

**Repeatedly Encountered Descriptions of Wrongdoing Seem More True but Less
Unethical: Evidence In a Naturalistic Setting**

Raunak M. Pillai¹, Lisa K. Fazio¹, and Daniel A. Effron²

¹ Department of Psychology and Human Development, Vanderbilt University

² Organisational Behaviour Subject Area, London Business School

Author Note

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1937963. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. The experiment was funded by a grant to the third author from the Center for the Science of Moral Understanding at UNC Chapel Hill.

All data, materials, and additional analyses are available online at the project's OSF site, along with our preregistration of the analyses and sample size: <https://osf.io/gn92m>.

Correspondence concerning this article should be addressed to Raunak Pillai, Department of Psychology and Human Development, Vanderbilt University, 230 Appleton Place, Nashville, TN 37203. Email: raunak.m.pillai@vanderbilt.edu

Abstract

When news about moral transgressions “goes viral,” the same person may repeatedly encounter identical reports about a wrongdoing. In a longitudinal experiment ($N = 607$ US-based MTurkers), we show that these repeated encounters can affect moral judgments. As participants went about their lives, we text-messaged them news headlines describing corporate wrongdoings (e.g., a cosmetics company harming animals). After 15 days, they rated these wrongdoings as less unethical than new wrongdoings. Extending prior laboratory research, these findings reveal that repetition can have a lasting effect on moral judgments in naturalistic settings, that affect plays a key role, and that increasing the number of repetitions generally makes moral judgments more lenient. Repetition also made fictitious descriptions of wrongdoing seem truer, connecting this *moral repetition effect* with past work on the *illusory truth effect*. The more times we hear about a wrongdoing, the more we may believe it – but the less we may care.

Keywords: moral judgement, illusory truth, repetition, desensitization, affect, exposure, transgressions; truth

Statement of Relevance

Smartphones frequently interrupt our lives with news about others' wrongdoing. As we compulsively check social media or receive news alerts throughout the day, we may repeatedly encounter the same "viral" report about the same wrongdoing. Our experiment suggests that these repeated encounters can make news of the wrongdoing seem truer, but the wrongdoing itself seem less wrong. To mimic the experience of repeated exposure to viral news stories, we texted participants real and fictitious news headlines about corporate wrongdoings (e.g., a cosmetics company harming animals) as participants went about their lives for two weeks. Receiving these texts led participants to judge these wrongdoings as less unethical but more true. These results reveal how repeated exposure to viral content can affect what we believe and what we condemn. As modern technology inundates us with information about the latest scandal, we may perceive wrongdoings as less wrong and falsehoods as less false.

Repeatedly Encountered Descriptions of Wrongdoing Seem More True but Less Unethical: Evidence In a Naturalistic Setting

Smartphones frequently interrupt our lives with reports about others' wrongdoing. As we receive news alerts or compulsively check social media, we encounter videos of people licking food at grocery stores (Garcia, 2019), rumors about children smuggled in cabinets (Dickson, 2020), and descriptions of other real and fictional wrongdoings. Morally outrageous content often "goes viral" on social media (Brady et al., 2017), making users likely to encounter the same description of the same wrongdoing repeatedly. How might this repeated exposure affect moral judgments?

Theoretically, repeated exposure could make a wrongdoing seem *more* unethical by increasing its salience (Mrkva & Van Boven, 2020) or by implying it is infamous (Jacoby et al., 1989; Weaver et al., 2007). However, in laboratory experiments, repetition makes wrongdoings seem *less* unethical (Effron, 2022; Effron & Raj, 2020). This *moral repetition effect* occurs because repeatedly reading about wrongdoings diminishes anger (*an affective-desensitization mechanism*; Dijksterhuis & Smith, 2002), and less anger means less severe moral judgments (Avramova & Inbar, 2013).

However, it is unclear whether the moral repetition effect occurs outside of brief lab studies. The present longitudinal experiment tests whether sending news headlines about wrongdoings to people's smartphones, during approximately two weeks of their daily lives, can make those wrongdoings seem less unethical days later. In addition to testing the robustness, longevity, and generalizability of the moral repetition effect, our experiment addresses theoretical questions about how repetition affects judgement.

The Moral Repetition Effect in Naturalistic Settings

Two limitations of prior research raise questions about whether the moral repetition effect will occur in naturalistic settings. First, prior research repeatedly exposed participants to descriptions of several wrongdoings without interruption, with only minutes separating each exposure (Effron, 2022; Effron & Raj, 2020). However, in real life, repeatedly-encountered descriptions of wrongdoings are usually interspersed with other content (e.g., other posts on social media), people are distracted, and days may separate each exposure. These factors may prevent people from becoming affectively desensitized to the wrongdoings, eliminating the moral repetition effect.

Second, participants in past experiments rated the morality of wrongdoings just a few minutes after first reading about them (Effron, 2022; Effron & Raj, 2020). However, in real life, people often express moral judgments days after reading about wrongdoings. Repetition may not affect moral judgments after such long delays. Even if people became affectively desensitized to a wrongdoing, they could become “re-sensitized” over time (Rankin et al., 2009), particularly if they next read about it in a different setting (e.g., home versus work; Vervliet et al., 2013). Yet, there are reasons to expect repetition *would* affect moral judgments after long delays.

Desensitization to unpleasant stimuli can endure for days (Ferrari et al., 2020) or longer (Wolpe, 1961), and repetition affects other judgments (e.g., truth, liking, fame) weeks after exposure (Bartlett et al., 1991; Henderson et al., 2021; Seamon et al., 1983).

In sum, it is important to test whether the moral repetition effect occurs in naturalistic settings because there are compelling theoretical reasons to predict it does not. Our research addresses two other key theoretical questions.

Does the Number of Views Matter?

First, in the real world, the number of times people encounter a description of a wrongdoing will vary. Does the size of the moral repetition effect increase with each exposure? The affective-desensitization mechanism predicts yes. The more times people read about a wrongdoing, the weaker their affective reactions should be (Rankin et al., 2009; Thompson & Spencer, 1966), and the less unethical they should judge the wrongdoing. However, the only prior experiment designed to test this prediction found no support (Effron, 2022, Experiment 4). Moral judgements became more lenient between the first and second time participants read about a wrongdoing, but not between the second and sixth.

This finding may seem inconsistent with the affective-desensitization mechanism. However, desensitization can exhibit a logarithmic trend where earlier exposures desensitize people more than later exposures (Rankin et al., 2009; Thompson & Spencer, 1966), making the marginal effect of later exposures on moral judgments small and hard to detect. Providing a more sensitive test of this logarithmic trend (Preacher et al., 2005), our experiment uses a wider range and a larger number of exposure conditions than previous research (i.e., 1, 2, 4, 8, and 16 exposures vs. 1, 2, and 6 exposures, respectively).

Does Repetition Have Conflicting Effects on Judgments of Morality vs. Truth?

Second, on social media, viral descriptions of wrongdoing may or may not be true (e.g., “fake news,” conspiracy theories; Dickson, 2020; Pennycook & Rand, 2021; Vosoughi et al., 2018). Repetition increases the perceived truth of true and false information alike (Dechêne et al., 2010; Pennycook et al., 2018). How might this *illusory truth effect* relate to the moral repetition effect? Intuitively, people should feel more upset about true wrongdoings than fictional ones. For example, people would presumably be more outraged if a flight attendant slapped a crying baby in real life than in a novel. Because negative affect drives moral

judgements (Greene & Haidt, 2002), true wrongdoings might receive more condemnation than fictional ones. Thus, by making wrongdoings seem truer, repetition might elicit *harsher* moral judgments, *inhibiting* the moral repetition effect.

However, the moral repetition effect and the illusory truth effect may instead be compatible. To maintain belief in a “just world” (Hafer & Bègue, 2005), people may be motivated to rationalize wrongdoings – particularly those that seem real. Rationalizing wrongdoings should reduce negative affect (Wilson & Gilbert, 2008), resulting in less severe moral judgements. Thus, by making wrongdoings seem truer, repetition might elicit *more lenient* moral judgments, *amplifying* the moral repetition effect.

The Present Research

The present research tests whether repetition in a naturalistic setting affects moral judgments. As participants went about their lives over 15 days, we text-messaged news headlines describing various wrongdoings to their phones, varying the number of times we texted different headlines. On day 16, participants rated the unethicity and truth of these same wrongdoings alongside wrongdoings we had never texted them. We predicted that the wrongdoings we had repeatedly texted would seem less unethical than new wrongdoings (the moral repetition effect), and that reductions in anger would mediate this effect (the affective-desensitization mechanism). Addressing an alternative explanation, we also examined whether repetition would make wrongdoings seem less unusual and thus less unethical (Lindström et al., 2018). This *norm-perception* mechanism received inconsistent support in prior research (Effron, 2022). We also examined whether the size of the moral repetition effect would increase with each exposure following a logarithmic pattern. Finally, we tested whether repeated wrongdoings would seem truer (the illusory truth effect), and how truth perceptions would relate to moral judgments.

In sum, we advance understanding of the moral repetition effect by resolving competing predictions about (a) whether it occurs outside certain artificial conditions, (b) whether the number of repetitions matters, and (c) whether truth perceptions amplify or inhibit this effect.

Method

Open Practices

We pre-registered this experiment, and we report all conditions, measures, and data exclusions. Verbatim materials, data, analysis code, and the pre-registration document are posted on the Open Science Framework (OSF), <https://osf.io/gn92m/>.

This experiment received ethics approval under protocol REC722 at the London Business School.

Participants

Recruitment

Using the CloudResearch service (Litman et al., 2017), we targeted U.S.-based adult Amazon Mechanical Turk workers. To enhance data quality, we only recruited people on the CloudResearch “approved participants” list, (Eyal et al., 2021), and we blocked participants with duplicate IP addresses or who had completed previous related studies of ours.

Participants could earn up to \$10. Specifically, they received \$1 for signing up to receive 5 text messages per day over 15 days, plus \$.05 per text message they replied to (see below), plus a \$2.25 bonus if they responded to more than 95% of the messages. Additionally, participants who responded to more than 50% of messages were eligible for the final survey, which they received \$2 for completing.

Based on the attrition rate in a prior study using the present paradigm (Fazio et al., 2022), we pre-registered our expectation that recruiting 1,120 participants would yield a final sample of

approximately 500 people. However, after an initial round of data collection (April 17th, 2021-May 3rd, 2021), we had received only 295 complete responses. Hence, prior to analyzing any data, we recruited another 448 people (June 12th, 2021-June 28th, 2021), of whom 312 provided complete responses (see OSF page for more detail about recruitment and attrition). The final sample size across the two rounds of data collection was 607 people who provided 9,712 observations to the main dependent measure. Because all manipulations were within participants, the design prevents concerns about unequal attrition across conditions.

Demographics

Approximate summary data from CloudResearch indicated most participants were White (79.6%, 9.95% Asian, 7.38% Black, 2.57% Multiracial, 0.48% Native American), male (56.66%, 43.34% female), and that participants had a median birth decade of the 1980s (32-41 years old), range: 1940's to 2000's.

Exclusions

As pre-registered, we excluded ratings of any headlines that participants said they had looked up during the study, which resulted in the loss of 7 observations from 6 participants. Thus, our final sample contained 9,705 observations from 607 participants.

Statistical Power

A simulation-based sensitivity analysis conducted using simR (Green & MacLeod, 2016) suggested that a sample of 607 people provides 80% power to detect a repetition effect of at least 1.4 points on the 100-point unethicality scale, smaller than the average difference of 2.3 points in laboratory studies of the moral repetition effect (see OSF page for details and analysis code). Note that the *powersim* function in simR only allows two-tailed tests, while our pre-registered analysis for this effect specified a one-tailed test. Thus, this sensitivity analysis is a conservative

estimate and likely overestimates the smallest effect size detectable at 80% power with our design.

Design

Over 15 days, we exposed participants to eight false news headlines about wrongdoings. On the 16th day, each participant rated these same eight headlines (*repeated headlines*), plus eight unfamiliar ones (*new headlines*). Our pre-registered hypotheses predict that people will rate repeated headlines as less unethical than new headlines.

We also varied how many times participants viewed each of the eight repeated headlines. Participants viewed two of these headlines twice, another two headlines four times, another two headlines eight times, and the final two headlines 16 times. Our preregistered hypotheses predict that unethicality ratings would increase and truth ratings would decrease as a function of the logarithm of the total number of times a headline was viewed.

Finally, we varied the timing of participants' headline views (following Fazio et al., 2022). The repeated headlines that people viewed 16 times were texted once every day of the study. The repeated headlines that people viewed fewer than 16 times were either spread out evenly throughout the study (*spaced version*) or shown on consecutive days at the end of the experiment (*massed version*). For example, headlines with 4 views total were either shown on days 4, 8, 12, and 16 (spaced) or on days 13–16 (massed). Since this massed-spaced manipulation is confounded with recency of the items, we do not analyze this variable. Instead, we collapsed data across the spaced and massed versions. Our analyses below only consider the number of times participants viewed a headline in the experiment, regardless of whether those views were massed or spaced (e.g., headlines G & H, shown in Figure 1, were coded as “2 views” or “repeated” in our analyses, even though exposures to G were spaced and exposures to

H were massed). Logistically, spacing some items out throughout the texting phase ensures no day has more than 5 key text messages and the variation in repetition spacing more closely mimics variation in repetitions that people encounter in the real world.

Figure 1 summarizes the design. Importantly, we counterbalanced which specific headlines participants viewed 1, 2, 4, 8, or 16 times, and which repeated headlines were presented in the spaced versus massed versions. By rotating headlines through the 16 repetition schedules listed in each row of Figure 1, we created 16 different counterbalancing groups and randomly assigned participants to groups.

Materials

The headlines were adapted from 10 true and 16 false news headlines describing corporate wrongdoing (Effron, 2022, Experiments 1a, 1b, 4; see Table 1 for examples). We replaced references to specific companies and individuals with generic descriptors (e.g., “beverage company,” “tech CEO”), and condensed the original headlines to fit within the character limit of the study’s text-messaging platform. The 16 false headlines were the experiment’s key stimuli, which participants rated on the dependent measures. We selected this set of headlines because they had produced a robust and replicable moral repetition effect in the lab (Effron, 2022, Experiments 1a, 1b), and because, for our test of the illusory truth effect, we were interested in how much people would believe false headlines. The 10 true headlines were fillers, shown during the experiment’s texting phase (see below), but not rated on the dependent measures.

Procedure

Overview

Following Fazio et al. (2022), this experiment had two phases. During a 15-day *texting phase*, participants received text messages containing various headlines. Then, during a *rating phase* on day 16, participants completed an online survey in which they rated new and repeated headlines on the dependent measures.

Texting Phase

After consenting, participants signed up to receive text messages through remind.com, a platform that allowed us to send text messages to their phones on a pre-determined schedule. Then, every day for 15 days, we texted participants a different headline every two hours from 11am to 7pm CDT – a total of five headlines daily. The number of false headlines (our key stimuli) texted each day varied as a function of our repetition manipulation (see Figure 1). The remainder of the five daily headlines were filler (true) headlines, which ensured that everyone received the same number of texts each day. We randomized the order in which we texted the five headlines each day.

Each text message instructed participants to reply indicating their interest in the headline on a scale of 1 (low) to 6 (high). The purpose of this filler rating was to incentivize participants to read each headline; we paid them for responding to each text within 24 hours.

Rating Phase

On the 16th day, eligible participants (i.e., those who had responded to at least 50% of the text messages) received an emailed link to the final survey and had 48 hours to complete it. On this survey, participants rated 16 headlines – half repeated, half new, order randomized – on four single-item measures by typing a number on a 0–100 scale. Three measures (used in Effron, 2022) asked participants about the behavior described in the headline: how unethical they found it (main DV), how angry it made them feel (affect mediator), and how unusual they found it

(norm-perception mediator; 0 = *not at all*; 100 = *extremely*). The remaining measure (truth DV; adapted from 6-point scale used in Fazio et al., 2022) asked participants to rate the headline from 0 = *definitely false* to 100 = *definitely true*. We randomized the order in which participants provided these ratings, with the constraint that no other measures came between the anger and norm-perception measures (the two potential mediators), and that the truth measure always came first or last.

Following Fazio et al.'s illusory truth study (2022), we informed participants before the ratings that they may have encountered some of the headlines before. (Participants in prior studies on the moral repetition effect, and in some studies on the illusory truth effect, did not receive this information; e.g., Effron, 2022; Udry et al., 2022). Unlike Fazio et al. (2022), we did not provide participants with any information about the truth or falsity of the headlines, because all headlines were false.

Before the ratings, we also told participants that they should not look up any headlines. At the end of the survey, we asked participants if they had looked up headlines during the texting phase or the rating phase. Participants answering “yes” described the headlines they looked up. As noted above, we excluded their ratings of these headlines from analysis, reasoning that looking the headlines up would interfere with participants’ truth judgments.

Additionally, participants were asked what they believed the purpose of the study was. Finally, we debriefed participants, emphasizing that all headlines on the final survey were false.

Results

Analytic Approach

To account for the repeated-measures design, we computed multi-level models with the “lme4” package in R (Bates et al., 2015). Our mediation analyses use bias-adjusted estimates and

bootstrapped confidence intervals computed with the “boot” package in R (Canty & Ripley, 2021) using 2,000 bootstrap replicates. Because we had strong directional predictions, we pre-registered one-tailed significance tests to increase statistical power, but the results yield identical conclusions with two-tailed tests. Non-preregistered analyses are flagged as exploratory and use two-tailed tests.

Repeated Wrongdoings Seemed Less Unethical

Replicating the moral repetition effect in a naturalistic setting, participants rated repeated headlines as significantly less unethical than new headlines ($M_{\text{repeated}} = 79.40$, 95% confidence interval, or CI = [78.58, 80.22]; $M_{\text{new}} = 80.85$, 95% CI = [80.05, 81.65]; $d_z = 0.12$), $b = -1.43$, $SE = 0.57$, $t(14.53) = 2.51$, $p = .012$. This result is from a mixed-effects linear regression model predicting unethicality ratings from repetition (0 = new, 1 = repeated), with random intercepts and repetition slopes by-participant and by-headline. Mixed-effect models like this, that include the maximum number of random effects, are ideal because they minimize the Type 1 error rate while preserving power (Barr et al., 2013). Table 2 shows the results. (Before fitting the model, we passed it through the *buildmer* package in R (Voeten, 2019). The model converged, but if it had not, then *buildmer* would have systematically removed random effects until it achieved convergence).

The effect size observed in this experiment was about half as large as in laboratory experiments using similar stimuli and fewer total views. In the present experiment, repetition decreased unethicality judgements by 1.43 points, an effect size of $d_z = .12$; in prior experiments, repetition decreased unethicality judgments by 3 points, $d_z = .25$ (Effron, 2022, Experiments 1a and 1b).

Anger and Norm Perceptions Mediated the Moral Repetition Effect

To test our predictions about mechanism, we conducted a multi-level mediation analysis of the effects of repetition (0 = new, 1 = repeated) on unethicality judgments, with anger and norm perceptions as parallel mediators. Our pre-registered hypothesis was that anger would mediate the moral repetition effect above and beyond norm perceptions. We did not have a hypothesis about whether norm perceptions would also mediate this effect because prior studies produced inconsistent results (Effron, 2022). Our analysis used the multilevel model structure reported in Table 3, adding fixed paths from repetition to unethicality judgements via anger and norm perception.

We fit this multi-level mediation model using the approach recommended by Bauer et al. (2006), which allows the use of univariate modelling software while providing unbiased estimates of effects. Specifically, we began by formatting the data such that the DV (unethicality judgements) and both mediators (anger, norm perceptions) were placed under a single “stacked” variable. We used three dummy-coded (0, 1) selector variables to identify whether a given row of this stacked variable was the DV, the anger mediator, or the norm perception mediator. Then, we specified the mediation model as a univariate model with the stacked variable as the outcome. This univariate model consisted of the predictors for each pathway in the model, multiplied by the appropriate selector variable to ensure the model was being fit to the appropriate data. For instance, the “anger” selector variable was multiplied with the repetition term to estimate the effect of repetition on anger using only those rows in the stacked data frame that corresponded to anger judgements. The annotated code for processing and analyzing these data is available at the project’s OSF page. For a tutorial introducing this approach, see UCLA: Statistical Consulting Group (2014).

As predicted, we find evidence that anger mediated the moral repetition effect, above and beyond norm perceptions (see Figure 2). Specifically, people felt less angry about wrongdoings described in repeated headlines versus new headlines, and the less angry people felt, the less unethical they found the wrongdoings – a significant indirect effect, $b = -0.97$, bootstrapped 95% CI = [-1.61, -0.32]. We also found a significant, but numerically smaller, indirect effect through the norm perception measure, $b = -0.25$, bootstrapped 95% CI = [-0.46, -0.03]. That is, people rated wrongdoings as less unusual if they were described in repeated versus new headlines, and the less unusual the wrongdoings seemed, the less unethical they were rated.

Few Participants Guessed the Hypothesis

Finally, to address the possibility that repetition affected moral judgments simply because participants guessed the hypothesis (i.e., demand characteristics), we conducted an exploratory analysis of participants responses to the open-ended questions, “What do you think this study was about? What do you think we were trying to test?” Reducing concerns about hypothesis-guessing, only 3 of our 607 participants (0.49%) mentioned any connection between exposure to the wrongdoings and judgements of morality or unethicability, and an additional 14 (2.31%) mentioned desensitization or changes in feelings with exposure.

Additional Repetitions Had Diminishing Marginal Effects

We next examined if and how the moral repetition effect depended on the number of times participants viewed the headlines. We hypothesized that the effect’s magnitude would decrease with the number of views – a logarithmic effect where, for example, the effect of increasing the number of views from 1 to 2 would be larger than the effect of increasing it from 15 to 16. As Figure 3 shows, the results appear broadly consistent with the hypothesized logarithmic effect.

This effect was significant in a mixed-effects regression model predicting unethicality ratings from the natural log of the number of views, $b = -.80$, $SE = .23$, $t(605.25) = 3.47$, $p < .001$. As in the above analyses, the model's random-effect structure was determined by the “buildmer” function in R (see Table 3). As predicted, this model fit the data better than an alternative model that predicted unethicality ratings from the linear effect of the number of views (see OSF page). Together, these results suggest that successive repetitions have diminishing marginal effects on unethicality judgments.

We also compared the model reported in Table 3 to the model reported in Table 2, which predicts unethicality from repetition status (0 = new, 1 = repeated). Though not pre-registered, this comparison tests whether repetition continually decreases unethicality judgements (model in Table 3) or whether this decrease is all-or-none (model in Table 2), as Effron (2022) found. This comparison also favored the model reported in Table 3 predicting unethicality ratings from the natural logarithm of the number of views (see OSF page for details).

Finally, similar to the analysis reported in Figure 2, we conducted a multilevel mediation analysis examining the effects of number of views on unethicality judgments, with anger and norm perceptions as parallel mediators. We again found that anger and norm perceptions mediated the moral repetition effect, and bootstrapped estimates of the indirect effects revealed a larger indirect effect through anger than through norm perceptions (see OSF page for details).

Repeated Headlines Were Perceived as Truer

The results thus far demonstrate the moral repetition effect in a naturalistic setting. The results also replicate the illusory truth effect in a naturalistic setting. As in past research (Fazio et al., 2022), the more times participants viewed a headline in daily life, the more true they thought it was – a logarithmic effect whereby truth ratings increase sharply between initial increases in

number of views, and then plateau as the number of views increased further (see Figure 4). A pre-registered test of this logarithmic effect was significant in a mixed-effects model that followed the same analysis plan used to arrive at the model reported in Table 4, but with truth ratings as the dependent variable, $b = 1.49$, $SE = .39$, $t(20.18) = 3.79$, $p < .001$ (see Table 4). A model comparison favored the model predicting perceived truth from the natural logarithm of the number of views as compared to the raw number of views (see OSF page).

As shown in Table 4, one log-unit increase in the number of views resulted in about a 1.49-point increase in perceived truth. Note that while we replicated an effect of repetition on belief in this naturalistic setting, the effects were smaller than past research using more banal stimuli (trivia statements). In Fazio et al. (2022), truth ratings increased 0.25 of 6 points (proportional to 4.16 of 100 points) per unit increase in the natural logarithm of the number of views.

In addition to examining the effect of the number of views on belief, we also conducted an exploratory analysis predicting truth ratings from repetition (0 = new, 1 = repeated), as we had with the unethicality ratings. Again, replicating the illusory truth effect, we find that participants gave higher truth ratings to repeated, $M = 51.12$, than new items, $M = 47.33$ (see OSF page for full results).

Perceived Truth Also Mediated the Moral Repetition Effect

The results thus far demonstrate that repeatedly viewing descriptions of wrongdoing can make them seem less unethical (the moral repetition effect) and more true (the illusory truth effect). What is the relationship between these effects? Does the illusory truth effect *amplify* the moral repetition effect because truer headlines seem less unethical? Or does the illusory truth effect *inhibit* the moral repetition effect because truer headlines seem *more* unethical?

To explore these questions, we conducted a non-preregistered multi-level mediation analysis. This model was identical to the one reported in Figure 2, with the addition of perceived truth as a parallel mediator variable. The results were most consistent with the possibility that the illusory truth effect amplifies the moral repetition effect. That is, we observed a small but significant negative indirect effect from repetition to truth judgments and then to unethicality judgments, $b = -.08$, bootstrapped 95% CI = [-0.18, -0.02]), above and beyond the indirect effects through anger and norm perceptions discussed earlier (see Figure 5). A caveat is that the results do not allow us to definitively conclude that truer wrongdoings seemed less unethical, because the direct effect between from truth perceptions to unethicality judgments, controlling for anger and norm perceptions, was not significant (see Figure 5).

Discussion

When a report about a wrongdoing “goes viral,” people may repeatedly encounter it through social media or news alerts on their phones. Capturing this experience, our longitudinal experiment reveals that these repeated encounters to identical reports of a wrongdoing can make the wrongdoing seem less unethical (the moral repetition effect) and the news seem truer despite being false (the illusory truth effect).

Our results make three key theoretical contributions. First, they demonstrate that the moral repetition effect can occur outside the lab in a naturalistic setting, where people repeatedly encounter descriptions of wrongdoings over weeks, their encounters are separated by a day or more, they see content between encounters, and they face distractions. We also provide the first evidence that repetition can have a *lasting* effect on moral judgments – a day or more after last reading about it. These findings are important because some theorizing suggests that affective desensitization might not occur in these settings (Rankin et al., 2009), and thus that repetition

would have no effect on moral judgments. Future work is needed to understand whether distractions, delays, and other features of naturalistic settings *attenuate* the moral repetition effect, but in our study such features did not *eliminate* it.

Second, we resolve an apparent inconsistency. Theoretically, more repetitions should lead to less moral condemnation – but in prior work, six exposures to a wrongdoing did not affect moral judgments significantly more than two exposures (Effron, 2022). Our results clarify that the size of the moral repetition effect *does* depend on the number of repetitions, but that increasing the number of repetitions has a progressively smaller effect on moral judgments (more of a logarithmic relationship than a linear or null relationship). Prior work may have failed to observe such effects because as the number of repetitions increases, the marginal effect of each additional repetition gets smaller and harder to detect.

Third, the results offer the first evidence of a relationship between the illusory truth effect and the moral repetition effect. Although the illusory truth effect might plausibly have *inhibited* the moral repetition effect (i.e., by making wrongdoings seem truer, repetition might elicit *harsher* moral judgments), the results suggest that the illusory truth effect may *amplify* the moral repetition effect (i.e., by making wrongdoings seem truer, repetition may have elicited *more lenient* moral judgments). Future work should confirm the robustness of this mediation effect, especially because the direct effect of truth perceptions on moral judgements was not significant; see Figure 5), and because statistical mediation cannot demonstrate causation (Bullock et al., 2010). Still, we speculate that truth perceptions may constitute a novel mechanism behind the moral repetition effect. Perhaps to preserve belief in a “just world” (Hafer & Bègue, 2005), people are more motivated to rationalize real (vs. fictional) wrongdoings.

The results also suggest that two other parallel mechanisms may explain the moral repetition effect. As in prior research, repetition made people less angry about wrongdoings, and the less angry they felt, the less unethical they judged the wrongdoing (Effron, 2022). Multiple processes may drive this affective desensitization (e.g., habituation, rationalization; Dijksterhuis & Smith, 2002) and more work is needed to evaluate these possibilities. Repetition also made wrongdoings seem less unusual – and the less unusual they seemed, the less unethical they seemed. Although this mechanism received only inconsistent support in prior research and explained less variance than the affective-desensitization mechanism, it is consistent with the idea that people confuse familiarity with prevalence (Tversky & Kahneman, 1973; Weaver et al., 2007) and infer that more-prevalent behaviors are more moral (Lindström et al., 2018).

Finally, the findings address a limitation of past work on the illusory truth effect (Fazio et al., 2022) by showing the effect holds not only in naturalistic *settings* but also with naturalistic *stimuli* (morally-charged news versus trivia) that are more emotionally and personally salient.

Although repetition had only a small effect on moral judgements ($d_z = .12$), reliably estimating small effects is critical for understanding multidetermined processes like moral judgements (Götz et al., 2022). Moreover, while repeated headlines still received high unethicality ratings (79 out of 100), it is arguably impressive that repetition shifted moral judgments at all (Prentice & Miller, 1992) as participants judged severe wrongdoings (e.g., slapping a baby for crying; Table 1). Repetition is unlikely to make wrongdoings seem right, but it does appear to make them appear a little less wrong. We suspect that this small effect could be practically meaningful when aggregated across the millions of social media users who repeatedly read about the latest viral wrongdoing.

Additionally, although people in the real world, as in our experiment, may repeatedly encounter the same “viral” description of the same wrongdoing, they may also encounter *different* descriptions of the same wrongdoing. We predict that the moral repetition effect still occurs in those situations – presumably people would still become desensitized to the *wrongdoing* even though they would not become desensitized to the specific *description* of the wrongdoing – but this prediction awaits future tests. Finally, a limitation of our results is that they come from a convenience sample of U.S.-based participants.

Ultimately, our results reveal how repeated exposure to viral content can affect how we judge morality and truth. The more we hear about a wrongdoing, the more we may believe it – but the less we may care.

References

Avramova, Y. R., & Inbar, Y. (2013). Emotion and moral judgment. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4(2), 169–178.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>

Bartlett, J. C., Strater, L., & Fulton, A. (1991). False recency and false fame of faces in young adulthood and old age. *Memory & Cognition*, 19(2), 177–188. <https://doi.org/10.3758/BF03197115>

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, Articles*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>

Bauer, D. J., Preacher, K. J., & Gil, K. M. (2006). Conceptualizing and testing random indirect effects and moderated mediation in multilevel models: New procedures and recommendations. *Psychological Methods*, 11(2), 142–163. <https://doi.org/10.1037/1082-989X.11.2.142>

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. <https://doi.org/10.1073/pnas.1618923114>

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. <https://doi.org/10.1037/a0018933>

Canty, A., & Ripley, B. D. (2021). *boot: Bootstrap R (S-Plus) Functions*.

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 Psychology Review, 14*(2), 238–257. <https://doi.org/10.1177/1088868309352251>

Dickson, E. (2020, July 14). A Wayfair Child-Trafficking Conspiracy Theory Is Flourishing on TikTok, Despite It Being Completely False. *Rolling Stone*.
<https://www.rollingstone.com/culture/culture-news/wayfair-child-trafficking-conspiracy-theory-tiktok-1028622/>

Dijksterhuis, A., & Smith, P. K. (2002). Affective habituation: Subliminal exposure to extreme stimuli decreases their extremity. *Emotion, 2*(3), 203.

Effron, D. A. (2022). The moral repetition effect: Bad deeds seem less unethical when repeatedly encountered. *Journal of Experimental Psychology: General, 151*(10), 2562–2585.
<https://doi.org/10.1037/xge0001214>

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>

Eyal, P., David, R., Andrew, G., Zak, E., & Ekaterina, D. (2021). Data quality of platforms and panels for online behavioral research. *Behavior Research Methods*.
<https://doi.org/10.3758/s13428-021-01694-3>

Fazio, L. K., Pillai, R. M., & Patel, D. (2022). The effects of repetition on belief in naturalistic settings. *Journal of Experimental Psychology: General, 151*(10), 2604–2613.
<https://doi.org/10.1037/xge0001211>

Ferrari, V., Mastria, S., & Codispoti, M. (2020). The interplay between attention and long-term memory in affective habituation. *Psychophysiology*, 57(6), e13572.
<https://doi.org/10.1111/psyp.13572>

Garcia, S. E. (2019, July 6). Blue Bell Ice Cream Licker in Viral Video Is Found, Police Say. *New York Times*. <https://www.nytimes.com/2019/07/06/us/blue-bell-licker-lufkin-tx.html>

Götz, F. M., Gosling, S. D., & Rentfrow, P. J. (2022). Small effects: The indispensable foundation for a cumulative psychological science. *Perspectives on Psychological Science*, 17(1), 205–215. <https://doi.org/10.1177/1745691620984483>

Green, P., & MacLeod, C. J. (2016). 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/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. [https://doi.org/10.1016/S1364-6613\(02\)02011-9](https://doi.org/10.1016/S1364-6613(02)02011-9)

Hafer, C. L., & Bègue, L. (2005). Experimental research on just-world theory: Problems, developments, and future challenges. *Psychological Bulletin*, 131(1), 128–167.
<https://doi.org/10.1037/0033-2909.131.1.128>

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>

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.

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), 228–242. <https://doi.org/10.1037/xge0000365>

Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime. Com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, 49(2), 433–442.

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*, 118(6), 1118–1145. <https://doi.org/10.1037/pspa0000184>

Pennycook, G., Cannon, T. D., & Rand, D. G. (2018). Prior exposure increases perceived accuracy of fake news. *Journal of Experimental Psychology: General*, 147(12), 1865–1880. <https://doi.org/10.1037/xge0000465>

Pennycook, G., & Rand, D. G. (2021). The psychology of fake news. *Trends in Cognitive Sciences*, 25(5), 388–402. <https://doi.org/10.1016/j.tics.2021.02.007>

Preacher, K. J., Rucker, D. D., MacCallum, R. C., & Nicewander, W. A. (2005). Use of the extreme groups approach: A critical reexamination and new recommendations. *Psychological Methods*, 10(2), 178.

Prentice, D. A., & Miller, D. T. (1992). When small effects are impressive. *Psychological Bulletin*, 112(1), 160–164.

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., McSweeney, F. K., Wilson, D. A., Wu, C.-F., & Thompson, R. F. (2009). Habituation revisited: An updated and revised description

of the behavioral characteristics of habituation. *Neurobiology of Learning and Memory*, 92(2), 135–138. <https://doi.org/10.1016/j.nlm.2008.09.012>

Seamon, J. G., Brody, N., & Kauff, D. M. (1983). Affective discrimination of stimuli that are not recognized: II. Effect of delay between study and test. *Bulletin of the Psychonomic Society*, 21(3), 187–189. <https://doi.org/10.3758/BF03334682>

Thompson, R. F., & Spencer, W. A. (1966). Habituation: A model phenomenon for the study of neuronal substrates of behavior. *Psychological Review*, 73(1), 16.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), 207–232.

UCLA: Statistical Consulting Group. (2014, November 24). *How Can I Perform Mediation with Multilevel Data? (Method 2) | R FAQ*. Advanced Research Computing Statistical Methods and Data Analytics. <https://stats.oarc.ucla.edu/r/faq/how-can-i-perform-mediation-with-multilevel-data-method-2/>

Vervliet, B., Craske, M. G., & Hermans, D. (2013). Fear extinction and relapse: State of the art. *Annual Review of Clinical Psychology*, 9(1), 215–248. <https://doi.org/10.1146/annurev-clinpsy-050212-185542>

Voeten, C. C. (2019). *buildmer: Stepwise elimination and term reordering for mixed-effects regression*.

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151. <https://doi.org/10.1126/science.aap9559>

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. <https://doi.org/10.1037/0022-3514.92.5.821>

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>

Wolpe, J. (1961). The systematic desensitization treatment of neuroses. *Journal of Nervous and Mental Disease, 132*(3), 189–203.

Tables

Table 1*Sample Headlines*

Key (false)	Filler (true)
An electronics company manufactured headphones that exploded in a man's ear.	The CEO of an electric utility company was involved in a scheme to bribe an Ohio House Speaker.
A flight attendant slapped a 7-month-old baby in the face for crying during a flight.	A steel manufacturer falsified data on the strength and quality of steel products sold to customers.
A messaging app uses photo filtering tech to make a facial recognition database for the FBI, CIA.	A pharmacy chain falsely claimed their vitamin supplements can stave off colds.

Table 2*Mixed-Effects Model of Unethicality Ratings With Fixed Effects of Repetition and Random**Slopes and Intercepts by Participant and by Headline*

Fixed Effects	Estimate	SE	Df	t Value	p Value
Intercept	80.84	2.93	15.96	27.58	<.001
Repetition	-1.43	0.57	14.53	-2.51	.012
Random Effects	Variance	SD	Correlation		
Participant (Intercept)	160.54	12.67			
Participant (Repetition)	3.43	1.85	0.17		
Headline (Intercept)	131.48	11.47			
Headline (Repetition)	1.55	1.25	0.15		

Note. Model was fit to 9,705 unethicality ratings from 607 participants across 16 claims. Bolded values indicate significant effects. Correlation values in each row reflect correlations between the term and all preceding random effects terms in the same level (participant or headline) in order of appearance in the table. The *p* value for the repetition term reflects a one-tailed test, as pre-registered; the intercept term reflects a two-tailed test.

Table 3

Mixed-Effects Model of Unethicality Ratings With Fixed Effects of the Natural Logarithm of the Number of Views and Random Slopes and Intercepts by Participant and by Headline

Fixed Effects	Estimate	SE	Df	t Value	p Value
Intercept	80.82	2.96	16.11	27.36	<.001
Ln(Views)	-0.80	0.23	605.25	-3.47	<.001
Random Effects	Variance	SD	Correlation		
Participant (Intercept)	161.69	12.72			
Participant (Ln(Views))	0.46	0.68	0.24		
Headline (Intercept)	133.88	11.57			

Note. Model was fit to 9,705 unethicality ratings from 607 participants across 16 claims. Bolded values indicate significant effects. Correlation values in each row reflect correlations between the term and all preceding random effects terms in the same level (participant or headline) in order of appearance in the table. The *p* value for the natural logarithm of the number of views term reflects a one-tailed test, as pre-registered; the intercept term reflects a two-tailed test.

Table 4

Mixed-Effects Model of Truth Ratings With Fixed Effects of the Natural Logarithm of the Number of Views and Random Slopes and Intercepts by Participant and by Headline

Fixed Effects	Estimate	SE	Df	t Value	p Value
Intercept	47.92	3.94	15.61	12.15	<.001
Ln(Views)	1.49	0.39	20.18	3.79	<.001
Random Effects	Variance	SD	Correlation		
Participant (Intercept)	189.65	13.77			
Participant (Ln(Views))	14.60	3.82	-0.18		
Headline (Intercept)	242.17	15.56			
Headline (Ln(Views))	1.15	1.07	-0.34		

Note. Model was fit to 9,705 truth ratings from 607 participants across 16 claims. Bolded values indicate significant effects. Correlation values in each row reflect correlations between the term and all preceding random effects terms in the same level (participant or headline) in order of appearance in the table. The *p* value for the natural logarithm of the number of views term reflects a one-tailed test, as pre-registered; the intercept term reflects a two-tailed test.

Figure Captions

Figure 1

Presentation Schedule for Key Headlines

Note. Headlines D, F and H have massed presentation schedules, while headlines C, E, and G have spaced schedules. We counterbalanced which specific headline corresponded to each letter.

Figure 2

Mediation Analysis of the Effects of Repetition on Unethicality Judgements via Anger and Norm Perceptions

Note. Values indicated path coefficients. Solid lines indicate significant paths; dotted lines indicate non-significant paths.

[†] $p < .10$ * $p < .05$. ** $p < .01$. *** $p < .001$.

Figure 3

Unethicality Ratings by Number of Times Viewed, with Mixed-Effects Linear Regression Predictions

Predictions

Note. The x-axis indicates the total number of times an item was seen in the experiment, regardless of its exact presentation schedule (see Figure 1). In the figure, and all related analyses, the “number of views” includes all exposures to a headline—including exposure to the headline during the final rating phase. Solid points indicate mean unethicality ratings. Solid black error bars indicate standard errors of the mean. Solid black line indicates predicted unethicality rating from the mixed-effects regression model described in Table 3. Shaded region indicates 95% confidence interval of the prediction from the same model.

Figure 4

Truth Ratings by Number of Times Viewed, with Mixed-Effects Linear Regression Predictions

Note. The x-axis indicates the total number of times an item was seen in the experiment, regardless of its exact presentation schedule (see Figure 1). In the figure, and all related analyses, the “number of views” includes all exposures to a headline—including exposure to the headline during the final rating phase. Solid points indicate mean truth ratings. Solid black error bars indicate standard errors of the mean. Solid black line indicates predicted truth rating from the

mixed-effect regression model described in Table 4. Shaded region indicates 95% confidence interval of the prediction from the same model.

Figure 5

Mediation Analysis of the Effects of Repetition on Unethicality Judgements via Anger, Norm Perceptions, and Perceived Truth

Note. Values indicated path coefficients. Solid lines indicate significant paths; dotted lines indicate non-significant paths.

[†] $p < .10$ * $p < .05$. ** $p < .01$. *** $p < .001$.