

LBS Research Online

[D Efron](#)

The Moral Repetition Effect : Bad Deeds Seem Less Unethical When Repeatedly Encountered
Article

This version is available in the LBS Research Online repository: <https://lbsresearch.london.edu/id/eprint/2332/>

[Efron, D](#)

(2022)

The Moral Repetition Effect : Bad Deeds Seem Less Unethical When Repeatedly Encountered.

Journal of Experimental Psychology: General, 151 (10). pp. 2562-2585. ISSN 0096-3445

DOI: <https://doi.org/10.1037/xge0001214>

American Psychological Association

<https://psycnet.apa.org/fulltext/2022-54629-001.ht...>

Users may download and/or print one copy of any article(s) in LBS Research Online for purposes of research and/or private study. Further distribution of the material, or use for any commercial gain, is not permitted.

**The Moral Repetition Effect:
Bad Deeds Seem Less Unethical When Repeatedly Encountered**

Daniel A. Effron

London Business School

Author Note

After Experiment 1a, all experiments were pre-registered. Pre-registration documents, verbatim materials, and data are available at

https://osf.io/6xpq8/?view_only=eb17d8e867aa4fe7b5007baa17086859

I have no conflicts of interest to disclose.

I received funding from the Center for the Science of Moral Understanding at University of North Carolina Chapel Hill.

I thank Justine Murray, Stephanie Rodriguez, and the LBS Research Lab for research assistance.

Correspondence concerning this article should be addressed to the author, Organisational Behaviour Subject Area, London Business School, Regent's Park, London, United Kingdom, NW1 4SA. Email: deffron@london.edu

Some of the ideas and data in this article were presented at the annual meeting of the Society for Personality and Social Psychology in February, 2021, and the annual meeting of the Society for Judgment and Decision Making in December, 2020.

Abstract

Reports of moral transgressions can “go viral” through gossip, continuous news coverage, and social media. When they do, the same person is likely to hear about the same transgression multiple times. The present research demonstrates that people will judge the same transgression less severely after repeatedly encountering an identical description of it. I present seven experiments (six of which were pre-registered; 73,265 observations from 3,301 online participants and urban residents holding 55 nationalities). Participants rated fake-news sharing, real and hypothetical business transgressions, violations of fundamental “moral foundations,” and various everyday wrongdoings as less unethical and less deserving of punishment if they had been shown descriptions of these behaviors previously. Results suggest that affect plays an important role in this *moral repetition effect*. Repeated exposure to a description of a transgression reduced the negative affect that the transgression elicited, and less-negative affect meant less-harsh moral judgments. Moreover, instructing participants to base their moral judgments on reason, rather than emotion, eliminated the moral repetition effect. An alternative explanation based on perceptions of social norms received only mixed support. The results extend understanding of when and how repetition influences judgment, and they reveal a new way in which moral judgments are biased by reliance on affect. The more people that hear about a transgression, the wider moral outrage will spread; but the more times an individual hears about it, the less outraged that person may be.

Keywords: moral judgment, ethics, repetition, familiarity, desensitization, affect, mere exposure

Abstract word count: 233

The Moral Repetition Effect:

Bad Deeds Seem Less Unethical When Repeatedly Encountered

Dr. David Dao did not want to give up his seat on oversold United Flight 3411 from Chicago. He had already boarded the flight, and his patients needed his help in Louisville. But United needed the seat for its staff, and when Dao refused to leave the flight, security officers forcibly removed him, banging his head against a seat as they did (Goldstein, 2017). Videos of Dao, bleeding and unresponsive, being dragged from the plane, went viral on social media. In a single day, social media users shared one such video 87,000 times and viewed it 6.8 million times (Marotti & Zumbach, 2017).

Reports of moral transgressions can spread fast and far through gossip, continuous news coverage, and social media posts (Brady et al., 2017; Crockett, 2017; Fernandes et al., 2017). The wider these reports spread, the more people see it, and the more individuals there are who feel outrage. At the same time, the wider these reports spread, the more likely the same individual is to encounter the same description of the same transgression multiple times. With coverage of Dao saturating traditional and social media, a single person is likely to have come across a headline and photo depicting the Dao incident more than once. The present research examines whether people judge the morality of a transgression differently if they have seen a description of it before. Widespread coverage of a moral transgression will clearly stoke collective outrage – but how will repeated exposure to the same coverage of the same transgression affect an individual’s moral judgments?

The question of how repetition affects judgments has long interested social and cognitive psychologists (e.g., Allport & Lepkin, 1945; Brauer et al., 1995; Hasher et al., 1977; Schacter, 1987; Zajonc, 1968), and past research offers several reasons to predict that people would judge

a moral transgression more harshly after previously encountering a description of it. Repetition could make people think that the transgressor is more infamous (Jacoby, Woloshyn, et al., 1989), that concern about the transgression is more widespread (Weaver et al., 2007), or that the description is more truthful (Dechêne et al., 2010). Repetition could also intensify a person's initial negative reaction to a transgression by making the transgression more salient (see Mrkva & Van Boven, 2020).

However, I propose the opposite – that people will judge a transgression more *leniently* after previously encountering a description of it. By *transgression*, I mean any behavior that violates a moral norm or value (see Turiel, 1983). I hypothesize that the same transgression will seem less unethical and less deserving of punishment after prior exposure to a description of it. I call this phenomenon the *moral repetition effect*.

This hypothesis is based on the theory that feelings – not just reason – influence moral judgments (for a review, see Avramova & Inbar, 2013; Greene & Haidt, 2002; Haidt, 2001). When people learn about a transgression, they feel upset – in other words, they experience negative affect. And the more intensely they feel the negative affect, the more morally wrong they judge the transgression to be (Haidt, 2001; Haidt & Kesebir, 2010). However, people are not always aware of why they feel a certain way. They might think they feel upset about the transgression when they actually feel upset for an unrelated reason – a case of *affect misattribution* (see Clore et al., 2001; Schwarz, 2011). Thus, manipulating how people are feeling when they judge a transgression can influence how morally wrong they think the transgression is – even when the source of the feeling is completely unrelated to the transgression (Gawronski et al., 2018; Gummerum et al., 2016; Strohminger et al., 2011). For

example, watching a comedic video made people more likely to think it was appropriate to murder someone in order to save others' lives (Valdesolo & DeSteno, 2006).

Repetition may influence moral judgments in a similar way: by changing how people feel when they judge a transgression. The first time you see a headline about David Dao being dragged off the flight, you might feel intense *moral outrage* – a type of negative affect involving anger in response to a moral violation (Batson et al., 2007).¹ The second or third time you see the same headline, you might feel somewhat less outraged. More generally, people are likely to feel more intense negative affect when they first encounter a description of a moral transgression than when they subsequently encounter the same description. And because affect informs moral judgments, repeated exposure to the transgression would make it seem less unethical. In short, I hypothesized that an *affective mechanism* would explain the moral repetition effect.

The idea that people experience less-intense affective reactions to stimuli and events after an initial exposure – i.e., *desensitization* (Campbell et al., 2014) – has received wide support. In the lab, people express less-extreme evaluations of the same affect-arousing stimulus after an initial encounter with it (Leventhal et al., 2007), and they display less-angry facial expressions in response to a first video clip depicting a terrorist attack than to a fifth clip depicting a different terrorist attack (Hoffman & Kaire, 2020). Outside the lab, people appear to become desensitized to violence after repeated exposure to it in the media (Carnagey et al., 2007; Funk et al., 2003), and clinicians leverage insights about desensitization to treat phobias and trauma (Rothbaum & Schwartz, 2002; Wolpe, 1961). Known by other names as well – affective habituation (Dijksterhuis & Smith, 2002), affective adaptation (Wilson & Gilbert, 2008), and hedonic

¹ Moral outrage may also involve disgust (Salerno & Peter-Hagene, 2013).

adaptation (Frederick & Loewenstein, 1999) – desensitization thus occurs in a wide range of contexts.

As an alternative to the affective mechanism, a cognitive mechanism for the moral repetition effect could be that repetition changes people’s perceptions of social norms. Repeated information feels familiar, and people can mistake familiarity for prevalence. For example, when a particular opinion sounds familiar, people will infer that many others must hold this opinion – even if the real reason the opinion sounds familiar is that the same person expressed it multiple times (Weaver et al., 2007). By the same token, repeatedly encountering the same description of the same person committing the same wrongdoing could erroneously lead people to think that *many* people commit that wrongdoing. And the more common people think a behavior is, the less morally wrong they will find it (Lindström et al., 2018). For example, when judging the morality of selfishness (Lindström et al., 2018), tax evasion (Monroe et al., 2018; Welch et al., 2005), unfamiliar behaviors (Eriksson et al., 2015), and a series of transgressions ranging from office-supply-theft to rape (McGraw, 1985), people appear to confuse what people *actually* do with what people *ought* to do. In short, repeatedly encountering the same wrongdoing could make it seem more common, which could reduce how unethical it seems (i.e., a *norm-perception mechanism*).

The present experiments test both the affective mechanism and the norm-perception mechanism – which are not mutually exclusive – but I thought the norm-perception mechanism would be less likely to account for the moral-repetition effect. People confuse familiarity with prevalence (Weaver et al., 2007) – but what specifically will they think is prevalent after repeatedly encountering the same description of the same transgression? Repetition could make *the transgression itself* seem prevalent (“everybody’s doing it”), which as noted should make the

transgression seem *less* unethical (Lindström et al., 2018; Monroe et al., 2018). Alternatively, repetition could make *knowledge* of a single person's transgression seem prevalent ("everybody's heard about what he did"), which in turn could make the transgression seem *more* unethical. That is, "things frequently seen are assumed to be widely known" (Kwan et al., 2015, p. 50), and people may believe that the most morally outrageous transgressions are the ones that become most widely known (see Brady et al., 2019; Brady et al., 2017). Or perhaps repeated exposure to the same information about a transgression could increase how frequently the *transgressor* commits it ("he does this all the time"), leading people to judge the transgression as *more* unethical than a first-time offense (Kliemann et al., 2008). Because it was unclear *a priori* which of these possibilities was more likely, I considered my initial tests of the norm-perception mechanism exploratory, whereas the tests of the affective mechanism were confirmatory.

By testing the moral repetition effect, the present work aims to demonstrate a new phenomenon grounded in both classic and recent work on how repetition affects judgments. First, the moral repetition effect is related to – though distinct from – the classic *mere-exposure effect*, by which repeatedly exposing people to a simple stimulus – such as a novel shape or symbol – increases how much they like it (Zajonc, 1968, 2001). That transgressions seem less unethical after repeated exposure could be considered a "moral mere-exposure effect," a phenomenon that scholars have speculated about but not experimentally tested (Jarudi et al., 2008; Weeks et al., 2005). However, evidence that repetition affects moral judgments would challenge the literature's conclusion that the mere-exposure effect does not occur reliably for negative stimuli (Albrecht & Carbon, 2014; Brickman et al., 1972; Grush, 1976; Meskin et al., 2013; Mrkva & Van Boven, 2020; Reber, Schwarz, et al., 2004; Swap, 1977; cf. Zajonc et al., 1974). This prior work suggests that repetition should *not* improve evaluations of transgressions,

which are highly negative stimuli. Moreover, the mere-exposure effect is stronger when the stimuli are very simple, like polygons, as opposed to moderately complex, like photographs (Montoya et al., 2017). Prior work thus raises doubts about whether mere exposure would affect how people judge highly complex stimuli like moral behaviors. In short, the mere-exposure literature does not clearly predict the moral repetition effect.

Second, the present research builds on the recent finding that repeated exposure to a political fake-news headline makes it seem less unethical to publish and share, even when people recognize its falsity (Effron & Raj, 2020). This effect emerged even when controlling for how much participants liked the headline, which helped distinguish it from the mere-exposure effect (Effron & Raj, 2020, Experiment 4). This research, to my knowledge, provides the only prior demonstration that repetition can affect moral judgments. However, the authors' (untested) explanation for the effect predicts that repetition will only affect moral judgments of spreading misinformation, and *not* moral judgments of other transgressions. Thus, this prior research neither predicts nor explains the moral repetition effect.

More specifically, Effron and Raj drew a theoretical distinction between what people believe, and what feels intuitively true to them (see Effron, 2018; Newman et al., 2015; Shidlovski et al., 2014). Repeating a falsehood, they argued, can make it feel intuitively true (i.e., gives it a “ring of truthfulness;” p. 76), even without convincing people to believe it. Because intuition drives moral judgments (Greene & Haidt, 2002), making a falsehood feel intuitively truthful – even when people explicitly acknowledge its falsity – should reduce how unethical people think it is to spread. With its focus on truthfulness, this proposed explanation predicts that repetition will affect moral judgments of spreading falsehoods like fake news, but *not* that repetition will affect moral judgments of other wrongdoings. For example, repeatedly reading

coverage of the Dao incident could make the incident feel more intuitively truthful, but it is unclear how this feeling would make the incident seem any less unethical. By providing the first demonstration of the moral repetition effect, the present research reveals a phenomenon that is much broader than the one studied by Effron and Raj and that is driven by a different mechanism than the one they proposed. As an ancillary goal, two of the present experiments also aim to replicate Effron and Raj's findings and provide a first test of the mechanism they proposed.

Finally, the present research identifies a new way in which affect biases moral judgments. Affect can motivate people to punish transgressors, compensate victims, and protest injustice (Iyer et al., 2007; Nelissen & Zeelenberg, 2009) – but it can also lead to inconsistent, normatively problematic patterns of moral judgment (Bloom, 2017; Slovic, 2007). For example, people judge the same transgression more harshly when the victim or the transgressor is personally identifiable (Gino et al., 2010; Small & Loewenstein, 2005), or when the victim is a single individual as opposed to a group (Kogut & Ritov, 2005, 2015; see also Nordgren & McDonnell, 2011), because we experience stronger feelings in response to the behavior of identifiable individuals. The present research identifies a different bias. While holding features of the victim and transgressor constant, the studies test whether our reliance on affect can lead us to form different moral judgments of the same transgression simply because we have read about it before.

The Present Research

The present studies test whether repeated exposure to descriptions of a wide variety of transgressions affect moral judgments of those behaviors. Experiment 1a provides a first demonstration of this moral repetition effect, showing that prior exposure to the same news headlines about business transgressions makes those transgressions seem less unethical. It also

replicates the finding that fake news seems less unethical to share when previously encountered, but it challenges the theoretical explanation originally offered for this finding (Effron & Raj, 2020). Experiment 1b was a high-powered, direct replication of Experiment 1a. Experiment 2 finds that the moral repetition effect generalizes to a new stimulus set: brief descriptions of “everyday” transgressions. Experiment 3 finds that exposing people to violations of six “moral foundations” (Graham et al., 2013) can make the violations seem less unethical and reduced how much punishment people think the violators deserved. Experiments 4 and 5 find support for the affective mechanism. Specifically, Experiment 4 shows that people report less-negative affective reactions to transgressions they have read about before (vs. are seeing for the first time), and that these affective reactions mediate the moral repetition effect; Experiment 5 replicates this mediation finding, and also finds that the moral repetition effect is eliminated when participants are instructed to base their moral judgments on reason instead of emotion. The studies also test the norm-perception mechanism, but only find inconsistent support (Experiments 3 and 4). Finally, Experiment 5 finds that a single repetition is sufficient to produce the moral repetition effect, and that five repetitions did not measurably increase the size of this effect. Together, these studies provide robust, replicable support for the moral repetition effect and its affective mechanism.

Open Practices

After Experiment 1a, all experiments were pre-registered. All sample sizes and data-exclusion criteria were determined in advance of data analysis, and I report all measures and experimental conditions. Pre-registration documents, verbatim materials, data, and analysis code are available at https://osf.io/6xpq8/?view_only=eb17d8e867aa4fe7b5007baa17086859 (Effron, 2021).

The experiments received ethics approval under protocol REC608 at the author's institution.

Experiments 1a and 1b

Experiment 1a had two goals. First, it sought to conceptually replicate the finding that sharing a fake-news headline would seem less unethical when the headline was seen previously (Effron & Raj, 2020). Second, it sought to extend this finding by exploring whether a variety of business transgressions would seem less unethical after repeatedly seeing descriptions of them. Experiment 1b was a large-sample, pre-registered replication conducted to confirm Experiment 1a's findings.

As an ancillary goal, Experiment 1b also tests Effron and Raj (2020)'s explanation for why fake-news seems less unethical to share when previously encountered. As noted, Effron and Raj suggested that repetition affects moral judgments by making fake news feel intuitively truthful. If this explanation is correct, then repetition should have a larger effect on moral judgments of sharing fake news than on moral judgments of committing transgressions that are unrelated to truthfulness.

Method

Experiment 1b was pre-registered at <https://aspredicted.org/383je.pdf>

Design. Both experiments had a 2 (headline-type: previously-seen vs. new; within-participants) X 2 (condition: judge sharing vs. judge behavior; between-participants) design with 16 repeated measures.

Participants. Residents of a large U.K. city who were enrolled in a behavior lab's participant pool completed Experiment 1a. The target sample size was 300 people, but more timeslots were posted to guard against no-shows, which resulted in 331 participants. *By a priori*

decision, I dropped the six participants who failed a reading-comprehension question before the experiment began.² Four additional participants did not provide responses to the dependent measures. The final sample size was thus 321 people, who provided a total of 5,134 responses in the repeated-measures design. The 316 people who provided demographics (216 women, 99 men, 1 non-binary person; M age = 27 years, SD = 11) held citizenship in 55 different countries.

Experiment 1b recruited U.K. nationals on Prolific Academic, an online participant pool whose users are less familiar with research methods and more diverse than users of Amazon Mechanical Turk (Peer et al., 2017). To promote data quality, participants could not begin if they failed a reading-comprehension question or were using a mobile device. The target sample size was 1,000 people; of the 987 who passed the screeners to begin the study, 940 remained after applying pre-registered exclusions (non-UK or duplicate IP address, duplicate participant IDs). These participants provided 15,040 responses to the main dependent measure.

Statistical power. Based on the smallest effect-size previously observed in this paradigm (Effron & Raj, 2020), Experiment 1a was powered to detect $d_z = -.15$ (equivalent to $b = -1.39$ in that study) for the difference between the two within-participants conditions in each of the two between-participants conditions. A power simulation determined that Experiment 1a could detect this effect size 88.6% of the time given its final sample size, design, and two-tailed $\alpha = .05$. The simulation used the *powerSim* function in R (Green & MacLeod, 2015), specified random intercepts for participants and items, and used the parameters observed in Experiment 2 of Effron

² For the reading-comprehension question, participants were asked to select the word that describes something that you use to cook. The correct answer was “stove;” incorrect answers included “car,” “guitar,” and “painting.” Subsequent studies used equivalent comprehension questions. See Online Supplement for verbatim materials.

& Raj (2020). It was not feasible to recruit enough participants from Experiment 1a's participant pool to have adequate power for testing the interaction between the two independent variables.

Experiment 1b, by contrast, could recruit enough online participants to test this interaction. The *powerSim* function determined that, with 1,000 participants and 16 repeated measures, Experiment 1b had 87.80% power to detect an interaction that attenuated the effect size observed in Experiment 1a by 2/3 (i.e., a moral repetition effect of $d_z = -0.31$, $b = -3.15$ attenuated by an interaction coefficient of $b = 2.10$). (Note that this result somewhat underestimates power because, whereas *powerSim* only allows two-tailed significance tests, I pre-registered one-tailed tests for Experiment 1b). This power simulation used the parameters observed in Experiment 1a.

Stimuli. The stimuli were 16 fake-news headlines about business transgressions, gathered from <https://www.snopes.com/fact-check/category/business/> and lightly edited, that actually circulated online (e.g., “A United Airlines flight attendant slapped a 7 month-old baby in the face for crying during a flight”; see Table 1). A relevant photograph (e.g., of a United Airlines plane) appeared above each headline.

Procedure. Participants completed Experiment 1a in the laboratory, embedded in randomly-ordered, unrelated studies. Participants completed Experiment 1b online.

Familiarization phase. The experiments first familiarized participants with 8 headlines, randomly selected from the bank of 16. Participants learned that they would be “rating some news headlines that were published on social media.” Then participants saw the 8 headlines once in a randomized order and completed a filler rating of each, and then saw the same headlines again in a different randomized order and completed a different filler rating of each. The order of the filler ratings (how engaging and how boring they found the headline) was counterbalanced.

Distraction phase. Next, to create a brief delay, the experiment asked participants 11 filler questions (e.g., “Do you own a smart phone?”). To introduce this phase, the instructions simply read, “We would now like you to answer a few questions about yourself.”

Judgment phase. Following Effron and Raj (2020), the following message introduced the next phase of the experiment:

In this part of the study you will be asked to read a series of **fake** news headlines. The information in these headline is not real.

The headlines describe behavior by various companies and individuals. However, the companies and individuals in question did NOT actually engage in these behaviors.

Effron and Raj originally included this message to demonstrate that repeated exposure to fake-news headlines makes them seem less unethical to share — even when people do not believe the headlines (cf. Pennycook et al., 2018).

Next, participants saw all 16 headlines (order randomized) — 8 of which they had seen earlier in the experiment (*previously-seen headlines*) and 8 of which they had not (*new headlines*) — and rated each on the dependent measure.

Between-participants manipulation and dependent measure. For the dependent measure, participants indicated their moral judgments about the 16 headlines. The specific judgments requested depended on randomly-assigned, between-participants condition. In the *judge-sharing* condition, participants indicated how unethical it would be to share each headline online given that it is false (as in Effron & Raj, 2020). In the *judge-behavior* condition, they instead indicated how unethical the behavior described in the headline would be if it were true (e.g., how unethical it would be for a flight-attendant to slap a baby for crying). Participants responded on a scale from 0 = *not at all* to 100 = *extremely* by moving a slider that initially appeared at the midpoint.

Comprehension check. Finally, participants indicated whether all 16 headlines were true, all were false, or some were true and some were false (multiple choice). If they selected the latter answer, they saw all 16 headlines again in a randomized order and indicated which were true versus false.

Results and Discussion

Analytic approach. I submitted the dependent variable to a mixed regression model with fixed effects for headline type (1 = previously seen, 0 = new), condition (judge sharing = 1, judge behavior = 0), and the headline-type X condition interaction, plus random intercepts for participants and stimuli (see Table 2).³ Then I computed simple slopes for the effect of headline type in each condition. The text below reports two-tailed significance tests for Experiment 1a, and pre-registered one-tailed tests for Experiment 1b because that experiment tested pre-registered directional predictions (see Cho & Abe, 2013). In these and all subsequent experiments, I computed mixed models using Stata 16's *mixed* command, which uses maximum likelihood (ML) estimation and an independent variance-covariance structure, and assesses significance using *z*-tests (see StataCorp, 2019).

Repetition reduces condemnation of fake-news sharing. Conceptually replicating previous research on fake political news (Effron & Raj, 2020), participants rated fake business headlines as less unethical to share when the headlines had been shown earlier in the study than when they had not, in both Experiment 1a ($M_s = 72.45$ and 75.66 , $SD_s = 23.69$ and 21.67), $d_z = -0.31$, $b = -3.15$, $SE(b) = .85$, $z = 3.72$, $p < .001$, and in Experiment 1b ($M_s = 70.26$ and 72.98 ,

³ I pre-registered fixed effects for stimuli in all studies, but later realized that stimuli should be treated as a random effect to account for stimulus sampling (see Judd et al., 2012). The main text thus reports the results from models with random intercepts for stimuli, but the conclusions are identical when stimuli are instead treated as fixed effects.

$SDs = 25.74$ and 24.29), $d_z = -0.26$, $b = -2.62$, $SE(b) = .49$, $z = 5.35$, $p < .001$, for the simple slope of headline-type in the judge-sharing condition.

Repetition reduces condemnation of other wrongdoing. Going beyond prior work, participants rated business transgressions as less unethical to commit when described in previously-seen headlines than when described in new headlines, in both Experiment 1a ($Ms = 74.55$ and 77.28 , $SDs = 15.64$ and 14.15), $d_z = -0.21$, $b = -2.94$, $SE(b) = .84$, $z = 3.51$, $p < .001$, and in Experiment 1b ($Ms = 75.19$ and 78.79 , $SDs = 14.25$ and 13.16), $d_z = -0.26$, $b = -3.64$, $SE(b) = .48$, $z = 7.58$, $p < .001$, for the simple slope of headline-type in the judge-behavior condition.

No significant interaction effect. Based on the theory that repeating misinformation influences moral judgments by making the misinformation feel more intuitively truthful (Effron & Raj, 2020), I hypothesized that the effect of repetition on judgments of sharing fake news would be larger than the effect on judgments of committing transgressions. However, the results showed no support. That is, the statistical interaction between headline-type and condition was not significant in either study, $b = -.21$, $p = .862$ in Experiment 1a, and $b = 1.02$, $p = .930$ in Experiment 1b (see Table 2; recall that only Experiment 1b was powered sufficiently to detect a small interaction). The p -value in Experiment 1b approaches 1.00 because the pre-registered one-tailed test was for a coefficient with the opposite sign ($p = .140$ by a two-tailed test).

Belief in the fake news. Prior work found that repeatedly encountering fake-news articles makes them seem less unethical to share even when people recognize them as fake (Effron & Raj, 2020). Consistent with this finding, the results above emerged even though the experiment explicitly informed participants the headlines were false; the comprehension-check measure showed that participants correctly identified the headlines as false on most trials (90%

in Experiment 1a, and 92% in Experiment 1b); and the conclusions were identical when I reran the analyses without responses to headlines that participants had misidentified as true.

Experiment 2

Replicating and extending the finding that previously-encountered misinformation seems less unethical to share (Effron & Raj, 2020), Experiments 1a and 1b show that a variety of business transgressions seem less unethical when repeatedly encountered. To examine generalizability, Experiment 2 sought to replicate this moral repetition effect with a new set of stimuli: brief descriptions of “everyday” transgressions.

Method

Experiment 2 was pre-registered at <https://aspredicted.org/ir3hf.pdf>

Participants. I requested 250 complete responses from American participants on Prolific Academic. The *powerSim* function in R determined that with this sample size, 30 repeated measures, two-tailed $\alpha = .05$, Experiment 2 would have 85.5% power to detect a moral repetition effect half the size of what Experiment 1b observed. (Because I pre-registered a one-tailed test, this study’s actual power was higher; *powerSim* does not compute power for one-tailed tests).

Participants could only begin the study if they had not completed a prior study in this series and correctly answered a reading-comprehension question. Of the 261 people who began the study, I dropped the 4 people who met pre-registered exclusion criteria (duplicate IP address; no duplicate participant IDs or non-U.S. IPs were detected), and 9 people did not provide any responses to the dependent measure. The final sample was thus 248 people (M age = 31 years, $SD = 12$; 117 men, 126 women, and 5 non-binary individuals), who provided 7,252 responses.

Stimuli. The transgressions used as stimuli were single-phrase descriptions of morally relevant behaviors selected based on a pretest, in which 128 American Prolific Academic users

rated 150 “common social behaviors” from Fuhrman et al. (1989, p. 587). The main experiments retained the 30 behaviors that these participants rated as most morally bad (e.g., “Stole money and jewelry from the relatives he was living with,” “Hit a car and left the scene of the accident”; see OSM for details about the pretest and for a complete list of the behaviors).

Procedure. The procedure closely resembled Experiments 1a and 1b. The familiarization phase exposed participants twice to 15 transgressions, randomly selected from the bank of 30 and presented in randomized orders. Participants rated each transgression on Experiment 1a’s filler items. After a brief distractor task (15 filler items about, e.g., social media use), participants completed the judgment phase, indicating how “morally bad” they found each of the 30 transgressions, half of which they had seen previously and half of which they had not (0 = *not at all* to 100 = *extremely*). Participants responded to the dependent measure by moving a slider that initially appeared at the scale midpoint.

Results and Discussion

As predicted, participants rated the transgressions as less morally bad when they had encountered them earlier in the study ($M = 70.81$, $SD = 14.92$) than when they had not ($M = 73.06$, $SD = 14.47$), $d_z = -0.28$, $b = -2.29$, $SE(b) = .41$, $z = 5.58$, $p < .001$ by a pre-registered one-tailed test (and a two-tailed test as well; see Table 3). These results replicate the moral repetition effect, and show that it generalizes to “everyday” wrongdoings and not just to corporate wrongdoings described in fake-news articles.

Experiment 3

To further examine generalizability, Experiment 3 tested whether repeated exposure to transgressions that violated each of six “moral foundations” (Graham et al., 2013) would reduce not only moral condemnation, but also punitive sentiment. Experiment 3 also tested the norm-

perception mechanism: that repeatedly encountering the same transgression would make it seem more common, which could reduce how unethical it seems. To test this mechanism, Experiment 3 measured judgments of how unusual each violation was. At this point in the research process, I was agnostic about whether this mechanism would receive support, but my pre-registered prediction was that this mechanism would be insufficient to completely account for the moral repetition effect.

Method

Experiment 3 was pre-registered at <https://aspredicted.org/k69i2.pdf>

Design. Participants rated 28 moral violations, half of which they had seen earlier in the study and half of which were new. Experiment 3 thus had two within-participants conditions with 28 repeated measures.

Participants and statistical power. The target sample was 250 urban residents enrolled in a U.K. behavioral lab's participant pool. With 28 repeated measures, this sample size provides approximately the same statistical power as Experiment 2.

Posting extra time-slots to guard against no-shows, I received responses from 273 people. Following pre-registered exclusion criteria, I dropped the two people who failed a reading-comprehension question before beginning the experiment. I excluded an additional person from analysis because computer problems forced them to repeat the same part of the experiment multiple times. Thus, the final sample size was 270 people (162 women, 101 men, and 7 people who did not report gender; M age = 33 years, SD = 13) who held citizenship in 49 different countries and provided a total of 7,552 responses to the dependent measure.

Stimuli. The stimuli were 28 one-sentence descriptions of behaviors selected from the standardized Moral Foundations Vignettes (Clifford et al., 2015). Five of the foundations –

authority, fairness, liberty, loyalty, and sanctity – were violated by four behaviors each. The sixth foundation, care, was violated by four behaviors involving emotional harm to people, two behaviors involving physical harm to people, and two behaviors involving physical harm to animals (see Table 4 for examples, and https://osf.io/6xpg8/?view_only=eb17d8e867aa4fe7b5007baa17086859 for a complete list).

Procedure. Experiment 3’s procedure was adapted from Experiments 1a and 1b, and participants completed it in the lab in randomized order with two other studies for unrelated projects. During the *familiarization phase*, participants saw 14 of the moral violations, presented in randomized order, and rated how poorly versus well-written each was (filler item). Then they saw the same violations a second time (order randomized again) and rated how boring versus interesting each was (filler). The 14 violations were randomly selected from the list of 28, with the constraint that each participant see two violations of the “care” foundation involving physical harm and two involving emotional harm, plus two violations from each of the remaining 5 moral foundations.

Next, for the experiment’s *distraction phase*, participants answered 11 filler questions about the behavioral lab where they were completing the experiment (e.g., how many times they had visited in the past six months).

Finally, for the *judgment phase*, participants saw all 28 violations (order randomized) – half of which they had previously seen, and half of which were new. Participants answered three questions about each violation in the following order: “How unethical is the behavior you saw?”, “How much does that behavior deserve to be punished?”, and “How unusual is it for a

person to engage in that behavior?”⁴ Participants responded by moving sliders on a 0–100 scale, anchored at *Not at all unethical/No punishment/Not at all unusual* and *Extremely unethical/Extremely harsh punishment/Extremely unusual*. The slider’s initial position was the scale midpoint, labelled *Moderately unethical/Moderate punishment/Moderately unusual*. An exploratory measure at the end of the experiment asked participants where they would place their views about “social issues” on a spectrum ranging from *Very liberal* to *Very conservative* (7-point scale).

Guessing the hypothesis. At the end of the study, participants responded to an open-ended question to assess whether they could guess the hypothesis: What do you think the researchers were hoping to find in this study?⁵

Results

Analytic approach. I analyzed each of the three dependent measures with a mixed regression model containing a dummy code for condition (1 = previously-seen violation, 0 = new violation), random intercepts for participants, and random intercepts for items (see Footnote 2; Judd et al., 2012). Table 5 displays the results. I report pre-registered one-tailed significance tests for directional hypotheses (see Cho & Abe, 2013).

Moral judgments. As predicted, repeated exposure reduced moral condemnation. Participants rated previously-seen violations ($M = 67.73$, $SD = 13.45$) as less unethical than new violations ($M = 69.13$, $SD = 13.64$), $d_z = -0.15$, $b = -1.39$, $z = 2.53$, $p = .006$.

Punishment. Also as predicted, repeated exposure reduced how much punishment participants thought the moral violations deserved. Participants recommended less harsh

⁴ In a follow-up study, the *unusualness* measure correlated highly with an alternative measure of norm perceptions: a subjective estimate of the number of people who engaged in each behavior, repeated-measures $r = -.75$; see Online Supplement.

⁵ Two other open-ended questions asked whether they had noticed any typos and if they had any other comments.

punishment for violations they had seen earlier in the experiment ($M = 51.64$, $SD = 16.88$) compared to violations they had not ($M = 53.72$, $SD = 16.23$), $d_z = -0.21$, $b = -2.07$, $z = 3.69$, $p < .001$. Descriptively, the effect size was slightly larger for punishment than for moral judgments (a difference of $d_z = .06$), but this difference was not significant, as shown by largely overlapping 95% CIs for these effect-size estimates: $d_z = -.21$, $CI = [-.33, -.09]$ for punishment and $d_z = -.15$, $CI = [-.27, -.03]$.

Social norm perceptions. Exploratory analyses supported the possibility that a norm-perception mechanism could contribute to the moral repetition effect – but as predicted, this mechanism was insufficient to completely explain the effect. Specifically, participants thought it was less unusual for a person to commit the violations they had previously seen in the experiment ($M = 54.47$, $SD = 13.47$) compared to violations they were seeing for the first time ($M = 56.24$, $SD = 12.75$), $d_z = -0.19$, $b = -1.70$, $z = 3.06$, $p = .002$ (two-tailed) – and the less unusual they perceived it, the less unethical they thought it was, $b = .25$, $z = 23.01$, $p < .001$. This indirect effect was significant in a multilevel mediation analysis with random intercepts for participants and fixed effects for item, $b = -.43$, $z = 3.04$, $p = .002$.⁶ However, participants still thought that previously-seen violations were significantly less unethical and less deserving of punishment than new violations after statistically controlling for how unusual they found the violations (see Table 5), $ps = .039$ and $.001$, respectively. Thus, the idea that repeated exposure to the vignettes made the relevant violations seem more common received support, but was insufficient to completely explain the moral repetition effect.

⁶ I ran the multilevel models using Stata's *gsem* command, and then multiplied the *a* and *b* paths together using the *nlcom* command (see StataCorp, 2021). A model with random intercepts for item failed to converge within 45 minutes.

Exploratory tests of moderation. Exploratory analyses reported in the OSM found no reliable evidence that repeated exposure had different effects depending on (a) the specific moral foundation violated, (b) how politically liberal or conservative participants were on social issues, or (c) the interaction between participants' politics and whether a behavior violated one of the "binding" foundations theorized to be more important to conservatives than to liberals (i.e., authority, loyalty, and sanctity; Graham et al., 2009).

Few people guessed the hypothesis. Reducing concerns about demand characteristics, virtually no one guessed the hypothesis. A research assistant coded participants' free responses to the question, "What did you think the researchers were hoping to find in this study?" Only 4 out of the 270 participants (i.e., 1.48%) mentioned anything about repeatedly encountering the same descriptions of behavior, and the conclusions were identical when I repeated the analyses without these participants.

Discussion

Experiment 3 suggests that the moral repetition effect generalizes broadly across the six "moral foundations" thought to carve out the moral domain. Repetition not only made violations of these foundations seem less unethical, but also made them seem less deserving of punishment.

The results were also consistent with the norm-perception mechanism: Repetition made the violations seem more common, which in turn predicted less moral condemnation. The evidence for this mechanism, however, is limited in two respects. First, it emerged in exploratory analyses, so a replication would be desirable. Second, many of the transgressions in Experiment 3 were highly unusual (e.g., intentionally running over a squirrel; having sex with a frozen chicken; see Gray & Keeney, 2015), which could have made social norms particularly salient and seem especially relevant for forming a moral judgment. A stimulus set with less-unusual

transgressions would provide a stronger test of the norm-perception mechanism. The next experiment provides such a test.

Importantly, the moral repetition effect was still reliable even after accounting for the norm-perception mechanism, which raises the possibility that another, unmeasured mechanism is needed to fully account for the effect (Zhao et al., 2010). This finding is consistent with prior work in which participants rated previously encountered fake-news headlines as less unethical to share than unfamiliar fake-news headlines, even after accounting for how widely participants thought the headlines had already been shared (Effron & Raj, 2020, Study 4). The unmeasured mechanism, I suggest, is the affective mechanism described earlier. Experiment 4 offers a first test of this mechanism.

Experiment 4

Experiment 4 had four key goals. First, it aimed to replicate the moral repetition effect with real-world moral transgressions, rather than with the fake news or hypothetical transgressions used in the previous experiments. To this end, Experiment 4's stimuli were news headlines describing actual transgressions by organizations and leaders. Second, Experiment 4 tested whether norm perceptions would mediate the moral repetition effect in a stimulus set without the sort of bizarre transgressions that Experiment 3 included. Third, Experiment 4 tested the mediating role of affect. I predicted that previously-encountered transgressions would elicit less-intense affective reactions, which in turn would predict less moral condemnation.

Finally, Experiment 4 explored whether the size of the moral repetition effect depends on the number of exposures. Whereas the previous experiments repeated each stimulus either 0 or 2 times, Experiment 4 repeated the stimuli 0, 1, or 5 times. I predicted that (a) repetition would

make the transgressions seem less unethical (i.e., a replication of the moral repetition effect), and (b) this effect would be larger with five repetitions than with one repetition.

Method

I pre-registered this study at <https://aspredicted.org/vz33g.pdf>

Participants. Because I expected a single repetition could have a smaller effect on moral judgments than two repetitions did in the earlier experiments, I recruited a large sample, posting slots for 800 American participants on Prolific Academic. With 21 repeated measures, this sample provides over 99% power to detect a mean difference of $b = .5$ on a 100-point scale, according to a power simulation with *powerSim* that used the variance parameters observed in Experiment 1b.

Participants could only begin the study if they passed a reading-comprehension question, were not on a mobile device, and had not participated in a previous study in this project. Of the 881 people who began the study, 10 were excluded for having a duplicate or non-US IP address, or a duplicate participant ID, and 74 could not be analyzed because they exited the study without responding to the main dependent measure for any items. Thus, the final sample size was 797 people, who provided a total of 16,657 observations in this repeated-measures design. (I neglected to collect demographics).

Stimuli. The stimuli were 21 headlines describing real transgressions committed by corporations or leaders (e.g., “Patisserie Valerie overstated its financial position by £94 million”). The headlines were taken from real news sources and rephrased in the active voice. A relevant photograph accompanied each headline (e.g., a picture of Patisserie Valerie), making the stimuli resemble the type of news content that people regularly encounter on social media. Table

6 lists the stimuli's text; the accompanying photos are available at

https://osf.io/6xpg8/?view_only=eb17d8e867aa4fe7b5007baa17086859

Procedure. The procedure was similar to Experiments 1a–3, except this time I varied the number of repetitions. In the familiarization phase, participants first rated 14 headlines on a filler item. Then, they rated 7 of these 14 headlines four more times each on additional filler items. Thus, by the end of the familiarization phase, they had seen 1/3 of the 21 headlines five times, 1/3 of the headlines just once, and 1/3 of the headlines not at all. The specific headlines seen 5, 1, or 0 times was randomized between participants. The filler items asked how boring, engaging, and well-written each headline was, how many words it had, and how colorful the accompanying image was. The order of these filler items was randomized so that the specific items on which participants rated 7 versus 14 items varied.

Next, for the distraction phase, participants answered 10 filler questions (e.g., “Do you own a smart phone?”).

Finally, for the judgment phase, participants indicated how unethical they found the behavior described in each headline (dependent measure), how unusual they found the behavior, and how intense their emotional reaction to the behavior was. The second two items were potential mediators, and their order was randomized. Participants responded to each item by typing a number from 0 = *not at all* to 100 = *extremely*.

At the end of the study, participants responded to the open-ended questions from Study 3, including a prompt to guess the hypothesis.

Results

Analytic approach. I submitted the dependent variable (moral judgments) to a multi-level regression model with two planned, orthogonal contrasts for condition (*contrast1*: no-prior-

exposure condition coded -2 , other two conditions coded $+1$; *contrast2*: no-prior-exposure condition coded 0 , one-prior-exposure condition coded -1 , five-prior-exposure condition coded $+1$), and random intercepts for participants and items (see Footnote 2; Judd et al., 2012). The first contrast tests the prediction that prior exposure makes violations seem less unethical (i.e., the moral repetition effect). The second contrast tests whether this effect is larger after five exposures than after just one. As a robustness check, I reran the model controlling for norm perceptions (as in Experiment 3, but not pre-registered); the conclusions reported below were identical. Table 7 displays descriptive statistics in each condition, and Table 8 displays the regression results. I report pre-registered one-tailed tests of directional predictions.

The moral repetition effect replicated. Participants rated the transgressions as less unethical when they had seen them earlier in the study ($M = 74.83$, $SD = 17.19$ for transgressions previously shown one or five times) than when they had not ($M = 75.89$, $SD = 17.99$ for headlines not previously shown), $d_z = -0.09$, $b = -.37$, $SE(b) = .11$, $z = 3.20$, $p < .001$ for *contrast1*.

More exposures did not significantly increase the moral repetition effect. There was not sufficient evidence to support the prediction that transgressions would be rated as less unethical when they were previously seen five times ($M = 74.55$, $SD = 18.16$) compared to just once ($M = 75.22$, $SD = 18.59$), $b = -.26$, $d_z = -0.05$, $SE(b) = .20$, $z = 1.29$, $p = .098$ for *contrast2*.

One exposure was enough. Whereas the earlier experiments showed a moral repetition effect after two prior exposures to a transgression, an exploratory analysis revealed that just one exposure was sufficient in the present experiment. That is, participants rated transgressions they had previously seen once ($M = 75.22$, $SD = 18.59$) as less unethical than transgressions they had

not previously seen ($M = 75.89$, $SD = 17.99$), $d_z = 0.05$, $b = -.85$, $SE(b) = .40$, $z = 2.13$, $p = .033$ (two-tailed, because this analysis was not pre-registered) for a dummy code comparing these two conditions in a mixed-regression analysis that also included a dummy code comparing the five-exposure condition to the control condition, fixed effects for transgression, and random intercepts for participant.

Mediation: analytic approach. I next tested two potential mediators of the moral repetition effect: affective intensity and norm perception. Using the *gsem* function in Stata, I modeled paths from *contrast1* (which compares the two repetition conditions to the no-repetition condition; see above), to the two mediators in parallel (mean-centered), to unethicity judgments. As controls, the model also included the orthogonal contrast (i.e., *contrast2*), fixed effects for transgression, and random intercepts for participants.⁷ Then I computed indirect effects by multiplying the relevant *a* and *b* paths together using Stata's *nlcom* command (StataCorp, 2021). As above, I report pre-registered one-tailed tests of directional predictions, except where indicated.

Affective intensity, but not norm perceptions, mediated the moral repetition effect. As predicted, the results showed a significantly negative indirect effect through affective intensity, $b = -.26$, $z = 5.79$, $p < .001$. That is, people rated their affective reactions to the transgressions as less intense when they had seen them before compared to when they had not, and the weaker their affective reactions, the less they condemned the transgressions (see Figure 1's top path).

Contrary to predictions, the results did not show a significant indirect effect through norm perceptions, $b = -.02$, $z = 1.17$, $p = .123$. In contrast to Experiment 3, previously-seen

⁷ A model with random intercepts for item fail to converge within 45 minutes.

transgressions were not rated as significantly less unusual than new transgressions, although the more unusual a transgression was perceived as being, the more it was condemned (see Figure 1's bottom path). Note that the indirect effect through affective intensity reported above was 6.5 times larger than the indirect effect through norm perceptions, and this difference in the size of the indirect effects was significant, $b = .24$, $z = 4.78$, $p < .001$. The indirect effect through norm perceptions was still not significant when the affective-intensity measure was omitted from the model, $b = -.04$, $z = 1.35$, $p = .177$ (two-tailed because this analysis was not pre-registered).⁸

Few people guessed the hypothesis. As in Study 3, almost no one guessed the hypothesis. A research assistant's coding revealed that only 6 out of the 797 participants (i.e., 0.75%) mentioned anything about repetition when asked what they thought the researchers hoped to find. Experiment 4's conclusions were identical when these 6 participants were excluded from analyses. These findings reduce concern that demand characteristics explain the results.

Discussion

Experiment 4 advances our understanding of the moral repetition effect in several ways. First, it suggests that this effect not only occurs when people judge hypothetical transgressions and fake news (Experiments 1a–3), but also when they judge headlines about a wide range of actual transgressions by organizations and their leaders (see Table 6). Second, Experiment 4 found mediational evidence consistent with my claim that prior exposure to a transgression

⁸ As in Experiment 3, the direct effect of repetition on moral judgments was significant when controlling for norm-perception, $b = -.32$, $z = 2.94$, $p = .003$ for *contrast1*, reinforcing the idea that norm-perceptions cannot fully account for the moral repetition effect. By contrast, this direct effect dropped to non-significance when controlling for affective intensity instead of norm perception, $b = -.07$, $z = .69$, $p = .244$ (one-tailed), suggesting that affective intensity on its own can account for the moral repetition effect.

reduces the intensity of the negative affect it elicits, which in turn reduces how unethical the transgression seems.

Third, Experiment 4 failed to replicate Experiment 3's exploratory finding that norm perceptions mediated the moral repetition effect. That is, the results did not show that previously-encountered transgressions were less unusual than transgressions seen for the first time. One possibility is that the specific stimuli used in Experiment 3 (moral-foundations violations, like sex with a chicken) drew people's attention to unusualness in a way that Experiment 4's stimuli did not (e.g., sexual harassment; see Gray & Keeney, 2015). In any case, Experiment 4 does replicate Experiment 3's finding that the moral repetition effect occurs even when controlling for norm perceptions. Thus, norm perceptions may play a mediating role in some circumstances, but they do not appear to be necessary to explain the moral repetition effect.

Finally, Experiment 4 found that even one exposure to a transgression was sufficient to make it seem less unethical. Unexpectedly, there was insufficient evidence to conclude that five exposures increased the magnitude of this effect. If the moral repetition effect is larger with five repetitions compared to just one, it is not dramatically larger. Although it remains possible that more than five repetitions would influence moral judgments more than just one repetition, it appears that the marginal effect of each additional repetition falls off steeply after the first. This observation is consistent with prior work showing the number of stimulus repetitions has a diminishing marginal effect on judgments (e.g., DiFonzo et al., 2016; Montoya et al., 2017) and desensitization responses (Rankin et al., 2009; Thompson & Spencer, 1966).

Experiment 5: Turning Off the Moral Repetition Effect

Experiment 4 provides mediational evidence that affect plays a key role in the moral repetition effect. Experiment 5 sought to replicate this effect, and also to test the mediating role

of anger, a specific emotional component of moral outrage (Batson et al., 2007; Salerno & Peter-Hagene, 2013). A limitation of this approach is that, although prior work provides a strong basis for assuming that affect precedes moral judgments (e.g., Greene & Haidt, 2002; Haidt, 2001; Haidt & Kesebir, 2010), no measurement-of-mediation approach can empirically speak to the causal order of affective reactions and moral judgments (Bullock et al., 2010; Fiedler et al., 2011). To address this limitation, Experiment 5 employed a moderation-of-process design (see Spencer et al., 2005). If affect is involved in the moral repetition effect, then instructing participants to form moral judgments based on reason instead of emotion should attenuate it. Thus, Experiment 5 investigated a theoretically grounded approach to reducing people's susceptibility to the moral repetition effect.

Method

The pre-registration document for Experiment 5 is posted at

<https://aspredicted.org/kx48h.pdf>

Stimuli, participants, and statistical power. The stimuli were the same 30 descriptions of wrongdoings used in Experiment 2. For this reason, Experiment 5's sample size was determined based on the observed effect size in Experiment 2. An *a priori* power simulation with *powerSim* estimated that a sample of 500 participants with 30 stimuli would provide 94% power to detect a statistical interaction that completely attenuates the effect size observed in Experiment 2 at $\alpha = .05$, two-tailed. However, to ensure that Experiment 5 could detect even a partially attenuated interaction, I targeted a larger sample, posting slots for 750 participants on Prolific Academic. To begin the study, participant had to have identified as a UK national and resident on a prescreen survey, pass a reading comprehension check, not be on a mobile device, and not have completed a prior study in this program of research. After applying the preregistered

exclusion criteria (i.e., duplicate or non-UK IP addresses, or duplicate participant IDs) and discarding participants who could not be analyzed because they did not provide any responses to the dependent measure, 725 people remained (451 women, 265 men, 9 non-binary; M age = 37 years, $SD = 14$) from the 773 who began the study. The remaining participants provided 21,630 responses to the dependent measure in this repeated-measures design.

Procedure. Experiment 5 had a 2 (behavior type: previously seen vs. new; within-participants) X 2 (judgment basis: emotion vs. reason; between-participants) factorial design with 30 repeated measures. The familiarization phase exposed participants twice to 15 transgressions, randomly selected from the bank of 30 and presented in randomized orders. Participants rated each transgression on Experiment 2's filler items. After a brief distractor task (15 filler items about, e.g., demographics and social media use), participants were randomly assigned instructions to use either *emotion* or *reason* during the upcoming judgment phase (a manipulation adapted from Martel et al., 2020; Nordgren & Dijksterhuis, 2008). Specifically, participants in the *emotion* condition read the following (emphasis in original):

Many people believe that **emotion is the cornerstone of good moral judgment**. To determine what's right and wrong, we should **tune into our emotions** and **use our feelings** rather than relying on reason.

We will now show you a series of behaviours and ask you to rate how **morally bad** each behavior is.

Please assess the behaviours by **relying on emotion, rather than reason**. **Go with your first instinct** about each behaviour, **pay attention to the feelings** you have about it (like anger or outrage), and **don't think too hard** when forming your judgment about how morally bad it is.

By contrast, participants in the *reason* condition instead read:

Many people believe that **reason is the cornerstone of good moral judgment**. To determine what's right and wrong, we should **think rationally** rather than relying on our emotions.

We will now show you a series of behaviors and ask you to rate how **morally bad** each behaviour is.

Please assess the behaviors by **relying on reason, rather than emotion**. **Think very hard** about each behaviour, try to **ignore any feelings** you might have about it (like anger or outrage), and **generate a clear reason** about why the behaviour is – or is not – so morally bad.

Next, participants completed the judgment phase, providing ratings of each of the 30 transgressions – half of which they had seen at the beginning of the study, and half of which were new. In the emotion condition, participants first typed a number to indicate the intensity of their emotional reaction to the transgression (0 = *no emotional reaction*, 100 = *extremely intense emotional reaction*), how angry the transgression made them feel (0 = *not at all angry*, 100 = *extremely angry*), and how morally bad they thought it was (0 = *not at all morally bad*, 100 = *extremely morally bad*). The first two items were potential mediators of the moral repetition effect, and aimed to draw participants' attention to their emotions. The reason condition replaced these two measures with filler items that aimed to draw participants' attention to their reasoning: "Can you think of a clear reason why that behaviour is – or is not – so morally bad?" (*yes/no*) and "How carefully have you now thought about how morally bad that behaviour is?" (0 = *not at all carefully*, 100 = *extremely carefully*).

Participants in all conditions then responded to the dependent measure: "How morally bad is that behavior?" (0 = *not at all morally bad*, 100 = *extremely morally bad*). Finally, they responded to two manipulation checks (Martel et al., 2020), indicating how much they had used (a) emotion and (b) reason when judging how morally bad the transgressions were (1 = *not at all* to 5 = *a great deal*).

Results

I pre-registered and report one-tailed significance tests for directional predictions.

Manipulation check. Participants said they relied on emotion more in the emotion condition ($M = 3.99$, $SD = .88$) than in the reason condition ($M = 2.50$, $SD = 1.08$), $t(716) = 20.15$, $p < .001$. Conversely, they said they relied on reason more in the reason condition ($M = 4.04$, $SD = .76$) than in the emotion condition ($M = 2.78$, $SD = 1.08$), $t(717) = 18.22$, $p < .001$.

(Because of missing data, the *dfs* for these analyses are smaller than the total sample size would imply).

Reasoning eliminated the moral repetition effect. To test whether the instructions to rely on reason versus emotion moderated the moral repetition effect, I submitted the ratings of moral badness to a mixed regression model with random intercepts for participants and fixed-effect dummy codes for the specific wrongdoing, wrongdoing condition (1 = previously seen, 0 = new), instruction condition (1 = reason, 0 = emotion), and the wrongdoing X instructions interaction. With this coding, the hypothesis predicts a significant interaction with a positive coefficient, which is exactly what the results showed, $b = 1.54$, $z = 2.96$, $p < .002$ (see Table 9).

Decomposing this interaction revealed the predicted pattern of results. When participants were instructed to rely on emotion, they rated previously-seen transgressions as less unethical than new wrongdoings ($M_s = 70.12$ and 71.91 , respectively, $SD_s = 17.24$ and 17.37), replicating the moral repetition effect, $b = -2.00$, $SE(b) = .38$, $z = 5.29$, $p < .001$. The effect size, $d_z = -0.22$, was similar to the effect size observed in Experiment 2 with the same stimuli when participants received no instructions about whether to rely on emotion or reason, $d_z = -0.28$. By contrast, when Experiment 5's participants were instructed to rely on reason, the results did not show a significant moral repetition effect ($M_s = 69.69$ and 70.09 for previously-seen and new wrongdoings, respectively, $SD_s = 17.25$ and 16.66), $b = -.45$, $SE(b) = .36$, $z = 1.25$, $p = .212$, $d_z = -.05$. These results suggest that affective processing plays an important role in the moral repetition effect.

Affective intensity and anger mediated the moral repetition effect. In the emotion condition, the moral repetition effect was significantly mediated by affective intensity (replicating Experiment 4) and also, separately, by anger. (Recall that the reason condition did

not assess these mediators to avoid focusing people on their emotions). That is, people reacted with less affective intensity and less anger to a wrongdoing if they had seen it previously, and the less affective intensity and anger they experienced, the less morally bad they thought the wrongdoing was (see Figure 2). These results are from two multilevel mediation analyses – one for each mediator – computed using the *gsem* command in Stata, with random intercepts for participant and fixed effects for item,⁹ mean-centering the mediator, and computing indirect effects by multiplying the *a* and *b* paths together with the *nlcom* command (StataCorp, 2021). As predicted, these analyses revealed a significantly negative indirect effect for both affective intensity, $b = -2.30$, $SE(b) = .23$, $z = 9.87$, $p < .001$, and anger, $b = -2.30$, $SE(b) = .25$, $z = 9.08$, $p < .001$. Suggesting that affective intensity and anger could each account for all the variance in the moral repetition effect, the direct effect of repetition on moral judgments was no longer significant after controlling for either affective intensity, $b = .31$, $SE(b) = .31$, $z = 1.01$, $p = .844$, or anger, $b = .33$, $SE(b) = .29$, $z = 1.14$, $p = .873$ (the *p*-values approach 1.00 for these one-tailed tests because the direct effect was in the opposite direction as the moral repetition effect).

Discussion

These results provide evidence that affect plays a causal role in the moral repetition effect. When participants were instructed to use their emotions to inform their moral judgments, the effect replicated, and was mediated by affective intensity in general (as in Experiment 4) and by anger in particular. By contrast, when participants were instructed to ignore their emotions, the moral repetition effect was eliminated. This moderation-of-process finding addresses ambiguity about the causal order of affect and moral judgments (Spencer et al., 2005).

⁹ A model with random intercepts for item failed to converge within 45 minutes.

General Discussion

Five experiments suggest that people judge bad deeds more leniently when they have encountered them before. Specifically, repeatedly encountering the same description of a bad deed reduced how unethical people thought it was to commit, and reduced how much punishment they thought the perpetrator deserved. This moral repetition effect was robust, replicating in several high-powered, pre-registered studies. It also generalized across populations and stimuli, emerging among participants from 55 nations judging transgressions that included spreading fake news, engaging in corporate wrongdoing, violating the six “moral foundations” (Graham et al., 2013), and committing “everyday” bad deeds. The effect replicated not only when the stimuli were short descriptions of hypothetical wrongdoing, but also when they were real news headlines about actual wrongdoing.

The results also suggest that an affective mechanism can explain the moral repetition effect. Experiments 4 and 5 found support for a statistical mediation model in which previously encountered transgressions elicited less negative affect in general and less anger in particular, which in turn led people to judge these transgressions as less morally wrong. Experiment 5 also found that instructing people to base their moral judgments on reason rather than emotion eliminated the moral repetition effect. This affective mechanism is consistent with the well-established finding that people become desensitized to aversive stimuli and negative experiences after an initial exposure (e.g., Bradley et al., 1993; Campbell et al., 2014; Dijksterhuis & Smith, 2002; Nelson & Meyvis, 2008).

Lower-level or higher-level cognitive processes could explain why participants became desensitized to the moral transgressions. On one hand, prior exposure to a moral transgression provides an opportunity to understand why it occurred and to minimize its consequences – in

other words, to “explain it away” (Wilson & Gilbert, 2008). Thus, perhaps participants felt less angry about the same transgression when they had seen it before because they had had time to think about it. On the other hand, higher-level cognitive processing is not needed to explain desensitization. A variety of non-human organisms habituate to aversive stimuli when repeatedly presented (Boisseau et al., 2016; Thompson, 2009; Thompson & Spencer, 1966). People rate negative words as less negative after being subliminally primed with them (Dijksterhuis & Smith, 2002), they show less activity in the amygdala in response to an emotional face if they have seen it before (Wright et al., 2001), and they show a weaker startle reflex in response to an unpleasant photo if it has been previously presented (Bradley et al., 1993). Thus, participants’ negative affective response to a moral transgression may have ebbed after an initial exposure even without taking the time to “explain away” the transgression. If this explanation is correct, then the moral repetition effect should still occur even when people are under cognitive load during the exposure phase – a prediction that future research should test. In practice, it seems likely that lower-level and higher-level cognitive processes would both contribute to the desensitization observed in the moral repetition effect.

Alternative Explanations

Norm perception. The present research did not find consistent support for the possibility that a norm-perception mechanism could explain the moral repetition effect. When participants judged the morality of violating the six moral foundations – violations that included a number of highly unusual behaviors, like having sex with a frozen chicken – norm perceptions significantly but only partially mediated the moral repetition effect (Experiment 3). That is, prior exposure to these violations made them seem less unusual, which in turn made them seem less morally wrong. By contrast, when participants judged the morality of committing real-world business

transgressions, norm perceptions did not significantly mediate the moral repetition effect because exposure did not make the transgressions seem any less unusual (Experiment 4). One explanation for these inconsistent results could be that exposure makes wrongdoing feel more familiar, but that people interpret familiarity differently in different contexts (Corneille et al., 2020; Schwarz, 2004; Unkelbach & Greifeneder, 2013). People may be more likely to interpret familiarity as a cue to prevalence when the question of prevalence is salient. In Experiment 3, the presence of highly unusual behaviors may have made participants think more about prevalence throughout the study. Future research is needed to understand whether and when the norm-perception mechanism reliably contributes to the moral repetition effect. Such research might try measuring norm perceptions with different items (e.g., estimates of the number of people who commit a behavior; see Footnote 4). The present research suggests, however, that an affective mechanism can explain the moral repetition effect above and beyond a norm-perception mechanism.

Demand characteristics. It is unlikely that the moral-repetition effect emerged simply because participants guessed the hypothesis. Experiments 3 and 4 showed that a very small percentage of participants guessed or came close to guessing the hypothesis (i.e., < 2%), and the results remained unchanged when these participants were excluded.

Other alternative explanations. Other known consequences of repetition would struggle to account for why repeated exposure to bad deeds affected moral judgments. Repeatedly encountering the same stimulus can get boring (Berlyne, 1970; Stang, 1975) – but whereas boredom should result in more-negative evaluations of stimuli, the moral repetition effect demonstrates more-positive evaluations of bad deeds after repetition. Stimuli can become more salient after repeated encounters (Mrkva & Van Boven, 2020) – but whereas salience makes

interpersonal judgments more extreme (McArthur & Solomon, 1978; Sadler & Tesser, 1973), repetition in the present studies made judgments *less* extreme.

The Role of Fluency

Prior work has explained the effect of repetition on judgment by noting that repeated stimuli are easier to cognitively process – that is, they feel more *fluent* (e.g., Graf et al., 2018; Jacoby, Woloshyn, et al., 1989; Schwarz et al., 2007; Whittlesea et al., 1990). For example, scholars generally agree that fluency explains why repeated statements seem truer than new statements (see Dechêne et al., 2010). There is some evidence that fluency can affect moral judgments: Transgressions can seem more unethical when described in disfluent versus fluent font (i.e., hard-to-read vs. easy-to-read; Laham et al., 2009). In theory, fluency could play a role in the moral repetition effect in two respects: by explaining (a) why repeated transgressions (sometimes) seem less unusual (i.e., the norm-perception mechanism), and (b) why repeated transgressions elicit less-negative affect (i.e., the affective mechanism). Thus, rather than offering an alternative mechanism for the moral repetition effect, fluency could explain why the two proposed mechanisms operate (a “mechanism behind the mechanisms”).

Fluency could underlie the norm-perception mechanism because “fluency implies frequency, which in turn implies consensus” (Alter & Oppenheimer, 2009, p. 227). That is, people could hold a *naïve theory* (Schwarz, 2004) that if information about a behavior feels easy to process, then lots of people must be performing that behavior. A similar, fluency-based explanation has been offered for why repeatedly encountering the same opinion makes people think it is widely held, even if it was actually the same individual who voiced the opinion on every encounter (Weaver et al., 2007).

Fluency could also underlie the affective mechanism. According to the *hedonic marker hypothesis*, fluency feels good (Reber, Schwarz, et al., 2004; Reber et al., 1998; Winkielman et al., 2012; Winkielman et al., 2003). Thus, because repetition increases fluency, repeatedly encountering the same stimulus should increase positive affect. This positive affect, in turn, could be misattributed to the stimulus itself, making judgments of it more positive (Bornstein & D'Agostino, 1994; Bornstein & D'Agostino, 1992).

However, fluency's potential role in the moral-repetition effect has been challenging to capture empirically. I tried the two approaches. First, I administered a validated, subjective measure of fluency: self-report ratings of how easy a stimulus feels to process (Graf et al., 2018). However, a pilot study found ceiling effects on this measure (i.e., most people rated the descriptions of transgressions as extremely easy to process). Next, I tried an objective measure of fluency: reaction time (Reber, Wurtz, et al., 2004; Unkelbach & Stahl, 2009; Wurtz et al., 2008). A supplemental study, Experiment S1, replicated the moral repetition effect ($d_z = -.25, p < .001, N = 251$ with 30 repeated measures; see OSM), but found no evidence that it was mediated by reaction time (see also DiFonzo et al., 2016). These results do not rule out a fluency mechanism, but they do illustrate that fluency's role is hard to pinpoint using currently accepted methods.

Theoretical Contributions

The findings expand understanding of when and how repetition affects judgment. We know that repeated statements seem truer (e.g., Dechêne et al., 2010; Fazio et al., 2015; Pennycook et al., 2018), repeated names seem more famous (e.g., Buchli, 2019; Jacoby, Kelley, et al., 1989; Jacoby, Woloshyn, et al., 1989; Topolinski & Strack, 2010), and repeated symbols seem more likeable (e.g., Montoya et al., 2017; Zajonc, 1968). The present research is the first,

however, to demonstrate that repetition can influence moral judgments of transgressions across the moral domain.

Other scholars have suggested that repetition might affect moral judgments in this way, but never tested this idea directly. Specifically, Jarudi et al. (2008) proposed a “moral mere-exposure effect” to explain why consumers rate common household objects, like refrigerators, as morally positive. Weeks et al. (2005) invoked the same idea to explain why business leaders who report having faced certain ethically fraught situations in the past – compared to those who report not having faced such situations – rate ethical violations in those situations as more acceptable. However, because they lacked an experimental manipulation, neither of these studies could assess whether exposure has a causal influence on moral judgments.

In demonstrating a causal relationship between repetition and moral judgments, the present research addresses key theoretical questions in the literatures on cognitive, social, and moral psychology about when repetition affects judgments. As noted, the classic mere-exposure effect shows that repeated exposure to symbols, words, and pictures increases how much people like these stimuli – unless the stimuli have negative associations (for a review, see Reber, Schwarz, et al., 2004). For example, when non-Chinese speakers were repeatedly exposed to a Chinese character, they liked it more – but only if they were told nothing about its meaning or told that it had a positive meaning, and not if they were told that it had a negative meaning (Mrkva & Van Boven, 2020). By contrast, the present research finds that morally negative behaviors seemed less morally wrong after repeated exposure. One explanation is that desensitization played a more important role in the present experiments than in a typical mere-exposure study. “A United Airlines flight attendant slapped a 7-month-old baby in the face for crying during a flight” (Experiments 1a and 1b) and other morally relevant behaviors offer more

potential for desensitization than an unfamiliar Chinese character because they presumably elicit more affect to begin with (Dijksterhuis & Smith, 2002). The present findings thus illustrate how a different understanding of repetition effects can emerge when examining affectively potent stimuli like descriptions of moral transgressions as opposed to the affectively weak stimuli typically used in mere-exposure paradigms (see Montoya et al., 2017).

Moving beyond the classic mere-exposure effect, the present findings also shed new light on recent research about repetition and moral judgments. Experiments 1a and 1b replicated and extended the finding that misinformation seems less unethical to spread when repeatedly encountered, but they challenge the explanation that repetition made misinformation feel more intuitively truthful (Effron & Raj, 2020). This explanation cannot account for why bad deeds unrelated to misinformation showed a moral repetition effect, and why this effect was not measurably different than the effect of repetition on moral judgments of misinformation (Experiments 1a and 1b). By contrast, this pattern of results is consistent with the affective mechanism for the moral repetition effect. Thus, fake news may seem less unethical to share when repeatedly encountered because people feel less angry about the falsity of such news if they have seen it before. Effron and Raj's findings appear to be a specific example of a general moral repetition effect. In short, the present research challenges classic and contemporary assumptions about when and how repetition will affect judgments.

More broadly, the research reveals a new way in which our moral judgments are biased by their reliance on affect. The results suggest that because people react with less affective intensity to a transgression they have read about before, they judge that transgression as less unethical. Other examples of how affect biases moral judgment stem from *compassion collapse*, “a general phenomenon of diminished affective sensitivity towards groups of people in need of

help” (Cameron & Payne, 2011, p. 2). Because of compassion collapse, people will donate more money to help a single victim than a group of victims (Kogut & Ritov, 2005; Västfall et al., 2014), and will punish perpetrators less severely when the crime harms more people (Nordgren & McDonnell, 2011). Relatedly, people think wrongdoing is worse when the victim is personally identifiable and thus evokes stronger affective reactions (Gino et al., 2010; see also Small & Loewenstein, 2003; Small et al., 2007). Along with this prior work, the moral repetition effect supports the idea that affect is an important driver of moral judgments (Greene & Haidt, 2002), but a driver that can lead to inconsistency and bias (Bloom, 2017).

However, the moral repetition effect is unlikely to involve the same processes as these previous examples of affective bias in moral judgment (see Cameron, 2017). One reason compassion collapse occurs is that people think more abstractly about groups than individuals (Kogut & Ritov, 2015); by contrast, it is not obvious that prior exposure to the same wrongdoing would make people perceive the wrongdoing more abstractly. Another reason is that people engage in motivated empathy suppression to avoid feeling overwhelmed when confronted by multiple victims. For example, people skilled at regulating their emotions reacted with less affective intensity to eight malnourished child victims of civil war than to just one (Cameron & Payne, 2011). However, it is not obvious that the descriptions of wrongdoing used in the present studies held the potential to emotionally overwhelm participants after just one or two exposures (e.g., a student copying a classmate’s answers in Experiment 3). Finally, repeated exposure to the same transgression does not make individual victims any more identifiable (cf. Gino et al., 2010). Thus, although driven by affect, the moral repetition effect is distinct from other known affective biases on moral judgment.

Limitations and Future Directions

Effect size. The effects observed under controlled experimental conditions were robust and consistent, but not large. Across studies, the standardized effect size was $d_z = .22$ (see Figure 3, top panel), which corresponds to the 35th percentile for effect sizes in social psychology and is equivalent to the median effect size in research on gender differences (Lovakov & Agadullina, 2021). In absolute terms, across studies, the moral repetition effect shifted judgment by 2.3 points (see Figure 3, bottom panel). This shift is small when compared against the full range of possible scale responses (from 0 = not at all unethical to 100 = completely unethical), but less small when compared against the range of responses that the stimuli could be expected to elicit. These stimuli were wrongdoings selected to be highly unethical (e.g., a flight attendant slapping a baby for crying; a CEO committing sexual harassment), so – as should be expected – the bulk of responses fell in a narrow range at the scale’s top end (e.g., across studies, the majority of ratings were greater than 81). In short, repetition did not lead people to judge wrongdoings as morally acceptable, but it did reduce how unethical people thought the wrongdoings were.

These small effects are important because, as discussed, they advance theory about how repetition affects moral judgments. Arguably, it is also impressive that simply reading the same description of a severe transgression more than once can have any impact on moral judgments (Prentice & Miller, 1992). Moreover, in the real world, small effects could be important when amplified across large populations. For example, viral social media content or continuous news reports about transgressions can be viewed by hundreds of millions of people. If repeated exposure to such content reduced the moral condemnation expressed by the average user by 1/5 of a standard deviation, as it did in the present experiments, it could potentially have meaningful consequences for decisions to share the content further or to call for the transgressor’s punishment. Future research is needed to understand whether, when amplified across the many

millions of people who may repeatedly encounter a report about a transgression, this small but consistent effect has measurable consequences outside the lab.

Potential boundary conditions. Does repeated exposure to transgressions, in and of itself, ever make those transgressions seem *more* unethical? Perhaps some transgressions' severity only becomes apparent upon reflection, and perhaps multiple exposures facilitate such reflection. Although thinking about negative events typically makes people feel better through a process of sense-making (Wilson & Gilbert, 2008), future research may identify exceptions where the more you understand an event, the more awful you realize it is.

Like participants in the present experiments, people in the real world sometimes repeatedly encounter identical content from the same source. For example, social media users may repeatedly see the same headline about the Dao incident because the same friend posts it on both Facebook and Twitter, because it repeatedly appears in their news feed as an advertisement for the publication that ran the headline, or simply because they check the same news source's homepage multiple times in the same day. However, unlike participants in the present experiments, people in the real world sometimes repeatedly encounter the same content from *different* sources. For example, several different friends on social media may post the same headline about the Dao incident, or different news outlets may run the same picture or video of that incident. Future research should test the moral repetition effect when each repetition comes from a different source. On one hand, repetition could still reduce moral condemnation in such cases. Previously hearing about a transgression should dampen people's affective reaction to it, regardless of whether they have heard about it from the same source or from a different source. On the other hand, it is possible that the moral repetition effect would be attenuated or reversed in such cases. Upon hearing about the same transgression from multiple friends or news outlets,

people could infer that the transgression is receiving widespread attention because it is sparking widespread outrage. As a result, people might condemn it more harshly.

This reversal of the moral repetition effect could be particularly likely to occur when people repeatedly encounter *condemnation* – and not just a description – of the same transgression from multiple sources. For example, a social-media user might notice that multiple friends have not only posted the video of David Dao, but also called for a boycott of United Airlines. In such cases, repetition could *increase* moral condemnation via three mechanisms: The more people you hear condemn the transgression, the more you could *feel* the transgression is wrong (via emotional contagion; Centola & Macy, 2007), *think* the transgression is wrong (via persuasion; Harkins & Petty, 1981), or merely *say* the transgression is wrong (because of conformity pressures; Bond, 2005; Latané, 1981). Additional research is needed to test these predictions. It would also be interesting to explore what happens when repetition is accompanied by competing commentary – such as if some commentators express outrage about a transgression whereas others minimize its severity.

Whereas the present studies repeatedly exposed people to the same description of the same person committing the same transgression, people in the real world sometimes encounter different descriptions of different people committing similar transgressions. For example, a person might read three different news stories, each one about how a different public figure broke COVID-19 lockdown rules. One possibility is that such contexts amplify the moral repetition effect by strengthening the norm-perception mechanism. A person is more likely to infer that lockdown violations are common after reading about multiple officials violating lockdown than after repeatedly reading about the same incident in which one politician violated lockdown — and if such violations seem common, they may not seem so morally wrong

(Lindström et al., 2018). A different possibility is that such contexts attenuate the moral repetition effect by making a given violation seem symptomatic of a larger problem. When just one official has violated lockdown, people may experience only moderate outrage because they infer that the official exercised poor judgment; when a dozen officials have violated lockdown, people may experience strong outrage, inferring that we live in a society where lawmakers feel above the law. Future research is thus needed to test how repeated exposure to the same transgression affects moral judgments when each exposure involves a different person committing the same transgression.

Other questions ripe for future research include whether the moral repetition effect depends on the medium that depicts the transgression (e.g., text versus video) or the delay between first encountering a transgression and subsequently judging it. Given that other effects of repetition on judgment emerge after delays of weeks or even months, the moral repetition effect could as well (Bacon, 1979; Brown & Nix, 1996; Henderson et al., 2021; Henkel & Mattson, 2011; Nadarevic & Erdfelder, 2014).

It will also be interesting to understand how the moral repetition effect changes with the number of repetitions. Experiment 4 did not find a measurable difference between the effect size produced by one versus five repetitions, but it is possible that more repetitions would produce a larger effect. Other effects of repetition on judgment increase logarithmically in size with the number of repetitions (e.g., DiFonzo et al., 2016; Hassan & Barber, 2021; Montoya et al., 2017), and desensitization effects tend to be logarithmic as well (Rankin et al., 2009; Thompson & Spencer, 1966); the moral repetition effect may follow the same pattern. However, it is also possible that the transgressions used as stimuli were severe enough on average that repetition could only reduce unethicity judgments up to a point. If so, more repetitions might only

increase the effect size with stimuli that elicit a wider range of moral judgments (e.g., behaviors that could seem either ethical or unethical). Another possibility is that the relationship between the number of repetitions and moral judgments is U-shaped, such that after a certain point, more repetitions start to make transgressions seem *more* unethical. Perhaps, for example, people would start feeling annoyed after hearing about the same transgression 20 or more times, misattribute this negative affect to transgression itself, and thus rate the transgression as more unethical. This finding would be consistent with the general idea that different psychological processes can operate at different numbers of repetitions (e.g., Berlyne, 1970; Pillai & Fazio, 2021; Stang, 1974). Future research should thus test larger numbers of repetitions.

Practical Implications

The results cast new light on the relationship between moral outrage and the viral spread of content on social media. When content depicting a person's wrongdoing appears on social media, moral outrage may fuel its viral spread (Brady et al., 2019; Crockett, 2017). And the more people who encounter the content, the higher the number of outraged individuals will be, with potentially severe consequences for the wrongdoer. Yet the wider the content spreads, the more likely any one individual is to repeatedly encounter it, and the less outraged that individual may become. Thus, even as collective outrage grows, a given individual's outrage could decline.

This possibility creates a conundrum for journalists, activists, and whistleblowers who wish to alert the public to wrongdoing. On the one hand, the more people there are who know about the wrongdoing, the more likely it is that the wrongdoer will be held accountable. Voicing concerns about the transgression repeatedly will increase the likelihood that one's voice is heard. On the other hand, the present results suggest that the more times any one individual hears about

the wrongdoing, the less inclined they will be to condemn and punish the wrongdoer. The same techniques that stir up collective outrage could ironically increase individual apathy.

The moral repetition effect has other concerning implications. By making transgressions seem less unethical, repetition could contribute to the normalization of wrongdoing in organizations and society. Once a novel transgression becomes familiar, people may be less inclined to condemn it. Perhaps repeatedly hearing the same description of the same wrongdoing could even reduce individuals' compunction about committing that wrongdoing themselves. In this sense, the moral repetition effect could contribute to *moral disengagement* – a suite of psychological processes that allow people to violate their moral standards without sanctioning themselves (Bandura, 1999; Moore, 2015).

At the same time, the moral repetition effect could also have desirable effects. Moral outrage can fuel conflict, political polarization, and intolerance (Crockett, 2017; Skitka et al., 2004). By reducing such outrage in response to a particular event, the moral repetition effect could allow “cooler heads to prevail.” Additionally, by normalizing once-taboo behaviors, the moral repetition effect could facilitate social reforms that many will regard as moral progress. A majority of Americans from both major political parties now support same-sex marriage: 70% overall, up from just 27% in 1996 (McCarthy, 2021). There are numerous reasons for this attitudinal sea-change (e.g., Tankard & Paluck, 2017); perhaps one is that some people's initial moral outrage about same-sex marriage simply faded as the idea became more familiar.

Coda

The same wrongdoing seems less unethical after repeatedly reading the same description of it. This moral repetition effect is robust, replicable, generalizes broadly across the moral domain, and occurs because repetition dampens people's emotional reactions to wrongdoing.

Cartoonist Roz Chast (2010) depicts the news cycle with an illustration of a man repeatedly remarking, “What’s new? ... WHAT? ... OH, NO! ... Oh, well. ... What’s new?” Indeed, moral outrage at wrongdoing seems to wane over the course of the news cycle. Although this cycle of outrage may seem to simply reflect a short collective attention span, the present research suggests a different possibility – a transgression may seem a bit less reprehensible when we have heard about it before.

Research Context

The present research emerged from my program of work on when and why people grant others a license to transgress (e.g., Effron, 2018; Effron, Kakkar, et al., 2018; Effron & Monin, 2010; Effron, O’Connor, et al., 2018) and the observation that social media often repeatedly exposes people to same piece of “viral” content about others’ wrongdoing. After my colleague I and found that repeatedly encountering a fake-news article can make it seem less unethical to share (Effron & Raj, 2020), I began the present project to understand why. The initial results suggested a broader phenomenon – driven by a different mechanism – than the one we had described in the 2020 paper. Subsequent studies confirmed that the moral repetition effect occurs with transgressions across the moral domain (not just fake-news sharing), and the results converged on an affective mechanism. In addition to testing boundary conditions of the moral repetition effect, future work in this research stream will examine other processes that shape moral judgments in our digital world.

References

- Albrecht, S., & Carbon, C.-C. (2014). The Fluency Amplification Model: Fluent stimuli show more intense but not evidently more positive evaluations. *Acta Psychologica, 148*, 195-203.
- Allport, F. H., & Lepkin, M. (1945). Wartime rumors of waste and special privilege: why some people believe them. *The Journal of Abnormal and Social Psychology, 40*(1), 3.
- Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychological Review, 13*(3), 219-235. <https://doi.org/10.1177/1088868309341564>
- Avramova, Y. R., & Inbar, Y. (2013). Emotion and moral judgment. *Wiley Interdisciplinary Reviews: Cognitive Science, 4*(2), 169-178.
- Bacon, F. T. (1979). Credibility of repeated statements: Memory for trivia. *Journal of Experimental Psychology: Human Learning and Memory, 5*(3), 241.
- Bandura, A. (1999). Moral disengagement in the perpetration of inhumanities. *Personality and Social Psychology Review, 3*(3), 193-209.
- Batson, C. D., Kennedy, C. L., Nord, L.-A., Stocks, E. L., Fleming, D. Y. A., Marzette, C. M., Lishner, D. A., Hayes, R. E., Kolchinsky, L. M., & Zenger, T. (2007). Anger at unfairness: Is it moral outrage? *European Journal of Social Psychology, 37*(6), 1272-1272.
- Berlyne, D. E. (1970). Novelty, complexity, and hedonic value. *Perception & Psychophysics, 8*(5), 279-286.
- Bloom, P. (2017). Empathy and its discontents. *Trends in Cognitive Sciences, 21*(1), 24-31.
- Boisseau, R. P., Vogel, D., & Dussutour, A. (2016). Habituation in non-neural organisms: Evidence from slime moulds. *Proceedings of the Royal Society B: Biological Sciences, 283*(1829), 20160446.
- Bond, R. (2005). Group size and conformity. *Group Processes & Intergroup Relations, 8*(4), 331-354. <https://doi.org/10.1177/1368430205056464>
- Bornstein, R. F., & D'Agostino, P. R. (1994). The attribution and discounting of perceptual fluency: Preliminary tests of a perceptual fluency/attributional model of the mere exposure effect. *Social Cognition, 12*(2), 103-128.
- Bornstein, R. F., & D'Agostino, P. R. (1992). Stimulus recognition and the mere exposure effect. *Journal of personality and social psychology, 63*(4), 545.
- Bradley, M. M., Lang, P. J., & Cuthbert, B. N. (1993). Emotion, novelty, and the startle reflex: habituation in humans. *Behavioral neuroscience, 107*(6), 970.
- Brady, W. J., Crockett, M., & Van Bavel, J. J. (2019). The MAD Model of Moral Contagion: The Role of Motivation, Attention, and Design in the Spread of Moralized Content Online. *Perspectives on Psychological Science, 17*45691620917336.
- Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., & Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences, 114*(28), 7313-7318.
- Brauer, M., Judd, C. M., & Gliner, M. D. (1995). The effects of repeated expressions on attitude polarization during group discussions. *Journal of personality and social psychology, 68*(6), 1014.

- Brickman, P., Redfield, J., Harrison, A. A., & Crandall, R. (1972). Drive and predisposition as factors in the attitudinal effects of mere exposure. *Journal of Experimental Social Psychology*, 8(1), 31-44.
- Brown, A. S., & Nix, L. A. (1996). Turning lies into truths: Referential validation of falsehoods. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(5), 1088.
- Buchli, D. R. (2019). False fame effects emerge following an internal context shift and are abolished with contextual reinstatement. *Memory*, 27(8), 1024-1033.
- Bullock, J. G., Green, D. P., & Ha, S. E. (2010). Yes, but what's the mechanism? (Don't expect an easy answer). *Journal of personality and social psychology*, 98(4), 550-558.
- Cameron, C. D. (2017). Compassion collapse: Why we are numb to numbers? In E. M. Seppälä, E. Simon-Thomas, S. L. Brown, M. C. Worline, C. D. Cameron, & J. R. Doty (Eds.), *The Oxford Handbook of Compassion Science* (Vol. 261, pp. 271). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780190464684.013.20>
- Cameron, C. D., & Payne, B. K. (2011). Escaping affect: How motivated emotion regulation creates insensitivity to mass suffering. *Journal of personality and social psychology*, 100(1), 1.
- Campbell, T., O'Brien, E., Van Boven, L., Schwarz, N., & Ubel, P. (2014). Too much experience: A desensitization bias in emotional perspective taking. *Journal of personality and social psychology*, 106(2), 272-285. <https://doi.org/10.1037/a0035148>
- Carnagey, N. L., Anderson, C. A., & Bushman, B. J. (2007). The effect of video game violence on physiological desensitization to real-life violence. *Journal of Experimental Social Psychology*, 43(3), 489-496.
- Centola, D., & Macy, M. (2007). Complex contagions and the weakness of long ties. *American Journal of Sociology*, 113(3), 702-734.
- Chast, R. (2010, November 15). The News Cycle. *The New Yorker*. <https://www.cartooncollections.com/cartoon?searchID=CC133907>
- Cho, H.-C., & Abe, S. (2013). Is two-tailed testing for directional research hypotheses tests legitimate? *Journal of Business Research*, 66(9), 1261-1266. <https://doi.org/https://doi.org/10.1016/j.jbusres.2012.02.023>
- Clifford, S., Iyengar, V., Cabeza, R., & Sinnott-Armstrong, W. (2015). Moral foundations vignettes: A standardized stimulus database of scenarios based on moral foundations theory. *Behavior Research*, 47, 1178-1198. <https://doi.org/10.3758/s13428-014-0551-2>
- Clore, G. L., Gasper, K., & Garvin, E. (2001). Affect as information. In *Handbook of affect and social cognition* (pp. 121-144). Erlbaum.
- Corneille, O., Mierop, A., & Unkelbach, C. (2020). Repetition increases both the perceived truth and fakeness of information: An ecological account. *Cognition*, 205, 104470.
- Crockett, M. J. (2017). Moral outrage in the digital age. *Nature Human Behaviour*, 1(11), 769-771.
- Dechêne, A., Stahl, C., Hansen, J., & Wänke, M. (2010). The truth about the truth: A meta-analytic review of the truth effect. *Personality and Social Psychological Review*, 14(2), 238-257. <https://doi.org/10.1177/1088868309352251>
- DiFonzo, N., Beckstead, J. W., Stupak, N., & Walders, K. (2016). Validity judgments of rumors heard multiple times: The shape of the truth effect. *Social Influence*, 11(1), 22-39.
- Dijksterhuis, A., & Smith, P. K. (2002). Affective habituation: Subliminal exposure to extreme stimuli decreases their extremity. *Emotion*, 2(3), 203-214. <https://doi.org/10.1037/1528-3542.2.3.203>

- Effron, D. A. (2018). It could have been true: How counterfactual thoughts reduce condemnation of falsehoods and increase political polarization. *Personality and Social Psychology Bulletin*, 44(5), 729-745. <https://doi.org/10.1177/0146167217746152>
- Effron, D. A. (2021). *Supplemental materials for The Moral Repetition Effect*. https://osf.io/6xpq8/?view_only=eb17d8e867aa4fe7b5007baa17086859
- Effron, D. A., Kakkar, H., & Knowles, E. D. (2018). Group cohesion benefits individuals who express prejudice, but harms their group. *Journal of Experimental Social Psychology*, 79, 239-251. <https://doi.org/10.1016/j.jesp.2018.08.002>
- Effron, D. A., & Monin, B. (2010). Letting people off the hook: When do good deeds excuse transgressions? *Personality and Social Psychology Bulletin*, 36, 1618-1634. <https://doi.org/10.1177/0146167210385922>
- Effron, D. A., O'Connor, K., Leroy, H., & Lucas, B. J. (2018). From inconsistency to hypocrisy: When does “saying one thing but doing another” invite condemnation? *Research in Organizational Behavior*, 38, 61-75. <https://doi.org/10.1016/j.riob.2018.10.003>
- Effron, D. A., & Raj, M. (2020). Misinformation and morality: Encountering fake-news headlines makes them seem less unethical to publish and share. *Psychological Science*, 31(1), 75-87. <https://doi.org/10.1177/0956797619887896>
- Eriksson, K., Strimling, P., & Coultas, J. C. (2015). Bidirectional associations between descriptive and injunctive norms. *Organizational Behavior and Human Decision Processes*, 129, 59-69.
- Fazio, L. K., Brashier, N. M., Payne, B. K., & Marsh, E. J. (2015). Knowledge does not protect against illusory truth. *Journal of Experimental Psychology: General*, 144(5), 993-1002. <https://doi.org/10.1037/xge0000098>
- Fernandes, S., Kapoor, H., & Karandikar, S. (2017). Do we gossip for moral reasons? The intersection of moral foundations and gossip. *Basic and Applied Social Psychology*, 39(4), 218-230.
- Fiedler, K., Schott, M., & Meiser, T. (2011). What mediation analysis can (not) do. *Journal of Experimental Social Psychology*, 47(6), 1231-1236.
- Frederick, S., & Loewenstein, G. (1999). Hedonic adaptation. In D. Kahneman, E. Diener, & N. Schwarz (Eds.), *Well-being: The foundations of hedonic psychology* (pp. 302-329). Russel Sage Foundation.
- Fuhrman, R. W., Bodenhausen, G. V., & Lichtenstein, M. (1989). On the trait implications of social behaviors: Kindness, intelligence, goodness, and normality ratings for 400 behavior statements. *Behavior Research Methods, Instruments, & Computers*, 21(6), 587-597.
- Funk, J. B., Buchman, D. D., Jenks, J., & Bechtoldt, H. (2003). Playing violent video games, desensitization, and moral evaluation in children. *Journal of Applied Developmental Psychology*, 24(4), 413-436.
- Gawronski, B., Conway, P., Armstrong, J., Friesdorf, R., & Hütter, M. (2018). Effects of incidental emotions on moral dilemma judgments: An analysis using the CNI model. *Emotion*, 18(7), 989.
- Gino, F., Shu, L. L., & Bazerman, M. H. (2010). Nameless + harmless = blameless: When seemingly irrelevant factors influence judgment of (un)ethical behavior [Peer Reviewed]. *Organizational Behavior and Human Decision Processes*, 111(2), 93-101. <https://doi.org/10.1016/j.obhdp.2009.11.001>

- Goldstein, M. (2017, December 20). Biggest travel story of 2017: The bumping and beating of Dr. David Dao. *Forbes*.
<https://www.forbes.com/sites/michaelgoldstein/2017/12/20/biggest-travel-story-of-2017-the-bumping-and-beating-of-doctor-david-dao/?sh=7e0113ddf61f>
- Graf, L. K., Mayer, S., & Landwehr, J. R. (2018). Measuring processing fluency: One versus five items. *Journal of Consumer Psychology*, 28(3), 393-411.
- Graham, J., Haidt, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S., & Ditto, P. H. (2013). Moral foundations theory: The pragmatic validity of moral pluralism. *Advances in experimental social psychology*, 47, 55-130.
- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology*, 96(5), 1029.
<https://doi.org/0.1037/a0015141>
- Gray, K., & Keeney, J. E. (2015). Impure or just weird? Scenario sampling bias raises questions about the foundation of morality. *Social Psychological and Personality Science*, 6(8), 859-868.
- Green, P., & MacLeod, C. J. (2015). SIMR: An R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution*, 7(4), 493-498.
<https://doi.org/https://doi.org/10.1111/2041-210X.12504>
- Greene, J., & Haidt, J. (2002). How (and where) does moral judgment work? *Trends in Cognitive Sciences*, 6(12), 517-523.
- Grush, J. E. (1976). Attitude formation and mere exposure phenomena: A nonartificial explanation of empirical findings. *Journal of personality and social psychology*, 33(3), 281.
- Gummerum, M., Van Dillen, L. F., Van Dijk, E., & López-Pérez, B. (2016). Costly third-party interventions: The role of incidental anger and attention focus in punishment of the perpetrator and compensation of the victim. *Journal of Experimental Social Psychology*, 65, 94-104.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108, 814-834.
- Haidt, J., & Kesebir, S. (2010). Morality. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.), *Handbook of Social Psychology* (5th ed., pp. 797-832). Wiley.
- Harkins, S. G., & Petty, R. E. (1981). The multiple source effect in persuasion: The effects of distraction. *Personality and Social Psychology Bulletin*, 7(4), 627-635.
- Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior*, 16(1), 107-112.
[https://doi.org/10.1016/S0022-5371\(77\)80012-1](https://doi.org/10.1016/S0022-5371(77)80012-1)
- Hassan, A., & Barber, S. J. (2021). The effects of repetition frequency on the illusory truth effect. *Cognitive Research: Principles and Implications*, 6(1), 1-12.
- Henderson, E. L., Simons, D. J., & Barr, D. J. (2021). The trajectory of truth: A longitudinal study of the illusory truth effect. *Journal of Cognition*, 4(1), 1-23.
<https://doi.org/10.5334/joc.161>
- Henkel, L. A., & Mattson, M. E. (2011). Reading is believing: The truth effect and source credibility. *Consciousness and Cognition*, 20(4), 1705-1721.
- Hoffman, A. M., & Kaire, J. (2020). Comfortably numb: Effects of prolonged media coverage. *Journal of Conflict Resolution*, 64(9), 1666-1692.

- Iyer, A., Schmader, T., & Lickel, B. (2007). Why Individuals Protest the Perceived Transgressions of Their Country: The Role of Anger, Shame, and Guilt. *Personality and Social Psychology Bulletin*, 33(4), 572-572.
- Jacoby, L. L., Kelley, C., Brown, J., & Jasechko, J. (1989). Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. *Journal of personality and social psychology*, 56(3), 326.
- Jacoby, L. L., Woloshyn, V., & Kelley, C. (1989). Becoming famous without being recognized: Unconscious influences of memory produced by dividing attention. *Journal of Experimental Psychology: General*, 118(2), 115.
- Jarudi, I., Kreps, T., & Bloom, P. (2008). Is a refrigerator good or evil? The moral evaluation of everyday objects. *Social Justice Research*, 21(4), 457-469.
- Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of personality and social psychology*, 103(1), 54-69. <https://doi.org/10.1037/a0028347>
- Kliemann, D., Young, L., Scholz, J., & Saxe, R. (2008). The influence of prior record on moral judgment. *Neuropsychologia*, 46(12), 2949-2957.
- Kogut, T., & Ritov, I. (2005). The “identified victim” effect: An identified group, or just a single individual? *Journal of Behavioral Decision Making*, 18(3), 157-167.
- Kogut, T., & Ritov, I. (2015). Target dependent ethics: discrepancies between ethical decisions toward specific and general targets. *Current Opinion in Psychology*, 6, 145-149.
- Kwan, L. Y. Y., Yap, S., & Chiu, C.-y. (2015). Mere exposure affects perceived descriptive norms: Implications for personal preferences and trust. *Organizational Behavior and Human Decision Processes*, 129, 48-58. <https://doi.org/https://doi.org/10.1016/j.obhdp.2014.12.002>
- Laham, S. M., Alter, A. L., & Goodwin, G. P. (2009). Easy on the mind, easy on the wrongdoer: Discrepantly fluent violations are deemed less morally wrong. *Cognition*, 112(3), 462-466.
- Latané, B. (1981). The psychology of social impact. *American Psychologist*, 36, 343-356.
- Leventhal, A. M., Martin, R. L., Seals, R. W., Tapia, E., & Rehm, L. P. (2007). Investigating the dynamics of affect: Psychological mechanisms of affective habituation to pleasurable stimuli. *Motivation and Emotion*, 31(2), 145-157. <https://doi.org/10.1007/s11031-007-9059-8>
- Linden, A. (2021). *RMCORR: Stata module to compute a correlation for data with repeated measures*. In *Statistical Software Components, Boston College Department of Economics* (Version S458971) <https://ideas.repec.org/c/boc/bocode/s458971.html>
- Lindström, B., Jangard, S., Selbing, I., & Olsson, A. (2018). The role of a “common is moral” heuristic in the stability and change of moral norms. *Journal of Experimental Psychology: General*, 147(2), 228.
- Lovakov, A., & Agadullina, E. R. (2021). Empirically derived guidelines for effect size interpretation in social psychology. *European Journal of Social Psychology*, 51(3), 485-504. <https://doi.org/10.1002/ejsp.2752>
- Marotti, A., & Zumbach, L. (2017, April 11, 2017). Video shows United Airlines' passenger dragged off plane. *Chicago Tribune*. <https://www.chicagotribune.com/business/ct-united-drag-passenger-0411-biz-20170410-story.html>

- Martel, C., Pennycook, G., & Rand, D. G. (2020). Reliance on emotion promotes belief in fake news. *Cognitive Research: Principles and Implications*, 5(1), 1-20.
- McArthur, L. Z., & Solomon, L. K. (1978). Perceptions of an aggressive encounter as a function of the victim's salience and the perceiver's arousal. *Journal of personality and social psychology*, 36(11), 1278.
- McCarthy, J. (2021, June 8). Record-high 70% in U.S. support same-sex marriage. *Gallup*. <https://news.gallup.com/poll/350486/record-high-support-same-sex-marriage.aspx>
- McGraw, K. M. (1985). Subjective probabilities and moral judgments. *Journal of Experimental Social Psychology*, 21(6), 501-518.
- Meskin, A., Phelan, M., Moore, M., & Kieran, M. (2013). Mere exposure to bad art. *British Journal of Aesthetics*, 53(2), 139-164.
- Monroe, A. E., Dillon, K. D., Guglielmo, S., & Baumeister, R. F. (2018). It's not what you do, but what everyone else does: On the role of descriptive norms and subjectivism in moral judgment. *Journal of Experimental Social Psychology*, 77, 1-10. <https://doi.org/10.1016/j.jesp.2018.03.010>
- Montoya, R. M., Horton, R. S., Vevea, J. L., Citkowicz, M., & Lauber, E. A. (2017). A re-examination of the mere-exposure effect: The influence of repeated exposure on recognition, familiarity, and liking. *Psychological Bulletin*, 143(5), 459-498.
- Moore, C. (2015). Moral disengagement. *Current Opinion in Psychology*, 6, 199-204.
- Mrkva, K., & Van Boven, L. (2020). Salience theory of mere exposure: Relative exposure increases liking, extremity, and emotional intensity. *Journal of personality and social psychology*.
- Nadarevic, L., & Erdfelder, E. (2014). Initial judgment task and delay of the final validity-rating task moderate the truth effect. *Consciousness and Cognition*, 23, 74-84. <https://doi.org/https://doi.org/10.1016/j.concog.2013.12.002>
- Nelissen, R. M., & Zeelenberg, M. (2009). Moral emotions as determinants of third-party punishment: Anger, guilt and the functions of altruistic sanctions. *Judgment and Decision Making*, 4(7), 543.
- Nelson, L. D., & Meyvis, T. (2008). Interrupted consumption: Disrupting adaptation to hedonic experiences. *Journal of Marketing Research*, 45(6), 654-664.
- Newman, E. J., Garry, M., Unkelbach, C., Bernstein, D. M., Lindsay, D. S., & Nash, R. A. (2015). Truthiness and falsiness of trivia claims depend on judgmental contexts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(5), 1337-1348. <https://doi.org/10.1037/xlm0000099>
- Nordgren, L. F., & Dijksterhuis, A. (2008). The devil is in the deliberation: Thinking too much reduces preference consistency. *Journal of Consumer Research*, 36(1), 39-46.
- Nordgren, L. F., & McDonnell, M.-H. M. (2011). The scope-severity paradox: Why doing more harm is judged to be less harmful. *Social Psychological and Personality Science*, 2(1), 97-102.
- Oppenheimer, D. M. (2008). The secret life of fluency. *Trends in Cognitive Sciences*, 12(6), 237-241.
- Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70, 153-163.

- Pennycook, G., Cannon, T., & Rand, D. G. (2018). Prior exposure increases perceived accuracy of fake news. *Journal of Experimental Psychology: General*, *147*, 1865–1880. <https://doi.org/10.1037/xge0000465>
- Pillai, R. M., & Fazio, L. K. (2021). The effects of repeating false and misleading information on belief. *Wiley Interdisciplinary Reviews: Cognitive Science*, *12*(6). <https://doi.org/doi.org/10.1002/wcs.1573>
- Prentice, D. A., & Miller, D. T. (1992). When small effects are impressive. *Psychological Bulletin*, *112*(1), 160.
- Rankin, C. H., Abrams, T., Barry, R. J., Bhatnagar, S., Clayton, D. F., Colombo, J., Coppola, G., Geyer, M. A., Glanzman, D. L., & Marsland, S. (2009). Habituation revisited: An updated and revised description of the behavioral characteristics of habituation. *Neurobiology of Learning and Memory*, *92*(2), 135-138.
- Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing Fluency and Aesthetic Pleasure: Is Beauty in the Perceiver's Processing Experience? *Personality and Social Psychology Review*, *8*(4), 364-382.
- Reber, R., Winkielman, P., & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science*, *9*(1), 45-48. <https://doi.org/10.1111/1467-9280.00008>
- Reber, R., Wurtz, P., & Zimmermann, T. D. (2004). Exploring “fringe” consciousness: The subjective experience of perceptual fluency and its objective bases. *Consciousness and Cognition*, *13*(1), 47-60.
- Rothbaum, B. O., & Schwartz, A. C. (2002). Exposure therapy for posttraumatic stress disorder. *American journal of psychotherapy*, *56*(1), 59-75.
- Sadler, O., & Tesser, A. (1973). Some effects of salience and time upon interpersonal hostility and attraction during social isolation. *Sociometry*, 99-112.
- Salerno, J. M., & Peter-Hagene, L. C. (2013). The interactive effect of anger and disgust on moral outrage and judgments. *Psychological Science*, *24*(10), 2069-2078.
- Schacter, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*(3), 501.
- Schwarz, N. (2004). Metacognitive experiences in consumer judgment and decision making. *Journal of Consumer Psychology*, *14*(4), 332-348.
- Schwarz, N. (2011). Feelings-as-information theory. In P. A. M. V. Lange, A. Kruglanski, & E. T. Higgins (Eds.), *Handbook of theories of social psychology* (Vol. 1, pp. 289-308). Sage.
- Schwarz, N., Sanna, L. J., Skurnik, I., & Yoon, C. (2007). Metacognitive experiences and the intricacies of setting people straight: Implications for debiasing and public information campaigns. In M. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 39, pp. 127-161). [https://doi.org/10.1016/S0065-2601\(06\)39003-X](https://doi.org/10.1016/S0065-2601(06)39003-X)
- Shidlovski, D., Schul, Y., & Mayo, R. (2014). If I imagine it, then it happened: The implicit truth value of imaginary representations. *Cognition*, *133*(3), 517-529.
- Skitka, L. J., Bauman, C. W., & Mullen, E. (2004). Political tolerance and coming to psychological closure following the September 11, 2001, terrorist attacks: An integrative approach. *Personality and Social Psychology Bulletin*, *30*(6), 743-756.
- Slovic, P. (2007). “If I look at the mass I will never act.” Psychic numbing and genocide. *Judgment and Decision Making*, *2*, 29-85.
- Small, D. A., & Loewenstein, G. (2003). Helping a victim or helping the victim: Altruism and identifiability. *Journal of Risk and uncertainty*, *26*(1), 5-16.

- Small, D. A., & Loewenstein, G. (2005). The devil you know: The effects of identifiability on punishment. *Journal of Behavioral Decision Making*, 18(5), 311-318.
- Small, D. A., Loewenstein, G., & Slovic, P. (2007). Sympathy and callousness: The impact of deliberative thought on donations to identifiable and statistical victims. *Organizational Behavior and Human Decision Processes*, 102(2), 143-153.
- Spencer, S. J., Zanna, M. P., & Fong, G. T. (2005). Establishing a causal chain: Why experiments are often more effective than mediational analyses in examining psychological processes. *Journal of personality and social psychology*, 89(6), 845.
- Stang, D. J. (1974). Methodological factors in mere exposure research. *Psychological Bulletin*, 81(12), 1014.
- Stang, D. J. (1975). Effects of "mere exposure" on learning and affect. *Journal of personality and social psychology*, 31(1), 7.
- StataCorp. (2019). *Stata 16 Base Reference Manual*. StataCorp LLC.
- StataCorp. (2021). Example 42g — One- and two-level mediation models (multilevel). In *Stata 17 Base Reference Manual* (pp. 437-445). StataCorp LLC.
- Strohinger, N., Lewis, R. L., & Meyer, D. E. (2011). Divergent effects of different positive emotions on moral judgment. *Cognition*, 119(2), 295-300.
- Swap, W. C. (1977). Interpersonal attraction and repeated exposure to rewarders and punishers. *Personality and Social Psychology Bulletin*, 3(2), 248-251.
- Tankard, M. E., & Paluck, E. L. (2017). The effect of a Supreme Court decision regarding gay marriage on social norms and personal attitudes. *Psychological Science*, 28(9), 1334-1344. <https://doi.org/10.1177/0956797617709594>
- Thompson, R. F. (2009). Habituation: A history. *Neurobiology of Learning and Memory*, 92(2), 127.
- Thompson, R. F., & Spencer, W. A. (1966). Habituation: a model phenomenon for the study of neuronal substrates of behavior. *Psychological Review*, 73(1), 16.
- Topolinski, S., & Strack, F. (2010). False fame prevented: Avoiding fluency effects without judgmental correction. *Journal of personality and social psychology*, 98(5), 721.
- Turiel, E. (1983). *The development of social knowledge*. Cambridge University Press.
- Unkelbach, C., & Greifeneder, R. (2013). A general model of fluency effects in judgment and decision making. In C. Unkelbach & R. Greifeneder (Eds.), *The experience of thinking: How the fluency of mental processes influences cognition and behaviour* (pp. 11-32). Psychology Press.
- Unkelbach, C., & Stahl, C. (2009). A multinomial modeling approach to dissociate different components of the truth effect. *Consciousness and Cognition*, 18(1), 22-38.
- Valdesolo, P., & DeSteno, D. (2006). Manipulations of emotional context shape moral judgment. *Psychological Science*, 17, 476-477.
- Västfall, D., Slovic, P., Mayorga, M., & Peters, E. (2014). Compassion fade: Affect and charity are greatest for a single child in need. *PloS one*. <https://doi.org/10.1371/journal.pone.0100115>
- Weaver, K., Garcia, S. M., Schwarz, N., & Miller, D. T. (2007). Inferring the popularity of an opinion from its familiarity: A repetitive voice can sound like a chorus. *Journal of personality and social psychology*, 92(5), 821-833.
- Weeks, W. A., Longenecker, J. G., McKinney, J. A., & Moore, C. W. (2005). The role of mere exposure effect on ethical tolerance: A two-study approach. *Journal of Business Ethics*, 58(4), 281-294.

- Welch, M. R., Xu, Y., Bjarnason, T., Petee, T., O'Donnell, P., & Magro, P. (2005). "But everybody does it...": The effects of perceptions, moral pressures, and informal sanctions on tax cheating. *Sociological Spectrum*, 25(1), 21-52.
- Whittlesea, B. W., Jacoby, L. L., & Girard, K. (1990). Illusions of immediate memory: Evidence of an attributional basis for feelings of familiarity and perceptual quality. *Journal of Memory and Language*, 29(6), 716-732.
- Wilson, T. D., & Gilbert, D. T. (2008). Explaining away: A model of affective adaptation. *Perspectives on Psychological Science*, 3(5), 370-386. <https://doi.org/10.1111/j.1745-6924.2008.00085.x>
- Winkielman, P., Huber, D. E., Kavanagh, L., & Schwarz, N. (2012). Fluency of consistency: When thoughts fit nicely and flow smoothly. In B. Gawronski & F. Strack (Eds.), *Cognitive consistency: A fundamental principle in social cognition* (pp. 89-111). Guilford press.
- Winkielman, P., Schwarz, N., Fazendeiro, T., & Reber, R. (2003). The hedonic marking of processing fluency: Implications for evaluative judgment. In J. Munsch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion* (Vol. 189, pp. 217). Erlbaum.
- Wolpe, J. (1961). The systematic desensitization treatment of neuroses. *Journal of Nervous and Mental Disease*, 132, 189-203. <https://doi.org/10.1097/00005053-196103000-00001>
- Wright, C. I., Fischer, H., Whalen, P. J., McInerney, S. C., Shin, L. M., & Rauch, S. L. (2001). Differential prefrontal cortex and amygdala habituation to repeatedly presented emotional stimuli. *NeuroReport*, 12(2), 379-383.
- Wurtz, P., Reber, R., & Zimmermann, T. D. (2008). The feeling of fluent perception: A single experience from multiple asynchronous sources. *Consciousness and Cognition*, 17(1), 171-184.
- Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of personality and social psychology*, 9, 1-27. <https://doi.org/10.1037/h0025848>
- Zajonc, R. B. (2001). Mere exposure: A gateway to the subliminal. *Current Directions in Psychological Science*, 10(6), 224-228.
- Zajonc, R. B., Markus, H., & Wilson, W. R. (1974). Exposure effects and associative learning. *Journal of Experimental Social Psychology*, 10(3), 248-263.
- Zhao, X., Lynch, J. G., & Chen, Q. (2010). Reconsidering Baron and Kenny: Myths and truths about mediation analysis. *Journal of Consumer Research*, 37(2), 197-206.

Tables

Table 1

Fake-News Headlines Used as Stimuli in Experiments 1a and 1b

A worker at a British dairy that produces Anchor Butter contaminated the company's products with HIV-infected blood
Apple manufactured an Airpod that exploded in a man's ear
Coca-Cola and Nestle are negotiating with the Brazilian government to privatize the Guarani Aquifer in South America.
Appliance manufacturer Dyson Ltd made a vacuum cleaner that severely injured a man's penis
Facebook vice president Jeff Rothschild said "In order to finalize the New World Order process we need a 3rd world war to exterminate 90% of the world's population."
Fisher-Price manufactured a "Tiny Toker" toyset featuring a toy bong and other marijuana paraphernalia
H&M created an "Equally Racist" collection featuring children of various ethnicities in offensive clothing
Heineken once published an advertisement that showed a toddler drinking a beer and boasted about having the youngest customers in the business
The L'Oreal cosmetics company severely injured a monkey at one of its testing facilities
A McDonald's in Birmingham stored horse meat in their freezers
PepsiCo adds the toxic preservative formaldehyde to their "Naked" brand juices and smoothies
The company that manufactured Nutella-brand spread has cancer-causing chemicals in the factories where Nutella is made
Snapchat is using its "Lenses" photo and video filtering technology to create a facial recognition database for use by law enforcement agencies
Through its app, TikTok allows unauthorized users to obtain the personal contact details of other users who have not made their details public
An Uber driver in Tampa, FL, attempted to kidnap a female passenger to sell to sex traffickers
A United Airlines flight attendant slapped a 7 month-old baby in the face for crying during a flight

Table 2

Regression Results for Experiments 1a and 1b

		Main Effects Model						Interaction Model					
		<i>b</i>	SE(<i>b</i>)	<i>z</i>	<i>p</i>	95% CI (<i>b</i>)		<i>b</i>	SE(<i>b</i>)	<i>z</i>	<i>p</i>	95% CI (<i>b</i>)	
Fixed effects													
Previously-seen headline	Experiment 1a	-3.04	0.59	-5.11	< .001	-4.21	-1.88	-2.94	0.84	-3.50	< .001	-4.58	-1.29
	Experiment 1b	-3.14	0.34	-9.16	<u>< .001</u>	-3.81	-2.47	-3.64	0.48	-7.58	< .001	-4.58	-2.70
Judge-sharing condition	Experiment 1a	-1.80	2.02	-0.89	.374	-5.76	2.17	-1.69	2.11	-0.80	.423	-5.83	2.45
	Experiment 1b	-5.37	1.25	-4.30	< .001	-7.81	-2.92	-5.88	1.29	-4.54	< .001	-8.42	-3.35
Interaction	Experiment 1a							-0.21	1.19	-0.18	.860	-2.55	2.13
	Experiment 1b							1.02	0.69	1.49	<u>.930</u>	-0.32	2.37
Intercept	Experiment 1a	77.37	2.77	27.89	< .001	71.94	82.81	77.32	2.79	27.71	< .001	71.85	82.79
	Experiment 1b	78.56	2.58	30.37	< .001	73.49	83.63	78.91	2.59	30.40	< .001	73.73	83.89
Random effects													
Participant intercept Variance	Experiment 1a	300.86	26.02			253.95	356.45	300.86	26.02			253.94	356.45
	Experiment 1b	338.17	16.88			306.65	372.93	338.17	16.88			306.65	372.93
Item intercept variance	Experiment 1a	89.23	32.07			44.11	180.47	89.22	32.07			44.11	180.46
	Experiment 1b	94.40	33.54			47.04	189.43	94.42	33.55			47.05	189.46
Residual Variance	Experiment 1a	453.95	9.27			436.14	472.49	453.95	9.27			436.14	472.49
	Experiment 1b	441.24	5.26			431.06	451.67	441.17	5.26			430.99	451.60

Note. Underlined *p*-values are pre-registered one-tailed tests of Experiment 1b’s directional predictions.

Table 3

Regression Results for Experiment 2

	<i>b</i>	SE(<i>b</i>)	<i>z</i>	<i>p</i>	95% CI (<i>b</i>)	
Fixed effects						
Previously-seen headline	-2.29	0.41	-5.58	<u>< .001</u>	-3.10	-1.49
Intercept	72.97	1.92	38.05	< .001	69.21	76.72
Random effects						
Participant intercept variance	189.32	17.96			157.20	228.00
Item intercept variance	84.87	22.26			50.77	141.90
Residual variance	304.67	5.16			294.73	314.96

Note. The results are from a mixed regression model with fixed effects for headline type (1 = previously seen, 0 = new) and random intercepts for participants and items. The conclusions were identical when I followed the pre-registered plan to treat item as a fixed effect; see Footnote 2. The underlined *p*-value is from a pre-registered one-tailed test, but it reached the same significance level with a two-tailed test.

Table 4

Examples of Experiment 3's Stimuli from Each Moral Foundation

Moral Foundation	Number of Stimuli	Example
Fairness	4	You see a student copying a classmate's answer sheet on a makeup final exam.
Liberty	4	You see a man telling his fiancé that she has to switch to his political party.
Authority	4	You see a girl ignoring her father's orders by taking the car after her curfew.
Loyalty	4	You see an employee joking with competitors about how bad his company did last year.
Sanctity	4	You see a man having sex with a frozen chicken before cooking it for dinner.
Care (emotional)	4	You see a teenage boy chuckling at an amputee he passes by while on the subway.
Care (physical, animal)	2	You see a woman swerving her car in order to intentionally run over a squirrel.
Care (physical, human)	2	You see a teacher hitting a student's hand with a ruler for falling asleep in class.

Note. The stimuli were selected from Clifford et al. (2015).

Table 5: Experiment 3’s Regression Results

Dependent Measure	Model for Main Analysis					Model Controlling for Unusualness Judgments						
	<i>b</i>	SE(<i>b</i>)	<i>z</i>	<i>p</i>	95% CI (<i>b</i>)	<i>b</i>	SE(<i>b</i>)	<i>z</i>	<i>p</i>	95% CI (<i>b</i>)		
Unethicality												
Fixed effects												
Previously-seen vignette	-1.39	0.55	-2.53	<u>.006</u>	-2.46	-0.31	-0.93	0.53	-1.76	<u>.039</u>	-1.96	0.10
Unusual							0.25	0.01	23.01	< .001	0.23	0.27
Intercept	69.13	2.64	26.16	< 0.001	63.95	74.31	54.97	2.61	21.03	< 0.001	49.85	60.09
Random effects												
Participant intercept variance	142.58	14.03			117.57	172.92	139.13	13.62			114.84	168.55
Item intercept variance	176.60	47.76			103.94	300.06	162.36	43.92			95.55	275.89
Residual Variance	564.62	9.37			546.54	583.29	526.39	8.74			509.53	543.80
Punishment												
Fixed effects												
Previously-seen vignette	-2.07	0.56	-3.69	<u>< 0.001</u>	-3.16	-0.97	-1.68	0.55	-3.08	<u>.001</u>	-2.75	-0.61
Unusual							0.22	0.01	19.39	< .001	0.20	0.24
Intercept	53.72	2.85	18.83	< 0.001	48.13	59.31	41.33	2.88	14.35	< 0.001	35.69	46.98
Random effects												
Participant intercept variance	228.53	21.52			190.01	274.87	230.16	21.58			191.53	276.59
Item intercept variance	199.74	53.98			117.60	339.23	192.81	52.10			113.53	327.44
Residual Variance	592.13	9.83			573.17	611.72	563.00	9.345			544.97	581.62
Unusual												
Fixed effects												
Previously-seen vignette	-1.70	0.56	-3.06	.002	-2.79	-0.61						
Intercept	56.20	2.91	19.29	< 0.001	50.49	61.91						
Random effects												
Participant intercept variance	128.90	12.92			105.91	156.88						
Item intercept variance	219.96	59.37			129.59	373.33						
Residual Variance	585.19	9.72			566.45	604.55						

Note. All *p*-values are from two-tailed significance tests. The main text reports pre-registered one-tailed tests of directional predictions. The results were identical when treating items as fixed effects instead of random effects, as pre-registered; see Footnote 2. Underlined *p*-values are from pre-registered one-tailed tests.

Table 6

Real Transgression News Headlines Used as Stimuli in Experiment 4

Airbnb farmed Craigslist to grow its listings
Apple deliberately slows down iPhones as they get older, forcing users to upgrade
Exxon has misled Americans on climate change for decades
FirstEnergy CEO involved in scheme to bribe Ohio House Speaker
Fitness app Polar exposes sensitive location details for thousands of users, including soldiers and secret agents
GE's CEO Jeff Immelt used two private planes, one empty, for business travel
Google Plus exposed private data of up to 500,000 users
Hillsborough Area Regional Transit Authority CEO Ben Limmer spent taxpayer money on personal expenses
Houston pastor defrauds elderly investors by selling them worthless bonds
Kobe Steel falsified data on the strength and quality of steel products sold to customers
Leslie Moonves, CEO of CBS corporation, allegedly assaulted six women
Listerine falsely claims their mouthwash is a substitute for floss
Married Intel CEO violates company policy with consensual, sexual relationship with employee
Microsoft Bing copies Google search results and presents them as their own
News of Worlds hacks phones to access exclusive information about celebrities and politicians
Olay misleads customers about the effects of its eye cream by airbrushing model
Patisserie Valeria overstated financial position by £94m
Salesforce hires fake protestors to protest a conference held by their biggest rivals, Siebel Systems
Uber books 5560 fake rides with rival Lyft to decrease driver availability
Valeant executive involved in conspiracy to cheat investors
Walgreens falsely claimed their Walborn vitamin supplements can stave off colds

Table 7

Descriptive Statistics for Each Measure in Each of Experiment 4's Within-Participants Condition.

	Unethical		Affectively intense		Unusual	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
0 prior exposures	75.89 ^a	(17.99)	42.21 ^a	(22.66)	42.32	(21.71)
1 prior exposure	75.22 ^b	(18.59)	40.78 ^b	(22.18)	41.73	(21.74)
5 prior exposures	74.55 ^b	(18.16)	39.02 ^c	(22.38)	41.93	(21.59)

Note. Means within the same column that have different letters differ at $p < .05$, two-tailed.

Table 8

Regression Results for Experiment 4

	<i>b</i>	SE(<i>b</i>)	<i>z</i>	<i>p</i>	95% CI (<i>b</i>)	
Fixed effects						
Repetition vs. no repetition	-0.37	0.11	-3.20	<u>< 0.001</u>	-0.59	-0.14
5 vs. 1 repetition						
Intercept	-0.26	0.20	-1.29	0.200	-0.65	0.13
Random effects						
Participant intercept variance	253.91	13.81			228.24	282.48
Item intercept variance	130.74	40.52			71.22	240.00
Residual variance	438.75	4.93			429.19	448.52

Note. The results are from a mixed regression model with fixed effects for headline type (1 = previously seen, 0 = new) and random intercepts for participants and items. The conclusions were identical when I followed the pre-registered plan to treat item as a random effect; see Footnote 2. The underlined *p*-value is from a pre-registered one-tailed test, but it remains at the same value (< .001) with a two-tailed test.

Table 9

Regression Results for Experiment 5


	Main Effects Model						Interaction Model					
	<i>b</i>	SE(<i>b</i>)	<i>z</i>	<i>p</i>	95% CI (<i>b</i>)		<i>b</i>	SE(<i>b</i>)	<i>z</i>	<i>p</i>	95% CI (<i>b</i>)	
Fixed effects												
Previously-seen headline	-1.19	0.26	-4.56	<u>< 0.001</u>	-1.70	-0.68	-2.00	0.38	-5.29	< 0.001	-2.73	-1.26
Deliberation condition	-1.22	1.23	-0.99	0.322	-3.64	1.20	-2.00	1.26	-1.58	0.114	-4.47	0.48
Interaction							1.54	0.52	2.96	<u>0.002</u>	0.52	2.57
Intercept	71.70	2.38	30.10	< 0.001	67.03	76.37	72.11	2.39	30.21	< 0.001	67.43	76.78
Random effects												
Participant intercept variance	263.25	14.50			236.31	293.26	263.27	14.50			236.32	293.28
Item intercept variance	146.05	37.84			87.89	242.69	146.08	37.85			87.91	242.75
Residual variance	367.46	3.60			360.48	374.58	367.31	3.60			360.33	374.42

Note. The underlined *p*-values are from pre-registered one-tailed tests of Experiment 5's directional predictions.

Figures

Figure 1


Affective Intensity, but not Norm Perceptions, Mediated the Moral Repetition Effect in Experiment 4.



Note. Solid lines indicate statistically significant paths. $^{***} p < .001$. The independent variable was coded -2 for new wrongdoings and $+1$ for 1 or 5 repetitions.

Figure 2

Mediation of the Moral Repetition Effect in Experiment 5's Emotion Condition. *Top panel:* Mediation by affective intensity. *Bottom panel:* Mediation by anger.




Note. Solid lines indicate statistically significant paths. $*** p < .001$


Figure 3

Effect Sizes for the Moral Repetition Effect

Top Panel: Standardized Effect Sizes



Bottom Panel: Raw Mean Differences



Note. Effect sizes are the difference in moral judgments between new and previously-seen transgressions. Thus, more-negative numbers indicate a larger repetition effect. The size in Experiment 5 is from the emotion condition, because the deliberation condition was intended to reduce the moral repetition effect. Participants made raw responses on a 100-point scale.