



ARTICLE

<https://doi.org/10.1057/s41599-019-0269-y>

OPEN

Social transmission favours the ‘morally good’ over the ‘merely arousing’

Joseph M. Stubbersfield¹ , Lewis G. Dean², Sana Sheikh^{2,3}, Kevin N. Laland⁴ & Catharine P. Cross²

ABSTRACT Moral stories are pervasive in human culture, forming the basis of religious texts, folklore, and newspaper articles. We used a linear transmission chain procedure to test three competing hypotheses: (1) that moral content in general is preferentially transmitted between individuals compared to non-moral content; (2) that negativity bias leads specifically to morally bad content being preferentially transmitted; and (3) that a bias towards pro-social information leads specifically to morally good content being preferentially transmitted. While we found no support for a bias for moral content in general, we did find that morally good content was transmitted with greater fidelity than neutral or morally bad content, with ratings of morally good content but not morally bad content predicting transmission. Moral content, therefore, appears to be particularly culturally potent when it describes the ‘virtuous’ rather than the ‘sinful’. A second study repeated the first but also tested the influence of physiological arousal on transmission by measuring the electrodermal activity of participants. This study also found that morally good content was transmitted with greater fidelity than neutral or morally bad content and that physiological arousal had a negative effect on transmission with more arousing material being less faithfully transmitted. These results suggest that the communication of content relating to moral virtue might serve to avoid negative impression formation and promote social bonding, and that this might partially explain the ubiquity of moral content in human culture.

¹Department of Anthropology, University of Durham, Durham, UK. ²School of Psychology and Neuroscience, University of St. Andrews, Scotland, UK.

³Cambridge Health Alliance, Harvard Medical School, Cambridge, USA. ⁴School of Biology, University of St. Andrews, Scotland, UK. Correspondence and requests for materials should be addressed to C.P.C. (email: cpc2@standrews.ac.uk)

Introduction

Information about morality—defined by Ayala (2010) as “value judgments concerning human behaviour” (p. 9016)—permeates human culture, in religious texts, folklore, fables and news stories. Moral information might be an explicit proclamation of what is moral or amoral, or a more implicit illustration of the moral norms of a social group. Morality is characterised by some authors as an adaptation, built upon emotions that are themselves adaptive, with specific moral codes emerging via gene-culture coevolution (Gintis et al., 2008; Graham et al., 2013; Norenzayan, 2014). While disagreement exists about the sequences of events and selection pressures underpinning the human capacity for morality, there is general agreement that the emergence and evolution of moral codes and norms is only possible through cultural transmission (Ayala, 2010; Haidt and Joseph, 2004; McNamara et al., 2019). The process by which moral content is embedded into culturally transmitted stories and artefacts is poorly understood, however. Here, we present two studies investigating cognitive content biases towards moral information.

A cognitive content bias is a predisposition that humans have for attending to, recalling, or re-producing, certain kinds of information. (Boyd and Richerson, 1985; Barrett and Nyhof, 2001; Henrich and McElreath, 2003). Evolutionary psychologists argue that such biases exist because they were adaptive in ancestral environments (e.g., tracking social relationships in complex groups), and that they modify the content and structure of cultural knowledge and artefacts, which makes them increasingly transmittable (Barrett and Nyhof, 2001; Laland and Brown, 2011; Mesoudi, 2016). Studies have demonstrated biases for a number of types of content: ecological information relevant to health and survival, hereafter *survival content* (Nairne, 2010; Stubbersfield et al., 2015); information relevant to social relationships and interaction, hereafter *social content* (Mesoudi et al., 2006; Stubbersfield et al., 2015), and information evoking an emotional response, hereafter *emotional content* (Eriksson and Coultas, 2014; Heath et al., 2001; Stubbersfield et al., 2017). A number of studies investigating biases in transmission have used the transmission chain, or serial reproduction, paradigm, in which experimental material is passed along a linear ‘chain’ of individuals. First developed by Bartlett (1932), this method allows researchers to assess which types of information are preserved with the greatest fidelity as they are passed along the chain, in turn revealing the biases in social transmission which influence the transmission and evolution of culture (Mesoudi and Whiten, 2008; Mesoudi et al., 2006). While evidence for a number of content biases in transmission has been documented (See Stubbersfield et al., 2018 for a review) a content bias for morally relevant content has yet to be experimentally examined.

Based on Haidt and Joseph’s (2004) moral foundations theory, one might expect a transmission bias for all morally relevant information, because any information relevant to the moral foundations would be salient. This hypothesis has not yet been directly tested but indirect evidence for a bias for moral content comes from the literature on impression formation. Wojciszke et al. (1998) for example, found that morality-related information played a more important role in global impression formation and was more cognitively accessible than competence-related information. Ybarra et al. (2001) found that participants were more sensitive to person-relevant information from the morality domain than the competence domain. Van Leeuwen et al. (2012) used a memory confusion paradigm and found that participants spontaneously categorised along a morality dimension but not along a competence dimension. Morality also influences the perception of specific actions: Pizarro et al. (2006) gave participants one of two vignettes—the first about a man intentionally

leaving a restaurant without paying and the second about a man who forgot to pay for his meal. Participants who read the vignette with the intentional—but not the accidental—moral transgression distorted the recalled amount to be larger than it was when asked for the size of the bill. Taken together, these studies suggest that moral content is particularly salient when forming impressions of other people and their actions.

Another known bias which could influence the transmission of moral content is a general negativity bias. Research on memory, perception, decision making and impression formation has suggested that negative entities (such as events, objects or personal traits) are more salient than their positive counterparts (Baumeister et al., 2001; Rozin and Royzman, 2001). It has been suggested that this bias is adaptive, because negative entities are likely to incur greater costs on an individual than positive entities incur benefits (Rozin and Royzman, 2001). A suggestion supported by Fessler et al. (2014), who found that participants were more credulous of negatively framed than positively framed information and that negative content was over-represented in a corpus of online urban legends and supernatural beliefs. Bebbington et al. (2016) also demonstrated a transmission advantage for negatively valenced content over positively valenced content, and found that ambiguous content is more likely to become negative than positive through transmission. Similarly, Walker and Blaine (1991) used a naturalistic “field-experiment” to demonstrate that rumours forecasting unpleasant consequences have an advantage in social transmission over rumours forecasting pleasant consequences. In addition, other studies have found a negativity bias in emotional expression in the arts (see Brand et al., 2019; Morin and Acerbi, 2017). Based on these findings it is feasible that content featuring immoral behaviour will be more salient and better transmitted than content featuring virtuous behaviour, because immoral agents could be perceived as more hazardous than moral agents are beneficial.

Alternatively, a positivity bias might shape the transmission of moral content. Experimental research has found a preference for choosing to transmit positively valenced vignettes over negatively valenced equivalents (van Leeuwen et al., 2018). Studies using “real world” data have found an advantage for positive content, with positively valenced messages being more frequently and more widely shared on social media than negative content (Ferrara and Yang, 2015a; Ferrara and Yang, 2015b; Fu et al., 2016) and urban legends featuring amusing content being more common than those featuring other, negative emotions (Stubbersfield et al., 2018). In addition, research has shown that children seek out and transmit content which supports a positive, pro-social, evaluation of their in-group over other types of information, including negative information about their out-group (Over et al., 2017). Therefore, the possibility of virtuous content being advantaged in transmission should also be considered. We therefore examine the transmission of both morally good content and morally bad content. In Study 1 we test three alternative hypotheses relating to the transmission of moral information in a linear transmission chain. In Study 2 we extend on Study 1 by including a measure of physiological arousal in order to examine the role of emotion in the transmission of moral information.

Study 1

The present research investigates the social transmission of moral and morally neutral (hereafter “non-moral”) content. We produced a series of vignettes, each with two versions, moral and non-moral. The moral version of each vignette features either a deliberately virtuous behaviour, which promotes or enhances

social life and interaction (hereafter “morally good”), or a deliberately morally transgressive act, which suggests selfishness or anti-sociality (hereafter “morally bad”). Non-moral versions of the same vignettes feature the same outcomes, but they are not brought about intentionally. For instance, in a vignette describing the damage of a person’s property, in the moral version it would be due to vandalism and in the non-moral version it would be due to accidental damage. A pre-test study (see Supplementary Material (SM) 1 and 2) validated the material by having participants ($N = 133$) rate the vignettes (see SM3) for moral content (good and bad) and other content that might influence recall or transmission (social information, gender stereotype consistency, survival information). Vignettes were used in the study if participants in the pre-test rated the moral version as significantly higher in either morally good or morally bad content than the non-moral version. Because moral content could have a transmission advantage simply because of its social and emotive nature (Graham et al., 2013; Pizarro, 2000), we address the question of whether a specific content bias exists for moral information, rather than moral content being favoured simply because it is social or emotive, by testing morality as a separate predictor of transmission alongside measures of social and emotive content.

The present study primarily seeks to examine three competing hypotheses:

1. There is a cognitive bias for transmitting general moral information. Consistent with research suggesting a moral information bias in the domains of impression formation and categorisation, moral information—both good and bad—may be preferentially transmitted over non-moral information (H1).
2. There is a cognitive bias for transmitting morally bad information. Consistent with research suggesting a negativity bias in transmission chains and recall, morally bad information may be preferentially transmitted over non-moral or morally good information (H2).
3. There is a cognitive bias for transmitting morally good information. Consistent with research suggesting a bias for transmitting positive information in experimental settings and on social media, morally good information may be preferentially transmitted over non-moral or morally bad information (H3).

Material and methods

Participants. Forty participants (32 female, 7 male, 1 other) aged 17 to 43 years ($M = 21.58$, $SD = 4.55$) took part. All participants gave their informed consent.

Materials. Vignettes were 19 to 62 words and contained 3 to 8 propositions (determined through propositional analysis (Kintsch, 1974)). See below for examples (bolded sections illustrate differences between versions and did not appear as such to participants).

Tyre–Non-moral version

Nigel returned to his bike after visiting his friends to find that it had a flat tyre. The tyre had been **punctured by a small thorn**. Now he would have to walk two miles home in the pouring rain. Nigel was so angry at the puncture he kicked a wall.

Tyre–Moral version

Nigel returned to his bike after visiting his friends to find that it had a flat tyre. The tyre had been **deliberately slashed by someone**. Now he would have to walk two miles home in the pouring rain. Nigel was so angry at the puncture he kicked a wall.

The results of the pre-test were used to determine ten vignettes most appropriate for analyses, i.e., those with fewest possible confounds between moral and non-moral versions. Participants in the pre-test also classified the vignettes according to the most prominent emotion (anger, disgust, elevation, gratitude, guilt/

shame, happiness), and rated the vignettes for emotional intensity, moral goodness, moral badness, survival information, social information, male stereotype consistency, and female stereotype consistency (see SM1–SM3).

Design. A linear transmission chain design was used (see Mesoudi and Whiten, 2008, for further details on this design and its use). This design was used as it allows for the examination of cumulative effects over generations, an advantage over a single generation design when considering recall and cultural transmission. As in previous research (Mesoudi and Whiten, 2004; Mesoudi et al., 2006), ten chains comprised of four participants, or ‘generations’ were used. The first participant in each of the ten chains received all ten vignettes. Participants received one of two sets of vignettes, with an equal mix of moral and non-moral vignettes. Half the participants received Set A, while half received Set B. Each set contained the same vignettes but with opposite versions, i.e., Set A would have the moral version of the ‘Tyre’ vignette, while Set B would contain the non-moral version, and so on. The second participant in a chain received all text generated by the first participant, and so on.

Procedure. Participants were presented with the vignettes on a computer screen. After reading five vignettes, they were asked to recall each one with a short prompt: e.g., “Please type in the box provided, as accurately as you can remember, the ‘Tyre’ story (the story about Nigel)”. Participants were informed that the product of their recall would be used as the material for the next participant in the chain.

Coding. Recalled material was coded for the presence of propositions found in the original version. Coding reliability was assessed by having an independent coder, blind to the hypothesis, code 10% of the material. The independent coder and experimenter were highly consistent ($r = 0.95$, $p < .001$). In cases of disagreement the first coder’s decision stood. Sensitivity tests were conducted to assess coder reliability (see below and SM10).

Statistical analysis. To test H1 a generalised linear mixed effects model (GLMM) was used to predict the proportion of original propositions correctly recalled, with moral vs. non-moral vignette version as a fixed effect, with nested random effects of vignette in participant, participant in generation, and generation in vignette set. To test H2 and H3 a second GLMM was constructed to predict proportion of original propositions correctly recalled, including participant age, participant gender, word count, number of propositions, emotion, emotional intensity, moral good score, moral bad score, survival information score, social information score, male stereotype consistency score, female stereotype consistency score and generation as fixed effects and the same nested random effects. Predictors were removed if doing so did not impair model fit, determined by Akaike information criterion (AIC; see Burnham and Anderson, 2002). No specific ΔAIC was used to determine predictor removal but all models with $\Delta AIC < 2$ relative to the best-fitting model were included in model averaging (see SM4 for AICs of each model produced).

Analyses was conducted using the lme4 package (Bates et al., 2014) in R versions 3.1.1 (R Core Team, 2014) and 3.2.2 (R Core Team, 2016) to fit all GLMMs, with multiple pair-wise comparisons conducted using the multcomp package (Hothorn et al., 2008). Model comparisons and averaging were performed, and relative importance measures (computed by the sum of AIC weights across all of the models where the variable occurs) determined using the MuMIn package (Bartoń, 2014). All non-categorical variables were centred on the mean.

Results. The moral proposition appearing in the original material was recalled in the majority cases (77%), and the moral proposition survived to the end of the majority of chains where it was present in the original material (60%). However, moral versions of vignettes were not transmitted with higher fidelity than non-moral versions (vs. generation-only model, $X^2_1 = .12$, $p > .05$), suggesting no general bias for moral content versus non-moral content in transmission. This was also found when just examining the recall of the key proposition which varied between versions, moral versions were not transmitted with higher fidelity than non-moral equivalents ($X^2_1 = 1.88$, $p > .05$). However, a higher rating for morally good information was an important predictor of transmission fidelity, while morally bad information was not: moral good score was retained as an effective predictor in all of the three best-fitting models, as determined by AICc (see Table 1). In addition, morally good score was a better predictor of recall for the key proposition than morally bad score ($X^2 = 5.14$, $p < .001$).

Vignette emotion, word count, and participant age and gender were also important predictors of transmission fidelity. Figure 1

Table 1 Fixed effects, AICc and ΔAIC of the three best fitting models (ΔAIC < 2) produced in analyses				
Model	Fixed effects	AICc	ΔAIC	
1	Good + Emotion + Male stereotype + Age + Gender + Words + Generation	446.58	0.00	
2	Good + Emotion + Age + Gender + Words + Generation	447.22	0.64	
3	Good + Emotion + Male stereotype + Female stereotype + Age + Gender + Words + Generation	448.22	1.64	




Fig. 1 Odds ratios with confidence intervals of fixed effects predicting transmission fidelity in the best fitting model in Study 1 (Model 1 in Table 1). Odds ratios are sorted from highest to lowest, with the highest at top. Values to the right of the dashed line indicate a positive effect, values to the left of the line indicate a negative effect. Produced using sjPlot (Lüdtke, 2016). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Reference categories are: anger for emotion, generation 1 for generation, and female for gender

shows the odds ratios for the fixed effects of the best fitting model. (For details of other models, see SM4).

Comparisons of vignettes with different emotions showed a number of significant differences. Happiness vignettes were transmitted with greater fidelity than those featuring elevation (Tukey's HSD corrected $z = 2.76$, $p < 0.05$) and gratitude ($z = 2.75$, $p < 0.05$). Disgust vignettes were transmitted with greater fidelity than those featuring guilt/shame ($z = 4.44$, $p < 0.001$), happiness ($z = 3.75$, $p < 0.01$), elevation ($z = 3.79$, $p < 0.01$) and gratitude ($z = 3.39$, $p < 0.01$). Anger vignettes were transmitted with greater fidelity than those featuring happiness ($z = 2.88$, $p < 0.05$), elevation ($z = 3.24$, $p < 0.05$) and gratitude ($z = 2.93$, $p < 0.05$) (all results based on best fitting model, see SM5 for table).

A multi-model averaging approach (see Burnham and Anderson, 2002; Grueber et al., 2011) was used to determine appropriate effect estimates (see Fig. 2). Morally good score had a positive effect on transmission fidelity (estimate = 0.80 ± 0.20 SE, $z = 3.92$). Higher male stereotype consistency score had a positive effect on transmission fidelity (estimate = 0.29 ± 0.28 SE, $z = 1.06$), while higher female stereotype consistency score had a small, negative effect on transmission fidelity (estimate = -0.06 ± 0.19 SE, $z = 0.19$). Participant age (estimate = 0.09 ± 0.04 SE, $z = 2.10$) and vignette word count (estimate = 0.04 ± 0.01 SE, $z = 3.41$) also had a small but consistent positive effect on transmission fidelity. Participant gender had an effect, with women recalling more propositions, on average, than other participants (estimate = -0.88 ± 0.44 SE, $z = 2.00$). Relative variable importance measures (see Fig. 2) suggest that, of the predictors, good score, emotion, gender, word count and age were the most important in determining recall. Furthermore, moral goodness score predicted recall in every generation (see Fig. 3).

Sensitivity tests. Sensitivity tests were conducted using data from the second coder to assess the robustness of results based on data from the original coder. As in the original results, moral versions of vignettes were not transmitted with higher fidelity than non-moral versions (vs. generation-only model, $X^2_1 = .05$, $p > .05$), suggesting the finding that there is no general bias for moral content versus non-moral content is robust. In addition, a higher rating for morally good information was also an important predictor in the sensitivity test; moral good was retained in all the best fitting models while morally bad was not ($\Delta AIC < 2$) and had a positive effect on transmission fidelity (model average estimate = 0.60 ± 0.20 SE, $z = 2.97$), suggesting this finding was




Fig. 2 Predictor effect size indicated by z value and relative variable importance (maximum value = 1) from the average model based on the three best fitting models in Study 1. See SM4 for a more complete report of model-averaged coefficients. ^aIndicates a categorical variable where mean z-value is presented




Fig. 3 Predicted probabilities for the effect of good score (from mean) on recall by generation, derived from the results of the best-fitting model in Study 1 (Model 1 in Table 1). Produced using sjPlot (Lüdtke, 2016)

also robust. For more details on the results of the sensitivity tests see SM10.

Discussion. The aim of this study was to test three alternative hypotheses regarding a transmission bias for moral information. We found no evidence to support H1 (there is a cognitive bias for transmitting general moral information) or H2 (there is a cognitive bias for transmitting morally bad information). In contrast, H3 (there is a cognitive bias for transmitting morally good information) was supported: higher morally good score predicted higher transmission fidelity, and this is true in all four ‘generations’.

Based on Haidt and Joseph’s (2004) moral foundations theory, one might expect a recall advantage and a transmission bias for all morally relevant information, because any information relevant to the moral foundations would be salient. One might further expect stories about morally corrupt action to be particularly culturally potent, given previous research suggesting that negative information is more salient and viewed as more credible than equivalent positive information (Baumeister et al., 2001; Fessler et al., 2014; Rozin and Royzman, 2001). This is not what we found. Rather, the transmission bias for moral information took the form of an advantage for morally good content. While this could be considered counter-intuitive given the extant literature suggesting a negativity bias in transmission experiments and cultural artefacts (Bebbington et al., 2016; Brand et al., 2019; Fessler et al., 2014; Morin and Acerbi, 2017), tales of the morally virtuous have also been found to be culturally successful in mythology (Bierlein, 2010). Further, as previously mentioned, analyses of social media show positively valenced content to be more successful than negatively valenced content (Ferrara and Yang, 2015a; Ferrara and Yang, 2015b; Fu et al., 2016), while children display a bias towards receiving and transmitting information portraying their in-group as pro-social (Dunham et al., 2011; Over et al., 2017) and adults display a preference for transmitting positive vignettes over negative ones (van Leeuwen et al., 2018). We therefore suggest that a valence-based transmission bias is dependent on a number of contextual factors that we discuss further in the general discussion.

Results regarding the other potential predictors are largely consistent with several key findings in the literature on transmission biases. We found that vignettes featuring male-

stereotypical content had high transmission, supporting Lyons and Kashima’s (2006) findings. We found, however, that while male stereotype consistency had a positive effect on transmission fidelity, female stereotype consistency did not. There are two possible explanations for this. First, out-group members are perceived to be more stereotypic than in-group members (Park and Rothbart, 1982). Because most of our participants were female, stories about men might have been more likely to be treated as out-group stories, leading male gender stereotypes to be more readily accessed. Second, perceived stereotype sharedness plays a crucial role in how communicable stereotype-consistent content is (Clark and Kashima, 2007), and gender stereotypes about men are more homogenous than those about women, increasing the chance that they will be consistently shared (Wood and Eagly, 2012). It is also worth noting that the effects of gender stereotypes were less important than other predictors in terms of predicting transmission fidelity.

We found that vignettes featuring anger were more effectively transmitted than vignettes featuring other emotions. This is consistent with previous research showing a transmission advantage for content featuring ‘activating’ emotions rather than ‘deactivating’ ones (Berger and Milkman, 2010). We also found that vignettes featuring disgust had higher transmission fidelity than a number of other emotions, which is consistent with research demonstrating an advantage for disgusting content in transmission (e.g., Eriksson and Coultas, 2014). Our effects of participant age were very small in size and do not allow conclusions to be drawn owing to the restricted range and strongly skewed distribution of participant age. With regards to word count positively predicting transmission fidelity, participants may have put more effort into recalling the longer vignettes because they expected them to be harder to remember, or their larger size may have made them a more memorable package than shorter vignettes. Finally, because all but eight of our participants were female, gender was not a priori a variable of interest, and there is little existing literature on gender and social transmission in adults, we refrain from drawing conclusions as to the influence of gender. One potential limitation of this study is that an independent coder coded only a subset of the material. To assess the reliability of the results sensitivity tests were conducted (see results section and SM10). These tests suggest the key findings are robust.

Study 2

Study 1 demonstrated that emotion played an important role in the transmission of moral content. It has long been understood that emotion plays a key role in our moral experience and the evolution of our moral sense (Gintis et al., 2008; Graham et al., 2013; Norenzayan, 2014; Pizarro, 2000). As mentioned previously, previous research in cultural transmission has found an emotional content bias (Eriksson and Coultas, 2014; Heath et al., 2001; Stubbersfield et al., 2017). These studies propose that the mechanism explaining emotional content bias is that more emotional content elicits a greater physiological response, resulting in enhanced selection, recall and transmission, however, their measures of emotional arousal are self-reported ratings of emotional content, rather than any physiological measure of arousal. To our knowledge only one study has directly examined an interaction between physiological arousal and cultural transmission, Berger (2011) found that a physiological excitatory state (induced through physical exercise) increased the sharing of information despite the arousal being incidental to the material being shared.

The present study seeks to replicate and extend on Study 1 by including a measure of physiological arousal: electrodermal

activity (EDA). EDA measurement has been extensively used to examine physiological arousal in response to emotional content (Boucsein, 2012). Study 2 also builds on Study 1 by including consequences in the narratives, either a reward for morally good actions or a punishment for morally bad actions. Previous research demonstrates that the consequences of an action play a key role in moral judgement and that moral scenarios involving consequences elicit unique activity in the brain compared to nonmoral scenarios (Schaich Borg et al., 2006). We therefore hypothesise that vignettes which include the consequences of a moral action will be more faithfully transmitted than those which do not. Based on the results of Study 1 and previous studies examining emotion and the transmission of narratives we will examine four hypotheses:

1. There is a cognitive bias for transmitting morally good information. Consistent with Study 1 and research suggesting a bias for transmitting positive information in experimental settings and on social media, morally good information will be preferentially transmitted over non-moral or morally bad information (H4).

2. There is a bias for transmitting more physiologically arousing content. Consistent with previous research suggesting a bias for more emotional content, content which evokes a stronger physiological response (as measured using EDA) will be more faithfully transmitted along a linear chain (H5).

3. There is a bias for transmitting narratives which feature a consequence for a moral action. Stories which present either a reward for a morally good action or punishment for morally bad action will be more faithfully transmitted than those which do not (H6).

4. Self-reported emotion ratings provide an adequate proxy for actual emotional arousal. In order to provide an assessment of previous research which has used self-report measures, we will compare the EDA measures with self-report measures (H7).

Materials and methods

Participants. Thirty-six participants (27 female, 9 male) aged 18 to 49 years ($M = 22.61$, $SD = 5.40$) took part. All participants gave their informed consent.

Materials. Vignettes were 30 to 68 words and contained 3 to 10 propositions (determined through propositional analysis (Kintsch, 1974)). As in Study 1, moral and non-moral versions were created. In addition, versions with consequences and without consequences were included. See below for examples (bolded sections illustrate differences between versions and did not appear as such to participants).

Smoothie-Non-moral version

Jackie's partner read an interview with a famous actor who drank urine for its supposed health benefits. **He decided to try it out on himself.** He added some urine to his breakfast smoothie. He didn't see a problem—he hadn't noticed any difference in the taste. Jackie felt sickened at the thought.

Smoothie-Non-moral version with consequence

Jackie's partner read an interview with a famous actor who drank urine for its supposed health benefits. **He decided to try it out on himself.** He added some urine to his breakfast smoothie. He didn't see a problem—he hadn't noticed any difference in the taste. Jackie felt sickened at the thought. **Jackie and her partner had a big fight about it.**

Smoothie-Moral version

Jackie's partner read an interview with a famous actor who drank urine for its supposed health benefits. **He decided to try it out on the kids.** He added some urine to their breakfast

smoothie. He didn't see a problem—they hadn't noticed any difference in the taste. Jackie felt sickened at the thought.

Smoothie-Moral version with consequences

Jackie's partner read an interview with a famous actor who drank urine for its supposed health benefits. **He decided to try it out on the kids.** He added some urine to their breakfast smoothie. He didn't see a problem—they hadn't noticed any difference in the taste. Jackie felt sickened at the thought. **Jackie and her partner had a big fight about it.**

As in Study 1, the results of a pre-test were used to determine the eight vignettes most appropriate for analyses, using the same survey as the Study 1 pre-test (see SM1, SM6, and SM7).

EDA equipment and measurement. EDA was measured using a Biopac MP36R system operating AcqKnowledge 4.4 software, sampling at 50 Hz, with a gain level of $\times 1000$ and a high pass filter of .05. EDA responses for each participant were identified as the maximum phasic skin conductance response measured while reading each vignette (no recordings were taken when typing or not reading). To correct for individual differences, response scores were range corrected, that is each score was calculated as a proportion of that participants total EDA range. Settings and measurement were informed by Braithwaite et al. (2013) and pilot studies conducted by the researchers.

Design. As in Study 1, a linear transmission chain design was used. As in previous research, three 'generations' were used (Barrett and Nyhoff, 2001; Nielson et al., 2012; Stubbersfield et al., 2015; Stubbersfield et al., 2017) across twelve chains. The first participant in each of the twelve chains received all eight vignettes. Participants received an equal mix of moral, non-moral, moral with consequences and non-moral with consequences.

Procedure. The procedure matched that of Study 1, with the addition of EDA measurement. Participants attached pre-gelled electrodes to their own palms with instruction from the researcher, who also checked they were attached appropriately. The researcher remained in the room to monitor the EDA readout and participant activity for actions which could influence the EDA measurement (i.e., coughing, excessive movement).

Coding. Recalled material was coded for the presence of propositions found in the original version. Coding reliability was assessed by having an independent coder, blind to the hypothesis, code 11% of the material. The independent coder and experimenter were highly consistent ($r = 0.92$, $p < .001$). In cases of disagreement the first coder's decision stood. Sensitivity tests were conducted to assess coder reliability (see below and SM10).

Statistical analysis. To test H4 and H5 analyses followed that of Study 1 using the same software and R packages. A GLMM was constructed to predict the proportion of original propositions correctly recalled, including participant age, participant gender, word count, number of propositions, moral good score, moral bad score, survival information score, social information score, male stereotype consistency score, female stereotype consistency score, emotion, EDA score, and generation as fixed effects, with nested effects of vignette in participant, participant in generation, and generation in vignette set. Predictors were removed if doing so did not impair model fit, determined by AIC. As before, no specific ΔAIC was used to determine predictor removal but all models with $\Delta AIC < 2$ relative to the best-fitting model were included in model averaging (see SM4 for AICs of each model produced). All non-categorical variables were centred on the mean.

Table 2 Fixed effects, AICc and ΔAIC of the five best fitting models (ΔAIC < 2) produced in the analysis

Model	Fixed effects	AICc	ΔAIC
1	Good + Emotion + EDA score + Gender + Words + Generation	352.98	0.00
2	Good + Emotion + Gender + Words + Generation	353.00	0.02
3	Good + Emotion + EDA score + Gender + Generation	353.01	0.04
4	Good + Emotion + EDA score + Gender + Words + Generation + Female stereotype	353.25	0.28
5	Good + Emotion + EDA score + Gender + Words + Generation + Female stereotype + Male Stereotype	354.26	1.29




Fig. 4 Odds ratios with confidence intervals of fixed effects predicting transmission fidelity in the best fitting model in Study 2 (Model 1 in Table 2). Odds ratios are sorted from highest to lowest, with the highest at top. Values to the right of the dashed line indicate a positive effect, values to the left of the line indicate a negative effect. Produced using sjPlot (Lüdtke, 2016). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Reference categories are: anger for emotion, generation 1 for generation, and female for gender

To test H6 a GLMM was constructed to predict the proportion of original propositions correctly recalled with moral category (categories being moral-no-consequences (M), moral-consequences (MC), non-moral-no-consequences (N), and non-moral-consequences (NC)) as a fixed effect, with the same random effects as the H4 and H5 GLMM. To test H7 a version of the best fitting model was created in which EDA score was replaced with emotional rating score, and the effect on model fit evaluated.

Results. As in Study 1, a higher rating for morally good information was an important predictor of transmission fidelity: moral good score was retained as an effective predictor in all of the five best-fitting models, as determined by AICc. EDA score was less important as a predictor of transmission fidelity, being retained in four of the five best-fitting models (see Table 2). The moral proposition appearing in the original material was recalled in the majority cases (77.08%), and the moral proposition survived to the end of the majority of chains where it was present in the original material (58.33%). Morally good score was also a better predictor of recall for the key proposition than morally bad score ($X^2 = 0.59$, $p < .001$).

Vignette emotion and participant gender were also important predictors of transmission fidelity. Figure 4 shows the odds ratios for the fixed effects of the best fitting model in Study 2 (Model 1 in Table 2). (For details of other models, see SM8).

Pairwise comparisons of vignettes featuring different emotions showed that disgust vignettes were transmitted more faithfully than vignettes featuring either gratitude (Tukey's HSD corrected $z = -3.43$, $p < 0.005$) or elevation ($z = -3.13$, $p < 0.01$). No other




Fig. 5 Predictor effect size indicated by z value and relative variable importance (maximum value = 1) from the average model based on the five best fitting models for Study 2. See SM8 for a more complete report of model-averaged coefficients. ^aIndicates a categorical variable where mean z-value is presented

significant differences between emotions were found (z s = -0.57 to 1.28 , p s > 0.05). (See SM5).

A multi-model averaging approach (see Study 1) was used to determine appropriate effect estimates (see Fig. 5). Morally good score had a positive effect on transmission fidelity (estimate = 0.38 ± 0.11 SE, $z = 3.56$) while EDA score had a less consistent but negative effect on transmission fidelity (estimate = -0.70 ± 0.66 SE, $z = 1.06$). Participant gender had an effect, with women recalling more propositions, on average, than other participants (estimate = -1.13 ± 0.33 SE, $z = 3.38$). Vignette word count had a small negative effect (estimate = -0.03 ± 0.03 SE, $z = 1.12$). Higher female stereotype consistency score had a small positive effect on transmission fidelity (estimate = 0.16 ± 0.29 SE, $z = 0.54$), while higher male stereotype consistency score had a small negative effect on transmission fidelity (estimate = -0.03 ± 0.12 SE, $z = 0.27$). Relative variable importance measure (see Fig. 5) suggest that, of the predictors, good score, emotion and gender were the most important in determining recall. Each of these variables was also retained in the best-fitting models from Study 1 with different participants and material. The effect of good score was again present in all generations, as it was in Study 1 (see Fig. 6).

To test H7 a version of the best fitting model was created using emotional rating score in place of EDA score. In this model emotional rating also had a negative effect on transmission fidelity (estimate = -0.57 ± 0.37 SE, $z = -1.52$). When these two models were compared, the model with emotional rating score proved to be a better fit to the data than the model using EDA score ($X^2 = 0.14$, $p < 0.001$), however, the difference in model fit is small ($\Delta AIC < 2$).

Consequences (Moral and Nonmoral vs Moral with Consequences and Nonmoral with Consequences) is a significant predictor of transmission fidelity (vs. generation-only model, $X^2_3 = 10.31$, $p = .02$), however, vignettes without consequences




Fig. 6 Predicted probabilities for the effect of good score (from mean) on recall by generation, derived from the results of the best-fitting model in Study 2 (Model 1 in Table 2). Produced using sjPlot (Lüdtke, 2016)

(moral and non-moral) were transmitted more faithfully than their no-consequences equivalents (estimate = 0.55, SE = 0.26, $z = 2.15$, $p < 0.05$). Multiple pair-wise comparisons show that this primarily driven by the difference between Moral and Nonmoral with Consequences vignettes ($z = -3.04$, $p < 0.05$), no other significant differences between vignette types were found (z s = -2.18 to -0.00 , p s > 0.05).

Sensitivity tests. Sensitivity tests were conducted using data from the second coder to assess the robustness of results based on data from the original coder. As in the original results, morally good information was an important predictor, being retained in all the best fitting models, while EDA Score was not ($\Delta AIC < 2$). In addition, morally good score again had a positive effect on transmission fidelity (model average estimate = 0.36 ± 0.10 SE, $z = 3.81$) while EDA score had a less consistent but negative effect on transmission fidelity (model average estimate = -0.24 ± 0.50 SE, $z = 0.48$), suggesting these findings are also robust. Consequences were not a significant predictor of transmission (vs generation only model, $X^2_3 = 6.63$, $p = .08$). Emotional rating score proved to be a better fit to the data than the model using EDA score ($X^2 = 5.38$, $p < 0.001$) but both were negative predictors of transmission fidelity. For more details on the results of the sensitivity tests see SM10.

Discussion. The aim of this study was to test four hypotheses regarding a transmission bias for morally good content and the influence of physiological arousal on transmission. Consistent with the findings of Study 1, we found evidence to support H4 (there is a cognitive bias for transmitting morally good information). As in Study 1, this was true across all 'generations' (See Fig. 6). We found no evidence to support H5 (there is a bias for transmitting more physiologically arousing content), in fact the opposite was found. We also found no evidence to support H6 (there is a bias for transmitting narratives which feature a consequence for a moral action). We did, however, find evidence in support of H7 (self-reported emotion ratings provide an adequate proxy for actual emotional arousal). Self-reported emotion ratings negatively predict recall, as EDA measures do, suggesting that it is appropriate to use self-report measures for emotional content in future research.

Previous research examining an emotional bias in cultural transmission has found that emotive information (as determined through self-report measures) has a positive influence on transmission fidelity (Eriksson and Coultas, 2014; Heath et al., 2001; Stubbersfield et al., 2017). The explanation given for this finding is that emotive information is physiologically arousing,

which makes it more memorable and more likely to be selected for transmission. Our finding that physiological arousal (as measured through EDA) had a negative effect on transmission fidelity is inconsistent with this proposed mechanism and therefore with this previous research. To examine this finding further, in the context of our study, another GLMM was constructed predicting EDA score rather than recall, but otherwise the same as those models constructed to predict recall. A full model found that morally good content had a negative effect on EDA score (estimate = -0.36 ± 0.21 SE, $z = 0.08$), suggesting the vignettes with a higher morally good score were less physiologically arousing than other vignettes.

We therefore propose that the role of emotion and physiological arousal in cultural transmission is not a direct one, but one that increases transmission by making content more salient when memorability is the primary determinant of transmission success. Clark and Kashima (2007) demonstrated that participants' knowledge of their recall being transmitted to another person in a transmission chain produced different results from chains where they were not aware. They argued that participants' awareness of transmission led to communicative intent, producing a different result relative to the recall-only chains. In our experiment, transmission to another participant was also known, likely leading to a combination of recall and communicative intent where memorability (and hence arousal) may not have been the most influential factor. We suggest that less arousing content did not have a transmission advantage in our study because it was less arousing, but rather because the content bias for morally good content was a more important determinant of transmission success than memorability. This proposal is supported by van Leeuwen et al. (2018), who found that participants more frequently chose to transmit positive, low-arousal vignettes over negative, high arousal ones when transmission was to strangers. Further research is required to examine fully the role of arousal, memory, communicative intent and audience perception in the transmission of moral content and suggestions for future research are discussed in the general discussion.

We found no support for the hypothesis that vignettes which featured a consequence for moral actions would be more faithfully transmitted. The results instead suggested that vignettes with consequences were less faithfully transmitted than those without consequences. However, pairwise comparison suggests that this finding was driven by the differences in transmission fidelity between Moral vignettes and Nonmoral with Consequences vignettes. A key finding in studies examining the cultural transmission of narratives has been the process of sense-making and stories being altered to fit the transmitters' schema (Bartlett, 1932; Lyons and Kashima, 2006), therefore, that Nonmoral with Consequences vignettes were not faithfully transmitted may be better explained by the apparent incongruity between the non-moral, unintentional action having some form of punishment or reward as a consequence of it, than stories without consequences having some form of advantage in transmission. An appropriate interpretation of this result, therefore, is that in this case the inclusion of consequences had no clear direct effect, positive or negative, on the transmission of moral information.

We found that vignettes featuring female-stereotypical content had high transmission, supporting Lyons and Kashima's (2006) findings. We found, however, that while female stereotype consistency had a positive effect on transmission fidelity, male stereotype consistency did not. This finding is inconsistent with the finding of Study 1. It is again worth noting that, as in Study 1, the effects of gender stereotypes were less important than other predictors in terms of predicting transmission fidelity. It is

possible, therefore, that both studies give valid but noisy estimates of the same, small, ‘true’ effect of stereotype consistency. We also found that vignettes featuring disgust had higher transmission fidelity than a number of other emotions, which is consistent with research demonstrating an advantage for disgusting content in transmission (e.g., Eriksson and Coultas, 2014). Unlike Study 1, word count had a negative effect on transmission fidelity, with longer vignettes being more poorly transmitted. However, in Study 2, word count was confounded with the presence of consequences. This is likely to explain the inconsistency between Study 1 and Study 2 with regard to the effects of word count. Because all but nine of our participants were female, gender was not a priori a variable of interest, and there is little existing literature on gender and social transmission in adults, we refrain from drawing conclusions as to the influence of gender. To assess the reliability of the results, sensitivity tests were again conducted (see results section and SM10), and again found that the key findings are robust.

General discussion

The primary purpose of the studies presented here was to examine the transmission of moral information, while examining the influence of physiological arousal, emotion, and other predictors on that transmission. While we found no evidence of a general bias for moral information in transmission, in two studies with separate participants and materials we found an advantage for information rated as morally good. This could be considered counter-intuitive in the context of the negativity bias literature (Baumeister et al., 2001; Brand et al., 2019; Fessler et al., 2014; Morin and Acerbi, 2017; Rozin and Royzman, 2001), however, we now consider a possible functional reason for the preferential sharing of morally good content.

Stories of morally good actions might be preferentially transmitted because they appear to reinforce a shared moral norm, thereby serving a socially connective function between potential affiliates. Clark and Kashima (2007) argued that the transmission advantage for stereotype-consistent information arose because sharing stereotypes facilitated a perceived bond between participants. Our participants were aware that the products of their recall would be read and recalled by another participant. It may be that participants unconsciously enhanced their transmission of the morally good stories in order to convey a positive impression of themselves and their social norms to the next individual in the chain. Socially connective content may be more likely to be preferentially transmitted when building new social connections is important (i.e., receivers are unknown). Morality is important in impression formation and categorisation (van Leeuwen et al., 2012; Wojciszke et al., 1998; Ybarra et al., 2001), and impression formation is particularly prone to negativity bias (Lupfer et al., 2000; Rozin and Royzman, 2001). Participants may therefore be selectively transmitting positively valenced descriptions of morally good actions to avoid the contaminating effect of negatively valenced descriptions of morally bad actions on the impressions formed of them by strangers. We provide support for this hypothesis in SM9, which presents a supplementary study using the same material in which we find no advantage for moral information (whether good or bad) in individual recall. This supports the suggestion that our main result was driven by communicative intent rather than the memorability of morally good information. Further, an impression-driven mechanism behind the preferential transmission of morally good information also can also explain the finding that the inclusion of consequences had no effect (in Study 2). If morally good information is being preferentially transmitted to avoid the contaminating effect of morally bad

information, then whether that moral action is rewarded or not is irrelevant.

Our interpretation is supported by previous research which suggests that the expression of self-identity plays an important role in moral expression (Crockett, 2017). It is particularly supported by van Leeuwen et al. (2018), who found that participants more frequently chose to transmit positive, low-arousal vignettes over negative, high arousal ones when transmission was to strangers, suggesting that seeking social connection might influence transmission preferences. Further, unlike the negativity bias found in trends in emotional expression in art (see Brand et al., 2019; Morin and Acerbi, 2017), examination of trends in the cultural salience of moral terms shows no clear, linear trend (Wheeler et al., 2019), suggesting that the transmission of moral information may be particularly dependant on contextual factors. Further research is required to examine fully the role of arousal, memory, communicative intent and audience perception in the transmission of moral content. A limitation of the current study is the use of the vignettes’ original properties as a predictor of recall along a transmission chain and future research examining the transmission of moral information could address this. A purely selection-based paradigm with variable audiences could be used to assess effects of communicative intent and audience perception (similarly to van Leeuwen et al., 2018) or the relative effects of memorability and communicative intent could be examined through comparing recall-only and transmission-aware chains (as in Clark and Kashima, 2007). Future research could also allow participants to deliberately alter text (as in Stubbersfield et al., 2018) to further assess the influence of communicative intent on the transmission of moral information.

Our results show that the preferential transmission of morally good content is not simply the result of a bias for social content. It is also distinct from a bias associated for emotional content, as indicated by separable effects in our models. With regard to emotions, there were few clear patterns, except for the more unambiguously ‘activating’ emotions having an advantage over other, less ‘activating’ emotions. However, that EDA score had a negative effect on transmission fidelity suggests that arousal may not be the mechanism behind this finding. Our results regarding EDA suggest that the role of emotion in social transmission is that it interacts with other, perhaps more important factors including the context in which transmission takes place and the nature of informational content. In the case of Study 2, morally good information was more transmissible, likely because of the reasons discussed above. While our results did suggest that self-reported emotional ratings of content are a useful measure, we would encourage further research using physiological measures to properly assess the mechanisms involved in the transmission of emotional material.

Our studies provide the first evidence of a content bias in social transmission that specifically favours content related to morality. Further, our results suggest that the content bias specifically favours morally good information. We propose that this bias likely functions as a *transmitter bias*, favouring the transmission to another individual of content that may serve a socially connective purpose, rather than as a *receiver bias*, favouring the encoding and recall of content which may be advantageous at the individual level. The selective communication of content relating to moral virtue might avoid negative impression formation and promote social bonding, and this might partially explain the success of morality-related content in human culture. It should be noted that while the bias for moral content can be distinguished from a bias for social or emotional content in the lab, these biases are unlikely to operate completely separately in real-world contexts. Moral information is often emotionally charged, social, and survival-relevant.

Data availability

The datasets generated and analysed during the current study are available in the Dataverse repository: <https://doi.org/10.7910/DVN/LADWHL>. Data produced from the EDA measurement are not publicly available due to their large size but are available from the corresponding author on reasonable request.

Received: 28 February 2019 Accepted: 16 May 2019

Published online: 04 June 2019

References

- Ayala FJ (2010) The difference of being human: Morality. *Proc Natl Acad Sci* 107 (Suppl.2):9015–9022. <https://doi.org/10.1073/pnas.0914616107>
- Barrett JL, Nyhof MA (2001) Spreading non-natural concepts: the role of intuitive conceptual structures in memory and transmission of cultural materials. *J Cogn Cult* 1(1):69–100. <https://doi.org/10.1163/156853701300063589>
- Bartlett FC (1932) *Remembering*. Macmillan, Oxford
- Bartón K (2014) MuMIn: multi-model inference. R package version 1.10.5. <http://CRAN.R-project.org/package=MuMIn>
- Bates D, Maechler M, Bolker B, Walker S (2014) lme4: linear mixed-effects models using Eigen and S4. R package version 1.1-7. <http://CRAN.R-project.org/package=lme4>
- Baumeister RF, Bratslavsky E, Finkenauer C, Vohs KD (2001) Bad is stronger than good. *Rev Gen Psychol* 5(4):323. <https://doi.org/10.1037/1089-2680.5.4.323>
- Bebbington K, MacLeod C, Ellison TM, Fay N (2016) The sky is falling: evidence of a negativity bias in the social transmission of information. *Evol Hum Behav* 38(1):92–101. <https://doi.org/10.1016/j.evolhumbehav.2016.07.004>
- Berger J (2011) Arousal increases social transmission of information. *Psychol Sci* 22 (7):891–893. <https://doi.org/10.1177/0956797611413294>
- Berger J, Milkman K (2010) Social transmission, emotion, and the virality of online content. Wharton Research Paper 106. <http://opim.wharton.upenn.edu/~kmilkman/Virality.pdf>
- Bierlein JF (2010) *Parallel myths*. Ballantine Books, New York
- Boucsein W (2012) *Electrodermal activity*, 2nd edn. Springer Science+Business Media, New York. <https://doi.org/10.1007/978-1-4614-1126-0>
- Boyd R, Richerson PJ (1985) *Culture and the evolutionary process*. University of Chicago, University of Chicago Press, Chicago
- Braithwaite JJ, Watson DG, Jones R, Rowe M (2013) A guide for analysing electrodermal activity (EDA) & skin conductance responses (SCRs) for psychological experiments. *Psychophysiology* 49(1):1017–1034
- Brand CO, Acerbi A, Mesoudi A (2019) Cultural evolution of emotional expression in 50 years of song lyrics. Preprint at <https://osf.io/preprints/socarxiv/3j6wx/>.
- Burnham KP, Anderson DR (2002) *Model selection and multimodel inference: a practical information-theoretic approach*. Springer Science and Business Media, New York
- Clark AE, Kashima Y (2007) Stereotypes help people connect with others in the community: a situated functional analysis of the stereotype consistency bias in communication. *J Personal Soc Psychol* 93(6):1028–1039. <https://doi.org/10.1037/0022-3514.93.6.1028>
- Crockett MJ (2017) Moral outrage in the digital age. *Nat Hum Behav* 1 (11):769–771. <https://doi.org/10.1038/s41562-017-0213-3>
- Dunham Y, Baron AS, Carey S (2011) Consequences of “minimal” group affiliations in children. *Child Dev* 82(3):793–811. <https://doi.org/10.1111/j.1467-8624.2011.01577.x>
- Eriksson K, Coultas JC (2014) Corpses, maggots, poodles and rats: emotional selection operating in three phases of cultural transmission of urban legends. *J Cogn Cult* 14(1–2):1–26. <https://doi.org/10.1163/15685373-12342107>
- Ferrara E, Yang Z (2015a) Measuring emotional contagion in social media. *PLoS ONE* 10(11):e0142390. <https://doi.org/10.1371/journal.pone.0142390>
- Ferrara E, Yang Z (2015b) Quantifying the effect of sentiment on information diffusion in social media. *Peer J Comput Sci* 1:e26. <https://doi.org/10.7717/peerj-cs.26>
- Fessler DM, Pisor AC, Navarrete CD (2014) Negatively-biased credulity and the cultural evolution of beliefs. *PLoS ONE* 9(4):e95167. <https://doi.org/10.1371/journal.pone.0095167>
- Fu P, Lin Z, Lin H, Yuan F, Wang W, Meng D (2016) Quantifying the effect of sentiment on topic evolution in chinese microblog. In: Li F, Shim K, Zheng K, Liu G (eds) *Web technologies and applications*. APWeb 2016. Lecture Notes in Computer Science, Vol. 9931. Springer, Cambridge
- Gintis H, Henrich J, Bowles S, Boyd R, Fehr E (2008) Strong reciprocity and the roots of human morality. *Soc Justice Res* 21(2):241–253. <https://doi.org/10.1007/s11211-008-0067-y>
- Graham J, Haidt J, Koleva S, Motyl M, Iyer R, Wojcik SP, Ditto PH (2013) Moral foundations theory: The pragmatic validity of moral pluralism. *Advances in Experimental Social Psychology* 47:55–130. <https://doi.org/10.1016/B978-0-12-407236-7.00002-4>
- Grueber CE, Nakagawa S, Laws RJ, Jamieson IG (2011) Multimodel inference in ecology and evolution: challenges and solutions. *J Evolut Biol* 24(4):699–711. <https://doi.org/10.1111/j.1420-9101.2010.02210.x>
- Haidt J, Joseph C (2004) Intuitive ethics: how innately prepared intuitions generate culturally variable virtues. *Daedalus* 133(4):55–66. <https://doi.org/10.1162/0011526042365555>
- Heath C, Bell C, Sternberg E (2001) Emotional selection in memes: the case of urban legends. *J Personal Soc Psychol* 81(6):1028–1041. <https://doi.org/10.1037/0022-3514.81.6.1028>
- Henrich J, McElreath R (2003) The evolution of cultural evolution. *Anthropology* 12(3):123–135. <https://doi.org/10.1002/evan.10110>
- Hothorn T, Bretz F, Westfall P (2008) Simultaneous inference in general parametric models. *Biom J* 50(3):346–363. <https://doi.org/10.1002/bimj.200810425>
- Kintsch W (1974) *The Representation of meaning in memory*. Erlbaum, Oxford
- Laland KN, Brown GR (2011) *Sense and nonsense: evolutionary perspectives on human behaviour*. Oxford University Press, Oxford
- Lüdtke D (2016) sjPlot: Data visualization for statistics in social science. R package version 2.2.0. <https://CRAN.R-project.org/package=sjPlot>
- Lupfer MB, Weeks M, Dupuis S (2000) How pervasive is the negativity bias in judgments based on character appraisal? *Personal Soc Psychol Bull* 26 (11):1353–1366. <https://doi.org/10.1177/0146167200263004>
- Lyons A, Kashima Y (2006) Maintaining stereotypes in communication: Investigating memory biases and coherence-seeking in storytelling. *Asian J Soc Psychol* 9(1):59–71. <https://doi.org/10.1111/j.1467-839X.2006.00184.x>
- McNamara RA, Willard AK, Norenzayan A, Henrich J (2019) Weighing outcome vs. intent across societies: How cultural models of mind shape moral reasoning. *Cognition* 182:95–108. <https://doi.org/10.1016/j.cognition.2018.09.008>
- Mesoudi A (2016) Cultural evolution: a review of theory, findings and controversies. *Evolut Biol* 43(4):481–497. <https://doi.org/10.1007/s11692-015-9320-0>
- Mesoudi A, Whiten A (2004) The hierarchical transformation of event knowledge in human cultural transmission. *J Cogn Cult* 4(1):1–24. <https://doi.org/10.1348/000712605X85871>
- Mesoudi A, Whiten A (2008) The multiple roles of cultural transmission experiments in understanding human cultural evolution. *Philos Trans R Soc B* 363 (1509):3489–3501. <https://doi.org/10.1098/rstb.2008.0129>
- Mesoudi A, Whiten A, Dunbar R (2006) A bias for social information in human cultural transmission. *Br J Psychol* 97(3):405–423. <https://doi.org/10.1348/000712605X85871>
- Morin O, Acerbi A (2017) Birth of the cool: a two-centuries decline in emotional expression in Anglophone fiction. *Cogn Emot* 31(8):1663–1675. <https://doi.org/10.1080/02699931.2016.1260528>
- Nairne JS (2010) Adaptive memory: evolutionary constraints on remembering. *Psychol Learn Motiv* 53:1–32. [https://doi.org/10.1016/S0079-7421\(10\)53001-9](https://doi.org/10.1016/S0079-7421(10)53001-9)
- Nielson M, Cucchiari J, Mohamedally J (2012) When the transmission of culture is child's play. *PLoS ONE* 7:e34066. <https://doi.org/10.1371/journal.pone.0034066>
- Norenzayan A (2014) Does religion make people moral? *Behaviour* 151 (2–3):365–384. https://doi.org/10.1163/9789004263888_018
- Over H, Eggleston A, Bell J, Dunham Y (2017) Young children seek out biased information about social groups. *Dev Sci* 21(3):e12580. <https://doi.org/10.1111/desc.12580>
- Park B, Rothbart M (1982) Perception of out-group homogeneity and levels of social categorization: Memory for the subordinate attributes of in-group and out-group members. *J Personal Soc Psychol* 42(6):1051–1068. <https://doi.org/10.1037/0022-3514.42.6.1051>
- Pizarro D (2000) Nothing more than feelings? The role of emotions in moral judgment. *J Theory Soc Behav* 30(4):355–375. <https://doi.org/10.1111/1468-5914.00135>
- Pizarro DA, Laney C, Morris EK, Loftus EF (2006) Ripple effects in memory: judgments of moral blame can distort memory for events. *Mem Cogn* 34 (3):550–555. <https://doi.org/10.3758/BF03193578>
- R Core Team (2014) *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/>
- R Core Team (2016) *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/>
- Rozin P, Royzman EB (2001) Negativity bias, negativity dominance, and contagion. *Personal Soc Psychol Rev* 5(4):296–320. https://doi.org/10.1207/S15327957PSPR0504_2
- Schaich Borg J, Hynes C, Van Horn J, Grafton S, Sinnott-Armstrong W (2006) Consequences, action, and intention as factors in moral judgments: an fMRI investigation. *J Cogn Neurosci* 18(5):803–817. <https://doi.org/10.1162/jocn.2006.18.5.803>
- Stubbersfield JM, Flynn EG, Tehrani JJ (2018) Cognitive evolution and the transmission of popular narratives: a literature review and application to

- urban legends. *Evolut Stud Imaginative Cult* 1(1):121–136. <https://doi.org/10.26613/esic/1.1.20>
- Stubbersfield JM, Tehrani JJ, Flynn EG (2015) Serial killers, spiders and cybersex: social and survival information bias in the transmission of urban legends. *Br J Psychol* 106(2):288–307. <https://doi.org/10.1111/bjop.12073>
- Stubbersfield JM, Tehrani JJ, Flynn EG (2017) Chicken tumours and a fishy revenge: evidence for emotional content bias in the cumulative recall of urban legends. *J Cogn Cult* 17(1–2):12–26. <https://doi.org/10.1163/15685373-12342189>
- Stubbersfield J, Tehrani J, Flynn E (2018) Faking the news: intentional guided variation reflects cognitive biases in transmission chains without recall. *Cult Sci J* 10(1):54–65. <https://doi.org/10.5334/csci.109>
- van Leeuwen F, Park JH, Penton-Voak IS (2012) Another fundamental social category? Spontaneous categorization of people who uphold or violate moral norms. *J Exp Soc Psychol* 48(6):1385–1388. <https://doi.org/10.1016/j.jesp.2012.06.004>
- van Leeuwen F, Parren N, Miton H, Boyer P (2018) Individual choose-to-transmit decisions reveal little preference for transmitting negative or high-arousal content. *J Cogn Cult* 18(1–2):124–153. <https://doi.org/10.1163/15685373-12340018>
- Walker CJ, Blaine B (1991) The virulence of dread rumours: a field experiment. *Lang Commun* 11(4):291–297. [https://doi.org/10.1016/0271-5309\(91\)90033-R](https://doi.org/10.1016/0271-5309(91)90033-R)
- Wheeler MA, McGrath MJ, Haslam N (2019) Twentieth century morality: the rise and fall of moral concepts from 1900 to 2007. *PLoS ONE* 14(2):e0212267. <https://doi.org/10.1371/journal.pone.0212267>
- Wojciszke B, Bazinska R, Jaworski M (1998) On the dominance of moral categories in impression formation. *Personal Soc Psychol Bull* 24(12):1251–1263. <https://doi.org/10.1177/01461672982412001>
- Wood W, Eagly AH (2012) Biosocial construction of sex differences and similarities in behaviour. *Advances in Experimental Social Psychology* 46: 55–123. <https://doi.org/10.1016/B978-0-12-394281-4.00002-7>
- Ybarra O, Chan E, Park D (2001) Young and old adults' concerns about morality and competence. *Motiv Emot* 25(2):85–100. <https://doi.org/10.1023/A:1010633908298>

Acknowledgements

We would like to thank Claire Shirley and Carmen Granito for their contributions as independent coders of the material and the John Templeton Foundation for financial support (Grant #40128 to K.N.L.).

Additional information

The online version of this article (<https://doi.org/10.1057/s41599-019-0269-y>) contains supplementary material, which is available to authorized users.

Competing interests: The authors declare no competing interests.

Reprints and permission information is available online at <http://www.nature.com/reprints>

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019