Reputation fuels moralistic punishment that people judge to be questionably merited

Jillian J. Jordan¹ and Nour S. Kteily²

¹Harvard Business School, Soldiers Field Road, Boston, MA 02163

² Kellogg School of Management, Northwestern University,
2211 Campus Dr, Evanston, IL 60208

Abstract

Critics of outrage culture allege that virtue signaling fuels morally questionable punishment. But does reputation actually have the power to motivate punishment that people see as ambiguously deserved? Across four studies (total $n = 9,587$), among both liberals and conservatives, we find evidence that the answer is yes. In Studies 1-2, we use a vignette paradigm to demonstrate that even in scenarios where subjects judge punishment to be questionably merited, they often expect punishing to confer reputational benefits. Across a range of such scenarios featuring politicized moral transgressions, many subjects expected punishers to be evaluated positively by co-partisans (and especially more ideologically-minded co-partisans). Furthermore, this expectation sometimes held even for individuals who *personally* questioned the merits of punishment. In Studies 3-4, we use a behavioral paradigm to investigate the motivational force of reputation in ambiguous situations. To this end, we measure decisions to punish alleged sexual harassment (among liberal subjects) and anti-male discrimination (among conservatives). In conditions where punishment was judged to be morally questionable, subjects nonetheless used punishment to boost their reputations, punishing more frequently when their behavior was public than private. In fact, when approximately equating the strength of reputational incentives, reputation boosted punishment to a similar degree in conditions where punishment was seen as ambiguously vs. unambiguously deserved (Study 3). Furthermore, reputation drove punishment even among individuals with *personal* reservations about its morality (Study 4, featuring liberal subjects). Together, these results highlight the power of reputation and have implications for debates surrounding virtue signaling and outrage culture.

Keywords: Outrage, Signaling, Ideology, Moralistic punishment, Reputation

Reputation fuels moralistic punishment that people judge to be questionably merited

In October 2017, film producer Harvey Weinstein was accused, by more than eighty women, of sexual harassment, assault, and rape. Following these allegations, Weinstein was fired from his production company, denounced by politicians whom he had supported, left by his wife, and eventually sentenced to prison (Ransom, 2020). And while severe, these outcomes were widely seen as appropriate in light of the strong evidence of Weinstein's serious wrongdoing.

Yet people do not always agree that moralistic punishment is clearly deserved. Consider the case of Al Franken, a Democratic U.S. senator who, starting in November 2017, faced a number of sexual misconduct allegations that left, for many, a more equivocal impression. An initial allegation was clearly corroborated (a photo showed Franken pretending to grope a sleeping woman), but the severity of wrongdoing that the photo depicted was widely debated. Eventually, seven more women accused Franken of misconduct, but their allegations were less clearly corroborated. For example, an unnamed former congressional aide claimed that Franken tried to forcibly kiss her and declared that doing so was his "right as an entertainer"; Franken called this allegation "categorically not true" (Caygle, 2017).

In December 2017, Franken was pushed to resign from his post as senator. However, the morality of this decision remains controversial among Democrats and Republicans alike (Mayer, 2019). In other words, the case for Franken's punishment was seen as ambiguous: while some supported it, others were unsure that it was deserved—or even convinced that it was not. The Franken example also highlights two distinct reasons that punishment might be seen as ambiguously merited: people might question whether alleged acts of wrongdoing truly occurred,

and they also might question whether alleged wrongdoing is severe enough to merit the relevant punishment.¹

Thus, contemporary events highlight that when a transgression occurs, sometimes the proposed or enacted punishments are widely seen as justified: Beyond Weinstein, consider the case of Derek Chauvin, the American police officer who sparked worldwide racial-justice protests after killing Black civilian George Floyd (Hernández & Mueller, 2020). Yet in other cases, there is far less consensus about whether the proposed or enacted punishments are morally merited: Beyond Franken, consider the case of Justine Sacco, a public relations executive who was subjected to so much online outrage after writing a tweet that many saw as racist that she became the #1 worldwide trend on Twitter—raising questions about the proportionality of this outcome (Ronson, 2015). In this paper, we explore these more ambiguous contexts. In particular, we examine the role that *reputation* plays in driving punishment that people see as ambiguously merited.

Reputation as a driver of moralistic punishment

Generally speaking, we know that people have an appetite for outrage—and that this appetite can be fueled by reputation. People are frequently motivated to condemn and punish wrongdoers (Hofmann et al., 2018), even when they have not personally been harmed (Fehr & Fischbacher, 2004; Henrich et al., 2006; McAuliffe et al., 2015; Turillo et al., 2002), and such punishment plays a key role in upholding social norms and sustaining cooperative behavior

¹Although the #MeToo movement started in the United States (and the Franken case is a U.S. example), the phenomenon of moral ambiguity surrounding the merits of moralistic punishment is a global phenomenon. For example, while many see France's #MeToo movement as reflecting moral progress, it has also received substantial backlash (e.g., it has been characterized as a “witch hunt” and a threat to “sexual freedom”). Furthermore, #MeToo cases from France have involved ambiguity both with respect to the veracity of allegations (e.g., French film star Gérard Depardieu was accused of rape but denied the allegations) and the wrongness of undisputed allegations (e.g., French TV executive Eric Brion was accused of sexual harassment and admitted to making the relevant comments but argued that they did not constitute harassment) (Onishi, 2021).

(Boyd & Richerson, 1992; Mathew & Boyd, 2011; Spring et al., 2018). Yet punishing is also costly (Balafoutas et al., 2014; Dreber et al., 2008; Nikiforakis, 2008): it can take time, effort, and resources, elicit retaliation, and risk making the punisher appear overly aggressive.

One reason to pay these costs is that punishing can also make us look *good* in the eyes of others. Indeed, we know from the robust literatures on impression management, ingratiation, and reputation (Barasch et al., 2014; Baumard et al., 2013; Boyd & Richerson, 1989; Emler, 1990; Gordon, 1996; Jones, 1965; Jones & Wortman, 1973; Leary & Kowalski, 1990; Levine et al., 2017) that the desire to be perceived positively by others has a profound impact on social behavior. More specifically, reputation and impression management concerns can shape our moral judgements, emotions, and actions (Berman & Silver, 2022; Chaudhry & Loewenstein, 2019; Critcher et al., 2013; Effron & Conway, 2015; Everett et al., 2016; Merritt et al., 2012; Silver & Shaw, 2018), and in particular drive moralistic punishment. Punishing wrongdoers can earn us social rewards (Ohtsuki et al., 2009; Raihani & Bshary, 2015) and signal to others that we can be trusted not to commit wrongdoing ourselves (Barclay, 2006; Hok et al., 2019; Horita, 2010; Jordan et al., 2016; Jordan & Rand, 2017, 2019; Nelissen, 2008). And consequentially, people punish more when doing so can boost their reputations (Jordan et al., 2016; Jordan & Rand, 2019; Kurzban et al., 2007).

Can reputation fuel punishment that is judged to be ambiguously deserved?

While we know that reputation can motivate punishment, existing evidence comes from economic game experiments investigating punishment of selfishness (Jordan et al., 2016; Jordan & Rand, 2019; Kurzban et al., 2007) in which subjects were given no reason to feel uncertain about whether the alleged selfishness actually occurred. Moreover, this research has been framed around the premise that subjects see the selfish acts in question as deserving of the relevant

104 punishments (although this assumption has not been directly tested). Thus, previous work has
105 focused on punishment that is presumed to be seen as clearly merited. Consequently, while it is
106 clear that reputation can drive people to punish, we do not know whether—and to what extent—
107 reputation has this power in contexts where people see punishment as ambiguously deserved.

108 Our paper focuses on addressing this open question, which we believe has important
109 societal implications. To date, the literature has mostly treated the influence of reputation on
110 punishment as another example of the power of reputation to inspire socially beneficial behavior.
111 When focusing on punishment that is seen as *unambiguously* merited, this framing makes sense.
112 Just like reputation can benefit society by motivating direct acts of cooperation (e.g., donations
113 to charity), reputation can motivate people to punish individuals like Weinstein—an outcome
114 that has the positive consequence of deterring conduct that is seen as clearly immoral. But if
115 reputation also has the power to drive punishment that people judge to be questionably merited,
116 it would highlight the potential for reputation to inspire outcomes that are *not* widely seen as
117 socially beneficial.

118 Indeed, this suggestion has animated contemporary debates surrounding “outrage
119 culture.” Critics have argued that, as a consequence of incentives to engage in “virtue signaling”
120 (i.e., the conspicuous showcasing of one’s moral credentials for reputational gain), we now live
121 in an age of heightened outrage (Haidt & Rose-Stockwell, 2019; Tosi & Warmke, 2016). In
122 particular, critics suggest that society has become too willing to harshly punish alleged
123 wrongdoers in ever-more ambiguous cases (Pletka, 2020)—trends that they see as imposing
124 morally unjustified costs on alleged wrongdoers (Herzog, 2018), chilling social discourse
125 (Schlosser, 2015), weakening social movements by reducing group cohesion (Pengelly, 2019),
126 and creating psychological fragility (Lukianoff & Haidt, 2015) by encouraging people to

127 pathologize everyday experiences (i.e., “concept creep” (Haslam, 2016)). Furthermore, critics
128 claim that these trends are amplified by rising ideological polarization (Bakshy et al., 2015a;
129 Barberá & Rivero, 2015) and resulting “echo chambers”, in which people can allegedly expect
130 their audiences to look favorably upon even over-zealous punishment (Haidt & Rose-Stockwell,
131 2019; Sunstein, 2019).

132 In contrast, others see the trend to punish a broader swathe of transgressions positively—
133 as necessary to hold deserving wrongdoers accountable, and as activism that can redress past
134 harm and change norms for the better (Scott, 2018; Spring et al., 2018). From this perspective,
135 insofar as reputation motivates people to hold wrongdoers like Franken accountable, it is
136 positively impacting society. Regardless of where one lands in these debates, however, it is clear
137 that our willingness to punish in more ambiguous contexts has wide-reaching consequences.

138 **Present work**

139 In this paper, we thus investigate the power of reputation to fuel moralistic punishment
140 that people judge to be questionably merited. For reputation to have this power, two key
141 psychological conditions must be met. First, in ambiguous situations where people judge
142 punishment to be questionably merited, people must expect punishing to boost their reputations.
143 And second, this expectation must actually drive people to punish. Yet as outlined below,
144 existing research leaves open the question of whether, and when, each of these conditions are
145 likely to hold. Thus, across four studies (total $n = 9,587$), we (i) investigate whether people
146 expect punishing to confer reputational benefits in ambiguous situations (Studies 1-2) and (ii)
147 evaluate the motivational force of this expectation with respect to punishment behavior (Studies
148 3-4). Together, these inquiries speak to contemporary social debates and shed novel theoretical
149 light on important psychological questions.

In ambiguous situations, do people expect punishing to boost their reputations?

We begin by asking whether people expect punishing to confer reputational benefits in ambiguous situations. In general, highlighting attitudinal similarities with others is an effective way to earn their approval (Byrne, 1961; Byrne et al., 1967). Thus, to the extent that people expect their audiences to support a given punishment, they should expect punishing to improve their reputations—even if the situation is ambiguous. The question, then, becomes: how do people think about whether others will approve of punishment in ambiguous situations?

On the one hand, ambiguous situations by definition create a lack of consensus that punishment is deserved. For example, even liberal Americans debated the merits of Senator Franken's resignation, and many were unconvinced that it was deserved (Mayer, 2019). Therefore, in ambiguous situations, many people may feel unsure about how punishment will be perceived by others—and thus question its reputation value. Furthermore, such doubts may be especially prevalent among individuals who *personally* feel unsure that punishment is morally merited. People tend to show egocentric biases when considering the perspectives of others, and are inclined to expect others to see things the way they do (Epley et al., 2004). Thus, individuals who themselves harbor reservations about the morality of a particular punishment may be less likely to expect others to approve of that punishment.

On the other hand, the lack of consensus created by ambiguous situations also means that some people *will* approve of punishment. Therefore, some potential punishers in ambiguous situations may see their audience as likely to fall into this category (and thus believe that punishing is likely to confer reputational benefits). Whether or not a potential punisher holds this belief is likely to depend not only on their own individual perspective, but also on the information that they have about their audience. For example, one might be more likely to expect

173 punishing Franken to boost their reputation in the eyes of an audience who is more strongly
174 liberal or feminist. And more generally, in contexts where people have diagnostic information to
175 suggest that their audience is likely to support punishment, they should see punishment as having
176 greater reputation value. Furthermore, the expectation that punishing will confer reputational
177 benefits in ambiguous situations may sometimes extend even to individuals who personally
178 question whether punishment is merited, insofar as some people elude egocentric bias and
179 recognize that their audience may evaluate the case for punishment differently than they do.

180 In Studies 1-2, we use a vignette paradigm to ask: in situations where punishment is
181 judged to be questionably merited, do subjects expect punishing to confer reputational benefits?
182 Furthermore, how does the answer depend on the information that subjects have about the
183 relevant audience's ideology, and thus their likely attitudes about punishment? And finally, does
184 the expectation that punishment will confer reputational benefits ever extend to subjects who
185 themselves question the morality of punishment? If so, it would highlight the power of
186 reputational incentives to create a tension between an individual's personal moral values and
187 their self-interested desire to be evaluated positively by others.

188 To preview our results, we find, across a variety of ambiguous situations, that many
189 people *do* expect punishment to confer reputational benefits, especially in the eyes of more
190 ideological audiences. Moreover, we find that this expectation can even extend to some
191 individuals who are *personally* unconvinced that punishment is deserved.

192 ***In ambiguous situations, does reputation have motivational force?***

193 Next, we turn to asking: in ambiguous situations where people expect punishing to confer
194 reputational benefits, how readily does this expectation drive punishment behavior? In other
195 words, what is the motivational force of reputation in ambiguous situations?

Consider that in *unambiguous* situations, punishment is widely judged to be merited; the main factor that limits people from punishing is merely that punishing can be costly. Thus, when reputation drives punishment, it serves to encourage most people to act in ways that *align* with their personal moral values. For example, Weinstein had a history of financially supporting liberal political figures. Insofar as these figures felt reputation-based pressure to denounce Weinstein for his transgressions, this pressure likely inspired punitive actions that most saw as morally merited (but may otherwise have been reluctant to enact, given their costs).

In contrast, in ambiguous situations, some people question the merits of punishment, creating an additional barrier to punishing. Thus, for reputation to drive these individuals to punish, it must motivate people to act in ways that are potentially *at odds* with their personal values. This may be a more difficult motivational challenge, and it is unclear whether reputation is up to the task. For example, imagine that Senator Franken has a colleague named Maria who is personally unsure that Franken deserves to be ousted, but nonetheless confidently expects that punishing Franken would, on net, be perceived positively by her audience. How readily will this reputational incentive increase Maria's likelihood of punishing?

To answer this question, it is helpful to consider the ways that Maria's moral evaluation of punishment might combine with reputational incentives to shape her behavior. To begin, there is good reason to expect Maria's moral reservations about punishment to influence her "baseline propensity" to punish (i.e., her propensity to punish in a situation where there is no potential for reputation gain). We know that many people have a strong sense of justice (Gollwitzer & van Prooijen, 2016; Turillo et al., 2002) and care deeply about their moral identities (Aquino & Reed, 2002). Therefore, if Maria is unsure that Franken deserves to be ousted, her baseline probability of punishing is likely to be lower than it would be if she saw punishment as clearly deserved.

Furthermore, given the previously-cited literature on reputation and punishment, there is also good reason to believe that, if Maria *did* see punishment as clearly deserved, reputational incentives would likely increase her probability of punishing.

But what we investigate—and what is much less clear—is whether moral evaluations about punishment constrain the motivating force of reputation. Beyond influencing her baseline propensity to punish, might Maria’s concerns about the morality of punishment also make her less sensitive, or even impervious, to the influence of reputation? For example, imagine that Maria expects punishing Franken to boost her approval rating with a key constituent group by ten percentage points. Will her moral reservations about punishment cause her to devalue those ten percentage points of reputational gain, such that they do less—or even nothing—to increase her baseline propensity to punish?

Consistent with this possibility, there is evidence that moral concerns can sometimes make people less sensitive to instrumental incentives. For example, research on moral convictions (Skitka, 2010) and “sacred values” (Tetlock et al., 2000) highlights that people believe that certain moral values should not be compromised or traded off with other considerations. In fact, offering instrumental incentives can—at least in certain settings—actually make people *less willing* to compromise their moral values (Ginges et al., 2007; Ginges & Atran, 2009). Therefore, as people move from seeing punishment as clearly merited to morally questionable, we might expect reputational incentives to become less effective at motivating punishment.

On the other hand, moral objections to punishment in ambiguous situations may not rise to the level of a conviction or sacred value. And we know that people *do* sometimes trade off their moral values against other considerations (including self-interest), especially when the

value in question is relatively less deeply-held and it is possible to provide a plausible justification to defend one's decision (Babcock & Loewenstein, 1997; DeScioli et al., 2014; Tetlock, 2002, 2003). Furthermore, previous research investigating reputational considerations in particular has demonstrated that the desire to be perceived positively by others can drive people to misrepresent their opinions in order to conform to those expressed by social partners (Jones, 1965). Together, this work demonstrates that our personal values do not always cause us to ignore countervailing incentives.

Thus, even if Maria is unsure that Franken deserves to be ousted (and is thus likely to show a lower baseline inclination to push for his resignation), the prospect of a ten-point bump to her approval rating might still increase her drive to punish—perhaps even *just as readily* as it would if she were an enthusiastic supporter of punishment (with a higher baseline drive to punish). That is, reputation may have the power to readily drive punishment, even in contexts where many people harbor reservations about its morality.

In Studies 3-4, we use a behavioral paradigm to ask: in ambiguous situations where people expect punishment to boost their reputations, how readily does this expectation drive people to punish? To preview our results, we find evidence that reputation *can* drive punishment in ambiguous situations. In fact, when approximately equating the strength of reputational incentives, we find evidence that reputation boosts punishment to a similar degree in situations where punishment is seen as ambiguously (vs. unambiguously) deserved. Furthermore, we find that the motivating effects of reputation on punishment in ambiguous situations can extend even to individuals who *personally* question its morality.

Transparency and Openness

All studies in this paper received approval from an Institutional Review Board. Furthermore, Studies 1-4 were all pre-registered (although our Study 3 pre-tests were not), and we adhered to all pre-registered sample sizes and exclusion criteria. However, the theoretical focus of this paper differs somewhat from the theoretical focuses of our pre-registrations, such that our paper addresses a somewhat different set of questions than we planned to address in our pre-registrations. Therefore, some of our pre-registered analyses are not reported in this paper, and some of the analyses reported in this paper are not pre-registered. In SM Section 6, we detail for each study (i) the ways that our reported analyses differ from our pre-registrations and (ii) the rationales underlying these differences.

We also follow the Journal Article Reporting Standards (JARS; Kazak, 2018), reporting all experimental conditions and measured variables for all studies, although there are some conditions and measures that we do not analyze in our paper. Additionally, for all studies, all materials, data, analysis scripts, and pre-registrations are publicly available on OSF: https://osf.io/2ptzq/?view_only=1da273b1968742018caa5bcd03b14798. Data were analyzed using STATA MP 16.1 for Mac.

Study 1

In Study 1, we use a hypothetical vignette paradigm to ask: in situations where punishment is seen as ambiguously merited, do people ever expect punishment to confer reputational benefits? And if so, can this expectation extend even to individuals who personally question the merits of punishment? In addressing these questions, we first focus on a particular source of ambiguity: uncertainty about whether an alleged transgression actually occurred.

Method***Design***

To these ends, we presented subjects with a series of vignettes. In each vignette, an Observer witnesses a Punisher punish a Transgressor for an alleged moral transgression. To manipulate *ambiguity*, we varied the perceived credibility of the allegation against the Transgressor, while holding constant the allegation's severity and the Punisher's punitive response. In this way, we sought to vary whether the Punisher's punishment of the Transgressor would be perceived by subjects as ambiguously vs. unambiguously deserved. Then, after presenting each vignette, we asked subjects how the Observer's impression of the Punisher would change in light of their interaction. In this way, we measured subjects' expectations regarding the reputation consequences of punishment.

To tie our work to contemporary debates surrounding outrage culture, all studies investigated moralistic punishment in the context of politicized moral issues (that often feature prominently in such debates). Our paper also focuses on reputation in the eyes of co-partisans—an audience that people interact preferentially with, are especially motivated to be perceived positively by, and are often aligned with on politicized moral issues.

For all studies in this paper, we recruited American subjects via Amazon Mechanical Turk (Mturk). Because Republicans are under-represented on Mturk, some of our studies recruited both Democrats and Republicans, while others recruited just Democrats. In Study 1, we specifically recruited Democrats, and assigned them to read vignettes featuring racism, sexism, and homophobia—transgressions that we expected Democrats to be especially sensitive to. We also always described the Observer as a Democrat, such that subjects were judging how punishing these transgressions would be perceived in the eyes of a co-partisan.

Sample

For all studies in this paper, we employed the platform CloudResearch to recruit Americans from Mturk (and used the “verify worker country location” feature to confirm subjects’ national origin). To form our final samples, we excluded duplicate responses from the same IP addresses or Mturk worker IDs as well as incomplete survey responses. In Study 1, we recruited a target of $n = 400$ Democrats; our final sample consisted of $n = 401$ Democrats ($M_{\text{age}} = 40.31$ years, 46.38% male, 66.33% with a college or graduate degree).²

Procedure

We began by asking subjects to report their age, gender, and preferred political party affiliation (Democrat or Republican); subjects who identified as Republicans were redirected to a different study. Each subject was then assigned to one of our two ambiguity conditions, and then read and evaluated four vignettes (in random order) that all corresponded to that condition.

After presenting each vignette, we measured subjects’ initial evaluation of it. Specifically, we collected two measures of the perceived reputation value of punishment, as well as a set of secondary dependent variables. Next, we presented the four vignettes again (in a new random order); this time, subjects answered a set of questions relevant to their own perceptions of the moral merits of punishment. Finally, subjects completed an exit survey (containing demographic and other questions that are not analyzed in our main text; see SM Section 1.2 for details), and then were debriefed.

² We note that, in Study 1 (but no other studies in this paper), we presented an attention check at the beginning of the study (i.e., before subjects were assigned to conditions), which we describe in SM Section 1.2. Following our pre-registration, our Study 1 sample excludes subjects who failed this attention check. (This design feature was present in Study 1 but no other studies because Study 1 was conducted during a time period in which we were more concerned about inattentiveness in the Mturk subject pool.)

Vignette content. Each vignette begins by introducing an individual who works at an organization, and has been accused by a coworker of wrongdoing. We refer to this individual as the “Transgressor” in this paper, although we did not use this terminology in our vignettes. Across our four vignettes, the Transgressor was accused of (i) making a racist comment about a coworker, (ii) making homophobic remarks about a gay politician, (iii) making a sexist comment about a coworker, or (iv) wearing a racist Halloween costume.

Next, in the *ambiguous* condition of each vignette, subjects learn that the Transgressor denied that the transgression actually occurred, and proposed an alternative explanation for why he or she was being accused. Specifically, in two vignettes, the Transgressor suggested that the accuser may have made a perception error, and in two vignettes, the Transgressor suggested that the accuser may have been motivated by a rivalry between the two of them. In contrast, in the *unambiguous* conditions, subjects learn that the Transgressor heard about the allegation but did not respond to it. Afterwards, in both conditions, subjects learn that another individual at the organization—who we refer to in this paper as “the Punisher”—chose to exclude the Transgressor from a social event on the basis of the allegation.

For example, in our Democrat vignette about homophobia, *Sam* is the Punisher and *Brett* is the Transgressor. The vignette begins by introducing Sam and Brett, and describing the allegation against Brett:

Sam works at a medium-sized organization. One of his coworkers is named Brett. Sam and Brett are both straight. Recently, Brett was accused of homophobia. A gay employee at the organization named Alex claimed that he overheard Brett speak in a homophobic way about a gay political candidate. Specifically, Alex claimed that he overheard Brett and a colleague discussing an upcoming political election while watching TV coverage about the field of

353 *candidates. One of the candidates is openly gay, and the TV coverage featured a conversation in*
354 *which this candidate and his husband were laughing and joking around. Alex reported that*
355 *during this TV coverage, Brett—who had previously expressed that he found this particular*
356 *candidate really annoying—remarked “ugh, not these fucking gays.”*

357 In the ambiguous condition, the vignette then continues: *Brett responded to this*
358 *allegation by denying it. He insisted that he is not homophobic and did not reference the*
359 *candidate’s sexual orientation, but instead simply said “ugh, not these fucking guys.”* In the
360 unambiguous condition, the vignette instead continues: *Brett heard about this allegation against*
361 *him but did not respond to it.* Next, in both conditions, the vignette continues: *Sam learned about*
362 *the allegation against Brett and Brett’s reaction to it. The next day, Sam approached Brett about*
363 *a barbecue he was organizing for colleagues, and had previously invited Brett to. Sam said to*
364 *Brett: “The way you talked about that politician was really homophobic and offensive—I think*
365 *it’s best if you don’t come to the barbecue.”*

366 Finally, subjects learn that they will be answering questions about a third coworker, the
367 Observer, who is always described as “a Democrat.” Subjects then learn that the Observer
368 overheard the Punisher punishing the Transgressor. For example, in the homophobia vignette,
369 the vignette concludes: *Today, we’d like you to answer a few questions about another person*
370 *named Anthony. Anthony is a Democrat, and also works with Sam and Brett. Anthony learned*
371 *about the allegation against Brett and Brett’s reaction to it, and also overheard when Sam told*
372 *Brett not to attend the barbecue.*

373 For full vignette texts (across Studies 1 and 2), see SM Sections 7-8.

374 **Measures.** After reading each vignette, subjects completed an initial evaluation of that
375 vignette. Our primary DVs were two questions pertaining to the reputation value of punishment.

In the context of our homophobia vignette, subjects were asked: *After Anthony overheard Sam tell Brett not to attend the barbecue, how did Anthony's overall impression of Sam change?* We then measured responses to this question in two ways. First, in our *point-estimate* reputation measure, subjects provided their “best guess” answer on a 0 to 100 sliding scale (0 = Anthony's impression of Sam became a lot worse, 50 = No change, 100 = Anthony's impression of Sam became a lot better). Second, in our *probability-distribution* reputation measure, subjects assigned a percentage probability to five potential answers to the above question. Specifically, subjects were asked to consider the possibility that Anthony's impression of Sam (1) “Became a lot worse”, (2) “Became a little worse”, (3) “Did not change”, (4) “Became a little better”, or (5) “Became a lot better”, and were required to allocate a total of 100 percentage points across these five options. We then computed the mean of each subject's probability distribution and treated this value as our second primary DV.³

After reading and completing an initial evaluation of all four vignettes, subjects evaluated each vignette for a second time (in a new random order). This time, we asked subjects to provide their own evaluations of three questions relevant to the moral merits of punishment. In particular, using 0 to 100 sliding scales, we asked subjects to rate (i) the *likelihood* that the allegation was true (e.g., in the context of our homophobia vignette, “How likely do you think it is that the allegation against Brett is true?”; 0 = Not likely at all, 50 = Somewhat likely, 100 = Very likely), (ii) the *appropriateness* of punishment (e.g., “How appropriate do you think it was for Sam to call Brett out?”; 0 = Very inappropriate, 50 = Somewhat appropriate, 100 = Very appropriate), and (iii) the extent to which punishment was *justified* (e.g., “How justified do you think Sam was

³After measuring our primary DVs, we also measured a series of secondary DVs investigating (i) subjects' uncertainty versus confidence about the reputation consequences of punishment, and (ii) their expectations about how punishers would be perceived on a set of more specific traits (e.g., political partisanship). See SM Section 1.1 for more details about these variables, which we do not analyze in this paper.

in telling Brett not to attend the barbecue?”; 0 = Not justified at all, 50 = Somewhat justified, 100 = Very justified). These items formed a reliable scale ($\alpha = .92$), and an exploratory factor analysis of the three items produced one-factor solution; thus, in our analyses, we average them together to form a composite measure of the perceived moral merits of punishment. See SM Section 2.2 for factor analysis results.

Results

All analyses in this paper use linear regression. In our primary analyses of Study 1, we aggregate across vignettes, treat each evaluation of one vignette by one subject as a distinct observation, and cluster our standard errors on subject to account for the non-independence of repeated observations from the same subject. However, we also confirm that all results hold significantly within each individual vignette; see SM Section 1.3 for details.

Did our ambiguous conditions create moral ambiguity?

We begin by investigating whether the ambiguous condition of Study 1 caused subjects to see the case for punishment as morally ambiguous—both in absolute terms, and relative to the unambiguous condition. To this end, we analyze the perceived moral merits of punishment. We find that subjects in the ambiguous condition indeed saw punishment as less merited ($M = 63.54$, $SD = 23.01$) than subjects in the unambiguous condition ($M = 78.34$, $SD = 19.51$), $b = -14.81$ [$-18.34, -11.28$], $B = -.33$, $t = -8.25$, $p < .001$, $n = 401$.

Of note, even in the ambiguous condition, mean ratings of the merits of punishment were above our scale midpoint (which corresponded to characterizing the allegations as “somewhat” likely to be true, and punishment as “somewhat” appropriate and justified). However, our results nonetheless suggest that the ambiguous condition created a lack of consensus among subjects that punishment was deserved. Supporting this conclusion, in Figure 1A we plot the distribution

420 of ratings of the merits of punishment across ambiguity conditions. This figure highlights that (i)
421 our ambiguous condition created substantial heterogeneity across subjects, (ii) subjects in this
422 condition frequently provided ratings around the scale midpoint, suggesting middling
423 perceptions of the moral merits of punishment, and (iii) a substantial minority of ratings were
424 below the scale midpoint, suggesting moral disapproval of punishment.

Figure 1. In ambiguous situations where many people question the merits of punishment, it is nonetheless common to expect punishment to confer reputational benefits. (A) Across the ambiguous conditions of Studies 1–2, there was a lack of consensus among subjects that punishment was deserved—both in absolute terms, and relative to our unambiguous conditions. We plot ratings of the moral merits of punishment as a function of ambiguity. **(B)** Subjects in our ambiguous conditions nonetheless frequently reported that punishment would confer reputational benefits, especially in the eyes more ideological audiences. We plot ratings of the reputation value of punishment (as measured by our point-estimate reputation measure) within our

ambiguous conditions, as a function of Observer ideology (which we held constant in Study 1, but manipulated in Study 2). (C) This expectation sometimes held even when subjects were *personally* unconvinced that punishment was merited. Panel C replicates Panel B, but restricts to observations for which subjects rated the moral merits of punishment at or below the scale midpoint. Across panels, we separately plot results from Study 1 (Democrats), 2a (Democrats), and 2b (Republicans).

In our ambiguous condition, did subjects expect punishment to confer reputational benefits?

Thus, the ambiguous condition of Study 1 created a lack of consensus among subjects about whether punishment was merited. Next, we ask: did subjects in this condition ever nonetheless expect punishment to confer reputational benefits (i.e., to improve Observers' impressions of Punishers)? To address this question, we turn to investigating reputation ratings.

We find that subjects in the ambiguous condition, on average, expected Punishers to earn reputational *benefits* in the eyes of the Observer. Supporting this claim, mean ratings of punishment's reputation value were significantly above the scale midpoint of 50 for our point-estimate measure ($M = 57.76$, $SD = 21.58$, $t = 6.57$, $p < .001$, $n = 194$), and above the midpoint of 3 for our probability-distribution measure ($M = 3.27$, $SD = .93$, $t = 4.90$, $p < .001$).

In Figure 1B, we plot the distribution of reputation ratings within the ambiguous condition (focusing on our point-estimate measure for reasons of brevity; we find qualitatively identical results for our probability-distribution measure). This figure highlights that subjects in our ambiguous condition frequently provided reputation ratings above the scale midpoint, reflecting the expectation that punishment would confer reputational benefits. More specifically, 61.73% of ratings were above the scale midpoint.

Thus, subjects in the ambiguous condition of Study 1 *did* often expect punishment to confer reputational benefits. We also note that, while our theoretical focus in this paper is on the perceived reputation value of punishment in our *ambiguous* condition, Study 1 also revealed that

subjects in the unambiguous condition expected punishment to confer even *larger* reputational benefits; see SM Section 1.4 for more details about this comparison.

Next, we ask whether subjects in our ambiguous condition ever expected punishment to confer reputational benefits despite *personally* questioning the merits of punishment. In Figure 1C, we again plot the distribution of reputation ratings within our ambiguous condition (and again focus on our point-estimate measure). This time, however, we restrict to the subset of observations for which subjects reported questioning punishment's morality, as indicated by ratings of the moral merits of punishment at or below the scale midpoint. This subset constituted 29.26% of observations in the ambiguous condition ($n = 227$ observations across 108 unique subjects). And for 34.80% of this subset, subjects provided above-midpoint reputation ratings. Thus, even when subjects harbored *personal* reservations about the merits of punishment, they nonetheless sometimes expected punishing to confer reputational benefits.

Discussion

In sum, Study 1 provides evidence that, across a range of ambiguous situations, subjects frequently expect punishment to confer reputational benefits. Furthermore, this expectation sometimes holds even among individuals who are personally unsure that punishment is deserved.

Study 2

In Study 2, we replicate and extend these results in three important ways. First, we recruited both Democrats (Study 2a) and Republicans (Study 2b). Second, we shifted our theoretical focus to a different source of moral ambiguity. Whereas Study 1 focused on ambiguity surrounding whether an alleged transgression actually occurred, in Study 2 we considered ambiguity surrounding whether a transgression that definitely occurred was severe enough to merit the punishment it received. To this end, we sought to vary the perceived

offensiveness of a transgression while holding constant the punishment it received (and always making clear that the transgression did actually occur). Study 2 thus allows us to ask: when it is unclear that a transgression is severe enough to merit a particular punishment, do subjects still frequently expect that punishment to confer reputational benefits?

Third, Study 2 investigated how information about the Observer's ideology would influence the perceived reputation value of punishment. In Study 1, we provided *some* such information by always describing the Observer as a Democrat. However, we did not give subjects any further information about the Observer's likely leaning with respect to the moral domain at hand. For example, in the context of our homophobia vignette, subjects did not learn anything about the Observer's attitude towards gay rights, or receive any information about the strength of the Observer's partisanship. In Study 2, we explored the impact of such information.

Method

Design

In Study 2a, we again recruited Democrats to read vignettes about racism, sexism, and homophobia. In Study 2b, we recruited Republicans to read about a different set of transgressions, that we expected conservatives to be especially sensitive to: disrespect towards religion, veterans, and America.

Study 2 employed a two-by-two, between-subjects design, in which, in addition to manipulating *ambiguity*, we also manipulated the Observer's *ideology*. In all conditions, we described the Observer as a Democrat (in Study 2a) or a Republican (in Study 2b). And in our "less ideological" conditions, mirroring Study 1, this was the only description of the Observer that we provided. In our "more ideological" conditions, however, we described the Observer as a

strong Democrat or Republican, and also provided other information suggesting that they were relatively ideologically-minded with respect to the relevant moral domain.

Samples

We recruited a target of $n = 1600$ subjects in each of Studies 2a-b. Our final samples consisted of $n = 1591$ Democrats in Study 2a ($M_{\text{age}} = 36.60$ years, 44.63% male, 57.01% with a college or graduate degree) and $n = 1566$ Republicans in Study 2b ($M_{\text{age}} = 40.97$ years, 45.40% male, 52.62% with a college or graduate degree).

Procedure

As in Study 1, we began by asking subjects to report their age, gender, and preferred political party affiliation (Democrat or Republican). Based on their reported affiliation, we directed subjects to Study 2a or 2b. The procedure for Studies 2a-b was then identical to the procedure for Study 1, with the following exceptions.

First, we adapted our Study 1 vignettes and ambiguity manipulation. Specifically, in Study 2a, we modified our Study 1 vignettes so that the Punisher always witnesses the Transgressor behave in a way that the Punisher believes is offensive. In this way, our modified vignettes make clear to subjects that the transgression in question definitely occurred.

In the *unambiguous* condition of each Study 2a vignette, the Transgressor commits a transgression that is similar to the alleged transgression from Study 1, and is designed to be perceived by subjects as overtly offensive. Thus, across our four vignettes in Study 2a, the Transgressor (i) makes a racist comment about a coworker, (ii) makes homophobic remarks about a gay politician, (iii) makes a sexist comment about a coworker, or (iv) mentions having recently worn a racist Halloween costume. In the *ambiguous* condition of each vignette, the

Transgressor commits a version of these same transgressions that was designed to be perceived by subjects as less overtly offensive.

In Study 2b, we created a set of corresponding vignettes tapping moral issues that we expected conservatives to be particularly sensitive to. Specifically, we created a set of vignettes in which a Transgressor (i) makes a negative comment about religion to a child, (ii) behaves disrespectfully during a Veteran's Day ceremony, (iii) makes a joke about the September 11th attacks, or (iv) derogates an American flag. For each vignette, we likewise sought to manipulate the perceived offensiveness of the Transgressor's behavior.

For example, in both ambiguity conditions, our homophobia vignette (featured in Study 2a) began as follows: *Last week, Sam was having coffee with his coworker, Brett. While chatting, they began to discuss an upcoming political election and the field of candidates, including a specific candidate who is openly gay. Sam and Brett are both straight.*

In the unambiguous condition, the vignette continued: *During their conversation, Brett made a joke about the candidate being gay, stating that he was really starting to "get behind him" and then saying "that's what he said!". Then, after laughing at his own joke, Brett said "no, but seriously, it's pretty wild that a fag is running."*

In the ambiguous condition, the vignette instead continued: *During their conversation, Brett said that he was honestly very impressed by the field of candidates and could see himself voting for any of them. Then, he made a joke about the candidate being gay, stating that he was really starting to "get behind him" and then saying "that's what he said!".*

Next, the Punisher responds by verbally condemning the Transgressor. The Punisher always uses identical language across both ambiguity conditions. In the homophobia vignette, the vignette continues: *Sam responded by calling Brett out, saying "that's really homophobic*

550 *and offensive.*” Thus, in Study 2, the relevant punitive response is not social exclusion (like in
551 Study 1) but rather verbal condemnation. We note that we nonetheless continue to use the terms
552 “Punisher” and “punishment” throughout Study 2, both for consistency with Study 1 and because
553 we see verbal condemnation as an important (although often relatively less severe) way that
554 people “punish” each other in daily life.

555 After learning that the Punisher verbally condemned the Transgressor, subjects learn that
556 they will be answering questions about a third coworker, the Observer. In the *less ideological*
557 condition, the Observer is simply described as either “a Democrat” (in Study 2a) or “a
558 Republican” (in Study 2b). In the *more ideological* condition, the Observer is described as “a
559 strong Democrat” or “a strong Republican.” Furthermore, we provide some more specific
560 information designed to suggest that the Observer is relatively ideologically-minded with respect
561 to the moral domain of the transgression (see below for an example). Subjects then learn that the
562 Observer overheard the interaction between the Transgressor and the Punisher.

563 In the less ideological condition of the homophobia vignette, the vignette continues:
564 *Today, we’d like you to answer a few questions about another person named Anthony. Anthony*
565 *is a Democrat. Anthony also works with Sam and Brett. While Sam and Brett were having their*
566 *coffee conversation, Anthony (who, out of view, was standing in line for a coffee) happened to*
567 *overhear them. So Anthony ended up hearing what Brett said about the candidate, and hearing*
568 *Sam call Brett out.* In the more ideological condition, the Observer is instead described as
569 follows: *Anthony is a strong Democrat who strongly supports LGBT rights.*

570 Studies 2a-b employed the same measures as Study 1, with the exception that we
571 modified two of our three questions relevant to the moral merits of punishment to align with our
572 modified ambiguity manipulation. Specifically, instead of asking subjects about how likely the

transgression was to have occurred, we asked them how *offensive* the transgression was (e.g., in the context of the homophobia vignette, “How offensive do you think what Brett said was?” 0 = Not offensive at all, 50 = Somewhat offensive, 100 = Very offensive). Furthermore, instead of asking subjects how justified the Punisher was, we instead asked them how *proportionate* the punishment was (e.g., “Do you think that Sam’s approach to calling Brett out was too mild, proportionate, or too harsh?” 0 = Way too mild, 50 = Proportionate, 100 = Way too harsh). We also measured the perceived *appropriateness* of punishment, as in Study 1. Again, our three items formed a reliable scale (Democrats $\alpha = .78$, Republicans $\alpha = .81$), and exploratory factor analyses of the items produced one-factor solution for both Democrats and Republicans; thus, we average them together (after reverse-scoring the item assessing perceived proportionality) to form a composite measure. See SM Section 2.2 for factor analysis results.

Results

We analyze Studies 2a-b as we analyzed Study 1 (again aggregating across vignettes and confirming that all results hold significantly within each individual vignette in SM Section 1.3).

Did our ambiguous conditions create moral ambiguity?

We again begin by investigating the perceived moral merits of punishment. We find that subjects in the ambiguous conditions indeed saw punishment as less merited (Democrats $M = 58.76$, $SD = 15.03$; Republicans $M = 56.86$, $SD = 15.86$) than subjects in the unambiguous conditions (Democrats $M = 71.81$, $SD = 12.11$; Republicans $M = 70.01$, $SD = 13.68$), Democrats $b = -13.05$ $[-14.11, -11.99]$, $B = -.43$, $t = -24.17$, $p < .001$, $n = 1591$; Republicans $b = -13.15$ $[-14.34, -11.97]$, $B = -.41$, $t = -21.73$, $p < .001$, $n = 1566$. We also plot the distribution of ratings of the merits of punishment across ambiguity conditions in Figure 1A.

These analyses, and the distributions plotted in Figure 1A, highlight that the ambiguous conditions of Study 2 successfully created perceived moral ambiguity—both in absolute terms, and relative to our unambiguous conditions. Mirroring Study 1, there was again a lack of consensus among subjects in our ambiguous conditions that punishment was morally merited, this time owing to uncertainty about whether the transgressions in question were severe enough to merit the punishment they received.

In our ambiguous conditions, did subjects expect punishment to confer reputational benefits?

When faced with this ambiguity, did subjects expect punishment to be perceived positively? And how did the answer depend on Observer ideology?

Turning to our reputation measures, we again find that subjects in our ambiguous conditions did, on average, expect Punishers to earn reputational benefits—both in the eyes of more and less ideological Observers. Specifically, in the less ideological condition, ratings were significantly above the scale midpoint of 50 for our point-estimate reputation measure (Democrats $M = 58.14$, $SD = 21.85$, $t = 10.14$, $p < .001$, $n = 379$; Republicans $M = 59.10$, $SD = 22.35$, $t = 11.00$, $p < .001$, $n = 409$), and above the midpoint of 3 for our probability-distribution reputation measure (Democrats $M = 3.29$, $SD = .96$, $t = 7.71$, $p < .001$; Republicans $M = 3.30$, $SD = .99$, $t = 7.76$, $p < .001$). Similarly, in the more ideological condition, ratings were significantly above the scale midpoints both for our point-estimate (Democrats $M = 67.02$, $SD = 21.92$, $t = 20.57$, $p < .001$, $n = 407$; Republicans $M = 64.58$, $SD = 24.01$, $t = 15.71$, $p < .001$, $n = 377$) and probability-distribution (Democrats $M = 3.76$, $SD = .94$, $t = 20.18$, $p < .001$; Republicans $M = 3.55$, $SD = 1.02$, $t = 13.15$, $p < .001$) reputation measures.

Furthermore, within our ambiguous conditions, reputation ratings were significantly higher in the more (vs. less) ideological condition, both for our point-estimate (Democrats $b =$

8.88 [6.62, 11.14], $B = .20$, $t = 7.71$, $p < .001$, $n = 786$; Republicans $b = 5.48$ [3.04, 7.92], $B = .12$, $t = 4.41$, $p < .001$, $n = 786$) and probability-distribution (Democrats $b = .46$ [.36, .57], $B = .24$, $t = 8.65$, $p < .001$; Republicans $b = .26$ [.15, .37], $B = .13$, $t = 4.54$, $p < .001$) measures. Thus, describing the Observer as more ideological with respect to the relevant moral domain served to increase the perceived reputation value of punishment.

In Figure 1B, we plot the distribution of reputation ratings within our ambiguous conditions (again focusing on our point-estimate measure for reasons of brevity; we find qualitatively identical results on our probability-distribution measure). This figure highlights that (i) subjects frequently provided ratings above the scale midpoint, reflecting the expectation that punishment would confer reputational benefits, and (ii) this expectation was especially common when the Observer was more ideological. More specifically, 64.25% of Democrat ratings and 66.50% of Republican ratings were above the scale midpoint in the less ideological condition, and 79.05% of Democrat ratings and 74.87% of Republican ratings were above the scale midpoint in the more ideological condition.

Thus, subjects in our ambiguous conditions frequently expected punishment to confer reputational benefits, especially in the eyes of more ideological audiences. Furthermore, while our focus in this paper remains on the perceived reputation value of punishment in our *ambiguous* condition, we again show in SM Section 1.4 that subjects in our unambiguous condition expected punishment to confer even *larger* reputational benefits.

Finally, we turn to asking: how frequently did subjects in the ambiguous condition of Study 2 expect punishment to confer reputational benefits, despite *personally* questioning the morality of punishment? In Figure 1C, we plot the distribution of reputation ratings (on our point-estimate measure) within the ambiguous conditions of Study 2, restricting to the subset of

observations for which subjects provided ratings of the moral merits of punishment at or below the scale midpoint. This subset constituted 30.34% of Democrat observations ($n = 954$ observations across 495 unique subjects) and 33.08% of Republicans observations ($n = 1040$ observations across 520 unique subjects) within our ambiguous conditions.

Looking to this subset of observations, we find that subjects expected punishment to confer reputational benefits a meaningful fraction of the time. In the less ideological condition, we observe above-midpoint ratings on our point-estimate reputation measure for 43.14% of the subset of Democrat observations and 45.83% of the subset of Republican observations. And in the more ideological condition, we observe above-midpoint reputation ratings for 60.75% of the subset of Democrat observations and 62.11% of the subset of Republican observations. Thus, even when subjects reported personal reservations about the merits of punishment, they nonetheless expected punishing to confer reputational benefits a substantial proportion of the time—especially when the Observer was described as more ideological.

Discussion

Study 2 replicates and extends the results of Study 1. In the ambiguous conditions of both Studies 1 and 2, we successfully created a lack of consensus among subjects that punishment was deserved. Yet subjects in these conditions nonetheless frequently expected punishing to look good in the eyes of co-partisans. And in Study 2, describing the co-partisan audience as more ideological with respect to the relevant moral domain served to heighten this expectation.

Studies 1 and 2 also reveal that despite the human tendency towards egocentric biases in perspective taking (Epley et al., 2004), even individuals who *personally* question the merits of punishment sometimes expect punishment to confer reputational benefits. In this way, our results

highlight the potential for reputation to create a tension between an individual's personal moral values and their self-interested desire to be seen positively by others.

Importantly, Studies 1-2 provide evidence for these conclusions across a variety of politicized moral domains (e.g., racism, disloyalty towards America) and among both liberal and conservative subjects. And our conclusions appear robust to the source of moral ambiguity: subjects expected punishment to look good both when it was unclear that (i) an alleged transgression actually occurred (in Study 1) and (ii) a transgression that clearly occurred was severe enough to merit its punishment (in Study 2).

Study 3

In Study 3, we turn to asking: when people expect punishment to confer reputational benefits in ambiguous situations, are they willing to *act* on this expectation? To address this question, we investigate whether reputational incentives drive punishment behavior in ambiguous situations, where many individuals judge punishment to be morally questionable.

We also compare the motivational force of reputation across ambiguous vs. unambiguous situations. Putting reputation aside, we would expect people to differ in their baseline proclivities to punish across ambiguous vs. unambiguous situations, on the basis of their personal moral evaluations of punishment. But will moral reservations about punishment also serve to limit the power of reputational incentives to increase proclivities to punish? To address this question, we investigate the relative power of reputation to increase rates of punishment in situations where punishment is seen as questionably vs. clearly merited, when approximately equating the strength of perceived reputational incentives.

Method***Design***

To these ends, we designed a behavioral paradigm in which subjects made punishment decisions that they believed had real material costs, and meaningful social consequences. Our paradigm drew inspiration from the #MeToo movement, a hotbed for debates surrounding outrage culture and an archetypal context in which the case for punishment may be more or less ambiguous. Furthermore, our paradigm simultaneously incorporated both sources of ambiguity that we investigated in Studies 1-2—creating a strong test of whether reputation can fuel punishment in contexts where ambiguity is present.

In Study 3a, Democrat subjects read about alleged sexual harassment by a university professor. In Study 3b, Republican subjects read about alleged anti-male discrimination by university administrator (following her decision to discipline a male student who was accused of sexual assault). Furthermore, subjects in both studies read about a group of organizers who were attempting to punish the alleged transgressor (i.e., the professor or the administrator) via a relatively extreme punitive strategy. Our key DV was whether subjects decided to punish the alleged transgressor by donating money to support the organizers.

Studies 3a-b each employed a two-by-two, between-subjects design, in which we manipulated two key factors. First, we manipulated whether punishment had reputational consequences. To this end, we employed a version of the classic “Dictator Game” from behavioral economics. In particular, we told all subjects that another Mturk worker—who we refer to in this paper as “the Decider”—would receive an endowment of money (\$1) and decide how much of it to allocate to the subject (i.e., in the language of a Dictator Game, the Decider

played the role of “Dictator”). We then either told subjects that their punishment decisions would be observable to the Decider (*public* condition) or remain private (*private* condition).

Second, we endeavored to manipulate whether the organizers’ punitive approach would be perceived by subjects as clearly merited (*unambiguous* condition) or more questionably deserved (*ambiguous* condition). To this end, we presented subjects with a news article describing the allegations and varied both (i) how strong the evidence was that wrongdoing occurred and (ii) how severe the alleged wrongdoing was (while holding constant organizers’ punitive approach). Importantly, while subjects were led to believe that the allegations and news article, group of punishing organizers, and Decider were all real, they were actually fictitious.

We used this approach to investigate the power of reputation to drive punishment (by comparing rates of punishment in public vs. private), both in our ambiguous and unambiguous conditions. Critically, however, in order to create a maximally informative comparison between our ambiguous and unambiguous conditions, we sought to create a similarly strong reputation manipulation across ambiguity conditions. To this end, we attempted to ensure that subjects in the public versions of our ambiguous and unambiguous conditions expected punishing to look similarly good in the eyes of the Decider (such that making punishment public introduced comparable reputational incentives in both conditions).

Studies 1-2 highlight a potential challenge for this aim. Our analyses of Studies 1-2 focused on the finding that subjects in our ambiguous conditions expected punishment to confer reputational benefits. However, we also found that subjects in our unambiguous conditions expected punishment to confer even *larger* reputational benefits. We thus reasoned that subjects in the unambiguous (vs. ambiguous) conditions of Study 3 might likewise expect the Decider to

730 react more positively to punishment—posing a potential challenge for our goal of creating
731 comparable reputational incentives across conditions.

732 Yet by documenting the influence of audience ideology on the perceived reputation value
733 of punishment, Study 2 points to a potential solution to this challenge. Subjects in Study 2
734 expected punishment to look better in the eyes of more (vs. less) ideological audiences. Thus, we
735 reasoned that by describing the Decider as more ideological in the ambiguous (vs. unambiguous)
736 conditions of Study 3, we could cause subjects in both ambiguity conditions to expect
737 punishment, if observed, to look similarly good. To this end, in Study 3a we described the
738 Decider as “a Democrat who supports the #MeToo movement” in the ambiguous condition, but
739 simply “a Democrat” in the unambiguous condition. And in Study 3b we described the Decider
740 as “a Republican who is a committed member of the Men’s Rights Movement” in the ambiguous
741 condition, but simply “a Republican” in the unambiguous condition.

742 *Pre-tests*

743 A set of pre-tests suggest that this approach caused subjects in the public versions of our
744 ambiguous and unambiguous conditions to expect punishment to be perceived similarly (and
745 positively) by their Decider. Specifically, in our pre-tests (which were not pre-registered), we
746 assigned subjects to the above-described “ambiguous + more ideological” or “unambiguous +
747 less ideological” conditions, and then measured their expectations about the reputation value of
748 punishment in two ways. (We also pre-tested other conditions that we did not draw on in Studies
749 3a-b and thus do not describe here; see SM Section 3 for details.)

750 All pre-test subjects were assigned to the public condition of our behavioral paradigm,
751 and thus were told that the Decider would learn whether or not they punished before deciding
752 how much money to share with them. Then, subjects predicted, in random order, the number of

cents that the Decider would share if they did vs. did not punish; we calculated each subject's predicted financial gain from punishing as the difference between these numbers. Additionally, subjects predicted whether the Decider would evaluate them more positively if they did vs. did not punish, on a 1-9 Likert scale (1 = Much more positively if I do NOT donate, 5 = No difference, 9 = Much more positively if I DO donate).

As illustrated in Figure 2A, our pre-tests suggest that by describing the Decider as more ideological in the ambiguous than unambiguous condition, it is possible to create similarly-strong reputational incentives for punishment across ambiguity conditions. In our Democrat pre-test (target $n = 400$ across four total conditions; final $n = 194$ across the two conditions described here), subjects expected punishment to result in similar financial gains across conditions (unambiguous + less ideological $M = 19.90\text{¢}$, $SD = 30.13$, ambiguous + more ideological $M = 23.40\text{¢}$, $SD = 24.56$, $b = 3.50$ [-4.31, 11.32], $B = .06$, $t = 0.88$, $p = .378$, $n = 194$). They also expected punishment to be evaluated similarly (unambiguous + less ideological $M = 6.98$, $SD = 2.07$, ambiguous + more ideological $M = 7.16$, $SD = 1.67$, $b = .18$ [-.36, .71], $B = .05$, $t = 0.66$, $p = .509$). Similarly, in our Republican pre-test (target $n = 350$ across three total conditions; final $n = 231$ across the two conditions described here), subjects expected punishment to result in similar financial gains across conditions (unambiguous + less ideological $M = 16.61\text{¢}$, $SD = 31.29$, ambiguous + more ideological $M = 16.72\text{¢}$, $SD = 31.49$, $b = .12$ [-8.03, 8.26], $B = .002$, $t = 0.03$, $p = .978$). They also expected punishment to be evaluated similarly (unambiguous + less ideological $M = 6.60$, $SD = 2.06$, ambiguous + more ideological $M = 6.78$, $SD = 2.20$, $b = .18$ [-.37, .74], $B = .04$, $t = 0.65$, $p = .514$).

Thus, across both pre-tests, we found no evidence of meaningful differences between conditions on either of our measures. However, the absence of evidence for differences between

conditions is not definitive evidence of absence. The confidence intervals on the coefficients reported above reveal the upper bounds of differences between conditions that are plausible in light of our results (e.g., for our Likert scale measure, the upper bounds are about three quarters of a scale point on a nine-point scale for both Democrats and Republicans). Thus, our pre-test results suggest that the perceived reputation value of punishment was approximately equated across conditions, but do not rule out the possibility of some differences. For more details about the samples, methods, and results of our pre-tests, see SM Section 3.

In sum, then, Studies 3a-b employed a two-by-two, between-subjects design in which we crossed our reputation manipulation (public vs. private) with our ambiguity manipulation (ambiguous vs. unambiguous). Furthermore, in order to approximately equate the strength of our reputation manipulation across ambiguity conditions, we described the Decider as more ideological in the ambiguous condition and less ideological in the unambiguous condition.⁴ This intentional asymmetry in our design⁵ allowed us to purchase on our theoretical question of interest. Specifically, it allowed us to investigate the power of *similarly-strong* reputational incentives to drive punishment—both in situations where punishment is seen as clearly merited (unambiguous condition) and questionably deserved (ambiguous condition).

⁴We note that Study 3b also included a third set of public vs. private conditions, in which subjects were assigned to the ambiguous condition but the Decider was described as less ideological (i.e., simply as “a Republican”). These conditions were designed to address a theoretical question outside the scope of this paper, so we do not discuss them further.

⁵On the surface, the fact that we described the Decider as more ideological in the ambiguous (vs. unambiguous) conditions of Study 3 may seem like a “confound”. However, our design is *not* confounded with respect to our theoretical constructs of interest. Rather, by combining two design elements (i.e., a manipulation of ambiguity and a manipulation of Decider ideology), Study 3 successfully varies the perceived ambiguity of the moral case for punishment while holding the reputational incentives for punishment approximately constant. For a detailed discussion of this point, see SM Section 2.7.

794 ***Samples***

795 In each of Studies 3a-b, we recruited a target of $n = 1600$ subjects. Our final samples
796 consisted of $n = 1558$ Democrats in Study 3a (mean age = 35.38, 48.97% male, 54.04% with a
797 college or graduate degree) and $n = 1565$ Republicans in Study 3b (mean age = 38.74, 46.13%
798 male, 54.82% with a college or graduate degree).

799 ***Procedure***

800 We began by asking subjects to report their age, gender, and political affiliation (Democrat
801 or Republican). Based on their reported affiliation, we directed subjects to Study 3a (Democrats)
802 or 3b (Republicans). We then (i) introduced subjects to an allegation of wrongdoing, and explained
803 that they would have the opportunity to punish the alleged transgressor, (ii) told subjects that they
804 would also participate in an economic game with another Mturk worker (i.e., the Decider), (iii)
805 presented subjects with a full news article detailing the relevant allegations, and then (iv) measured
806 subjects' punishment decisions (and other dependent variables). Finally, subjects completed an
807 exit survey (containing demographic and other questions that are not analyzed in the main text;
808 see SM Section 2.3 for details), and then were debriefed.

809 **Punishment opportunity.** To introduce subjects to their punishment opportunity, we
810 began by presenting subjects with a news headline about alleged wrongdoing. Democrats in all
811 conditions of Study 3a saw the headline *Biologist Joseph Pringle Accused of Sexual Harassment*.
812 Republicans in the ambiguous condition of Study 3b saw the headline *Feminist Princeton Dean*
813 *accused of Anti-Male Discrimination following the Discipline of Basketball Captain Over Sexual*
814 *Assault Accusation*; in the unambiguous condition of Study 3b, the headline was identical except
815 that the word "Discipline" was replaced with the word "Expulsion."

We also presented subjects with a statement by the group of punishing organizers, which outlined a relatively extreme punitive strategy. For Democrats, the statement started as follows: *Joseph Pringle's treatment of women is unacceptable, and the fact that he is a prominent and celebrated biologist should not protect him from his deeply harmful behavior. Despite the very credible allegations against him, Columbia University has taken no action and he continues to hold power over the female students and staff working in his laboratory. This is a disgrace.*

For Republicans, the statement instead started: *In order to pursue her feminist agenda, Elizabeth Cartland has chosen to throw justice out the window and destroy an innocent man's life. Tyler Jones has done nothing wrong and yet has been forced to leave Princeton. It is unacceptable that Cartland's role as the Dean of Students has given her the power to discriminate against men with impunity. Despite the very strong allegations against Cartland, Princeton University has taken no action. This is a disgrace, and will allow her to continue pursuing fictional complaints against whomever she wants.*

The statement then concluded as follows, for Democrats/Republicans: *We are marching with megaphones outside of [his/her] office and personal residence in order to expose [his/her] [behavior/ bias and discrimination] to [his/her] friends, family, neighbors, and professional network, and mount pressure on [Columbia/Princeton] to take action against [him/her]. We have been continuously marching for over a week since [the allegations first came out/ Jones revealed the egregious process through which he was forced to leave Princeton], and will not stop until [he/Cartland] is held accountable for [his/her] disgusting behavior.*

Next, we introduced subjects to their punishment opportunity. We told subjects that the group of punishing organizers was soliciting online donations. Furthermore, we told subjects that they would receive 30 cents, and could either keep the money for themselves as bonus or donate

it to the group of punishers. Throughout this paper, we use the term “punishment” to refer to the decision to donate to the punishing organizers.

Introduction of economic game. Next, we introduced subjects to the economic game involving the Decider (which, as noted above, was a version of the classic “Dictator Game”). We described the game to subjects as “the Sharing Game”, and told subjects that another Mturk worker (i.e., the Decider, who we described to subjects as “Player 1”) would receive \$1 and choose how much, if anything, to share with them.

As outlined above, in Study 3a, we described the Decider as “a Democrat who supports the #MeToo movement” in the ambiguous condition, and “a Democrat” in the unambiguous condition. In Study 3b, we described the Decider as “a Republican who is a committed member of the Men’s Rights Movement” in the ambiguous condition, and “a Republican” in the unambiguous condition. Furthermore, all subjects were also told that *they* would be described to the Decider as a Democrat (in Study 3a) or a Republican (in Study 3b).

In our *public* conditions, we used the economic game with the Decider to create reputational incentives for punishment. To this end, we told subjects that the Decider would find out whether they punished before deciding how much to share with them. Specifically, subjects read that “Before deciding how much to share with you in the Sharing Game, Player 1 will see the news headline and organizer statement that you saw, read the full news article that you will read, learn that you saw these materials, and find out whether you decided to donate 30 cents to help hold [biologist Joseph Pringle/dean Elizabeth Cartland] accountable.” In contrast, in the *private* condition, subjects were told that their punishment decisions had no relevance for the Sharing Game. Specifically, subjects read that “The Sharing Game component of this HIT⁶ is

⁶ Note that “HIT” is a term used to describe a task, such as a study, on Mturk.

unrelated to the component in which you will decide whether to donate 30 cents to help hold [biologist Joseph Pringle/dean Elizabeth Cartland] accountable. Unlike you, Player 1 will NOT learn about the news article, group of organizers, or your donation decision.”

Thus, all subjects learned that they would participate in an economic game with a co-partisan Decider, but only subjects in the *public* condition believed that their punishment decisions would have reputational consequences. In this way, our design allows us to isolate the influence of reputation, while holding constant the effects of introducing an economic game and highlighting a shared co-partisan identity between the subject and the Decider.

After describing the economic game, we presented subjects with a series of questions about the game payoff structure, the Decider’s ideology, and the relationship between subjects’ punishment decisions and the game.⁷ Next, to enhance the salience of our observability manipulation, we presented subjects with screenshot(s) that purportedly illustrated the economic game from the perspective of the Decider. In the private condition, subjects saw just one screenshot (that made no mention of their punishment decision). In contrast, in the public condition, subjects saw two screenshots, in random order: one that would purportedly be shown if they chose *to* punish (informing the Decider of that decision), and one that would purportedly be shown if they chose *not* to punish (informing the Decider of *that* decision).

Presentation of news article. Next, we presented subjects with a full news article about the relevant allegations. Before doing so, we reminded subjects in the public conditions that the Decider would also read the article (and learn that they had read the article); subjects in the private conditions received no such reminder. As described above, we used the news article to

⁷ Per our pre-registration, our primary analyses of all studies include data from all subjects, but we find similar results when restricting to subjects who correctly answered our set of comprehension questions; see SM Sections 2-4 for details.

manipulate whether the organizers' punitive approach would be perceived as clearly merited (*unambiguous* condition) or more questionably deserved (*ambiguous* condition). To this end, we varied important details surrounding the credulity and severity of the relevant allegations.

For Democrats, in the *unambiguous* condition, the harassment allegations against the professor were very severe, and seemed very credible (the most severe accusation was of attempted rape; the professor did not deny the accusations; there were six accusers). In the *ambiguous* condition, the allegations were relatively less severe and seemed less credible (the most severe accusation was of relatively more minor unwanted touching; the professor denied the accusations; there were two accusers, both of whom had potential ulterior motives; sources vouched for the professor's character).

For Republicans, in the *unambiguous* condition, the discrimination allegations against the dean suggested severe bad-faith misconduct (the dean was accused of encouraging a female student to describe a sexual encounter with the college's basketball captain that she consented to but later regretted, and then blatantly mischaracterizing the student's story in order to get the captain expelled; a whistleblower who worked with the dean alleged that she was clearly motivated by anti-male bias; the dean did not defend herself). In the *ambiguous* condition, it was much less clear that the dean had acted in bad faith or that the basketball captain was actually treated unfairly (the dean forced the captain to take a one-year leave of absence following a real complaint against him; experts held a split opinion regarding whether the sexual encounter described in the complaint constituted sexual assault; the accuser later asked the dean to drop the complaint but the dean encouraged her to persist; the dean defended her actions). For full news article texts, see SM Section 9.

Thus, for both Democrats and Republicans, the ambiguous (vs. unambiguous) condition was designed to convey that (i) it was less likely that wrongdoing occurred and (ii) the alleged wrongdoing was less severe. Yet subjects in both ambiguity conditions learned that the organizers were applying the same (relatively severe) punitive approach. Thus, in the ambiguous condition, we expected subjects to see the organizers' punitive strategy as less clearly merited.

Measures. Next, we turned to collecting data from subjects. To measure punishment decisions, we asked subjects to decide whether to donate their 30¢ to the group of punishing organizers. We also measured two additional dependent variables.

First, on the same page as we measured punishment, we measured subjects' continuous commitment to punishing by asking subjects to rate their agreement with the statement "I am strongly committed to supporting the group of organizers" on a 1-10 Likert scale (1 = Strongly disagree, 10 = Strongly agree). We find similar patterns when analyzing punishment decisions versus reported commitment to punishing, both across Study 3 and Study 4 (which also measured both variables). Thus, in the main text, we report analyses of punishment decisions, and in SM Sections 2.5 and 4.2, we report analyses of commitment to punishing.

Second, on a new page, we measured subjects' perceptions of the moral merits of punishment. To this end, we presented subjects with ten questions relevant to the merits of punishment. In hopes of eliciting genuine responses, we informed subjects in all conditions that their answers would not be shown to the Decider. Analyses of these questions are reported in the main text, and are used to assess the success of our ambiguity manipulation.

In particular, in random order, we asked subjects to rate (1) how confident they were that the allegations were true, (2) how bad of a person they thought the alleged transgressor was, (3) how immoral they thought the alleged transgressor's actions were, (4) how angry they were at

the alleged transgressor, (5) how much of an outrage they thought the alleged transgressor's behavior was, (6) how much they thought the alleged transgressor deserved to be punished, (7) how important they thought it was that the alleged transgressor was punished, (8) how comfortable they were with the punishing group's approach, (9) the extent to which they thought the punishing group's approach was proportionate and appropriate, and (10) the extent to which they thought the punishing group's approach was moral. Subjects rated each item using a 1-10 scale (1 = Not at all, 10 = Very). This ten-item scale was reliable (Democrats ($\alpha = .97$, Republicans $\alpha = .97$), and exploratory factor analyses of our scale items produced one-factor solutions for both Democrats and Republicans; we thus average our ten items to form a single composite measure of the perceived moral merits of punishment.⁸ See SM Section 2.2 for precise item wordings and factor analysis results.

Results

To analyze Studies 3a-b, we begin by asking whether subjects in our ambiguous conditions indeed saw punishment as ambiguously merited, both in absolute terms and relative to our unambiguous conditions. To this end, we analyze the perceived moral merits of punishment. In particular, we analyze "baseline" perceptions in our private conditions, where reputation concerns are off the table.⁹ If, at baseline, many subjects in our ambiguous conditions saw punishment as questionably merited, it becomes interesting to ask whether reputational incentives nevertheless motivated increased punishment in our public conditions.

⁸We note that, because of a programming error, Study 3a did not require participants to complete our personal moral evaluation scale, leaving us with missing data for some subjects for some scale items. While all $n = 1558$ subjects completed at least one scale item, our sample sizes for individual items ranged from $n = 1545$ to $n = 1558$, with $n = 1529$ subjects completing all ten scale items. In our analyses of composite personal moral evaluations, we include all subjects (and compute each subject's score as the average of all scale items that they completed). In our analyses of our scale's reliability and factor structure, we consider only the $n = 1529$ subjects who completed all scale items.

⁹ We also note that our Study 3-4 designs allow us to compare moral evaluations across our public vs. private conditions, in order to investigate whether reputational incentives influenced these evaluations; see SM Sections 2.6 and 4.3 for these analyses.

Figure 2B plots average baseline ratings of the moral merits of punishment as a function of ambiguity, within the private conditions of Studies 3a-b. This figure illustrates that subjects in our unambiguous conditions saw punishment as clearly merited, as evidenced by high mean ratings (on our ten-point scale: Democrats $M = 7.91$, $SD = 1.85$; Republicans $M = 7.35$, $SD = 1.94$). In contrast, subjects in our ambiguous conditions saw the case for punishment as much more ambiguous, as evidenced by mean ratings near the scale midpoint of 5.5 (Democrats $M = 5.13$, $SD = 2.48$; Republicans $M = 5.73$, $SD = 2.40$) and significant differences between conditions (Democrats $b = -2.79$ $[-3.09, -2.48]$, $B = -.54$, $t = -17.96$, $p < .001$, $n = 788$; Republicans $b = -1.63$ $[-1.93, -1.32]$, $B = -.35$, $t = -10.55$, $p < .001$, $n = 801$). We therefore find that our ambiguity manipulation successfully induced perceived moral ambiguity.

Figure 2 thus highlights that, at baseline, many subjects in our ambiguous conditions harbored personal reservations about the merits of punishment, and saw punishment as much less morally merited than their counterparts in our unambiguous conditions (per Figure 2B). Yet Figure 2 *also* highlights that subjects in the public versions of our ambiguous and unambiguous conditions both perceived substantial—and similarly strong—reputational incentives to punish (per Figure 2A, which draws on pre-test data). To what extent did these reputational incentives drive punishment? To answer this question, we turn to considering punishment *decisions*.

In Figure 2C, we plot punishment decisions across conditions in Studies 3a-b. Additionally, for each study, we use effect coding to predict punishment decisions as a function of observability (private = $-.5$, public = $.5$), ambiguity (unambiguous = $-.5$, ambiguous = $.5$), and their interaction. We find significant main effects of ambiguity, with lower rates of punishment in our ambiguous than unambiguous conditions (Democrats $b = -.27$ $[-.31, -.22]$, $B = -.28$, $t = -11.66$, $p < .001$; Republicans $b = -.15$ $[-.20, -.11]$, $B = -.16$, $t = -6.55$, $p < .001$). We also find

significant main effects of observability, with higher rates of punishment in public than in private (Democrats $b = .13$ [.08, .17], $B = .14$, $t = 5.63$, $p < .001$, $n = 1558$; Republicans $b = .09$ [.05, .14], $B = .10$, $t = 4.02$, $p < .001$, $n = 1565$). And finally, we find no significant interactions between ambiguity and observability (Democrats $b = .02$ [-.07, .11], $B = .01$, $t = .53$, $p = .593$, $n = 1558$; Republicans $b = .001$ [-.09, .09], $B = -.001$, $t = .03$, $p = .979$, $n = 1565$).

Next, to break down these results and directly address our theoretical questions, we consider simple effects. We begin by considering the simple effects of ambiguity within our private conditions, where reputation concerns were off the table. We find significant negative effects of ambiguity (Democrats $b = -.28$ [-.34, -.22], $B = -.29$, $t = -8.67$, $p < .001$, $n = 788$; Republicans $b = -.16$ [-.22, -.09], $B = -.16$, $t = -4.71$, $p < .001$, $n = 801$). Thus, in addition to decreasing baseline moral evaluations of punishment (as reported above), moral ambiguity *also* decreased baseline rates of punishment behavior. This result supports the hypothesis that moral evaluations of punishment shape baseline proclivities to punish, and provides further evidence that our ambiguity manipulation was psychologically meaningful to subjects.

Next, we turn to considering the simple effects of observability across our two ambiguity conditions. We start with the unambiguous conditions, in which subjects, at baseline, were relatively inclined to punish. We find that making punishment observable increased rates of punishment from this high baseline, both among Democrats (private rate = .41, public rate = .53, $b = .12$ [.05, .18], $B = .12$, $t = 3.66$, $p < .001$, $n = 804$) and Republicans (private rate = .37, public rate = .47, $b = .09$ [.03, .16], $B = .10$, $t = 2.82$, $p = .005$, $n = 783$). Given that most subjects in the private versions of our unambiguous conditions were personally supportive of punishment (as illustrated in Figure 2B), this finding underscores the power of reputation to encourage punitive actions that are widely seen as merited (and to motivate individuals to live up to their personal

values). Furthermore, it also serves to further validate our paradigm by demonstrating that subjects were responsive to the reputational incentives that our public conditions introduced.

Next, we turn to the ambiguous conditions, in which subjects showed much lower baseline rates of punishment in private, but yet faced similar reputation-based pressure to punish in public. Did reputation *also* drive punishment in these conditions? Yes: despite their lower baseline inclinations to punish, subjects in our ambiguous conditions responded readily to reputational incentives, such that making punishment observable robustly increased punishment among both Democrats (private rate = .13, public rate = .27, $b = .14$ [.08, .20], $B = .15$, $t = 4.29$, $p < .001$, $n = 754$) and Republicans (private rate = .22, public rate = .31, $b = .10$ [.03, .16], $B = .10$, $t = 2.86$, $p = .004$, $n = 782$). Because Figure 2B also illustrates that many subjects in the private versions of our ambiguous condition were *not* personally supportive of punishment, this finding highlights the power of reputation to drive punishment even in situations where many judge punishment to be questionably merited (and suggests that reputation may sometimes push individuals *away* from their personal values).

Finally, given that subjects in our ambiguous vs. unambiguous conditions started with very different baseline rates of punishment (in private) and were similarly responsive to reputation, we also find robust simple effects of ambiguity within our *public* conditions. When decisions were public, subjects in our ambiguous conditions were again much less likely to punish than their counterparts in our unambiguous conditions (Democrats $b = -.25$ [-.32, -.19], $B = -.27$, $t = -7.82$, $p < .001$, $n = 770$; Republicans $b = -.15$ [-.22, -.09], $B = -.16$, $t = -4.56$, $p < .001$, $n = 764$).

Figure 2. Reputation fuels punishment that many people see as morally questionable. In fact, when approximately equating the strength of reputational incentives, reputation is similarly capable of driving punishment in situations where punishment is seen as ambiguously vs. unambiguously deserved. (A) Describing the Decider as more ideological in the ambiguous (vs. unambiguous) conditions of Study 3 allows us to vary ambiguity while approximately equating the strength of reputational incentives. Shown are data validating this approach from our Study 3 pre-tests, in which all subjects were assigned to the public condition of our paradigm. We plot the perceived reputation value of punishment (as measured by the expected financial gains from punishing) as a function of ambiguity (and Decider ideology). **(B)** In private, many subjects in our ambiguous conditions reported reservations about the morality of punishment. We plot ratings of the moral merits of punishment (higher numbers designate the perception that punishment is more moral) as a function of ambiguity, across the private conditions of Studies 3a-b. **(C)** Despite their reservations, subjects in the ambiguous conditions of Studies 3a-b readily responded to our reputation manipulation. In fact, reputation had a comparable influence on punishment in our ambiguous vs. unambiguous conditions. Shown are punishment decisions as a function of ambiguity and observability. Error bars are 95% CIs.

Discussion

Consistent with previous research on justice and morality (e.g., Aquino & Reed, 2002; Gollwitzer & van Prooijen, 2016), Study 3 provides strong evidence that moral considerations shaped punishment behavior in our paradigm. At baseline (i.e., in the *private* conditions of Study 3, where reputation concerns were off the table), subjects in our ambiguous conditions reported meaningful moral reservations about punishment, and were much less morally supportive of punishment than their counterparts in our unambiguous conditions. They also were much less likely than their counterparts to actually punish, suggesting that morally relevant details about the alleged transgressions shaped subjects' baseline punitive behavior. Furthermore, these morally relevant details were no less influential within our *public* conditions, which introduced the potential for reputational gain. Together, then, we find that subjects punished at lower overall rates when the case for punishment was more ambiguous—despite the fact that the financial cost of punishment, and expected reputational gains from punishment in public, did not differ across our ambiguity conditions. We thus find clear evidence that, rather than being solely focused on the personal costs and benefits of punishing, subjects were influenced by moral considerations.

Yet despite the fact that moral considerations influenced overall rates of punishment, we found no evidence that moral ambiguity limited the motivational force of reputation. Consistent with previous research on the influence of reputation on punishment, which has focused on contexts in which punishment is presumed to be seen as justified, observability increased rates of punishment within our *unambiguous* conditions. Strikingly, however, making punishment decisions observable also increased rates of punishment within our *ambiguous* conditions, in which many subjects saw punishment as morally questionable. In fact, when approximately equating the strength of reputational incentives, our results suggest that the motivational force of

reputation is comparable across ambiguous vs. unambiguous situations. In other words, while moral considerations shaped punishment behavior, we find no evidence that moral reservations about punishment served to restrain people from responding to reputational incentives. Instead, moral considerations and reputation motives exerted independent influences on behavior.

Of note, like in the context of our pre-tests, our null interaction results are not definitive evidence of absence of an interaction between ambiguity and observability. Despite our relatively large sample sizes, precisely estimating interaction effects is particularly difficult (Giner-Sorolla, 2018); the confidence intervals on the reported interaction coefficients reveal the upper bounds of interaction effect sizes that are plausible in light of our results (e.g., for punishment, we observe interaction effect size upper bounds of about ten percentage points for both Democrats and Republicans). Thus, while Study 3 clearly shows that reputation can drive punishment in ambiguous contexts, our conclusions with respect to the *relative* power of reputation in ambiguous vs. unambiguous contexts remain more tentative.

In sum, Study 3 suggests that in situations where people anticipate a reputational upside to punishing, we should expect reputation to drive heightened punishment—even if there is a lack of consensus that punishment is deserved, and many individuals see the case for punishment as morally murky. This conclusion highlights the motivational force of reputation, and is all the more notable in light of the clear evidence that subjects in our paradigm *were* morally engaged, with moral considerations markedly affecting overall rates of punishment.

Study 4

While Study 3 clearly shows that reputation can increase average rates of punishment in ambiguous situations, however, it leaves open an interesting individual-level question. In particular, it leaves open the question of *who* responds to reputational incentives in ambiguous

situations. Even in situations where many people question the morality of punishment, it is likely that at least some people will see punishment as clearly merited. Figure 2B illustrates this point in the context of Study 3: even in our ambiguous conditions, a substantial minority of subjects rated punishment as clearly morally merited. Thus, the positive effect of observability on punishment in the ambiguous conditions of Study 3—which held on average across subjects—could be explained in two different ways, both of which are interesting.

First, it is possible that the influence of reputation on punishment was driven exclusively by the subset of individuals who saw punishment as clearly merited. Alternatively, however, reputation might have the power to drive punishment even among individuals who *personally* harbor reservations about its morality. Regardless, Study 3 provides clear evidence that in ambiguous situations, the fact that many people question the morality of punishment does *not* prevent reputation from driving heightened rates of punishment—a conclusion that has important social implications. Yet the two aforementioned possibilities differ meaningfully from a psychological perspective. In Study 4, we thus ask: does the influence of reputation on punishment extend even to individuals who personally question its morality?

Method

To address this question, we recruited a large sample of subjects, and assigned them all to the ambiguous condition of our behavioral paradigm. Recruiting a large sample was critical for us to achieve sufficient power to draw inferences specifically about the subset of individuals who ultimately reported reservations about the morality of punishment. However, it was only feasible for us to do with Democrat subjects. Thus, Study 4 focuses exclusively on Democrats.

Specifically, in Study 4, we recruited a target of $n = 2500$ Democrats. Our final sample consisted of $n = 2415$ Democrats (mean age = 36.18, 42.28% male, 64.06% with a college or

graduate degree). We also measured ethnicity in Study 4 (but no other studies); our final sample was 73.46% White, 11.59% Black, 10.68% Asian, 1.08% American Indian or Alaska Native, .37% Native Hawaiian or Pacific Islander, 7.83% Hispanic/Latino, and 2.07% Multicultural (note that these categories were not mutually exclusive).

We employed a two-condition, between-subjects design in which we manipulated observability while assigning all subjects to our ambiguous condition and always describing the Decider as “a Democrat who supports the #MeToo movement.” Thus, Study 4 mirrors the ambiguous conditions of Study 3a, and again allows us to investigate the influence of (strong) reputational incentives on punishment that is judged to be questionably-merited. However, we modified the procedure of Study 4 in one critical way: we measured subjects’ initial moral evaluations of punishment *before* introducing our reputation manipulation.

Like in Study 3a, we began Study 4 by informing subjects that a university professor was accused of sexual harassment, via a news headline and a statement from the punishing group of organizers. However, we did *not* subsequently tell subjects that they would have the opportunity to make a punishment decision (by donating to the punishing organizers). Instead, we presented the full news article, and then measured subjects’ initial perceptions of the moral merits of punishment.

This design stood in contrast to Study 3a, which did not measure perceptions of the moral merits of punishment until the end of the study. And it allowed us to identify individuals with reservations about the morality of punishment, and ask: will *these* subjects act on reputational incentives for punishment? It also created an especially strong test of this question. By forcing subjects to report their reservations upfront—both to themselves and to the experimenter—we

1123 created a potentially strong consistency motive *not* to act on the reputational incentives for
1124 punishment that we later introduced.

1125 To measure subjects' initial perceptions of the moral merits of punishment, we presented
1126 subjects with a shorter version of our ten-item scale from Study 3a (using the items labeled 1, 2,
1127 4, 6, 8, and 9 in the Study 3 methods section). We also adjusted our scale in order to more
1128 explicitly identify individuals with reservations about the morality of punishment. In Study 3,
1129 our scale ranged from 1 to 10 and we only provided written labels at the endpoints (1 = "Not at
1130 all", 10 = "Very"). In Study 4, we modified the scale to range from 1 to 9, such that the midpoint
1131 would correspond to a whole number (5). We also labeled this midpoint with the phrase "I have
1132 reservations" (so that subjects were asked, for example, "How confident are you that the
1133 allegations against [the Professor] are true?" 1 = "Not at all", 5 = "I have reservations", 9 =
1134 "Very").

1135 This six-item scale was reliable ($\alpha = .94$) and an exploratory factor analysis produced a
1136 one-factor solution (see SM Section 2.2); we thus again averaged our items to form a composite
1137 measure of the perceived moral merits of punishment. Following our pre-registration, we then
1138 categorized subjects into four bins, based on their composite score on our scale: bin one (values
1139 ≤ 3), bin two (values > 3 and ≤ 5), bin three (values > 5 and ≤ 7), and bin four (values > 7) (see
1140 Figure 3). The first two bins, with values at or below the scale midpoint, categorize subjects who
1141 reported clear reservations about the morality of punishment.

1142 After measuring initial evaluations of punishment, we informed subjects about their
1143 punishment opportunity and the economic game involving the Decider (thus introducing our
1144 observability manipulation). We then measured punishment decisions, and commitment to
1145 punishing, as in Study 3a. Finally, we measured subjects' perceptions of the moral merits of

punishment for a second time; see SM Section 4.3 for more discussion of this DV, which we do not analyze here.

Results

Overall, subjects in Study 4 saw the case for punishment as highly ambiguous. Mean initial ratings of the moral merits of punishment were near our scale midpoint of 5 ($M = 5.10$, $SD = 1.92$), and 52% of subjects fell into bins one or two. Thus, mirroring the ambiguous conditions of Studies 3a-b, Study 4 successfully created an ambiguous situation. In this ambiguous situation, did reputation drive punishment—both overall, and specifically among individuals with clear reservations about its morality?

To answer this question, per our pre-registration, we report observability effects from models that control for subjects' (pre-treatment) initial evaluations of punishment; models without this control support identical conclusions. We begin by analyzing all subjects in Study 4. We find that overall, rates of punishment were higher in public (.34) than private (.25), $b = .11$ [.07, .14], $B = .12$, $t = 6.22$, $p < .001$, $n = 2415$, mirroring the overall effect of observability on punishment that we found in the ambiguous conditions of Studies 3a-b.

Next, we turn to asking: does this overall observability effect extend to individuals who reported moral reservations about punishment? In Figure 3, we plot the effect of observability on punishment within each bin. Starting with subjects who saw punishment as *more* clearly merited, we find significant observability effects both in the fourth bin (private rate = .51, public rate = .61, $b = .10$ [.01, .19], $B = .10$, $t = 2.08$, $p = .038$, $n = 424$), and the third bin (private rate = .36, public rate = .49, $b = .14$ [.06, .21], $B = .14$, $t = 3.78$, $p < .001$, $n = 741$). Strikingly, the influence of observability also extended to the second bin (private rate = .12, public rate = .23, $b = .11$ [.06, .16], $B = .15$, $t = 4.36$, $p < .001$, $n = 857$), revealing that reputation fueled punishment among

1169 subjects with clear reservations about its morality. In fact, observability even had an effect in the
 1170 first bin (private rate = .01, public rate = .07, $b = .06$ [.02, .10], $B = .14$, $t = 2.81$, $p = .005$, $n =$
 1171 393), among subjects who initially reported relatively clear *disapproval* of punishment.

1172

1173 **Figure 3. Evidence from Democrats reveals that reputation drives punishment, even among**
 1174 **individuals with reservations about its morality.** We plot punishment as a function of (i)
 1175 initial (i.e., pre-treatment) moral evaluations of punishment, and (ii) observability in Study 4. We
 1176 bin subjects on the basis of their initial composite ratings of the moral merits of punishment into
 1177 the ranges [1, 3], (3, 5], (5, 7], (7, 9]. Strikingly, the influence of reputation extends even to
 1178 individuals who, before reputational incentives were introduced, reported reservations about (or
 1179 even clear disapproval of) punishment (as indexed by ratings in bins 1 or 2).

1180
 1181

1182

Discussion

Study 4 bolsters our conclusion that reputation can fuel punishment that is judged to be ambiguously merited. In Study 4, reputation drove punishment in our ambiguous condition even among individuals who *personally* questioned the morality of punishment, and reported their reservations before learning that punishing could boost their reputations. Given the human desire for consistency, and the fact that consistency motives frequently shape behavior (Aronson et al., 1991; Bruneau et al., 2018, 2019; Gawronski, 2012; Stone & Fernandez, 2008), this result provides a striking demonstration of the power of reputation.

These results are particularly interesting given that the ambiguous conditions of our behavioral paradigm created uncertainty both over whether the relevant allegations were true, and whether they were severe enough to merit the organizers' punitive strategy. Thus, our results suggest that *neither* source of ambiguity gives rise to reservations that prevent people from acting on reputational incentives. And of note, the design of Study 4 allows us to bolster this claim by considering the specific types of reservations that subjects reported. In SM Section 4.4, we replicate our Study 4 analyses when binning subjects not on the basis of their *composite* ratings of the moral merits of punishment, but rather their responses to particular scale items asking about either (i) confidence that the allegations were true, or (ii) the proportionality and appropriateness of the organizers' punitive approach. Thus, we find that reputation can drive punishment among individuals with reservations pertaining to both sources of ambiguity.

Finally, we note that Studies 3 and 4 raise the interesting question of the extent to which subjects were *consciously* motivated by reputation when responding to our observability manipulations. In SM Section 5, we explore this question by plotting data from relevant

secondary measures. Broadly, these data suggest that subjects, across ambiguity conditions and bins, may have had relatively limited awareness of the influence of reputation on their behavior.

General Discussion

Across four studies (total $n = 9,587$), we find strong evidence that reputation can fuel moralistic punishment that subjects judge to be questionably merited. In Studies 1-2, we asked: in ambiguous situations, where there is a lack of consensus that punishment is merited, do people expect punishing to confer reputational benefits? And we found that subjects in such contexts frequently expect punishers to earn reputational benefits, especially in the eyes of more ideological audiences. Moreover, despite the human tendency towards egocentric biases in perspective taking (Epley et al., 2004), this expectation sometimes extended to individuals who were *personally* unsure that punishment was merited. We observed these patterns across a variety of politicized moral domains, among both Democrats and Republicans. They also held both when ambiguity reflected uncertainty that an alleged transgression actually occurred (Study 1, featuring only Democrats) as well as uncertainty that a transgression was severe enough to merit the relevant punishment (Studies 2a-b, featuring Democrats and Republicans).

Next, in Studies 3-4, we investigated the motivational force of reputation in ambiguous situations. We designed a behavioral paradigm to measure Democrats' willingness to punish a university professor accused of sexual harassment, and Republicans' willingness to punish a university administrator accused of anti-male discrimination. In the ambiguous conditions of this paradigm, we featured both aforementioned sources of uncertainty—causing many subjects to report reservations about the morality of punishing. Yet in Studies 3-4, subjects in our ambiguous conditions readily used punishment to boost their reputations, punishing more in public than private. In fact, in Study 3 we found no evidence that similarly-strong reputational

1228 incentives are any less capable of driving punishment in ambiguous (vs. unambiguous) situations
1229 (although we note that this inference, which relies on null effects, remains more tentative).
1230 Finally, in Study 4, we found evidence from Democrats that reputation drives punishment even
1231 among individuals who *personally* question the merits of punishment.

1232 **Theoretical and societal implications**

1233 Our results have important implications, both for theories of psychology and society.
1234 More specifically, our findings expand our understanding of the psychological power of
1235 reputation, as well as the breadth of its influence on social behavior. Previous research has
1236 documented the robust influence of reputation on behavior in the moral domain. Yet the focus
1237 has been on the power of reputation to fuel behaviors that are widely seen as morally good—
1238 such as direct acts of cooperation, or acts of punishment that are presumed to be seen by subjects
1239 as clearly justified. Thus, previous research has primarily made clear that reputation has the
1240 power to inspire socially beneficial behavior. And while our results do underscore this
1241 observation (most subjects in the unambiguous conditions of Studies 3a-b saw punishment as
1242 morally merited, and reputation increased their propensity to punish), we also find that reputation
1243 can drive behavior that is judged to be morally *questionable* (as evidenced by the robust
1244 influence of reputation on punishment in our ambiguous conditions across Studies 3-4).

1245 It was not obvious that reputation would have this power. Previous research has
1246 demonstrated that moral concerns can make people less responsive to instrumental incentives
1247 (Ginges et al., 2007; Ginges & Atran, 2009; Skitka, 2010; Tetlock et al., 2000). Thus, one might
1248 have expected subjects to be "principled", in the sense of becoming less responsive—or even
1249 entirely insensitive—to reputational incentives in contexts where they saw punishment as less
1250 clearly merited.

Furthermore, it is important to note that we *did* find strong evidence that subjects in our paradigm were influenced by moral considerations. In Studies 3a-b, subjects were overall much less likely to punish in our ambiguous conditions (in which moral reservations about punishment were common) than our unambiguous conditions (in which reservations were rare). These results suggest that morally relevant details about the relevant transgressions shaped subjects' baseline inclinations to punish, and continued to impact behavior even in our public conditions where reputational gains were on the table. Similarly, in Study 4, rates of punishment in our ambiguous condition were much lower among subjects who reported reservations.

Yet critically, highlighting the power of *reputation*, we found no evidence that moral reservations limit people from responding to reputational incentives for punishment in ambiguous situations. Instead, results from the ambiguous conditions of Studies 3-4 suggest that reservations make people less likely to punish, but do not prevent them from showing a robust uptick in punishment when reputation is on the line. In other words, reputation motives and moral considerations appear to exert independent influences on punishment behavior.

To make this conclusion more concrete, recall the motivating example of Senator Al Franken—and in particular, his hypothetical colleague, Maria, who personally questions whether Franken deserves to be ousted. Relative to a colleague who strongly supports Franken's forced resignation, our results suggest that Maria would likely have a much lower baseline inclination towards punishment. Yet our results *also* suggest that Maria would, on average, become meaningfully *more likely* to push for Franken's resignation if she expected doing so to boost her reputation. In fact, despite her lower baseline drive to punish, our results suggest that Maria might be *just as responsive* to reputational incentives as a colleague who enthusiastically supports punishment.

Our results thus demonstrate that, in addition to motivating socially beneficial behavior, reputation can motivate actions that many people see as morally ambiguous—and even push individuals away from their own moral values. In this way, our work implies that the scope of reputation’s influence on social behavior is far-reaching. And in doing so, it concords with evidence that personal values do not always prevent people from responding to countervailing incentives (Jones, 1965; Tetlock, 2002, 2003; Thompson, 2001) and contributes to a growing body of work highlighting the influence of self-interested incentives in the moral domain (Babcock & Loewenstein, 1997; DeScioli et al., 2014; Jordan & Rand, 2019; Melnikoff & Bailey, 2018; Melnikoff & Strohming, 2020).

This conclusion has implications for debates surrounding “outrage culture.” In recent years, it has been proposed that society has become increasingly willing to condemn and punish wrongdoers, including in contexts where there is less consensus that punishment is merited—a trend that some critics see as morally problematic, and others see as moral progress. Our results highlight that people are naturally much less inclined to punish in more ambiguous contexts. But they also support the proposal that “virtue signaling” motives may encourage people to treat ambiguous contexts more punitively.

Furthermore, our results suggest a potentially important role of audience ideology in this process. Within the ambiguous conditions of Study 2, subjects expected more ideological audiences to react more positively to punishment. Thus, in ideologically extreme contexts, virtue signaling may be especially likely to fuel punishment that is judged to be questionably merited. This conclusion highlights a potential consequence of rising extremity and polarization, and provides an empirical basis for suggestions that social media platforms may amplify outrage (Brady et al., 2017; Brady & Crockett, 2019; Crockett, 2017; Hawkins et al., 2018). By pairing

explicit reputation metrics (e.g., “like” buttons and retweets) with large, ideologically-polarized audiences (Bakshy et al., 2015b; Barberá & Rivero, 2015), social media platforms may be particularly well-equipped to encourage punishment in more ambiguous contexts.

Yet our results also hint at limits to the influence reputation may have on punishment in ambiguous situations. Studies 1-2 reveal that subjects expect punishment to confer larger reputational benefits when it is seen as unambiguously (vs. ambiguously) merited. Future research should investigate the generalizability of this finding, which concords with recent evidence that punishers are evaluated less positively when they are seen as disproportionately “piling on” (Sawaoka & Monin, 2018). If people reliably perceive stronger reputational incentives for punishment that is seen as clearly deserved, it would suggest that virtue signaling may preferentially fuel such punishment—and be less likely to encourage punishment that is seen as questionably merited.

Limitations and future directions

More generally, more work is needed to evaluate the generalizability of our conclusions. Importantly, we have supported many of our conclusions among both Democrats and Republicans and in the context of two distinct sources of ambiguity, and our vignette studies show consistent results across a variety of politicized moral domains. Furthermore, by designing studies involving detailed scenarios in specific moral domains, our work advances the literature on punishment and reputation, which has previously focused on punishment of selfishness in economic games. However, a limitation of our behavioral paradigm is that it focuses on just one scenario each for Democrat and Republican subjects. Moreover, Studies 1 and 4 exclusively recruited Democrat subjects, raising questions about the generalizability of their conclusions—including the key finding from Study 4 that even individuals with moral reservations about

punishment respond to reputational incentives. And while we endeavored (in both our vignette and behavioral paradigms) to create realistic scenarios rooted in the cultural zeitgeist surrounding outrage culture, our lab studies are inherently limited in their external validity.

For these reasons, our work serves to highlight that reputation can drive punishment that people see as ambiguously deserved, and that this conclusion can hold even for individuals who question its merits. Yet future work should continue to explore the influence of reputation on such punishment—across moral domains; across populations, including those that are not “WEIRD” (Henrich et al., 2010); and in analyses of real-world punishment behavior, including on the social media platforms that frequently animate debates surrounding outrage culture.

More specifically, an important question for future research is whether there are contexts in which moral concerns about punishment *do* limit the influence of reputation. In our studies, we found no evidence that moral concerns have this power. In fact, in Study 4, reputation drove punishment even among individuals who expressed *disapproval* of punishment, as evidenced by initial moral evaluations of punishment in the bottom quarter (i.e., “bin one”) of our scale. Yet it may have been rare for subjects in our paradigm—including those in this subgroup—to feel *morally convicted* (Skitka, 2010) that punishment was undeserved. Thus, it remains to be seen whether there are contexts in which people are more deeply committed in their opposition to punishment, to the point that they are less sensitive to the motivational thrust of reputation.

This prospect also raises another interesting question: will reputation have the same motivational force in contexts where punishment is likely to appeal to *counter*-partisan audiences (i.e., audiences whose political ideology differs from the potential punisher’s)? In our studies, subjects faced incentives to punish in ways that would appeal to *co*-partisan audiences. We made this design choice because people preferentially interact with co-partisans and are especially

1343 motivated to be perceived positively by them. However, people *do* sometimes face incentives to
1344 signal their virtue to members of the opposing political party. In such cases, do reputational
1345 incentives fuel punishment as readily?

1346 One might expect the answer to be yes, given that our results generally highlight the
1347 power of reputation. That is, given that reputation readily drove our subjects to enact punitive
1348 acts that they expected *co*-partisans to approve of (but saw as ambiguously-merited), we might
1349 likewise expect reputation to drive punitive acts that people expect *counter*-partisans to approve
1350 of (even when people do not personally support those acts). On the other hand, there are several
1351 reasons to question this prediction.

1352 First, in politicalized moral domains where *co*-partisans support punishment, people may
1353 be especially likely to experience moral objections to punishment that rise to the level of moral
1354 convictions. For example, a conservative individual who is deeply committed to opposing gay
1355 rights might be strongly opposed to punishing an act of homophobia (a decision that would
1356 appeal to *counter*-partisans, but violate a potentially sacred value). And her opposition to such
1357 punishment might be more constraining than the aversion she would feel to punishing an
1358 ambiguous act of alleged disrespect towards the military (a decision that would appeal to *co*-
1359 partisans, but she might see as questionably merited). Interestingly, then, reputation might be less
1360 effective at driving acts of punishment that appeal to *counter*-partisan audiences, insofar as (i)
1361 people are more likely to feel morally convicted that such punishment is unmerited and (ii) such
1362 convictions—unlike the moral reservations that our ambiguous stimuli elicited—are powerful
1363 enough to constrain the influence of reputation.

1364 Second, group-based motivations, like the desire to avoid satisfying *counter*-partisans or
1365 confirming their worldview (Frimer et al., 2017), could plausibly also limit our willingness to

ingratiate ourselves with counter-partisans. For example, a liberal individual with disdain for conservatives might be disinclined to support a punitive act that she believes her Republican co-worker would enthusiastically support—even if she does *not* strongly oppose the act and *does* face incentives earn her co-worker’s respect—simply in order to avoid the aversive feeling of making an outgroup member feel good (Cikara et al., 2011).

Finally, because reputational gain does not take place in a vacuum, incentives to ingratiate ourselves with counter-partisans may carry less weight. For example, the aforementioned liberal individual might stand to gain meaningfully from impressing her Republican co-worker. Yet punishing in a way that aligns with his conservative values could also carry a broader reputational risk: it could make her look *bad* in the eyes of her *liberal* co-workers, if they were to find out about her behavior. And given that people tend to preferentially interact (and value their relationships) with co-partisans, this risk may often loom large—limiting the power of reputation to drive people to impress counter-partisans. Future research should investigate each of these interesting possibilities, and more generally the question of how one’s connection to a reputational audience influences reputation’s motivational force.

Conclusion

In conclusion, we have conducted a systematic examination of the influence of reputation on punishment that is judged to be questionably merited. Together, our results provide strong evidence that reputation can fuel such punishment, and have important implications for both theories of psychology and contemporary social debates.

References

- 1387
1388 Aquino, K., & Reed, I. I. (2002). The self-importance of moral identity. *Journal of Personality*
1389 *and Social Psychology*, 83(6), 1423.
- 1390 Aronson, E., Fried, C., & Stone, J. (1991). Overcoming denial and increasing the intention to use
1391 condoms through the induction of hypocrisy. *American Journal of Public Health*, 81(12),
1392 1636–1638.
- 1393 Babcock, L., & Loewenstein, G. (1997). Explaining bargaining impasse: The role of self-serving
1394 biases. *Journal of Economic Perspectives*, 11(1), 109–126.
- 1395 Bakshy, E., Messing, S., & Adamic, L. A. (2015a). Exposure to ideologically diverse news and
1396 opinion on Facebook. *Science*, 348(6239), 1130–1132.
- 1397 Bakshy, E., Messing, S., & Adamic, L. A. (2015b). Exposure to ideologically diverse news and
1398 opinion on Facebook. *Science*, 348(6239), 1130–1132.
- 1399 Balafoutas, L., Grechenig, K., & Nikiforakis, N. (2014). Third-party punishment and counter-
1400 punishment in one-shot interactions. In *Economics Letters* (Vol. 122, Issue 2, pp. 308–
1401 310).
- 1402 Barasch, A., Levine, E. E., Berman, J. Z., & Small, D. A. (2014). Selfish or selfless? On the
1403 signal value of emotion in altruistic behavior. *Journal of Personality and Social*
1404 *Psychology*, 107(3), 393–413.
- 1405 Barberá, P., & Rivero, G. (2015). Understanding the political representativeness of Twitter users.
1406 *Social Science Computer Review*, 33(6), 712–729.
- 1407 Barclay, P. (2006). Reputational benefits for altruistic punishment. *Evolution and Human*
1408 *Behavior*, 27(5), 325–344. <https://doi.org/Doi.10.1016/J.Evolhumbehav.2006.01.003>

- 1409 Baumard, N., André, J.-B., & Sperber, D. (2013). A mutualistic approach to morality: The
1410 evolution of fairness by partner choice. *Behavioral and Brain Sciences*, 36(01), 59–78.
- 1411 Berman, J. Z., & Silver, I. (2022). Prosocial behavior and reputation: When does doing good lead
1412 to looking good? *Current Opinion in Psychology*, 43, 102–107.
- 1413 Boyd, R., & Richerson, P. J. (1989). The evolution of indirect reciprocity. *Social Networks*,
1414 11(3), 213–236.
- 1415 Boyd, R., & Richerson, P. J. (1992). Punishment allows the evolution of cooperation (or
1416 anything else) in sizeable groups. *Ethology and Sociobiology*, 13(3), 171–195.
1417 [https://doi.org/10.1016/0162-3095\(92\)90032-y](https://doi.org/10.1016/0162-3095(92)90032-y)
- 1418 Brady, W. J., & Crockett, M. J. (2019). How effective is online outrage. *Trends in Cognitive*
1419 *Sciences*, 23(2), 79.
- 1420 Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., & Van Bavel, J. J. (2017). Emotion shapes
1421 the diffusion of moralized content in social networks. *Proceedings of the National*
1422 *Academy of Sciences*, 114(28), 7313–7318.
- 1423 Bruneau, E., Kteily, N., & Falk, E. (2018). Interventions highlighting hypocrisy reduce collective
1424 blame of Muslims for individual acts of violence and assuage anti-Muslim hostility.
1425 *Personality and Social Psychology Bulletin*, 44(3), 430–448.
- 1426 Bruneau, E., Kteily, N., & Urbiola, A. (2019). A collective blame hypocrisy intervention
1427 enduringly reduces hostility towards Muslims. *Nature Human Behaviour*, 1–10.
1428 <https://doi.org/10.1038/s41562-019-0747-7>
- 1429 Byrne, D. (1961). Interpersonal attraction and attitude similarity. *The Journal of Abnormal and*
1430 *Social Psychology*, 62(3), 713.

- 1431 Byrne, D., Griffitt, W., & Stefaniak, D. (1967). Attraction and similarity of personality
1432 characteristics. *Journal of Personality and Social Psychology*, 5(1), 82.
- 1433 Caygle, H. (2017). *Another woman says Franken tried to forcibly kiss her*.
1434 [https://www.politico.com/story/2017/12/06/al-franken-accusation-sexual-harassment-](https://www.politico.com/story/2017/12/06/al-franken-accusation-sexual-harassment-2006-281049)
1435 [2006-281049](https://www.politico.com/story/2017/12/06/al-franken-accusation-sexual-harassment-2006-281049)
- 1436 Chaudhry, S. J., & Loewenstein, G. (2019). Thanking, apologizing, bragging, and blaming:
1437 Responsibility exchange theory and the currency of communication. *Psychological*
1438 *Review*, 126(3), 313.
- 1439 Cikara, M., Bruneau, E. G., & Saxe, R. R. (2011). Us and Them: Intergroup Failures of
1440 Empathy. *Current Directions in Psychological Science*, 20(3), 149–153.
1441 <https://doi.org/10.1177/0963721411408713>
- 1442 Critcher, C. R., Inbar, Y., & Pizarro, D. A. (2013). How quick decisions illuminate moral
1443 character. *Social Psychological and Personality Science*, 4(3), 308–315.
- 1444 Crockett, M. J. (2017). Moral outrage in the digital age. *Nature Human Behaviour*, 1(11), 769.
- 1445 DeScioli, P., Massenkoff, M., Shaw, A., Petersen, M. B., & Kurzban, R. (2014). Equity or
1446 equality? Moral judgments follow the money. *Proceedings of the Royal Society B:*
1447 *Biological Sciences*, 281(1797), 20142112.
- 1448 Dreber, A., Rand, D. G., Fudenberg, D., & Nowak, M. A. (2008). Winners don't punish. *Nature*,
1449 452(7185), 348–351.
- 1450 Effron, D. A., & Conway, P. (2015). When virtue leads to villainy: Advances in research on
1451 moral self-licensing. *Current Opinion in Psychology*, 6, 32–35.
- 1452 Emler, N. (1990). A social psychology of reputation. *European Review of Social Psychology*,
1453 1(1), 171–193.

- 1454 Epley, N., Keysar, B., Van Boven, L., & Gilovich, T. (2004). Perspective taking as egocentric
1455 anchoring and adjustment. *Journal of Personality and Social Psychology*, 87(3), 327.
- 1456 Everett, J. A., Pizarro, D. A., & Crockett, M. J. (2016). Inference of Trustworthiness From
1457 Intuitive Moral Judgments. *Journal of Experimental Psychology: General*.
- 1458 Fehr, E., & Fischbacher, U. (2004). Third-party punishment and social norms. *Evolution and*
1459 *Human Behavior*, 25(2), 63–87. [https://doi.org/10.1016/s1090-5138\(04\)00005-4](https://doi.org/10.1016/s1090-5138(04)00005-4)
- 1460 Frimer, J. A., Skitka, L. J., & Motyl, M. (2017). Liberals and conservatives are similarly
1461 motivated to avoid exposure to one another's opinions. *Journal of Experimental Social*
1462 *Psychology*, 72, 1–12. <https://doi.org/10.1016/j.jesp.2017.04.003>
- 1463 Gawronski, B. (2012). Back to the future of dissonance theory: Cognitive consistency as a core
1464 motive. *Social Cognition*, 30(6), 652–668.
- 1465 Giner-Sorolla, R. (2018, January 24). Powering Your Interaction. *Approaching Significance*.
1466 <https://approachingblog.wordpress.com/2018/01/24/powering-your-interaction-2/>
- 1467 Ginges, J., & Atran, S. (2009). Noninstrumental reasoning over sacred values: An Indonesian
1468 case study. *Psychology of Learning and Motivation*, 50, 193–206.
- 1469 Ginges, J., Atran, S., Medin, D., & Shikaki, K. (2007). Sacred bounds on rational resolution of
1470 violent political conflict. *Proceedings of the National Academy of Sciences*, 104(18),
1471 7357–7360.
- 1472 Gollwitzer, M., & van Prooijen, J.-W. (2016). Psychology of justice. In *Handbook of social*
1473 *justice theory and research* (pp. 61–82). Springer.
- 1474 Gordon, R. A. (1996). Impact of ingratiation on judgments and evaluations: A meta-analytic
1475 investigation. *Journal of Personality and Social Psychology*, 71(1), 54.

- 1476 Haidt, J., & Rose-Stockwell, T. (2019). The Dark Psychology of Social Networks. *The Atlantic*.
1477 [https://www.theatlantic.com/magazine/archive/2019/12/social-media-](https://www.theatlantic.com/magazine/archive/2019/12/social-media-democracy/600763/?fbclid=IwAR3aVENsaj-ndERsHRnsqT_rft7_Whkjr5lIx_c0rGjIfH_zlDSf06c0-Ao)
1478 [democracy/600763/?fbclid=IwAR3aVENsaj-](https://www.theatlantic.com/magazine/archive/2019/12/social-media-democracy/600763/?fbclid=IwAR3aVENsaj-ndERsHRnsqT_rft7_Whkjr5lIx_c0rGjIfH_zlDSf06c0-Ao)
1479 [ndERsHRnsqT_rft7_Whkjr5lIx_c0rGjIfH_zlDSf06c0-Ao](https://www.theatlantic.com/magazine/archive/2019/12/social-media-democracy/600763/?fbclid=IwAR3aVENsaj-ndERsHRnsqT_rft7_Whkjr5lIx_c0rGjIfH_zlDSf06c0-Ao)
- 1480 Haslam, N. (2016). Concept creep: Psychology's expanding concepts of harm and pathology.
1481 *Psychological Inquiry*, 27(1), 1–17.
- 1482 Hawkins, S., Yudkin, D., Juan-Torres, M., & Dixon, T. (2018). Hidden Tribes: A Study of
1483 America's Polarized Landscape. *Hidden Tribes*.
- 1484 Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world?
1485 *Behavioral and Brain Sciences*, 33(2–3), 61–83.
- 1486 Henrich, J., McElreath, R., Barr, A., Ensminger, J., Barrett, C., Bolyanatz, A., Cardenas, J. C.,
1487 Gurven, M., Gwako, E., Henrich, N., Lesorogol, C., Marlowe, F. W., Tracer, D., & Ziker,
1488 J. (2006). Costly punishment across human societies. *Science*, 312(5781), 1767–1770.
1489 <https://doi.org/10.1126/science.1127333>
- 1490 Hernández, J. C., & Mueller, B. (2020, June 1). Global Anger Grows Over George Floyd Death,
1491 and Becomes an Anti-Trump Cudgel. *The New York Times*.
1492 <https://www.nytimes.com/2020/06/01/world/asia/george-floyd-protest-global.html>
- 1493 Herzog, K. (2018). *Call-Out Culture Is a Toxic Garbage Dumpster Fire of Trash*. The Stranger.
1494 [https://www.thestranger.com/slog/2018/01/23/25741141/call-out-culture-is-a-toxic-](https://www.thestranger.com/slog/2018/01/23/25741141/call-out-culture-is-a-toxic-garbage-dumpster-fire-of-trash)
1495 [garbage-dumpster-fire-of-trash](https://www.thestranger.com/slog/2018/01/23/25741141/call-out-culture-is-a-toxic-garbage-dumpster-fire-of-trash)
- 1496 Hofmann, W., Brandt, M. J., Wisneski, D. C., Rockenbach, B., & Skitka, L. J. (2018). Moral
1497 punishment in everyday life. *Personality and Social Psychology Bulletin*, 44(12), 1697–
1498 1711.

- 1499 Hok, H., Martin, A., Trail, Z., & Shaw, A. (2019). When Children Treat Condemnation as a
1500 Signal: The Costs and Benefits of Condemnation. *Child Development*, 91(5), 1439–1455.
1501 <https://doi.org/10.1111/cdev.13323>
- 1502 Horita, Y. (2010). Punishers may be chosen as providers but not as recipients. *Letters on*
1503 *Evolutionary Behavioral Science*, 1(1), 6–9.
- 1504 Jones, E. E. (1965). Conformity as a tactic of ingratiation. *Science*, 149(3680), 144–150.
- 1505 Jones, E. E., & Wortman, C. B. (1973). *Ingratiation: An attributional approach*. General
1506 Learning Press Morristown, NJ.
- 1507 Jordan, J. J., Hoffman, M., Bloom, P., & Rand, D. (2016). Third-party punishment as a costly
1508 signal of trustworthiness. *Nature*, 530(7591), 473–476.
- 1509 Jordan, J. J., & Rand, D. (2017). Third-party punishment as a costly signal of high continuation
1510 probabilities in repeated games. *Journal of Theoretical Biology*, 421, 189–202.
- 1511 Jordan, J. J., & Rand, D. G. (2019). Signaling when nobody is watching: A reputation heuristics
1512 account of outrage and punishment in one-shot anonymous interactions. *Journal of*
1513 *Personality and Social Psychology*.
- 1514 Kazak, A. E. (2018). *Journal article reporting standards*.
- 1515 Kurzban, R., DeScioli, P., & O'Brien, E. (2007). Audience effects on moralistic punishment.
1516 *Evolution and Human Behavior*, 28(2), 75–84.
1517 <https://doi.org/10.1016/j.evolhumbehav.2006.06.001>
- 1518 Leary, M. R., & Kowalski, R. M. (1990). Impression management: A literature review and two-
1519 component model. *Psychological Bulletin*, 107(1), 34.
- 1520 Levine, E. E., Barasch, A., Rand, D. G., Berman, J. Z., & Small, D. A. (2017). *Signaling*
1521 *Emotion and Reason in Cooperation*.

- 1522 Lukianoff, G., & Haidt, J. (2015). The coddling of the American mind. *The Atlantic*, 316(2), 42–
1523 52.
- 1524 Mathew, S., & Boyd, R. (2011). Punishment sustains large-scale cooperation in prestate warfare.
1525 *Proceedings of the National Academy of Sciences*, 108(28), 11375–11380.
- 1526 Mayer, J. (2019). *The Case of Al Franken* | *The New Yorker*.
1527 <https://www.newyorker.com/magazine/2019/07/29/the-case-of-al-franken>
- 1528 McAuliffe, K., Jordan, J. J., & Warneken, F. (2015). Costly third-party punishment in young
1529 children. *Cognition*, 134, 1–10.
- 1530 Melnikoff, D. E., & Bailey, A. H. (2018). Preferences for moral vs. Immoral traits in others are
1531 conditional. *Proceedings of the National Academy of Sciences*, 115(4), E592–E600.
- 1532 Melnikoff, D. E., & Strohminger, N. (2020). The automatic influence of advocacy on lawyers
1533 and novices. *Nature Human Behaviour*, 1–7.
- 1534 Merritt, A. C., Effron, D. A., Fein, S., Savitsky, K. K., Tuller, D. M., & Monin, B. (2012). The
1535 strategic pursuit of moral credentials. *Journal of Experimental Social Psychology*, 48(3),
1536 774–777.
- 1537 Nelissen, R. (2008). The price you pay: Cost-dependent reputation effects of altruistic
1538 punishment. *Evolution and Human Behavior*, 29(4), 242–248. <https://doi.org/Doi>
1539 10.1016/J.Evothumbehav.2008.01.001
- 1540 Nikiforakis, N. (2008). Punishment and counter-punishment in public good games: Can we
1541 really govern ourselves? *Journal of Public Economics*, 92(1), 91–112.
- 1542 Ohtsuki, H., Iwasa, Y., & Nowak, M. A. (2009). Indirect reciprocity provides only a narrow
1543 margin of efficiency for costly punishment. *Nature*, 457(7225), 79–82. <https://doi.org/Doi>
1544 10.1038/Nature07601

- 1545 Onishi, N. (2021, April 8). Powerful Men Fall, One After Another, in France’s Delayed
1546 #MeToo. *The New York Times*.
1547 <https://www.nytimes.com/2021/04/08/world/europe/france-metoo-sandra-muller.html>
- 1548 Pengelly, M. (2019, April 6). Barack Obama warns progressives to avoid “circular firing squad.”
1549 *The Guardian*. [https://www.theguardian.com/us-news/2019/apr/06/barack-obama-](https://www.theguardian.com/us-news/2019/apr/06/barack-obama-progressives-circular-firing-squad-democrats)
1550 [progressives-circular-firing-squad-democrats](https://www.theguardian.com/us-news/2019/apr/06/barack-obama-progressives-circular-firing-squad-democrats)
- 1551 Pletka, D. (2020). Opinion | I never considered voting for Trump in 2016. I may be forced to
1552 vote for him this year. *Washington Post*. [https://www.washingtonpost.com/opinions/i-](https://www.washingtonpost.com/opinions/i-cant-stand-trump-but-democrats-may-force-me-to-vote-for-him/2020/09/14/1cf10518-f6c4-11ea-a275-1a2c2d36e1f1_story.html)
1553 [cant-stand-trump-but-democrats-may-force-me-to-vote-for-him/2020/09/14/1cf10518-](https://www.washingtonpost.com/opinions/i-cant-stand-trump-but-democrats-may-force-me-to-vote-for-him/2020/09/14/1cf10518-f6c4-11ea-a275-1a2c2d36e1f1_story.html)
1554 [f6c4-11ea-a275-1a2c2d36e1f1_story.html](https://www.washingtonpost.com/opinions/i-cant-stand-trump-but-democrats-may-force-me-to-vote-for-him/2020/09/14/1cf10518-f6c4-11ea-a275-1a2c2d36e1f1_story.html)
- 1555 Raihani, N. J., & Bshary, R. (2015). Third-party punishers are rewarded—but third-party helpers
1556 even more so. *Evolution*.
- 1557 Ransom, J. (2020, March 11). Harvey Weinstein’s Stunning Downfall: 23 Years in Prison. *The*
1558 *New York Times*. [https://www.nytimes.com/2020/03/11/nyregion/harvey-weinstein-](https://www.nytimes.com/2020/03/11/nyregion/harvey-weinstein-sentencing.html)
1559 [sentencing.html](https://www.nytimes.com/2020/03/11/nyregion/harvey-weinstein-sentencing.html)
- 1560 Ronson, J. (2015, February 12). How One Stupid Tweet Blew Up Justine Sacco’s Life. *The New*
1561 *York Times*. [https://www.nytimes.com/2015/02/15/magazine/how-one-stupid-tweet-](https://www.nytimes.com/2015/02/15/magazine/how-one-stupid-tweet-ruined-justine-saccos-life.html)
1562 [ruined-justine-saccos-life.html](https://www.nytimes.com/2015/02/15/magazine/how-one-stupid-tweet-ruined-justine-saccos-life.html)
- 1563 Sawaoka, T., & Monin, B. (2018). The Paradox of Viral Outrage. *Psychological Science*, 29(10),
1564 1665–1678. <https://doi.org/10.1177/0956797618780658>
- 1565 Schlosser, E. (2015, June 3). *I’m a liberal professor, and my liberal students terrify me*. Vox.
1566 <https://www.vox.com/2015/6/3/8706323/college-professor-afraid>

- 1567 Scott, S. (2018, February 1). In Defense of Call-out Culture. *City Arts Magazine*.
1568 <https://www.cityartsmagazine.com/defense-call-culture/>
- 1569 Silver, I. M., & Shaw, A. (2018). Pint-Sized Public Relations: The Development of Reputation
1570 Management. *Trends in Cognitive Sciences*, 22(4), 277–279.
1571 <https://doi.org/10.1016/j.tics.2018.01.006>
- 1572 Skitka, L. J. (2010). The psychology of moral conviction. *Social and Personality Psychology*
1573 *Compass*, 4(4), 267–281.
- 1574 Spring, V. L., Cameron, C. D., & Cikara, M. (2018). The Upside of Outrage. *Trends in Cognitive*
1575 *Sciences*, 22(12), 1067–1069. <https://doi.org/10.1016/j.tics.2018.09.006>
- 1576 Stone, J., & Fernandez, N. C. (2008). To practice what we preach: The use of hypocrisy and
1577 cognitive dissonance to motivate behavior change. *Social and Personality Psychology*
1578 *Compass*, 2(2), 1024–1051.
- 1579 Sunstein, C. R. (2019). We Need a Word for Destructive Group Outrage. *Bloomberg Opinion*.
- 1580 Tetlock, P. E. (2002). Social functionalist frameworks for judgment and choice: Intuitive
1581 politicians, theologians, and prosecutors. *Psychological Review*, 109(3), 451.
- 1582 Tetlock, P. E. (2003). Thinking the unthinkable: Sacred values and taboo cognitions. *Trends in*
1583 *Cognitive Sciences*, 7(7), 320–324.
- 1584 Tetlock, P. E., Kristel, O. V., Elson, S. B., Green, M. C., & Lerner, J. S. (2000). The psychology
1585 of the unthinkable: Taboo trade-offs, forbidden base rates, and heretical counterfactuals.
1586 *Journal of Personality and Social Psychology*, 78(5), 853.
- 1587 Thompson, L. (2001). The mind and heart of the negotiator. *Englewood Cliffs, New Jersey:*
1588 *Prentice Hall*.

- 1589 Tosi, J., & Warmke, B. (2016). Moral Grandstanding. *Philosophy & Public Affairs*, 44(3), 197–
1590 217. <https://doi.org/10.1111/papa.12075>
- 1591 Turillo, C. J., Folger, R., Lavelle, J. J., Umphress, E. E., & Gee, J. O. (2002). Is virtue its own
1592 reward? Self-sacrificial decisions for the sake of fairness. *Organizational Behavior and*
1593 *Human Decision Processes*, 89(1), 839–865.
- 1594

Supplementary Materials *for*

Reputation fuels moralistic punishment that people judge to be questionably merited

1. Studies 1-2	2
1.1. Description of secondary DVs	2
1.2. Description of exit survey and Study 1 attention check	2
1.3. Results within individual vignettes	3
1.4. Comparing reputation ratings in our ambiguous vs. unambiguous conditions	5
2. Studies 3a-b	5
2.1. Description of comprehension questions	5
2.2. Items used to measure the perceived moral merits of punishment, and factor analyses of these items	6
2.3. Description of exit survey	9
2.4. Results among comprehenders	10
2.5. Analyses of self-reported continuous commitment to punishing	10
2.6. Effects of observability on the perceived moral merits of punishment	11
2.7. Discussion of whether Studies 3a-b create a “confounded” design	11
3. Study 3 pre-tests	12
3.1. Additional design information	12
3.2. Additional procedure information	13
3.3. Results from Kolmogorov-Smirnov tests of the equity of distributions	15
3.4. Results among comprehenders	15
3.5. Results from additional experimental conditions	16
4. Study 4	16
4.1. Results among comprehenders	16
4.2. Analyses of self-reported continuous commitment to punishing	17
4.3. Effect of observability on post-treatment moral evaluations of punishment	17
4.4. Considering two sources of moral ambiguity	17
5. Reported reputation motives in Studies 3-4	18
6. Discussion of pre-registered analysis plans	20
7. Full vignette texts for Study 1	23
8. Full vignette texts for Study 2	26
9. Full news article texts for Studies 3-4	31

1. Studies 1-2

1.1. Description of secondary DVs

As noted in the main text, in Studies 1-2 we measured a series of secondary DVs investigating (i) subjects' uncertainty regarding the reputation consequences of punishment, and (ii) their expectations about how punishers would be perceived on a set of more specific traits. These DVs were designed to provide a richer picture of subjects' expectations regarding the reputation consequences of punishment. However, our focus in this paper is on the question of whether subjects in our ambiguous conditions expected punishment to have overall positive global reputational consequences (i.e., to confer reputational *benefits*). For this reason, we do not analyze these DVs in this paper. Instead, we plan to analyze our secondary DVs in future work, which will focus on comparing our ambiguous vs. unambiguous conditions with respect to subjects' expectations about both global and specific reputation consequences of punishment.

Here, however, we describe the secondary DVs that we collected in Studies 1-2. First, we measured self-reported uncertainty versus confidence about the reputation value of punishment. Specifically, subjects rated their agreement with four statements on a 0 to 100 sliding scale (0 = Strongly disagree, 50 = Neither agree nor disagree, 100 = Strongly agree). Two statements expressed confidence (e.g., in the context of our Study 1 vignette about homophobia: "Overall, I feel like I have a reasonably good sense of how Anthony reacted to Sam's decision to tell Brett not to attend the barbecue") and two expressed uncertainty (e.g., "I feel like I have no clear sense of what Anthony thought of Sam's decision to tell Brett not to attend the barbecue"); we randomized which set of statements subjects rated first.

Next, we presented subjects with a set of eight questions about some more specific reputation consequences of punishment. To illustrate these questions, in the context of our Study 1 vignette about homophobia, subjects were asked to further consider Anthony's impression of Sam, after Anthony overheard Sam tell Brett not to attend the barbecue. Subjects then rated, in random order, the extent to which Anthony thinks (1) Sam is politically liberal, (2) Sam always treats LGBT people respectfully, (3) Sam is a genuine person, (4) Sam is an overly sensitive person, (5) Sam is someone who rushes to judgement, (6) that by calling Brett out, Sam did something good for the world, (7) Sam is someone he would like to hang out with on the weekend, and (8) Sam is someone he would like to hire to work with him. Subjects answered each question with a 0 to 100 sliding scale (0 = Not at all, 50 = Somewhat, 100 = Very much). We note that item (1) was replaced with a rating of political conservatism in our Study 2b vignettes, and item (2) was adapted across vignettes to form a measure of respectful behavior in the relevant moral domain.

1.2. Description of exit survey and Study 1 attention check

As reported in the main text, after completing our manipulation check measures, subjects in Studies 1-2 completed a post-experimental survey. This survey contained a simple analogy question ("Dog is to puppy as cat is to ____", designed to check for inattentive subjects, or subjects who did not understand English) and a set of additional demographic questions (recall that subjects reported their age, gender, and political party affiliation at the start of the study; here they additionally reported their education, income, belief in god, strength of political party affiliation, general political conservatism or liberalism, conservatism or liberalism on social

issues, and conservatism or liberalism on economic issues). See survey materials on OSF for exact wording of each of these questions.

In the main text, we also report that in Study 1, we presented an attention check at the beginning of the experiment (i.e., before subjects were assigned to conditions) and excluded subjects who failed to pass the check. This attention check was presented after subjects reported their age, gender, and preferred political party affiliation, but before we presented our first vignette. To measure attention, we asked subjects to “Please carefully read the following story about a woman named Lisa” and then presented the following short vignette: “Lisa works at a Home Depot. At the store, Lisa's job is to work in customer service. Normally, Lisa works Tuesday-Sunday but does not work Mondays. However, last week Lisa's coworker Ben asked her to cover his Monday shift. So this Monday, Lisa has to work a 6-hour shift.” Then, on a new page, we asked subjects “What is Lisa's job at the Home Depot?” and provided the following answer choices: (i) Manager, (ii) Cashier, (iii) Stocker, (iv) Customer service, and (iv) It was not specified in the story. Correct answer: (iv).

1.3. Results within individual vignettes

In the main text, we report that all results from Studies 1-2 hold significantly within each individual vignette. Here, we expand on this claim. First, in Figure S1, we reproduce main text Figure 1 (which shows results from Studies 1-2) within each individual vignette. This figure highlights the robustness of our results across vignettes.

Furthermore, we note that all significance tests for Studies 1-2 that are reported in the main text are significant within each individual vignette. Specifically, for each individual vignette in Study 1, we find that (i) composite ratings of the moral merits punishment are significantly lower in our ambiguous condition than our unambiguous condition and (ii) within our ambiguous condition, mean reputation ratings are significantly above the scale midpoint for both of our reputation measures. Across all of these tests, all $ps \leq .003$.

Similarly, for each individual vignette in each of Studies 2a-b, we find that (i) relative to our unambiguous conditions, composite ratings of the moral merits of punishment are significantly lower in our ambiguous conditions; (ii) within our ambiguous conditions, mean reputation ratings are significantly above the scale midpoint for both of our reputation measures; and (iii) within our ambiguous conditions, ratings on both of our reputation measures are significantly higher in our “more ideological” condition than our “less ideological” condition. Across all of these tests, all $ps \leq .015$.

For reasons of brevity, we do not report results from each of these individual tests; however, we note that our OSF page for this paper includes scripts with code to reproduce them.

Finally, we note that our main text reports, for both Studies 1 and 2, that subjects in our ambiguous conditions who *personally* saw punishment as morally questionable (as indexed by ratings of the moral merits of punishment at or below the scale midpoint) nonetheless expected punishment to confer reputational benefits (as indexed by above-midpoint reputation ratings) a meaningful fraction of the time. While this claim was not supported by a significance test, we note that the relevant fraction can be observed, for each individual vignette, in Figure S1.

Figure S1. Results from Studies 1-2 within individual vignettes. Here we reproduce the results from main text Figure 1, within each individual vignette. In panel (A), we plot the perceived moral merits of punishment as a function of ambiguity. In panel (B), we plot ratings of the reputation value of punishment (as measured by our point-estimate reputation measure) within our ambiguous conditions, as a function of Observer ideology. In panel (C), we replicate Panel B, but restrict to observations for which subjects rated the moral merits of punishment at or below the scale midpoint. Across panels, we separately plot results from Studies 1 (Democrats), 2a (Democrats), and 2b (Republicans).

1.4. Comparing reputation ratings in our ambiguous vs. unambiguous conditions

In the main text, we note that while our theoretical focus is on the perceived reputation value of punishment in our *ambiguous* conditions, Studies 1-2 also revealed that subjects in our unambiguous conditions expected punishment to confer even *larger* reputational benefits. Here, we elaborate on this claim. Indeed, in Study 1 we find that reputation ratings were significantly higher in our unambiguous (vs. ambiguous) condition, both for our point-estimate ($b = 11.97$ [8.61, 15.33], $B = .27$, $t = 7.01$, $p < .001$, $n = 401$) and probability-distribution ($b = .51$ [.36, .66], $B = .27$, $t = 6.72$, $p < .001$) measures. Furthermore, in Studies 2a-b we find that, collapsing across our ideology conditions, reputation ratings were significantly higher in our unambiguous (vs. ambiguous) conditions, both for our point-estimate (Democrats $b = 11.70$ [10.04, 13.36], $B = .26$, $t = 13.85$, $p < .001$, $n = 1591$; Republicans $b = 9.82$ [7.88, 11.76], $B = .20$, $t = 9.92$, $p < .001$, $n = 1566$) and probability-distribution (Democrats $b = .46$ [.39, .54], $B = .24$, $t = 11.94$, $p < .001$; Republicans $b = .43$ [.34, .52], $B = .20$, $t = 9.55$, $p < .001$) measures.

2. Studies 3a-b

2.1. Description of comprehension questions

In the main text, we note that the procedure for Studies 3a-b involved presenting subjects with a series of comprehension questions after we introducing the economic game involving the Decider. Here, we describe these questions.

After introducing the economic game in Studies 3a-b, we presented subjects with four comprehension questions about the economic game and its relationship to their punishment decision. The first question read “*Imagine that Player 1 is deciding how much to share with to you. Which decision will result in Player 1 earning the highest payoff?*” and had the following answer choices: (i) *Player 1 deciding to share 0 cents*, (ii) *Player 1 deciding to share 50 cents*, (iii) *Player 1 deciding to share \$1*. Correct answer = (i).

The second question read “*Imagine that Player 1 is deciding how much to share with to you. Which decision will result in you earning the highest payoff?*” and had the following answer choices: (i) *Player 1 deciding to share 0 cents*, (ii) *Player 1 deciding to share 50 cents*, (iii) *Player 1 deciding to share \$1*. Correct answer = (iii).

In the third question, we asked subjects about the Decider’s ideology. In the “less ideological Decider” conditions, the third question read “*What political affiliation is Player 1?*” and had the following answer choices: (i) *Player 1 is a Democrat*, (ii) *Player 1 is a Republican*. Correct answers: (i) in Study 3a and (ii) in Study 3b. In the “more ideological Decider” conditions, the third question instead read “*What information do you know about Player 1’s political affiliation?*” and had the following answer choices: (i) *Player 1 is a Republican*, (ii) *Player 1 is a Democrat*, (iii) *Player 1 is a Democrat who supports the #MeToo movement* (in Study 3a) or *Player 1 is a Republican who is a committed member of the Men’s Rights Movement* (in Study 3b). Correct answers: either (ii) or (iii). (Note that we counted both of these two answers as correct because of ambiguity created by the question; while subjects were told that Player 1 supports #MeToo (in Study 3a) or that Player 1 is a committed member of the Men’s Rights Movement (in Study 3b), it is unclear whether this constitutes information about Player 1’s “political affiliation”.)

Finally, the fourth question read “*Is the Sharing Game related to the component of the HIT in which you will decide whether to donate 30 cents to help hold [biologist Joseph Pringle / dean Elizabeth Cartland] accountable?*” (with Study 3a referencing Pringle and Study 3b

referencing Cartland). In the private conditions, the answer choices were: (i) *No, it is unrelated: Player 1 will NOT learn about the news article, group of organizers, or my donation decision* and (ii) *Yes, it is related*, and the correct answer was (i). In the public conditions, the answer choices were: (i) *No, it is unrelated*, (ii) *Yes, it is related: before deciding how much to share with me, Player 1 will see all of the materials that I saw and will see, learn that I saw them, and find out whether I decided to donate 30 cents*, and the correct answer was (ii).

2.2. Items used to measure the perceived moral merits of punishment, and factor analyses of these items

Across each of Studies 1-4, we measured the perceived moral merits of punishment across a set of scale items. In the main text, we report that within each study, exploratory factor analyses of our scale items produced one-factor solutions. In this section, we provide results from these exploratory factor analyses, which used the principal-component factors method, for each study. These results consistently reveal that only one factor (with very high loadings for all scale items) was retained, and all other factors had Eigenvalues below 1.

Additionally, for Studies 3-4, we report the exact wording of our scale items used to measure the perceived moral merits of punishment. (For Studies 1-2, this information is reported in the main text.)

Study 1 (Democrats)

Factor analysis/correlation
Method: principal-component factors
Rotation: (unrotated)

Number of obs = 1,604
Retained factors = 1
Number of params = 3

Factor	Eigenvalue	Difference	Proportion	Cumulative
Factor1	2.58710	2.24663	0.8624	0.8624
Factor2	0.34046	0.26803	0.1135	0.9759
Factor3	0.07244	.	0.0241	1.0000

LR test: independent vs. saturated: $\chi^2(3) = 4409.09$ Prob> $\chi^2 = 0.0000$

Factor loadings (pattern matrix) and unique variances

Variable	Factor1	Uniqueness
likely	0.8715	0.2404
appropriate	0.9544	0.0890
justified	0.9574	0.0834

Study 2a (Democrats)

Factor analysis/correlation
Method: principal-component factors
Rotation: (unrotated)

Number of obs = 6,364
Retained factors = 1
Number of params = 3

Factor	Eigenvalue	Difference	Proportion	Cumulative
Factor1	2.20445	1.63295	0.7348	0.7348
Factor2	0.57149	0.34744	0.1905	0.9253
Factor3	0.22406	.	0.0747	1.0000

LR test: independent vs. saturated: $\chi^2(3) = 8047.35$ Prob> $\chi^2 = 0.0000$

Factor loadings (pattern matrix) and unique variances

Variable	Factor1	Uniqueness
offensive	0.9101	0.1718
appropriate	0.8916	0.2050
harsh	-0.7624	0.4188

Study 2b (Republicans)

Factor analysis/correlation
 Method: principal-component factors
 Rotation: (unrotated)

Number of obs = 6,264
 Retained factors = 1
 Number of params = 3

Factor	Eigenvalue	Difference	Proportion	Cumulative
Factor1	2.09705	1.43893	0.6990	0.6990
Factor2	0.65812	0.41329	0.2194	0.9184
Factor3	0.24483	.	0.0816	1.0000

LR test: independent vs. saturated: $\chi^2(3) = 6794.64$ Prob> $\chi^2 = 0.0000$

Factor loadings (pattern matrix) and unique variances

Variable	Factor1	Uniqueness
offensive	0.8904	0.2071
appropriate	0.8975	0.1945
harsh	-0.7061	0.5014

Study 3a (Democrats)

Subjects answered each of the below questions on a 1-10 scale (1 = Not at all, 10 = Very):

1. How confident are you that the allegations against Joseph Pringle are true?
2. How bad of a person do you think Joseph Pringle is?
3. How immoral do you think Joseph Pringle's actions were?
4. How angry are you at Joseph Pringle?
5. How much of an outrage is Joseph Pringle's behavior?
6. How much do you think Joseph Pringle deserves to be punished?
7. How important do you think it is that Joseph Pringle is punished?
8. How comfortable are you with the group of organizers' approach to holding Joseph Pringle accountable?
9. To what extent do you think the group of organizers' approach to holding Joseph Pringle accountable is proportionate and appropriate?
10. To what extent do you think the group of organizer's approach to holding Joseph Pringle accountable is moral?

Factor analysis/correlation
 Method: principal-component factors
 Rotation: (unrotated)

Number of obs = 1,529
 Retained factors = 1
 Number of params = 10

Factor	Eigenvalue	Difference	Proportion	Cumulative
Factor1	8.12305	7.45405	0.8123	0.8123
Factor2	0.66900	0.39691	0.0669	0.8792
Factor3	0.27209	0.03852	0.0272	0.9064
Factor4	0.23357	0.05879	0.0234	0.9298
Factor5	0.17478	0.03170	0.0175	0.9473
Factor6	0.14309	0.00365	0.0143	0.9616
Factor7	0.13944	0.02461	0.0139	0.9755
Factor8	0.11483	0.03112	0.0115	0.9870
Factor9	0.08371	0.03727	0.0084	0.9954
Factor10	0.04643	.	0.0046	1.0000

LR test: independent vs. saturated: $\chi^2(45) = 2.2e+04$ Prob> $\chi^2 = 0.0000$

Factor loadings (pattern matrix) and unique variances

Variable	Factor1	Uniqueness
personal_e-1	0.8869	0.2134
personal_e-2	0.9070	0.1773
personal_e-3	0.8960	0.1972
personal_e-4	0.8884	0.2108
personal_e-5	0.9168	0.1595
personal_e-6	0.9450	0.1070
personal_e-7	0.9427	0.1114
personal_e-8	0.8621	0.2567
personal_e-9	0.8822	0.2217
personal_-10	0.8821	0.2218

Study 3b (Republicans)

Subjects answered each of the below questions on a 1-10 scale (1 = Not at all, 10 = Very):

1. How confident are you that Elizabeth Cartland discriminates against men?
2. How bad of a person do you think Elizabeth Cartland is?
3. How immoral do you think Elizabeth Cartland's decision to take action against Jones was?
4. How angry are you at Elizabeth Cartland?
5. How much of an outrage is Elizabeth Cartland's behavior?
6. How much do you think Elizabeth Cartland deserves to be punished?
7. How important do you think it is that Elizabeth Cartland is punished?
8. How comfortable are you with the group of organizers' approach to holding Elizabeth Cartland accountable?
9. To what extent do you think the group of organizers' approach to holding Elizabeth Cartland accountable is proportionate and appropriate?
10. To what extent do you think the group of organizers' approach to holding Elizabeth Cartland accountable is moral?

Factor analysis/correlation
 Method: principal-component factors
 Rotation: (unrotated)

Number of obs = 1,565
 Retained factors = 1
 Number of params = 10

Factor	Eigenvalue	Difference	Proportion	Cumulative
Factor1	7.62284	6.69387	0.7623	0.7623
Factor2	0.92898	0.59997	0.0929	0.8552
Factor3	0.32901	0.06134	0.0329	0.8881
Factor4	0.26767	0.04956	0.0268	0.9149
Factor5	0.21812	0.03217	0.0218	0.9367
Factor6	0.18595	0.02586	0.0186	0.9553
Factor7	0.16009	0.03136	0.0160	0.9713
Factor8	0.12874	0.03778	0.0129	0.9841
Factor9	0.09095	0.02330	0.0091	0.9932
Factor10	0.06765	.	0.0068	1.0000

LR test: independent vs. saturated: $\chi^2(45) = 2.0e+04$ Prob> $\chi^2 = 0.0000$

Factor loadings (pattern matrix) and unique variances

Variable	Factor1	Uniqueness
personal_e-1	0.8745	0.2353
personal_e-2	0.8562	0.2669
personal_e-3	0.8843	0.2181
personal_e-4	0.8602	0.2600
personal_e-5	0.9075	0.1765
personal_e-6	0.9266	0.1415
personal_e-7	0.9227	0.1486
personal_e-8	0.8146	0.3364
personal_e-9	0.8390	0.2960
personal_~10	0.8379	0.2978

Study 4 (Democrats)

When providing their initial rating of the moral merits of punishment, subjects in Study 4 answered questions 1, 2, 4, 6, 8, and 9 from our set of Study 3a items.

Factor analysis/correlation Number of obs = 2,415
 Method: principal-component factors Retained factors = 1
 Rotation: (unrotated) Number of params = 6

Factor	Eigenvalue	Difference	Proportion	Cumulative
Factor1	4.60583	3.94086	0.7676	0.7676
Factor2	0.66497	0.43689	0.1108	0.8785
Factor3	0.22808	0.04614	0.0380	0.9165
Factor4	0.18194	0.01241	0.0303	0.9468
Factor5	0.16953	0.01989	0.0283	0.9751
Factor6	0.14965	.	0.0249	1.0000

LR test: independent vs. saturated: $\chi^2(15) = 1.4e+04$ Prob> $\chi^2 = 0.0000$

Factor loadings (pattern matrix) and unique variances

Variable	Factor1	Uniqueness
first_pers-1	0.8928	0.2028
first_pers-2	0.8924	0.2036
first_pers-3	0.8957	0.1977
first_pers-4	0.9203	0.1530
first_pers-5	0.8220	0.3244
first_pers-6	0.8291	0.3127

2.3. Description of exit survey

As reported in the main text, after completing our dependent measures, subjects in Studies 3a-b completed a post-experimental survey. In this survey, we began by asking subjects an open-ended question about how they made their decisions in the study. Next, we asked subjects a series of questions designed to investigate the extent to which they were consciously motivated by reputation concerns (including as compared to other motivations). Specifically, we asked subjects to rate the extent to which they made their decisions because they (i) personally felt it was truly the right decision, (ii) wanted to see themselves as a good person, (iii) wanted others to see them as a good person, and (iv) wanted the Decider to see them as a good person. Furthermore, subjects in the public condition additionally rated the extent to which they (i) would have behaved differently if the Decider were out of the picture, (ii) would have behaved differently if punishing did not cost money, (iii) decided purely on the basis of what would look good in the eyes of the Decider, and (iv) decided purely on the basis of what would earn them the most money. After answering these questions, we presented all subjects with an open-ended question about anything they'd like to share about the impression of the study.

Next, we asked subjects about their previous participation in related studies, as well as three questions about their beliefs that various components of the experiment (specifically, “the news article”, “the group of organizers”, and “Player 1 and the sharing game”). Next, we presented subjects with a simple analogy question (“Dog is to puppy as cat is to ____”, designed to check for inattentive subjects, or subjects who did not understand English) and a set of additional demographic questions (recall that subjects reported their age, gender, and political party affiliation at the start of the study; here they additionally reported their education, income, belief in god, strength of political party affiliation, general political conservatism or liberalism, conservatism or liberalism on social issues, and conservatism or liberalism on economic

issues). Finally, we presented subjects with a three-item Cognitive reflection task. See survey materials on OSF for exact wording of each of these questions.

We also note that Study 4 employed the same post-experimental survey as Studies 3a-b, with the exception that we added a measure of ethnicity to our set of demographic questions.

2.4. Results among comprehenders

As described in the main text, per our pre-registration, our primary analyses of Studies 3a-b include data from all subjects. However, we find similar results when restricting to subjects who correctly answered our set of comprehension questions. Here, we support this claim.

In Study 3a, 80% of subjects answered all four of our comprehension questions correctly. However, in the analyses reported here, we only restrict based on our first *three* questions. This choice reflects that our fourth question, which had different answer choices (and a different correct answer) in the public vs. private conditions, had a significantly lower pass rate in public than private ($b = -.03 [-.06, -.01]$, $B = -.07$, $t = -2.67$, $p = .008$, $n = 1558$); restricting to subjects who passed this question thus introduces a selection effect that can undermine random assignment.

Looking to the 83% of Study 3a subjects who passed all of our first three comprehension questions ($n = 1299$), when analyzing punishment decisions, we find a significant negative main effect of ambiguity ($b = -.29 [-.34, -.24]$, $B = -.31$, $t = -11.80$, $p < .001$), a significant positive main effect of observability ($b = .12 [.07, .17]$, $B = .13$, $t = 4.98$, $p < .001$), and no significant interaction between ambiguity and observability ($b = -.01 [-.10, .09]$, $B = -.003$, $t = -.10$, $p = .917$). Turning to simple effects, we find significant positive effects of observability both within the unambiguous ($b = .12 [.06, .19]$, $B = .13$, $t = 3.65$, $p < .001$, $n = 668$) and ambiguous ($b = .12 [.05, .19]$, $B = .13$, $t = 3.40$, $p < .001$, $n = 631$) conditions. We also find significant negative effects of ambiguity both within the private ($b = -.29 [-.35, -.22]$, $B = -.31$, $t = -8.37$, $p < .001$, $n = 665$) and public ($b = -.29 [-.36, -.22]$, $B = -.31$, $t = -8.32$, $p < .001$, $n = 634$) conditions.

In Study 3b, 74% of subjects answered all four comprehension questions correctly. The pass rate on our fourth question was not significantly different in public vs. private ($b = -.02 [-.05, .02]$, $B = -.02$, $t = -.98$, $p = .329$, $n = 1565$); thus, in our analyses of comprehenders, we restrict based on all four questions. Looking to the 74% of Study 3b subjects who passed all four questions ($n = 1156$), when analyzing punishment decisions, we find a significant negative main effect of ambiguity ($b = -.17 [-.22, -.12]$, $B = -.18$, $t = -6.38$, $p < .001$), a significant positive main effect of observability ($b = .14 [.09, .19]$, $B = .15$, $t = 5.23$, $p < .001$), and no significant interaction between ambiguity and observability ($b = -.03 [-.14, .07]$, $B = .02$, $t = -.58$, $p = .562$). Turning to simple effects, we find significant positive effects of observability both within the unambiguous ($b = .16 [.08, .23]$, $B = .17$, $t = 4.14$, $p < .001$, $n = 587$) and ambiguous ($b = .12 [.05, .20]$, $B = .13$, $t = 3.26$, $p = .001$, $n = 569$) conditions. We also find significant negative effects of ambiguity both within the private ($b = -.16 [-.23, -.08]$, $B = -.17$, $t = -4.21$, $p < .001$, $n = 607$) and public ($b = -.19 [-.26, -.11]$, $B = -.20$, $t = -4.80$, $p < .001$, $n = 549$) conditions.

2.5. Analyses of self-reported continuous commitment to punishing

As described in the main text, in addition to making punishment decisions, subjects in Studies 3a-b also reported their continuous commitment to punishing. As with punishment decisions, we used effect coding to investigate the main effects of ambiguity and observability, as well as their interaction, on this dependent variable.

Mirroring the patterns we observe when looking to punishment decisions, we find significant main effects of ambiguity: ambiguity decreased reported commitment to punishing (Study 3a: $b = -1.94 [-2.22, -1.66]$, $B = -.33$, $t = -13.69$, $p < .001$; Study 3b: $b = -1.05 [-1.33, -.77]$, $B = -.18$, $t = -7.41$, $p < .001$). We also find significant main effects of observability: making punishment observable increased reported commitment to punishing (Study 3a: $b = .66 [.38, .94]$, $B = .11$, $t = 4.63$, $p < .001$; Study 3b: $b = .30 [.02, .58]$, $B = .05$, $t = 2.11$, $p = .035$). Finally, we find no significant interactions between ambiguity and observability (Study 3a: $b = .46 [-.10, 1.01]$, $B = .04$, $t = 1.61$, $p = .107$; Study 3b: $b = -.26 [-.82, .30]$, $B = -.02$, $t = -.92$, $p = .359$).

2.6. Effects of observability on the perceived moral merits of punishment

As described in the main text, the designs of Studies 3a-b allow us to investigate the influence of observability on the perceived moral merits of punishment. In the main text, we merely used subjects' ratings of the merits of punishment to validate our ambiguity manipulations (by demonstrating that, within our private conditions, subjects rated punishment as less merited in our ambiguous conditions than our unambiguous conditions). However, we can also ask: did our reputation manipulation influence the perceived moral merits of punishment, such that ratings differed in the public vs. private conditions? Here, we report analyses that address this question.

To this end, we use effect coding to investigate the main effects of ambiguity and observability, as well as their interaction, on ratings of the moral merits of punishment. We find significant main effects of ambiguity: ambiguity decreased the perceived merits of punishment (Study 3a: $b = -2.58 [-2.80, -2.36]$, $B = -.51$, $t = -23.17$, $p < .001$; Study 3b: $b = -1.70 [-1.92, -1.48]$, $B = -.36$, $t = -15.16$, $p < .001$). In Study 3a, we also find a marginally significant positive main effect of observability ($b = .20 [-.02, .42]$, $B = .04$, $t = 1.80$, $p = .073$), and a marginally significant positive interaction between ambiguity and observability ($b = .42 [-.02, .85]$, $B = .04$, $t = 1.88$, $p = .061$). In contrast, in Study 3b, we find no significant main effect of observability ($b = -.03 [-.25, .19]$, $B = -.01$, $t = -.30$, $p = .764$) and no significant interaction between ambiguity and observability ($b = -.14 [-.58, .30]$, $B = -.02$, $t = -.64$, $p = .519$). Overall, then, we do not find strong evidence that our reputation manipulation influenced the perceived moral merits of punishment.

2.7. Discussion of whether Studies 3a-b create a “confounded” design

As described in the main text, in Studies 3a-b, we describe the Decider as more ideological in our ambiguous conditions than our unambiguous conditions. On the surface, this design may sound “confounded”. Here, however, we provide a detailed discussion of why our design is *not* actually confounded with respect to our theoretical constructs of interest.

As we show in Figure 2 of the main text, describing the Decider as more ideological in the ambiguous (vs. unambiguous) conditions of Study 3 allows us to (i) vary the perceived ambiguity of the moral case for punishment, but also (ii) hold constant the perceived strength of reputational incentives. Achieving these two aims is critical for our design to provide a valid test of our key theoretical question. Only by achieving these two aims can we ask: is the motivational force of similarly-strong reputational incentives any weaker in contexts where subjects perceive a more ambiguous moral case for punishment?

In contrast, if we had manipulated ambiguity *without* varying the observer's ideology, we *would* have introduced an important confound with respect to our theoretical constructs of interest. We find evidence that subjects perceive weaker reputational incentives when the moral

case for punishment is more ambiguous. For this reason, pairing our ambiguous vs. unambiguous conditions with equally ideological observers would have created a design in which we simultaneously varied both (i) the perceived ambiguity of the moral case for punishment, and (ii) the perceived strength of reputational incentives.

Thus, our design is *not* “confounded” in a way that is problematic. Rather, by combining two design elements (i.e., a manipulation of ambiguity and a manipulation of Decider ideology), our design successfully manipulates the theoretical construct that we *do* wish to vary (moral ambiguity) and holds approximately constant the construct that we do *not* wish to vary (reputational incentives).

Another way to think about this issue is: what is the conclusion that we are drawing from our design, and does the apparent “confound” provide an alternative explanation? In Study 3, we find that varying the ambiguity of the moral case for punishment does *not* reduce the influence of reputation on punishment behavior. More specifically, we find robust effects of observability on punishment within our ambiguous conditions (that are no smaller than the effects within our unambiguous conditions), despite the fact that subjects in our ambiguous conditions report less personal moral support for punishment. We interpret this as evidence that the power of reputation to drive punishment is *not* attenuated by ambiguous situations. Can our results be dismissed (and this conclusion thus challenged) by the fact that the ambiguous condition is paired with a more ideological Decider than the unambiguous condition?

We argue that the answer is “no”. One might argue: “more ideological Decider will serve to inflate reputation effects, because people expect more ideological observers to react more positively to punishment”. This is true—but the logic behind our design is that this fact serves to offset a counter-veiling force: holding Decider ideology constant, people would perceive *weaker* reputational incentives for punishment in the ambiguous (vs. unambiguous) condition of our paradigm.

Alternatively, one might argue: “introducing a more ideological Decider could prime subjects in the ambiguous condition to evaluate the alleged moral transgression more harshly”. This may be true, but the Decider is described as more ideological in both the public *and* private versions of our ambiguous conditions. Therefore, such an effect would simply serve to increase baseline moral support for punishment in the ambiguous condition, relative to the unambiguous condition. Yet as reported in the main text, we nonetheless find that at baseline (i.e., in private), subjects report much less personal moral support for punishment in the ambiguous condition than the unambiguous condition. Thus, we find clear evidence that our ambiguity manipulation swamped any countervailing effect of Decider ideology on personal moral support for punishment, such that subjects in our ambiguous condition saw punishment as less morally merited than subjects in our unambiguous condition (as intended). It is therefore striking that we nonetheless find no evidence that subjects in our ambiguous condition were any less responsive to reputational incentives than subjects in our unambiguous condition.

For all of these reasons, we believe that Study 3 is appropriately designed to answer our theoretical questions of interest and is not “confounded”.

3. Study 3 pre-tests

3.1. Additional design information

Here, we provide more information about the design of our Study 3 pre-tests. In the main text, we report that in our Democrat and Republican pre-tests, we assigned subjects to the “public” versions of the “ambiguous + more ideological” or “unambiguous + less ideological”

conditions that we employed in Studies 3a-b, and then measured subjects' expectations about the reputation value of punishment. Furthermore, we explain that our pre-tests served to provide evidence that by pairing our ambiguous conditions with more ideological Deciders, we were able to approximate equate the perceived reputation value of punishment across our ambiguous vs. unambiguous conditions.

In order to find a proper calibration (i.e., to determine *how* much more ideological the Decider should be in our ambiguous conditions), our pre-tests included more than just the two conditions that we describe in the main text and drew on in Studies 3a-b. Specifically, our pre-tests featured additional conditions in which we paired our ambiguous condition with a more ideological Decider, but described that Decider somewhat differently than we did in Studies 3a-b. We do not describe these conditions in the main text, because we did not draw on them in Studies 3a-b. Here, however, we provide more information about them.

In our Democrat pre-test (across all conditions, target $n = 400$, final $n = 397$, mean age = 37.02, 36.52% male, 64.23% with a college or graduate degree), subjects were either (i) assigned to the unambiguous condition, and told that the Decider was "a Democrat", (ii) assigned to the ambiguous condition, and told that the Decider was "a Democrat who supports the #MeToo movement", (iii) assigned to the ambiguous condition, and told that the Decider was "a strong Democrat who supports the #MeToo movement", or (iv) assigned to the ambiguous condition, and told that the Decider "supports the #MeToo movement". The first two conditions (i.e., conditions (i) and (ii); $n = 194$) are the conditions that we drew on in Study 3a and therefore focus on in the main text.

In our Republican pre-test (across all conditions, target $n = 350$, final $n = 343$, mean age = 39.23, 51.90% male, 50.15% with a college or graduate degree), subjects were either (i) assigned to the unambiguous condition, and told that the Decider was "a Republican", (ii) assigned to the ambiguous condition, and told that the Decider was "a Republican who is a committed member of the Men's Rights Movement", or (iii) assigned to the ambiguous condition, and told that the Decider was "a Republican who is a committed member of the Men's Rights Movement and believes that the #MeToo movement has gone too far". Similarly, we drew on the first two conditions ($n = 231$) in Study 3b and therefore focus on them in the main text.

3.2. Additional procedure information

Here, we provide additional procedure information about our Study 3 pre-tests. Our Study 3 pre-tests closely mirrored the public conditions of Studies 3a-b, with a few exceptions that we describe below.

First, because our Study 3 pre-tests featured a wider variety of Decider ideology descriptions, there were some corresponding changes to our third comprehension question about the economic game (in which subjects were asked about the Decider's ideology). Specifically, in our Democrat pre-test, in condition (i) the third question read "*What political affiliation is Player 1?*" and had the following answer choices: (i) *Democrat*, (ii) *Republican*, (iii) *Independent*. Correct answer: (i). In conditions (ii) and (iii), the question read "*What information do you know about Player 1's political affiliation?*" and the first two answer choices were always: "*Player 1 is a Republican*" and "*Player 1 is a Democrat*". In condition (ii), the third answer choice was "*Player 1 is a Democrat who supports the #MeToo movement*" and in condition (iii), the third answer choice was "*Player 1 is a strong Democrat who supports the #MeToo movement*". Like in Studies 3a-b, we counted both the second and third answer choices as correct (given the ambiguity created

by the question). Finally, in condition (iv), the question read “*What information do you know about Player 1?*” and the answer choices were “*Player 1 is male*”, “*Player 2 is female*”, and “*Player 1 supports the #MeToo movement*”. Only the third answer was counted as correct.

In our Republican pre-test, in condition (i), the third question read “*What political affiliation is Player 1?*” and had the following answer choices: (i) *Player 1 is a Republican*, (ii) *Player 1 is a Democrat*. Correct answer: (i). In conditions (ii) and (iii), the question read “*What information do you know about Player 1's political affiliation?*” and the first two answer choices were always: “*Player 1 is a Democrat*” and “*Player 1 is a Republican*”. In condition (ii), the third answer choice was “*Player 1 is a Republican who is a committed member of the Men's Rights Movement*” and in condition (iii), the third answer choice was “*Player 1 is a Republican who is a committed member of the Men's Rights Movement and believes that the #MeToo movement has gone too far*”. We counted both the second and third answer choices as correct (given the ambiguity created by the question).

Second, we note that after introducing our economic game involving the Decider, Studies 3a-b sought to enhance the salience of our observability manipulation by presenting subjects with screenshot(s) that purportedly illustrated the economic game from the perspective of the Decider. While we employed this approach in our Republican pre-test, our Democrat pre-test did not present any screenshots.

Third, we note that after subjects read the full news articles detailing the relevant allegations (i.e., the harassment allegations in Study 3a and the discrimination allegations in Study 3b), rather than asking subjects to make punishment decisions, we instead measured their expectations regarding the reputation value of punishment. Specifically, we told subjects that, before deciding whether to donate to the group of organizers, we would like them to make some guesses about Player 1 (i.e., the Decider). We also reminded subjects (i) of the ideology information about the Decider, and (ii) that the Decider would learn about their punishment decision (and more generally would be able to see all of the materials that they saw). We then measured their expectations about the reputation value of punishment.

Furthermore, only after collecting this data did we inform subjects that they would not actually be making punishment decisions. Thus, importantly, while subjects our Study 3 pre-tests merely rated their expectations regarding the reputation value of punishment and did *not* make punishment decisions, they were told that they would be making punishment decisions before we measured their expectations. We made this design choice so that our pre-tests could accurately capture the psychology of a decision-maker contemplating whether to punish.

As overviewed in the main text, we measured subjects' expectations regarding the reputation value of punishment in two ways. First, subjects predicted, in random order, the number of cents (between 0 and 100, measured in 10-cent increments) that the Decider would share with them if they (i) did, and (ii) did not, donate to the group of punishing organizers. Specifically, subjects were asked (i) “*If you choose NOT to donate 30 cents to the group of organizers, and thus NOT to help hold [Pringle/Cartland] accountable...How many cents, if any, do you think Player 1 will share with you?*” and (ii) “*If you choose TO donate 30 cents to the group of organizers, and thus TO help hold [Pringle/Cartland] accountable...How many cents, if any, do you think Player 1 will share with you?*”. We then computed each subject's predicted financial gain from punishing as the difference between these numbers.

Second, subjects answered the question “*Do you think Player 1 will see you more positively if you DO or do NOT donate?*” via a 1-9 Likert scale (1 = Much more positively if I do NOT donate, 5 = No difference, 9 = much more positively if I DO donate).

After subjects completed these guesses about the Decider they were paired with, we asked them to make a second set of guesses about a second and hypothetical Decider. Subjects were asked to imagine that this Decider did not get to read the full news article (but did get to see the news headline and organizer statement). This second set of guesses was designed to address a research question we do not discuss in this paper (and was measured after all DVs of interest); thus, they are not analyzed. We also note that before making their second set of guesses, subjects were asked to answer one comprehension question about the second set of guesses (in order to confirm that they understood that the hypothetical second Decider did not get to read the full news article).

After we measured this set of DVs, subjects completed a post-experimental survey. In this survey, we presented the same set of questions as in Studies 3a-b, with the exceptions that (i) we did not ask subjects our series of questions about the extent to which subjects were consciously motivated by reputation concerns (including as compared to other motivations), (ii) we did not ask subjects about previous participation in related studies, and (iii) in our Democrat pre-test but not our Republican pre-test, we included a Machiavellianism scale in our set of demographic questions.

3.3. Results from Kolmogorov-Smirnov tests of the equity of distributions

In main text, we report that our pre-tests suggest that subjects expected punishment to confer similar reputation benefits across our “ambiguous + more ideological” and “unambiguous + less ideological” conditions. To support this claim, we compare mean values on our two measures of the perceived reputation value of punishment (i.e., predicted financial gains from punishing and Likert scale ratings) across conditions, and report that we found no significant differences for either measure. Here, we note that using Kolmogorov-Smirnov tests of the equality of distributions, we also find no condition differences in the distributions of these measures. Specifically, in our Democrat pre-test, Kolmogorov-Smirnov tests reveal no condition differences for our financial gains ($D = .09, p = .868$) or Likert ($D = .09, p = .834$) measures. And in our Republican pre-test, Kolmogorov-Smirnov tests reveal no condition differences in the distributions for our financial gains ($D = .07, p = .965$) or Likert ($D = .14, p = .217$) measures.

3.4. Results among comprehenders

Here, we report our pre-test results, restricting to subjects who answered all comprehension questions correctly. Because all pre-test subjects were assigned to the public condition of our behavioral paradigm, there is no concern that pass rates on our fourth comprehension question might have differed across public vs. private conditions; we thus restrict based on all four comprehension questions.

In our Democrat pre-test, 77% of subjects answered all comprehension questions correctly. Among these subjects, across the two conditions we focus on in the main text, subjects predicted that punishment would result in similar financial gains (relative effect of “ambiguous + more ideological” condition, in cents: $b = 3.55 [-4.73, 11.82]$, $B = .07, t = 0.85, p = .398$), and be evaluated similarly (on our 9-point Likert scale, $b = .10 [-.48, .68]$, $B = .03, t = 0.35, p = .728$), $n = 150$. Using Kolmogorov-Smirnov tests of the equality of distributions, we also find no condition differences in the distributions of these measures ($D = .10, p = .871$ and $D = .06, p = 1.000$ for our first and second measures, respectively).

In our Republican pre-test, 77% of subjects answered all comprehension questions correctly. Among these subjects, across the two conditions we focus on in the main text, subjects

predicted that punishment would result in similar financial gains (relative effect of “ambiguous + more ideological” condition: $b = -1.61 [-10.49, 7.27]$, $B = -.03$, $t = -0.36$, $p = .721$), and be evaluated similarly ($b = .005 [-.59, .60]$, $B = .001$, $t = 0.02$, $p = .987$), $n = 186$. Using Kolmogorov-Smirnov tests of the equality of distributions, we also find no condition differences in the distributions of these measures ($D = .06$, $p = .991$ and $D = .13$, $p = .433$).

3.5. Results from additional experimental conditions

As noted above, while our Democrat and Republican pre-tests featured four and three conditions, respectively, in the main text we focus on the first two conditions from each pre-test (because we drew on these conditions in Studies 3a-b). In Tables S1a-b, however, we report the means and standard deviations for each of our two dependent measures across each of our conditions (for our Democrat and Republican pre-tests, respectively).

	Condition (i)	Condition (ii)	Condition (iii)	Condition (iv)
Dependent variable	Unambiguous + Democrat ($n = 100$)	Ambiguous + Democrat who Supports #MeToo ($n = 94$)	Ambiguous + Strong Democrat who Supports #MeToo ($n = 105$)	Ambiguous + Supports #MeToo ($n = 98$)
Predicted extent to which punishment will be perceived positively (on Likert scale)	$M = 6.98$, $SD = 2.07$	$M = 7.16$, $SD = 1.67$	$M = 7.04$, $SD = 2.13$	$M = 7.51$, $SD = 1.74$
Predicted financial gain from punishing (in cents)	$M = 19.90$, $SD = 30.13$	$M = 23.40$, $SD = 24.56$	$M = 16.19$, $SD = 25.62$	$M = 24.80$, $SD = 24.46$

Table S1a. Results across all conditions of our Democrat pre-test.

And second, in the table below, we report these results for our Republican pre-test.

	Condition (i)	Condition (ii)	Condition (iii)
Dependent variable	Unambiguous + Republican ($n = 112$)	Ambiguous + Republican who is a committed member of the Men's Rights Movement ($n = 119$)	Ambiguous + Republican who is a committed member of the Men's Rights Movement and believes #MeToo has gone too far ($n = 112$)
Predicted extent to which punishment will be perceived positively (on Likert scale)	$M = 6.60$, $SD = 2.06$	$M = 6.78$, $SD = 2.20$	$M = 6.51$, $SD = 2.37$
Predicted financial gain from punishing (in cents)	$M = 16.61$, $SD = 31.30$	$M = 16.72$, $SD = 31.49$	$M = 19.29$, $SD = 32.51$

Table S1b. Results across all conditions of our Democrat pre-test.

4. Study 4

4.1. Results among comprehenders

In Study 4, 79% of subjects answered all comprehension questions correctly. Here, we report our key results, restricting to subjects who passed our comprehension questions. However, as in other analyses, we do not restrict based on our fourth question, which had a significantly lower pass rate in public than private ($b = -.04 [-.06, -.02]$, $B = -.09$, $t = -4.35$, $p < .001$, $n = 2415$). Looking to the 82% of subjects who passed all of our first three comprehension questions, we find that overall, there was a significant positive observability effect on punishment ($b = .12 [.08, .15]$, $B = .13$, $t = 6.54$, $p < .001$), $n = 1982$. Next, we repeat these analyses among each bin. Looking to our punishment DV, we find significant positive observability effects across each of

bins one ($b = .05$ [.01, .09], $B = .12$, $t = 2.34$, $p = .020$, $n = 351$), two ($b = .11$ [.05, .16], $B = .15$, $t = 4.01$, $p < .001$, $n = 727$), three ($b = .18$ [.10, .26], $B = .18$, $t = 4.44$, $p < .001$, $n = 574$) and four ($b = .12$ [.01, .23], $B = .12$, $t = 2.22$, $p = .027$, $n = 330$).

4.2. Analyses of self-reported continuous commitment to punishing

Like with Studies 3a-b, in addition to making punishment decisions, subjects in Study 4 also reported their continuous commitment to punishing. Mirroring the patterns we observe when looking to punishment decisions, we find that overall, observability increased subjects' reported commitment to punishing ($b = .50$ [.34, .66], $B = .09$, $t = 6.32$, $p < .001$, $n = 2415$). When looking to each individual bin, we find significant observability effects on commitment to punishing in each of bins one ($b = .48$ [.20, .77], $B = .16$, $t = 3.33$, $p = .001$, $n = 393$), two ($b = .63$ [.37, .88], $B = .16$, $t = 4.87$, $p < .001$, $n = 857$), and three ($b = .43$ [.14, .73], $B = .10$, $t = 2.87$, $p = .004$, $n = 741$), as well as a marginally significant effect in bin four ($b = .42$ [-.01, .85], $B = .09$, $t = 1.91$, $p = .056$, $n = 424$).

4.3. Effect of observability on post-treatment moral evaluations of punishment

In Study 4, in addition to measuring subjects' initial (i.e., pre-treatment) perceptions of the moral merits of punishment (using items 1, 2, 4, 6, 8, and 9 from our Study 3a scale), we also collected a second, post-treatment measure of the perceived moral merits of punishment. To do so, we presented subjects with a second six-item scale. This scale used the remaining items from our Study 3a scale (i.e., items 3, 5, 7, and 10) plus two additional items ("How accurate do you think the allegations against [the Professor] are?" and "How fair and reasonable do you think the group of organizers and their approach are?").

Thus, our Study 4 design allows us to investigate the effects of observability on post-treatment moral evaluations of punishment. However, because we measured initial moral evaluations at the beginning of the study using a very similar scale, we did not necessarily expect our observability manipulation to have a substantial impact on final moral evaluations.

Overall, we find a significant positive effect of observability on post-treatment ratings of the moral merits of punishment in Study 4, $b = .07$ [.005, .13], $B = .02$, $t = 2.11$, $p = .035$, $n = 2415$. However, this effect is very small (note that the unstandardized coefficient of .07 scale points is in the context of a nine-point scale). Furthermore, when we look individually at each bin, we do not find a significant effect within any of bins one ($b = .09$ [-.10, .28], $B = .04$, $t = .98$, $p = .329$, $n = 393$), two ($b = .11$ [-.01, .22], $B = .05$, $t = 1.86$, $p = .064$, $n = 857$), three ($b = .08$ [-.04, .19], $B = .04$, $t = 1.32$, $p = .186$, $n = 741$), or four ($b = -.005$ [-.11, .10], $B = -.004$, $t = -.09$, $p = .927$, $n = 424$).

4.4. Considering two sources of moral ambiguity

As discussed in the main text, the ambiguous condition of our paradigm introduced two sources of ambiguity: uncertainty over whether the allegations were true, and uncertainty over whether the alleged transgressions were bad enough to merit the organizers' severe punitive strategy. However, subjects in the ambiguous condition of our paradigm nonetheless readily punished for reputational gain. Thus, the results from Studies 3-4 suggest that *neither* source of ambiguity gives rise to reservations that are sufficient to prevent people from responding to reputational incentives for punishment.

Interestingly, the design of Study 4 allows us to bolster this conclusion by considering the specific types of reservations that subjects reported. In particular, when analyzing Study 4, we

can bin subjects based on (i) their composite ratings of the moral merits of punishment, across our ten scale items, (ii) their responses to a single scale item about confidence that the allegations are true (item 1 from our 10-item scale), or (iii) their responses to a single item about whether the organizers' punitive approach is proportionate and appropriate (item 9 from our 10-item scale). Here, we show that regardless of which binning approach we use, we find comparable results. Thus, for both sources of ambiguity, we find that subjects with reservations are willing to act on reputational incentives.

In the main text, we bin subjects based on their composite ratings of the moral merits of punishment and find significant positive effects of observability on punishment across all four bins. When we instead bin subjects based on their responses to the scale item about confidence that the allegations are true, we likewise find significant positive observability effects across each of bins one ($b = .06$ [.02, .10], $B = .11$, $t = 2.76$, $p = .006$, $n = 564$), two ($b = .15$ [.09, .20], $B = .17$, $t = 5.28$, $p < .001$, $n = 866$), three ($b = .08$ [.01, .16], $B = .08$, $t = 2.18$, $p = .029$, $n = 662$) and four ($b = .13$ [.03, .24], $B = .14$, $t = 2.45$, $p = .015$, $n = 323$). And finally, when we bin subjects based on their responses to the scale item about the proportionality and appropriateness of the organizers' punitive approach, we again find significant positive observability effects across each of bins one ($b = .08$ [.04, .13], $B = .13$, $t = 3.71$, $p < .001$, $n = 678$), two ($b = .12$ [.06, .18], $B = .14$, $t = 3.85$, $p < .001$, $n = 715$), three ($b = .11$ [.03, .18], $B = .11$, $t = 2.89$, $p = .004$, $n = 663$) and four ($b = .13$ [.02, .23], $B = .13$, $t = 2.43$, $p = .016$, $n = 359$).

5. Reported reputation motives in Studies 3-4

In this section, we discuss the extent to which subjects in Studies 3-4 explicitly reported being driven by reputation motives. Overall, subjects in our Studies 3-4 were quite responsive to our observability manipulations, punishing at higher rates in public than private. But what was the psychological mechanism through which observability influenced punishment behavior? In particular, to what extent were subjects consciously motivated by reputation concerns?

As described in Section 2.3 of this document, at the end of Studies 3-4 we asked subjects a series of questions designed to investigate the extent to which they were consciously motivated by reputation concerns (including as compared to other motivations). We pre-registered some secondary analyses of these questions; however, instead of reporting these analyses, we have chosen (for reasons of brevity) to simply provide some descriptive plots of subjects' responses. Specifically, in Figures S2a-c, we plot responses to these questions for Studies 3a, 3b, and 4, respectively, in order to provide a descriptive impression of subjects' self-reported reputation motives.

The first rows of Figures S2a-c plot responses to the set of four questions presented to subjects in all conditions. In Figures S2a-b, we show results as a function of observability and ambiguity in Studies 3a-b, and in Figure S2c, we show results as a function of observability and subjects' pre-treatment ratings of the moral merits of punishment (binned as in our main text analyses) in Study 4.

These subplots highlight three key findings, which hold broadly across ambiguity conditions and bins. First, ratings of explicit reputation concerns (the two right-most subplots) are very low, both in absolute terms and as compared to ratings of other motivations (the two left-most subplots). In particular, subjects report primarily having made their decisions because they personally felt that it was the right decision, and also being somewhat driven by the desire to see themselves as a good person. Second, despite being low in absolute terms, ratings of reputation concerns *are* somewhat higher in public than private (although this is less true in

Study 3b than Studies 3a or 4), suggesting that subjects in the public conditions may have been somewhat conscious of acting on reputational incentives. Third, ratings of non-reputational motivations are *also* somewhat higher in public than private (although this effect is again weaker in Study 3b), suggesting that subjects in the public conditions may have incorrectly attributed some of their sensitivity to reputational incentives to other factors.

Together, these results suggest that subjects may have been relatively unaware that their punishment decisions were driven by reputation motives. And interestingly, this seems to be true even for subjects in the ambiguous condition, and even for subjects in this condition who reported reservations about the morality of punishment before the introduction of our observability manipulation.

However, it is of course also possible that subjects were consciously strongly motivated by reputation concerns, but did not want to admit to these concerns (e.g., because reputation concerns are self-interested rather than morally-motivated, and thus potentially less socially desirable). To explore this hypothesis, we also compared ratings of reputation motives to ratings of another motivation, the desire to earn money, that we anticipated subjects might also be hesitant to admit to (because, like reputation motives, monetary motives are self-interested rather than morally-motivated).

The second rows of Figures S2a-c plot responses to the four questions presented only to subjects in the public conditions; these questions presented two matched pairs of questions about motivations to (i) earn the most money possible vs. (ii) earn a positive reputation. These subplots again reveal low absolute ratings of explicit reputation motives, and also reveal that these ratings are even lower than ratings of explicit monetary motives. And these results again hold broadly across ambiguity conditions and bins. They thus lend some further support to the proposal that subjects may not have been especially motivated by explicit reputation concerns.

Of course, however, it is possible that subjects were simply less comfortable admitting to explicit reputation motives than explicit monetary motives. Overall, then, we simply interpret these results as providing suggestive evidence that subjects may have been sensitive to our observability manipulations without having been consciously motivated by reputation. Future research should further investigate this interesting possibility.

Figure S2a. Explicit reputation motives (and other motives) in Study 3a.

Figure S2b. Explicit reputation motives (and other motives) in Study 3b.

Figure S2c. Explicit reputation motives (and other motives) in Study 4.

6. Discussion of pre-registered analysis plans

With the exception of our Study 3 pre-tests, all studies in this paper were pre-registered. As described in the main text, for all of these studies, we adhered to our pre-registered sample sizes and exclusion criteria. However, the theoretical focus of our paper differs somewhat from the theoretical focuses of our pre-registrations, such that our paper addresses a somewhat different set of questions than our pre-registrations planned to address. Therefore, some of our pre-registered analyses are not reported in this paper, and some of the analyses reported in this paper are not pre-registered. In this section, we detail, for each study, the ways that our analyses deviate from our pre-registered analysis plans.

6.1. Study 1

Study 1 was pre-registered at <https://aspredicted.org/blind.php?x=d7ji9i>. Here, we list ways that our analyses of Study 1 differed from our pre-registration.

First, our exploratory factor analysis of our three questions relevant to the moral merits of punishment was not pre-registered, and we did not plan to combine these three items into a scale. However, given that our factor analysis produced a one-factor solution, and the three items formed a reliable scale (as reported in the main text), we did combine them to form a composite measure of the moral merits of punishment. Then, instead of conducting our planned manipulation check analyses (investigating the effects of ambiguity on each of our three items), we instead conducted a single analysis of the effect of ambiguity on our composite scale.

Second, in our Study 1 pre-registration, we planned to focus in our primary analyses on the effect of our ambiguity manipulation on the perceived reputation value of punishment. However, because our theoretical focus in this paper is on the perceived reputation value of punishment in our *ambiguous* conditions, we report this analysis only in SI Section 1.4. In our main text, we instead focus on analyzing reputation ratings within our ambiguous condition. To this end, we conduct analyses demonstrating that, in absolute terms, subjects in our ambiguous condition expected punishment to confer reputation *benefits* (by comparing reputation ratings to the scale midpoint). These analyses were pre-registered, albeit as secondary analyses. However, we also report (despite not having planned to do so) the frequency with which subjects in our ambiguous condition *personally* saw punishment as morally questionable (as indexed by appropriateness ratings at or below the scale midpoint) but nonetheless expected punishment to confer reputational benefits (as indexed by above-midpoint reputation ratings).

Second, we planned to conduct secondary analyses of our secondary DVs. These DVs were designed to provide a richer picture of subjects' expectations regarding the reputation consequences of punishment. However, our focus in this paper is on the question of whether subjects in our ambiguous conditions expected punishment to have overall positive global reputational consequences (i.e., to confer reputational *benefits*). For this reason, we do not analyze these DVs in this paper.

6.2. Studies 2a-b

Study 2a was pre-registered at <http://aspredicted.org/blind.php?x=5xp2rm> and Study 2b was pre-registered at <http://aspredicted.org/blind.php?x=zd5qn9>. Here, we list ways that our analyses of Studies 2a-b differed from these pre-registrations.

First, as described above in the context of Study 1, our exploratory factor analysis of our three questions relevant to the moral merits of punishment was not pre-registered, and we did not plan to combine these three items into a scale. Furthermore, in our Study 2 pre-registrations, we planned as a manipulation check analyses to simply compare ratings of offensiveness across ambiguity conditions. However, we instead investigated the effect of ambiguity on our composite measure of the perceived moral merits of punishment (mirroring Study 1).

Second, as in Study 1, we planned to focus in our primary analyses on the effects of our ambiguity manipulation on the perceived reputation value of punishment (including by reporting the main effect of ambiguity, its interaction with ideology, and simple effects of ambiguity at each level of ideology). However, because our theoretical focus in this paper is on the perceived reputation value of punishment in our *ambiguous* conditions, we do not focus much on the effects of ambiguity on reputation ratings (although we do report main effects in SI Section 1.4). Instead, we focus on analyzing reputation ratings within our ambiguous conditions, again comparing reputation ratings within our ambiguous conditions to the scale midpoint; unlike in Study 1, these analyses were not pre-registered in Study 2. Furthermore, as in Study 1, we report (despite not having planned to do so) the frequency with which subjects in our ambiguous

conditions *personally* saw punishment as morally questionable but nonetheless expected punishment to confer reputational benefits.

Third, as in Study 1, we planned to conduct secondary analyses of our secondary DVs, but do not analyze these DVs in this paper.

Fourth, we planned to account for repeated measures across vignettes by using repeated-measures ANOVAs that modeled vignette as a within-subject factor (although we did not plan to report effects of vignette). However, to create a consistent statistical approach across all analyses in this paper, we instead chose to analyze our data using linear regressions that account for repeated measures by clustering standard errors on subject. We note that we pre-registered this analysis plan in Study 1, and that its regression-based approach is congruent with the rest of our paper (which consistently reports analyses from linear regressions rather than ANOVAs). We also note, however, that both the ANOVA and regression-based analysis approaches produce very similar (and qualitatively identical) results.

Fifth, in addition to analyzing each vignette individually, we planned to conduct a set of secondary analyses looking only at the first vignette that each subject evaluated. We find that all of our results remain significant when restricting to the first vignette; however, for reasons of brevity, we do not report this analysis.

Finally, we planned to conduct a secondary analysis repeating our analyses after excluding subjects who failed to correctly answer a simple analogy question. Across all experiments in this paper, rates of incorrect responses to this question were very low (always less than 3%) and our key results are qualitatively unchanged by excluding subjects with incorrect responses; thus, for brevity we do not report this secondary analysis.

6.3. Study 3a

Study 3a was pre-registered at <http://aspredicted.org/blind.php?x=hq3be5>. Here, we list ways that our analyses of Study 3a differed from our pre-registration.

First, neither our exploratory factor analysis (of our ten moral merits of punishment scale items) nor our manipulation check analysis (of the effect of ambiguity on composite ratings of the moral merits of punishment, within the private conditions of Study 3a) were pre-registered.

Second, in our analyses of punishment decisions, we planned to report simple effects of observability, but not simple effects of ambiguity. However, for completeness, we report both types of simple effects.

Third, for reasons of brevity, we report our planned analyses of self-reported continuous support for punishment and the perceived moral merits of punishment only in this SM and not in the main text.

Forth, we planned to report, and do report, a secondary analysis of Study 3a that restricts to subjects who correctly answered our set of comprehension questions. However, we planned to restrict to subjects who answered all four comprehension questions correctly. Instead, we restrict only on the basis of our first three comprehensions. This choice reflects that our fourth question, which had different answer choices (and a different correct answer) in the public vs. private conditions, had a significantly lower pass rate in public than private; restricting on the basis of this question thus introduces a selection effect that can undermine random assignment. We also note that for reasons of brevity, we only extended our analyses of comprehenders to our punishment decision DV (but not our continuous support for punishment or perceived moral merits of punishment DVs).

Fifth, in our pre-registrations, we planned to conduct a few secondary analyses that, for brevity, we do not report results for. Specifically, we planned to investigate potential moderation of our effects by political ideology or CRT scores, and to test whether our observability effects are robust to controlling for reported explicit reputation motives, and conduct various statistical analyses of reported explicit reputation motives. However, with respect to this last set of analyses, we note that Figures S2a-c plot some relevant data.

Finally, as in Study 2, we planned to repeat our analyses after excluding subjects who failed to correctly answer a simple analogy question, but do not report these analyses.

6.4. Study 3b

Study 3b was pre-registered at <https://aspredicted.org/blind.php?x=8a38p7>. Our analyses of Study 3b differed from our pre-registration in the same ways that we describe above for Study 3a, with three exceptions.

First, with respect to our analyses of self-reported continuous support for punishment and the perceived moral merits of punishment, our Study 3b pre-registration planned to report the simple effects of observability within each ambiguity condition. However, for reasons of brevity, we do not report these simple effects (although we do follow our plan to report the main effects of observability and ambiguity, as well as the interaction between these factors).

Second, Study 3b did not pre-register secondary analyses of comprehenders. However, for consistency with Study 3a, we nonetheless report such analyses.

Third, as reported in the main text, Study 3b included a third set of public vs. private conditions, in which subjects were assigned to the ambiguous condition but the Decider was described as less ideological (i.e., the Decider was described simply as “a Republican”). These conditions were designed to address a theoretical question outside the scope of this paper, so we do not describe them further or analyze their results. However, our pre-registrations planned analyses of these conditions.

6.5. Study 4

Study 4 was pre-registered at: <https://aspredicted.org/blind.php?x=e3bc2w>. Here, we list ways that our analyses of Study 4 differed from our pre-registration.

First, neither our exploratory factor analysis (of our six initial perceived moral merits of punishment scale items) nor our analyses of Study 4 that consider two sources of ambiguity (reported in SM Section 4.3) were pre-registered.

Second, like in Study 3a, our secondary analyses of comprehenders restrict only on the basis of our first three comprehension questions (because pass rates for our fourth question differed significantly across our public vs. private conditions).

Third, like in Studies 3a-b, we do not report planned secondary analyses involving reported explicit reputation motives, or exclusions of subjects who failed to an analogy question.

7. Full vignette texts for Study 1

Here, we report full texts for each of our Study 1 vignettes, across experimental conditions.

7.1. Racist comment (Vignette 1)

Rita works at a medium-sized organization. Two of her coworkers are named **Emily** and **Ashanti**. Whereas Rita and Emily are White, Ashanti is Black.

Recently, Ashanti accused Emily of racism towards her. Specifically, Ashanti claimed that she overheard Emily say that she was sure that Ashanti was only hired because of affirmative action, and that she didn't like having "her type" around the office.

Ambiguous condition: Emily responded to this allegation by denying it. She insisted that she is not racist and would never speak that way about Ashanti. She also said that she wasn't sure why Ashanti would say something like this, but maybe it was to make her look bad because they were both vying for the same upcoming promotion.

Unambiguous condition: Emily heard about this allegation against her, but didn't say anything to dispute it.

Rita learned about Ashanti's allegation against Emily and Emily's reaction to it. The next day, Rita approached Emily about a picnic she was organizing for colleagues, and had previously invited Emily to. Rita said to Emily: "What you said about Ashanti was really racist and offensive—I think it's best if you don't come to the picnic".

Today, we'd like you to answer a few questions about another person named **Veronica**. Veronica is a Democrat, and also works with Rita, Emily, and Ashanti. Veronica learned about the allegation against Emily and Emily's reaction to it, and also overheard when Rita told Emily not to attend the picnic.

7.2. Homophobia (Vignette 2)

Sam works at a medium-sized organization. One of his coworkers is named **Brett**. Sam and Brett are both straight.

Recently, Brett was accused of homophobia. A gay employee at the organization named Alex claimed that he overheard Brett speak in a homophobic way about a gay political candidate.

Specifically, Alex claimed that he overheard Brett and a colleague discussing an upcoming political election while watching TV coverage about the field of candidates. One of the candidates is openly gay, and the TV coverage featured a conversation in which this candidate and his husband were laughing and joking around. Alex reported that during this TV coverage, Brett—who had previously expressed that he found this particular candidate really annoying—remarked "ugh, not these fucking gays".

Ambiguous condition: Brett responded to this allegation by denying it. He insisted that he is not homophobic and did not reference the candidate's sexual orientation, but instead simply said "ugh, not these fucking guys".

Unambiguous condition: Brett heard about this allegation against him but did not respond to it.

Sam learned about the allegation against Brett and Brett's reaction to it. The next day, Sam approached Brett about a barbecue he was organizing for colleagues, and had previously invited Brett to. Sam said to Brett: "The way you talked about that politician was really homophobic and offensive—I think it's best if you don't come to the barbecue".

Today, we'd like you to answer a few questions about another person named **Anthony**. Anthony is a Democrat, and also works with Sam and Brett. Anthony learned about the allegation against Brett and Brett's reaction to it, and also overheard when Sam told Brett not to attend the barbecue.

7.3. Sexism (Vignette 3)

David works at a medium-sized organization. One of his coworkers is named **Josh**.

Recently, Josh was accused of sexism. An employee at the organization named Monica claimed that she overheard Josh and a colleague discussing a presentation she had just given, and that Josh spoke about her in a sexist way.

Specifically, Monica claimed that she overheard Josh say that he hadn't been paying any attention to her presentation because her skirt was "slutty and distracting" and that she is "kind of a bitch".

Ambiguous condition: Josh responded to this allegation by denying it. He insisted that he is not sexist and would never speak that way about Monica. He also said that he wasn't sure why Monica would make something like this up, but maybe it was related to a dispute that they recently had over a work project.

Unambiguous condition: Josh heard about this allegation against him, but didn't do anything to address it.

David learned about the allegation against Josh and Josh's reaction to it. The next day, David approached Josh about a game night he was organizing for colleagues, and had previously invited Josh to. David said to Josh: "What you said about Monica was really sexist and offensive—I think it's best if you don't come to the game night".

Today, we'd like you to answer a few questions about another person named **Ken**. Ken is a Democrat, and also works with David and Josh. Ken learned about the allegation against Josh and Josh's reaction to it, and also overheard when David told Josh not to attend the game night.

7.4. Racist costume (Vignette 4)

Please carefully read the following story. Then, we will ask you to answer some questions about it.

Sarah works at a medium-sized organization. One of her coworkers is named **Frank**. Recently, Frank was accused of racism. A Black employee at the organization named Damian claimed that he saw Frank at a Halloween party at a club, and Frank was wearing a racist costume.

Specifically, Damian claimed that he saw Frank (who is White) dressed up as Stevie Wonder (who is Black)—and that Frank had painted his face black, used makeup to thicken his lips, and worn a fake black nose with exaggerated flared nostrils.

Ambiguous condition: Frank responded to this allegation by denying it. He insisted that he is not racist and wasn't even out at a club on the night in question, and that Damian must have mistaken him for somebody else.

Unambiguous condition: Frank heard about this allegation against him but didn't say anything about it.

Sarah learned about the allegation against Frank and Frank's reaction to it. The next day, Sarah approached Frank about a happy hour she was organizing for colleagues, and had previously invited Frank to. Sarah said to Frank: "The Halloween costume you wore was really racially insensitive and offensive—I think it's best if you don't come to the happy hour".

Today, we'd like you to answer a few questions about another person named **Michael**. Michael is a Democrat, and also works with Sarah and Frank. Michael learned about the allegation against Frank and Frank's reaction to it, and also overheard when Sarah told Frank not to attend the happy hour.

8. Full vignette texts for Study 2

Here, we report full texts for each of our Study 2 vignettes, across experimental conditions.

8.1. Racist comment (Democrat Vignette 1)

Last week, **Rita** was having coffee with her coworker, **Emily**. While chatting, they began to discuss another employee, **Ashanti**, who was recently hired. Whereas Rita and Emily are White, Ashanti is Black. Rita mentioned to Emily that she recently had a fun conversation with Ashanti and thought that she seemed cool.

Ambiguous condition: Emily agreed, and added that Ashanti also seemed really smart, so she definitely didn't think she was hired because of affirmative action or anything like that.

Unambiguous condition: Emily said that she hadn't talked to Ashanti yet, but that she didn't like having "her type" around the office, and was sure that she was only hired because of affirmative action.

Rita responded by calling Emily out for this comment and its implications, saying, “that’s really racist and offensive”.

Today, we’d like you to answer a few questions about another person named **Veronica**.

More ideological condition: Veronica is a strong Democrat who strongly supports the Black Lives Matter movement.

Less ideological condition: Veronica is a Democrat.

Veronica also works with Rita, Emily, and Ashanti. While Rita and Emily were having their coffee conversation, Veronica (who, out of view, was standing in line for a coffee) happened to overhear them. So Veronica ended up hearing what Emily said about Ashanti, and hearing Rita call Emily out.

8.2. Homophobia (Democrat Vignette 2)

Last week, **Sam** was having coffee with his coworker, **Brett**. While chatting, they began to discuss an upcoming political election and the field of candidates, including a specific candidate who is openly gay. Sam and Brett are both straight.

Ambiguous condition: During their conversation, Brett said that he was honestly very impressed by the field of candidates and could see himself voting for any of them. Then, he made a joke about the candidate being gay, stating that he was really starting to “get behind him” and then saying “that’s what he said!”.

Unambiguous condition: During their conversation, Brett made a joke about the candidate being gay, stating that he was really starting to “get behind him” and then saying “that’s what he said!”. Then, after laughing at his own joke, Brett said “no, but seriously, it’s pretty wild that a fag is running”.

Sam responded by calling Brett out, saying “that’s really homophobic and offensive”.

Today, we’d like you to answer a few questions about another person named **Anthony**.

More ideological condition: Anthony is a strong Democrat who strongly supports LGBT rights.

Less ideological condition: Anthony is a Democrat.

Anthony also works with Sam and Brett. While Sam and Brett were having their coffee conversation, Anthony (who, out of view, was standing in line for a coffee) happened to overhear them. So Anthony ended up hearing what Brett said about the candidate, and hearing Sam call Brett out.

8.3. Sexism (Democrat Vignette 3)

Last week, **David** was having coffee with his coworker, **Josh**. While chatting, they began to discuss a female coworker of theirs, **Monica**, who recently gave a presentation that they attended.

Ambiguous condition: After David asked Josh a question about the content of Monica's presentation, Josh commented that he felt bad but that he had honestly been a bit distracted by her short skirt, which he said was "kind of an interesting choice for the office" but that he "wasn't complaining about it".

Unambiguous condition: After David asked Josh a question about the content of Monica's presentation, Josh commented that he hadn't been paying any attention to Monica because her "slutty short skirt" was too distracting, and that she "is generally kind of a bitch anyways".

David responded by calling Josh out, saying "that's really objectifying and sexist".

Today, we'd like you to answer a few questions about another person named **Ken**.

More ideological condition: Ken is a strong Democrat and considers himself a strong feminist ally.

Less ideological condition: Ken is a Democrat.

Ken also works with David, Josh, and Monica. While David and Josh were having their coffee conversation, Ken (who, out of view, was standing in line for a coffee) happened to overhear them. So Ken ended up hearing what Josh said about Monica, and hearing David call Josh out.

8.4. Racist costume (Democrat Vignette 4)

Shortly after Halloween, **Frank** was having coffee with his coworker, **Sarah**. Frank and Sarah had both attended Halloween parties over the weekend, so they were discussing their costumes and swapping pictures. Sarah told Frank about her costume first, which was a bumble bee.

Ambiguous condition: Next, Frank (who is White) explained to Sarah that he had dressed up as Stevie Wonder (who is Black). He also showed her a picture of his costume, in which he wore a wig with dreadlocks to match Stevie's hairstyle.

Unambiguous condition: Next, Frank (who is White) explained to Sarah that he had dressed up as Stevie Wonder (who is Black). He also showed her a picture of his costume. In the costume, he painted his face black, and thickened his lips and painted them bright red using makeup. He also wore a fake black nose with exaggerated flared nostrils, and a gold teeth grill (even though Stevie Wonder doesn't wear grills).

Sarah responded by calling Frank out, saying "that costume is really racially insensitive and offensive".

Today, we'd like you to answer a few questions about another person named **Michael**.

More ideological condition: Michael is a strong Democrat who strongly supports the Black Lives Matter movement.

Less ideological condition: Michael is a Democrat.

Michael also works with Frank and Sarah. While Frank and Sarah were having their coffee conversation, Michael (who, out of view, was standing in line for a coffee) happened to overhear them. So Michael ended up hearing about Frank's costume, and hearing Sarah call Frank out.

8.5. Religion (Republican Vignette 1)

Last weekend, **David** attended a BBQ function organized by his company for employees and their families. At the event, David had a conversation with his co-worker **Josh**. While David and Josh were chatting, they were approached by another co-worker's seven-year-old daughter. She had wandered over to them while her parents were getting more food.

David and Josh began talking to this seven-year-old girl. During their conversation, she mentioned that she and her parents had attended church that morning, and then asked David and Josh if they believe in God.

Ambiguous condition: Josh responded by explaining that he has some doubts about whether God exists, and then telling the girl that when she got older, "she would have to think carefully about it for herself."

Unambiguous condition: Josh responded by telling the girl that "everything they tell you in church is a total lie" and that "people who believe in God are really stupid".

After Josh made these comments, David called him out, saying, "making negative comments about religion to somebody else's kid is really disrespectful and inappropriate".

Today, we'd like you to answer a few questions about another person named **Nate**.

More ideological condition: Nate is a strong Republican and a religious Christian.

Less ideological condition: Nate is a Republican.

Nate also works with David and Josh. At the BBQ, he happened to be seated behind David and Josh (out of their view). So Nate ended up overhearing the exchange between the seven-year-old girl and David and Josh, and hearing David call Josh out.

8.6. Veteran's day (Republican Vignette 2)

Last November, Veterans day fell on a Sunday. That Sunday, **Jason** attended church and sat next to **Steve**, another member of his congregation. The pastor explained that they would start the day's service with a short Veterans day ceremony. The ceremony began with the national anthem.

Ambiguous condition: During the anthem, Jason noticed that Steve kept whispering to his wife instead of paying attention.

Unambiguous condition: During the anthem, Steve kept loudly cracking jokes and laughing with a few of his friends sitting behind him. Then, later in the ceremony, Steve got a phone call, and instead of silencing it, got up to answer it (loudly saying "hang on a minute bro"), and then caused a major disruption as he squeezed out past others in the pew. A few minutes later, he squeezed back past everyone into his seat, again disrupting them while they paid their respects during the ceremony.

After the service ended, Jason called Steve out for his behavior, saying, "that was really disrespectful and offensive".

Today, we'd like you to answer a few questions about another person named **Ryan**.

More ideological condition: Ryan is a strong Republican who cares strongly about veterans' rights.

Less ideological condition: Ryan is a Republican.

Ryan also attends the same congregation as Jason and Steve. When Jason called Steve out for his behavior, Ryan (who, out of their view, happened to be walking behind them) overheard their conversation. So Ryan ended up hearing Jason describe Steve's behavior and call Steve out.

8.7. September 11 (Republican Vignette 3)

Last week, **Jackie** and her co-worker **Emma** were having lunch in their company's break room.

Ambiguous condition: During their lunch, Emma mentioned that she had recently overheard a joke at a coffee shop about the 9/11 attacks.

Emma said that she knew it wasn't great to joke about the attacks, but that she honestly found herself chuckling a little bit, and then told Jackie the joke. "What's Al Qaida's favorite football team?" she asked, and then answered "The New York Jets!"

Unambiguous condition: During their lunch, the company PA made an announcement, asking for a company-wide moment of silence to remember the victims of September 11th, 2001.

After the moment of silence ended, Emma told Jackie a joke about the 9/11 attacks. “What is the New York City Fire Department's favorite song?” she asked, and then answered “It’s Raining Men!” Afterwards, she laughed at her own joke and then said, “but seriously, kind of absurd that they’re imposing a moment of silence on everyone nearly two decades after the attacks”.

In response, Jackie called Emma out, saying, “that’s really disrespectful and offensive”.

Today, we’d like you to answer a few questions about another person named **Samantha**.

More ideological condition: Samantha is a strong Republican and deeply patriotic.

Less ideological condition: Samantha is a Republican.

Samantha also works with Jackie and Emma. While Jackie and Emma were having their lunch conversation, Samantha (who, out of view, was also in the break room) happened to overhear them. So Samantha ended up hearing what Emma said, and hearing Jackie call Emma out.

8.8. Flag (Republican Vignette 4)

Last weekend, **Rachel** and her co-worker **Becky** both attended a sporting event as part of a work function. Since their company purchased a large block of seats, each employee received an American Flag as a “thank you” as they exited the stadium.

Ambiguous condition: Upon receiving her flag, Becky turned to Rachel and said, “I’m not much of a flag person” and then dismissively handed the flag back to the person who gave it to her.

Unambiguous condition: Upon receiving her flag, Becky turned to Rachel and said, “I’m not much of a flag person” and then dismissively threw it on the floor. Then, a moment later, she picked it up and said “Ha, actually this is perfect because I just threw out the old rag I keep in my car for when I need something to wipe up a mess with.”

Afterwards, Rachel called Becky out, saying, “disrespecting the flag like that is really offensive”.

Today, we’d like you to answer a few questions about another person named **Caroline**.

More ideological condition: Caroline is a strong Republican who is deeply patriotic.

Less ideological condition: Caroline is a Republican.

Caroline also works with Rachel and Becky. She also attended the sporting event, and was walking behind Rachel and Becky (out of their view) as they were leaving. So Caroline ended up observing Rachel and Becky’s exchange over the flag, and hearing Rachel call Becky out.

9. Full news article texts for Studies 3-4

Here, we report full texts for each of our news articles from Studies 3-4, across ambiguity conditions.

9.1. Articles for Democrats (Study 3a and Study 4)

Ambiguous condition

Biologist Joseph Pringle Accused of Sexual Harassment

Joseph Pringle, the celebrated Columbia University biologist and Goldberg Prize winner known for his foundational research on immune functioning, established a university-wide graduate scholarship in September. The scholarship was designed to “recruit and retain the most talented women applicants.”

But in interviews, two women who worked for Dr. Pringle at Columbia (as his lab manager and graduate student, respectively) have described encounters when the biologist, now 63, treated them inappropriately.

In 2016, during her first week as Dr. Pringle’s lab manager, Jessica Sampora, now 26, said that the biologist invited her to his apartment to celebrate her new job. When she arrived, she said, he offered her a glass of wine, removed his pants and then asked her to undress. She declined, left the apartment and said nothing because, she said, she was worried about holding her job.

However, questions have been raised about the motivations behind Ms. Sampora’s allegation. In 2017, Columbia chose not to renew Ms. Sampora’s contract as a lab manager after she made several costly errors in the laboratory—a decision that Dr. Pringle, along with several other biology faculty members, supported. Before leaving her former department, Ms. Sampora was vocal about her disagreement with this decision.

Gabrielle Jordan, 27, also alleged that Dr. Pringle crossed a line. Ms. Jordan said she had been warned about Dr. Pringle’s advances after joining his laboratory as a graduate student in 2015. “A colleague said, ‘He will open huge doors for you professionally, but beware that he likes young women,’” she said. Then, while talking with her at the department holiday party two months later, she said, Dr. Pringle rubbed his hand down her lower back and onto her bottom, Ms. Jordan recalled. Like Ms. Sampora, Ms. Jordan said that she was too afraid to report the incident to anyone. She also declined to name the colleague who warned her about Dr. Pringle.

However, an anonymous source reported to the Columbia Daily Spectator, the Columbia University newspaper, that Gabrielle Jordan “had a reputation” in the biology department for “flirting with everyone and creating constant personal drama,” but declined to elaborate further.

Many former students and collaborators of Dr. Pringle have also issued statements expressing their surprise at the allegations, and describing their positive relationships with Dr. Pringle. “Dr. Pringle was my PhD advisor for six years, starting in 2010, and was an extremely dedicated mentor and advocate of my career. He has always treated me with respect. I cannot recall a time

that he ever made me feel uncomfortable in any way, or crossed any professional boundaries. And I have never seen any evidence of him treating other women this way. I am honestly completely shocked by these allegations,” wrote Dr. Jeanette Frank, a former student of Dr. Pringle’s who is now a professor at University of California, Berkeley.

For his part, Dr. Pringle has denied the allegations. “I had purely professional working relationships with Ms. Sampora and Ms. Jordan. The incidents they have described absolutely did not happen, and have not been corroborated by any evidence. Moreover, I was not even in attendance during the 2015 department holiday party. I am cooperating with Columbia University and enthusiastically invite a full investigation of my behavior.”

Unambiguous condition

Biologist Joseph Pringle Accused of Sexual Harassment

Joseph Pringle, the celebrated Columbia University biologist and Goldberg Prize winner known for his foundational research on immune functioning, established a university-wide graduate scholarship in September. The scholarship was designed to “recruit and retain the most talented women applicants.”

But many women who have worked for Dr. Pringle at Columbia as lab managers or graduate students have described encounters when the biologist, now 63, treated them inappropriately.

Four have described incidents in which they were invited to Dr. Pringle’s private apartment, where he exposed himself or asked them to undress, according to interviews. A fifth woman said in an interview that Dr. Pringle grabbed her underwear through her dress at a department holiday party, and a sixth described an incident with Dr. Pringle in which she said she had to flee his home after he forcefully pulled her onto a bed.

In 2016, during her first week as Dr. Pringle’s lab manager, Jessica Sampora, now 26, said that the biologist invited her to his apartment to celebrate her new job. When she arrived, she said, he offered her a glass of wine, removed his pants and then asked her to undress. She declined, left the apartment and said nothing because, she said, she was worried about holding her job.

Gabrielle Jordan, 27, said she had been warned about Dr. Pringle’s advances after joining his laboratory as a graduate student in 2015. “A couple of people had said, ‘He will open huge doors for you professionally, but beware that he loves attractive women,’” she said. Then, while talking with her at the department holiday party two months later, she said, Dr. Pringle moved his hand from the small of her back down to her behind, where he played with her thong underwear through her dress. “He started to roll my underwear around in his fingers,” Ms. Jordan recalled. Like Ms. Sampora, Ms. Jordan said that she was too afraid to report the incident to anyone.

Sara Clint described her experience with Dr. Pringle in 2011, when she was a visiting student at Columbia. Dr. Pringle invited Ms. Clint to a dinner party “for his laboratory” at his apartment, she said, but she turned out to be the only guest. After dessert, Ms. Clint said that Dr. Pringle forcefully tried to kiss her, and she went to leave. “He grabbed me from the back with both of his

arms and started pulling me backward,” Ms. Clint said. “I twisted and pulled away from him, and he grabbed one of my arms and started dragging me down the hallway toward the bedroom.”

“He pushes me on the bed and lays down on top of me while I’m twisting and pushing him away and saying, ‘No, no, no,’” she continued. “I was pretty aggressive about telling him no, but he wasn’t listening.” She said she finally broke free, ran to her car and locked all the doors. “Then he was right at the window: ‘Come on, come back in,’ he said. ‘I got the car started, got to the bottom of the long driveway, stopped and just sat there in my car crying and shaking.’”

She initially didn’t report the incident, she said, because she felt afraid. In 2014, however, she told her boyfriend Steve Honnald on a cross-country road trip they took together. “We put 7,000 miles on a rental, and she told me her story about Dr. Pringle,” Mr. Honnald confirmed in an interview. “I was horrified.”

9.2. Articles for Republicans (Study 3b)

Ambiguous condition

Feminist Princeton Dean accused of Anti-Male Discrimination following Discipline of Basketball Captain over Sexual Assault Accusation

Elizabeth Cartland, the Dean of Student Affairs at Princeton University, is being accused of gender discrimination following the discipline of former basketball captain Tyler Jones. In 2018, Cartland, 53, initiated disciplinary action against Jones, 22, reporting that he had sexually assaulted another Princeton student. The ultimate result was that Jones, a star power forward and Ivy League student-athlete, was found responsible for sexual assault and forced to take a one-year leave of absence from Princeton—an outcome that he says was unfair, but Cartland claims was justified.

A source at Princeton, who requested to remain anonymous, reported that in her capacity as Dean, Cartland organized an event for students to discuss the #MeToo movement, and its implications for students and culture at Princeton. After the event, a sophomore student, whose name has not been released to the public and who is identified as Jane Doe in some internal Princeton documents, spoke with Cartland in private. She told Cartland about an October 2017 sexual encounter that she had with Jones, which she felt uncomfortable about.

Cartland asked Doe if she had consented to the encounter, and Doe indicated that Jones was very pushy, and she had used nonverbal cues to indicate her discomfort and tried to physically move away several times. However, she also acknowledged that she “was not that physically forceful” and couldn’t remember whether she verbally asked to stop their encounter.

When asked, independent experts recently consulted to comment on the case had split opinions on whether the encounter that Doe described met Princeton’s definition of sexual assault. Moreover, Doe did not ask Cartland to file a complaint against Jones. However, the Princeton source said that Cartland discovered that a similar—albeit slightly more severe—report had been made about Jones over two years ago (but the student filing the report chose not to take action).

Cartland ultimately decided to file a complaint stating that Jones had sexually assaulted Doe—and the complaint resulted in Jones being required to take a one-year leave from school. A few weeks after the complaint was filed, Doe asked Cartland to drop the complaint, but Cartland encouraged her to proceed with it, explaining that Jones was a threat to campus safety.

Jones has accused Cartland of discriminating against men by pursuing weak sexual assault complaints against them. Cartland describes herself as a feminist on her personal Twitter account, including via a positive review for the book “A History of Women in America: From Founding Mothers to Feminists, How Women Shaped the Life and Culture of America”.

“Amid increased societal pressure to focus on issues of sexual misconduct, Ms. Cartland and Princeton University have chosen to prioritize an agenda over the facts. Despite my innocence, they targeted me in order to demonstrate that the university is tough on men who ‘victimize’ female students,” Jones said in a statement to the media.

Cartland has defended herself, arguing that women tend to question and downplay the culpability of men who ignore their discomfort, citing research documenting that repeat offenders—such as, in her view, people with records like Jones’—commit most of the sexual assaults on college campuses, and noting that the university teaches all incoming students about the importance of “affirmative consent”. Despite Jones’ allegations, Cartland has not been disciplined by Princeton and continues to serve as the Dean of Students.

Unambiguous condition

Feminist Princeton Dean accused of Anti-Male Discrimination following Expulsion of Basketball Captain over Sexual Assault Accusation

Elizabeth Cartland, the Dean of Student Affairs at Princeton University, is being accused of gender discrimination following the expulsion of former basketball captain Tyler Jones. In 2018, Cartland, 53, allegedly coerced a sophomore student into discussing a consensual sexual encounter she had with Jones, and then repeatedly lied about key details of her story in order to get Jones expelled for sexual assault.

As a result, Jones—a star power forward and academic award-winning student-athlete—was forced out of the university during the spring semester of his senior year. According to several sources, Jones did nothing wrong, but fell victim to Cartland’s radical feminist agenda and blatant bias against men. Cartland describes herself as a “strong feminist” on her personal Twitter and Facebook accounts, where she recently posted a positive review for the book “How to Date Men When You Hate Men”.

In January of 2018, in her capacity as Dean, Cartland organized an event for students in support of the #MeToo movement. Throughout the event, Cartland repeatedly invited students to talk to her about “any sexual experiences they had at Princeton that had made them uncomfortable”. A sophomore, whose name has not been released and who is identified as Jane Doe in internal Princeton documents, chose to speak with Cartland.

Doe was interviewed for this story. She says that she told Cartland about an October 2017 sexual encounter that she had with Jones, which she had later come to regret. Doe explained to Cartland that she had consented to the encounter—but later, after talking with a friend, came to feel that Jones had pushed things along too quickly. Doe says that she stressed that she did not blame Jones, because she had done nothing to communicate that she wanted to slow down. However, Cartland told Doe that Jones crossed a line with his “piggish” male behavior, and that his “toxic masculinity” was “rampant among jocks” and “promoting rape culture”. Cartland then pressured Doe to file a sexual assault complaint against Jones.

“Dean Cartland told me that even if Tyler’s behavior may not have technically been assault, charging men with assault whenever there’s even a hint of an issue is necessary to enact change. When I told her that I wasn’t comfortable misrepresenting the facts, she told me that Tyler had previously been accused of assaulting others. I only later learned that this was a lie, and that he had no previous record. Dean Cartland told me that if I didn’t file a complaint, I would be responsible when Tyler hurt others in the future. I was scared, so I agreed”, Doe said. According to Doe, Cartland then took formal notes on their conversation that completely misrepresented Doe’s story. Cartland falsely stated that Doe had reported Jones “raping and abusing” her, and that Doe had been crying throughout their sexual encounter.

According to a whistleblower at the Princeton Office of Sexual Assault Prevention (OSAP), whose name has been kept private for this story, Cartland used her authority to pressure OSAP staff against investigating Jones’ side of the story. When Doe also contacted OSAP staff directly to tell them that she felt uncomfortable with Cartland’s complaint and that she was not raped, Cartland deleted all records of Doe’s conversations with OSAP staff. Jones was ultimately found responsible for “penetration without consent” and expelled.

“I handled over a hundred cases before Cartland was hired, and I have never seen investigations hampered like this. Cartland has thrown justice out the window to pursue her radical anti-male agenda. Look, as a woman, I care deeply about the safety of female students. But what Cartland is doing just isn’t right. She’s ruining people’s lives and it makes me sick to watch it happen. These are complete witch-hunts and it’s honestly shocking that she’s gotten away with them. I know that many others in my office feel exactly the same way,” the whistleblower said.

In a statement to the media, Jones spoke out about the harm he has suffered. “Dean Cartland has chosen to ignore the facts in order to punish me simply because I am male, and a prominent figure in campus athletics. As a result, I’ve been forced out of my community, lost out on the degree I worked so hard for, and had my reputation irreparably damaged,” he said.

Members of Jones’ community have also attested to his moral character. Jones is a regular church-goer, and volunteers every weekend as a “little helper” in a retirement home. He regularly helps Dorothy Brown, 86, picking up groceries for her and keeping her company. “Tyler is a true light in my life and the lives of many others. I don’t know what I’d do without him. God bless him, and protect him from evildoers who seek him harm”, Brown said in an interview.

Despite the allegations against her, Cartland has not been disciplined by Princeton and continues to serve as the Dean of Students. According to university records, since Cartland began as Dean the number of male students expelled over sexual assault charges has quadrupled, and is now much higher than at peer schools like Harvard and Yale.