* 21–34. http://dx.doi.org/10.1080/15534510.2011.640199

Barden, J., Rucker, D. D., & Petty, R. E. (2005). "Saying one thing and doing another": Examining the impact of event order on hypocrisy judgments of others. *Personality and Social Psychology Bulletin, 31,* 1463–1474. http://dx.doi.org/10.1177/0146167205276430

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2014). *lme4: Linear mixed-effects models using Eigen and S4*. R package version 1.1–6. http://CRAN.R-project.org/package=lme4

Bird, F., & Waters, J. (1989). The moral muteness of managers. *California Management Review, 32,* 73–88. http://dx.doi.org/10.2307/41166735

Blader, S. L., & Rothman, N. B. (2014). Paving the road to preferential treatment with good intentions: Empathy, accountability and fairness. *Journal of Experimental Social Psychology, 50,* 65–81. http://dx.doi.org/10.1016/j.jesp.2013.09.001

Bravata, D. M., & Olkin, I. (2001). Simple pooling versus combining in meta-analysis. *Evaluation & the Health Professions, 24,* 218–230. http://dx.doi.org/10.1177/01632780122034885

Broockman, D., & Kalla, J. (2016). Durably reducing transphobia: A field experiment on door-to-door canvassing. *Science, 352,* 220–224. http://dx.doi.org/10.1126/science.aad9713

Brown, M. E., Treviño, L. K., & Harrison, D. A. (2005). Ethical leadership: A social learning perspective for construct development and testing. *Organizational Behavior and Human Decision Processes, 97,* 117–134. http://dx.doi.org/10.1016/j.obhdp.2005.03.002

Capehart, J. (2012, October 16). Obama must make Romney own flip-flops. *The Washington Post*. Retrieved from https://www.washingtonpost.com/blogs/post-partisan/post/mitt-romneys-flip-flop-character-flaw/2012/10/16/6b9158de-17ad-11e2-8792-cf5305eddf60_blog.html

Cha, S. E., & Edmondson, A. C. (2006). When values backfire: Leadership, attribution, and disenchantment in a values-driven organization. *The Leadership Quarterly, 17,* 57–78. http://dx.doi.org/10.1016/j.leaqua.2005.10.006

Christy, A. G., Kim, J., Vess, M., Schlegel, R. J., & Hicks, J. A. (2017). The reciprocal relationship between perceptions of moral goodness and knowledge of others' true selves. *Social Psychological & Personality Science*. Advance online publication. http://dx.doi.org/10.1177/1948550617693061

Cialdini, R. B., Cacioppo, J. T., Bassett, R., & Miller, J. A. (1978). Low-ball procedure for producing compliance: Commitment then cost. *Journal of Personality and Social Psychology, 36,* 463–476. http://dx.doi.org/10.1037/0022-3514.36.5.463

Cialdini, R. B., Reno, R. R., & Kallgren, C. A. (1990). A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology, 58,* 1015–1026. http://dx.doi.org/10.1037/0022-3514.58.6.1015

Cialdini, R. B., & Trost, M. R. (1998). Social influence: Social norms, conformity and compliance. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., Vol. 1 and 2, pp. 151–192). New York, NY: McGraw-Hill.

Cole Wright, J., Cullum, J., & Schwab, N. (2008). The cognitive and affective dimensions of moral conviction: Implications for attitudinal and behavioral measures of interpersonal tolerance. *Personality and Social Psychology Bulletin, 34,* 1461–1476. http://dx.doi.org/10.1177/0146167208322557

Conger, J. A., & Kanungo, R. N. (1998). *Charismatic leadership in organizations*. Thousand Oaks, CA: Sage.

Effron, D. A., & Miller, D. T. (2015). Do as I say, not as I've done: Suffering for a misdeed reduces the hypocrisy of advising others against it. *Organizational Behavior and Human Decision Processes, 131,* 16–32. http://dx.doi.org/10.1016/j.obhdp.2015.07.004

Forsyth, D. R. (1980). A taxonomy of ethical ideologies. *Journal of Personality and Social Psychology, 39,* 175–184. http://dx.doi.org/10.1037/0022-3514.39.1.175

Fredrickson, G. M. (2008). *Big enough to be inconsistent: Abraham Lincoln confronts slavery and race*. Cambridge, MA: Harvard University Press. http://dx.doi.org/10.4159/9780674033733

Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology, 96,* 1029–1046. http://dx.doi.org/10.1037/a0015141

Grant, A. (2015, November 14). The virtue of contradicting ourselves. *The New York Times*. Retrieved from http://www.nytimes.com/2015/11/15/opinion/sunday/the-virtue-of-contradicting-ourselves.html

Greenberg, J., Pyszczynski, T., Solomon, S., Rosenblatt, A., Veeder, M., Kirkland, S., & Lyon, D. (1990). Evidence for terror management theory II: The effects of mortality salience on reactions to those who threaten or bolster the cultural worldview. *Journal of Personality and Social Psychology, 58,* 308–318. http://dx.doi.org/10.1037/0022-3514.58.2.308

Hardy, C. L., & Van Vugt, M. (2006). Nice guys finish first: The competitive altruism hypothesis. *Personality and Social Psychology Bulletin, 32,* 1402–1413. http://dx.doi.org/10.1177/0146167206291006

Hastorf, A. H., & Cantril, H. (1954). They saw a game: A case study. *The Journal of Abnormal and Social Psychology, 49,* 129–134. http://dx.doi.org/10.1037/h0057880

Hedges, L. V. (1981). Distribution theory for Glass's estimator of effect size and related estimators. *Journal of Educational Statistics, 6,* 107–128.

Heine, S. J., Proulx, T., & Vohs, K. D. (2006). The meaning maintenance model: On the coherence of social motivations. *Personality and Social Psychology Review, 10,* 88–110. http://dx.doi.org/10.1207/s15327957pspr1002_1

hypocrisy. (2014). *In Merriam-Webster.com*. Retrieved September 20, 2014, from http://www.merriam-webster.com/dictionary/hypocrisy

James, F. (2014). 4 takeaways from Hillary Clinton's "Fresh Air" interview. Retrieved December 16, 2015, from http://www.npr.org/sections/itsallpolitics/2014/06/13/321773040/4-takeaways-from-hillary-clintons-fresh-air-interview

Johnson, B. T., & Eagly, A. H. (2000). Quantitative synthesis of social psychological research. In H. T. Reis & C. M. Judd (Eds.), *Handbook of research methods in social and personality psychology* (pp. 496–528). New York, NY: Cambridge University Press.

Kaiser, C. (2012). Gay marriage support: Obama's most courageous move. Retrieved June 23, 2016, from http://www.cnn.com/2012/05/09/opinion/kaiser-obama-same-sex-marriage/index.html

Kay, A. C., Gaucher, D., Peach, J. M., Laurin, K., Friesen, J., Zanna, M. P., & Spencer, S. J. (2009). Inequality, discrimination, and the power of the status quo: Direct evidence for a motivation to see the way things are as the way they should be. *Journal of Personality and Social Psychology, 97,* 421–434. http://dx.doi.org/10.1037/a0015997

Knight, P. A. (1984). Heroism versus competence: Competing explanations for the effects of experimenting and consistent management. *Organizational Behavior & Human Performance, 33,* 307–322. http://dx.doi.org/10.1016/0030-5073(84)90026-6

Kreps, T. A., & Monin, B. (2011). "Doing well by doing good"? Ambivalent moral framing in organizations. *Research in Organizational Behavior, 31,* 99–123. http://dx.doi.org/10.1016/j.riob.2011.09.008

Kreps, T. A., & Monin, B. (2014). Core values versus common sense: Consequentialist views appear less rooted in morality. *Personality and Social Psychology Bulletin, 40,* 1529–1542. http://dx.doi.org/10.1177/0146167214551154

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2016). lmerTest: Tests in linear mixed effect models. R package version 2.0–33. http://CRAN.R-project.org/package=lmerTest

Lakens, D., & Etz, A. J. (2017). Too true to be bad: When sets of studies with significant and non-significant findings are probably true. *Social Psychological and Personality Science*. Advance online publication. http://dx.doi.org/10.1177/1948550617693058

Levine, S. (2015). George H. W. Bush: I have "mellowed" on gay marriage. Retrieved June 30, 2016, from http://www.huffingtonpost.com/entry/george-hw-bush-gay-marriage_us_563b78b9e4b0307f2cac4609

Luttrell, A., Petty, R. E., Briñol, P., & Wagner, B. C. (2016). Making it moral: Merely labeling an attitude as moral increases its strength. *Journal of Experimental Social Psychology, 65,* 82–93. http://dx.doi.org/10.1016/j.jesp.2016.04.003

McAllister, D. J. (1995). Affect- and cognition-based trust as foundations for interpersonal cooperation in organizations. *Academy of Management Journal, 38,* 24–59. http://dx.doi.org/10.2307/256727

Medcof, J. W., & Evans, M. G. (1986). Heroic or competent? A second look. *Organizational Behavior and Human Decision Processes, 38,* 295–304. http://dx.doi.org/10.1016/0749-5978(86)90002-6

Miller, D. T., & Ratner, R. K. (1998). The disparity between the actual and assumed power of self-interest. *Journal of Personality and Social Psychology, 74,* 53–62. http://dx.doi.org/10.1037/0022-3514.74.1.53

Miller, Z. J. (2015, February 10). Axelrod: Obama misled nation when he opposed gay marriage in 2008. *Time*. Retrieved from http://time.com/3702584/gay-marriage-axelrod-obama/

Mullen, E., & Nadler, J. (2008). Moral spillovers: The effect of moral violations on deviant behavior. *Journal of Experimental Social Psychology, 44,* 1239–1245. http://dx.doi.org/10.1016/j.jesp.2008.04.001

Murray, S. L., Holmes, J. G., & Griffin, D. W. (1996). The benefits of positive illusions: Idealization and the construction of satisfaction in close relationships. *Journal of Personality and Social Psychology, 70,* 79–98. http://dx.doi.org/10.1037/0022-3514.70.1.79

Newman, G. E., Bloom, P., & Knobe, J. (2014). Value judgments and the true self. *Personality and Social Psychology Bulletin, 40,* 203–216. http://dx.doi.org/10.1177/0146167213508791

Norman, S. M., Avolio, B. J., & Luthans, F. (2010). The impact of positivity and transparency on trust in leaders and their perceived effectiveness. *The Leadership Quarterly, 21,* 350–364. http://dx.doi.org/10.1016/j.leaqua.2010.03.002

Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology, 45,* 867–872. http://dx.doi.org/10.1016/j.jesp.2009.03.009

Paunesku, D., Akhtar, O., & Tormala, Z. L. (2013). Weak > strong: The ironic effect of argument strength on supportive advocacy. In S. Botti & A. Labroo (Eds.), *Advances in consumer research* (Vol. 41). Duluth, MN: Association for Consumer Research.

Pew Research Center. (2015). Changing attitudes on gay marriage. Retrieved from http://www.pewforum.org/2015/07/29/graphics-slideshow-changing-attitudes-on-gay-marriage/

Preacher, K. J., & Hayes, A. F. (2004). SPSS and SAS procedures for estimating indirect effects in simple mediation models. *Behavior Research Methods, Instruments, & Computers, 36,* 717–731. http://dx.doi.org/10.3758/BF03206553

Pronin, E., Puccio, C., & Ross, L. (2002). Understanding misunderstanding: Social psychological perspectives. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 636–665). New York, NY: Cambridge University Press. http://dx.doi.org/10.1017/CBO9780511808098.038

R Core Team. (2014). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. http://www.R-project.org/

Reich, T., & Tormala, Z. L. (2013). When contradictions foster persuasion: An attributional perspective. *Journal of Experimental Social Psychology, 49,* 426–439. http://dx.doi.org/10.1016/j.jesp.2013.01.004

Revelle, W. (2014). psych: Procedures for personality and psychological research. Version 1.4.5., Evanston, IL: Northwestern University. http://CRAN.R-project.org/package=psych

Rice, R. W., Instone, D., & Adams, J. (1984). Leader sex, leader success, and leadership process: Two field studies. *Journal of Applied Psychology, 69,* 12–31. http://dx.doi.org/10.1037/0021-9010.69.1.12

Ross, L., & Ward, A. (1996). Naive realism: Implications for social conflict and misunderstanding. In T. Brown, E. Reed, & E. Turiel (Eds.), *Values and knowledge*. Mahwah, NJ: Erlbaum.

Rothstein, H. R., Sutton, A. J., & Borenstein, M. (2006). *Publication bias in meta-analysis: Prevention, assessment and adjustments*. Hoboken, NJ: Wiley.

Schimmack, U. (2012). The ironic effect of significant results on the credibility of multiple-study articles. *Psychological Methods, 17,* 551–566. http://dx.doi.org/10.1037/a0029487

Skitka, L. J., Bauman, C. W., & Sargis, E. G. (2005). Moral conviction: Another contributor to attitude strength or something more? *Journal of Personality and Social Psychology, 88,* 895–917. http://dx.doi.org/10.1037/0022-3514.88.6.895

Skitka, L. J., & Morgan, G. S. (2014). The social and political implications of moral conviction. *Advances in Political Psychology, 35,* 95–110. http://dx.doi.org/10.1111/pops.12166

Skitka, L. J., & Mullen, E. (2002). Understanding judgments of fairness in a real-world political context: A test of the value protection model of justice reasoning. *Personality and Social Psychology Bulletin, 28,* 1419–1429. http://dx.doi.org/10.1177/014616702236873

Skitka, L. J., Washburn, A. N., & Carsel, T. S. (2015). The psychological foundations and consequences of moral conviction. *Current Opinion in Psychology, 6,* 41–44. http://dx.doi.org/10.1016/j.copsyc.2015.03.025

Staw, B. M., & Ross, J. (1980). Commitment in an experimenting society: A study of the attribution of leadership from administrative scenarios. *Journal of Applied Psychology, 65,* 249–260. http://dx.doi.org/10.1037/0021-9010.65.3.249

Stone, J., & Fernandez, N. C. (2008). To practice what we preach: The use of hypocrisy and cognitive dissonance to motivate behavior change. *Social and Personality Psychology Compass, 2,* 1024–1051. http://dx.doi.org/10.1111/j.1751-9004.2008.00088.x

Tetlock, P. E., Kristel, O. V., Elson, S. B., Green, M. C., & Lerner, J. S. (2000). The psychology of the unthinkable: Taboo trade-offs, forbidden base rates, and heretical counterfactuals. *Journal of Personality and Social Psychology, 78,* 853–870. http://dx.doi.org/10.1037/0022-3514.78.5.853

Van Vugt, M., Hogan, R., & Kaiser, R. B. (2008). Leadership, followership, and evolution: Some lessons from the past. *American Psychologist, 63,* 182–196. http://dx.doi.org/10.1037/0003-066X.63.3.182

Van Zant, A. B., & Moore, D. A. (2015). Leaders' use of moral justifications increases policy support. *Psychological Science, 26,* 934–943. http://dx.doi.org/10.1177/0956797615572909

Wuthnow, R. (1993). *Acts of compassion*. Princeton, NJ: Princeton University Press.

Received December 17, 2015
Revision received April 11, 2017
Accepted April 19, 2017 ■