Goal inference in moral narratives

Judy S. Kim[1,2]*, Clara Colombatto[3]*, & M.J. Crockett[1,2]

[1] University Center for Human Values

[2] Department of Psychology

Princeton University

[3] Department of Experimental Psychology

University College London

Author Note

Correspondence address: Judy Kim (judyseinkim@gmail.com), M.J. Crockett (mj.crockett@princeton.edu), Department of Psychology, Princeton University, Peretsman-Scully Hall 310, Princeton, NJ, USA, 08540.

Declaration of interests: None

Abstract

We often form beliefs about others based on narratives they tell about their own moral actions. When constructing such moral narratives, narrators balance multiple goals, such as conveying accurate information about what happened ('informational goals') and swaying audiences' impressions about their moral characters ('reputational goals'). Here, we ask to what extent audiences' detection of narrators' reputational goals guide or prevent them from making moral character judgments intended by narrators. Across two pre-registered experiments, "audiences" read narratives written by real "narrators" about their own moral actions. Each narrator was incentivized to write about the same action twice while trying to appear like a morally good or bad person (positive and negative reputational goals). Audiences detected narrators' reputational goals with high accuracy and made judgments about moral character that aligned with narrators' goals. However, audiences were more vigilant toward positive than negative reputational goals, requiring more evidence of high informational goals. These results demonstrate how audiences' inferences of reputational goals can both support and hinder narrators: accurate goal recognition increases the chance that audiences will make judgments intended by narrators, but inferred positive reputational goals can lead to vigilance about accuracy. More generally, this provides a novel approach to studying how moral information about people is transmitted through naturalistic narratives.

Keywords: narrative, testimony, reputational goals, morality, pragmatics

Introduction

Much of our everyday communication revolves around people and their moral characters (Yerkovich, 1977; Wiessner, 2014; Baumeister, Zhang, & Vohs, 2004; Dunbar, 2004; Pasupathi & Wainryb, 2010; Bietti, Tilston, & Bangerton, 2018; Li & Tomasello, 2021; Schlenker, 2003; Banerjee, Heyman, & Lee, 2020). People share with others information about their past actions, decisions, and motivations, painting themselves as morally good or bad. These *moral narratives* are often unverifiable, describing facts that audiences could not have witnessed firsthand (e.g., imagine stories you hear when catching up with a friend, listening to a court hearing, or scrolling through social media). How do audiences make sense of other people's narratives about their moral characters? At times, they might interpret and trust a narrative just as intended by the narrator. After all, accurate comprehension and acceptance of other people's testimony allows valuable transmission of knowledge (Harris & Koenig, 2006; Harris et al., 2018; Mahr & Csibra, 2020; Hills, 2013; Sperber, 2001). At other times, however, audiences might unintentionally fail to take in a narrator's intended meaning or refuse to do so out of mistrust (Grice, 1979; Yus, 1999; Sperber & Wilson, 1986; Sperber et al., 2010; Mercier, 2020).

While most research on testimony has focused on audience comprehension and trust of simple utterances and propositions, we focus on testimony in the form of *narratives*. Narratives typically describe concrete events, actions, and mental states (e.g., Currie, 2010; Kim & Crockett, 2022; Lombrozo & Aronowitz, 2020; Genette, 1980). In contrast to a mere story which represents the facts of "what happened", a narrative provides a subjective version of the story, reflecting the goals of its narrator (Currie, 2010; Dahlstrom, 2019; Fraser, 2021; Genette, 1980; Kim & Crockett, 2022; Bietti, Tilston, & Bangerton, 2018). In the case of moral narratives, narrators often have *reputational goals* of appearing morally good or bad to others (e.g., Schlenker, 2003; Sedikides, Hoorens & Dufner, 2015; Banjeree, Heyman & Lee, 2020; Brown & Levinson, 1987). A narrator with a positive reputational goal might highlight the positive

outcomes of their actions while omitting negative ones (and vice versa for a negative reputational goal). In this way, narrators with different goals might select different but truthful details that they think will steer audiences toward the desired beliefs (e.g., about the narrator's own moral character).

In this paper, we examine comprehension and trust of moral narratives by focusing on the role of audiences' inferences of narrators' intentions. While an utterance might not explicitly state the intended evaluation, it might make an *intention* to transmit it more or less overt (e.g., Brown & Levinson, 1978; Horn, 2010; Davis, 2016a; Haugh, 2009; Camp, 2018). In the case of moral narratives, narrators might vary in how clearly they signal their desire to lead audiences to particular beliefs about moral character (i.e., make their reputational goals recognizable). For instance, a narrator who only presents positive information about themselves might make their positive reputational goal more obvious than another narrator who presents a more balanced image. Similarly, narrators might vary in how much they commit to the truth of their evaluation (e.g., signaling low vs. high confidence about their own belief that they are a good person), making their reputational goal more or less clear (cf., Reboul, 2017; Mazzarella et al., 2018; Vullioud et al., 2017; Meibauer, 2014). Whether it is wise for narrators to make their goals overt should depend on how the detection of those goals influences audiences' acceptance of testimony, which is what we will explore here (Sperber et al., 2010; Reboul, 2017; Sperber & Mercier, 2011; Bietti, Tilston, & Bangerton, 2018; Mazzarella et al., 2018).

On the one hand, recognition of a narrator's intentions can help align audience beliefs to what the narrator desires, avoiding miscommunication (Sperber & Wilson, 1986; Sperber et al., 2010). All utterances have multiple possible meanings, making uncertainty and noise inherent to communication. This is most evident in cases of indirect speech, where narrators do not explicitly state what they mean, but count on audiences inferring their non-literal meaning (Brown & Levinson, 1987; Searl, 1975; Yus, 1999). For example, someone who tells their friend

that their poem is "not *terrible*" likely intends for the friend to recognize that they are softening a negative opinion in order to be polite, but also runs the risk that the friend will make a literal interpretation and misunderstand it as "the opposite of terrible" (Yoon et al., 2020). Moral narratives, as we have defined them, similarly allow varying interpretations – even when narrators intend for audiences to make particular assertions such as "I'm a good person", they do not merely state them, but demonstrate and provide support for them through the selective presentation of details. If a narrator can clearly signal that their intention is for the audience to come to a particular interpretation (e.g., a reputational goal of appearing not too terrible), there may be greater likelihood of audiences ending up with that interpretation (e.g., believe the narrator is not too terrible).

Note that according to this account, audiences do not need to accept the narrator's intended meaning upon recognizing it. Rather, the idea is that intended meaning recognition can provide extra signal which should be considered for interpreting the meaning in the first place. Upon recognizing the narrator's intention, audiences can disagree with it or reject it as untrustworthy (Sperber & Mercier, 2017; Sperber et al., 2010). In fact, an overt goal of influencing the audience's beliefs might itself arouse suspicion and mistrust. The possibility of deception is inherent to testimonial communication, leading audiences to guard against it with epistemic vigilance (Sperber et al., 2010; Keonig & Harris, 2004; Mercier, 2020). A crucial component of epistemic vigilance involves "mind-reading" the narrator to infer whether or not their intention is to be informative, helpful, and truthful, or whether there are ulterior motives that prevent them from being fully informative (Sperber et al., 2010; Mascaro & Sperber, 2009; Mazzarella et al., 2018; Harris et al., 2018). In other words, audience acceptance of testimony often depends on their inferences about narrators' *informational goals*, which may be inferred to be low when potentially conflicting goals are detected. For instance, audiences often perceive "humblebrags", or brags accompanied by statements downplaying the accomplishment (e.g., "I

won this award but I don't deserve it"), to be inauthentic attempts to appear humble (e.g., Steinmetz, Sezer & Sedikides, 2017). They infer that the humblebragger's goal of appearing humble conflicts with their informational goal, preventing them from honestly expressing their pride.

Reputational goals may similarly backfire if they are inferred as indicative of or potentially motivating low informational goals (e.g., Crant, 1996; Nguyen, Seers, & Hartman, 2008). Work on self-presentation has shown that audiences condemn people who are *revealed* to be deceptive for the sake of their reputational goals (e.g., Lafreniere et al., 2014; Sedikides, Hoorens & Dufner, 2015). However, less is known about whether audiences perceived reputational goals themselves as signaling potential deception. In principle, reputational goals are nondiagnostic about whether narrators are telling the truth. For example, a narrator may attempt to appear morally good (i.e., have a positive reputational goal), either because they actually believe themselves to be that good and want to tell the truth (high informational goal), or because they are distorting the truth to appear better than they actually are (low informational goal). However, a large body of work suggests that people frequently self-enhance in order to appear *better* than they believe they are (Steimetz, Sezer & Sedikides, 2017; Tversky & Marsh, 2004; Bourdage, Roulin & Tarraf, 2018). Whether audiences exercise vigilance toward moral narratives may depend on prior beliefs about how likely others are to truthfully or deceptively appear morally good or bad.

In the current study, we examine how audiences infer reputational and informational goals in moral narratives and how such inferences influence their acceptance of narrators' intended meaning. We probed audience goal inferences in narratives from an experiment where "narrators" were incentivized to appear like a morally good or bad person while writing about a moral event from their own lives (Colombatto et al., under review). We presented these moral narratives to two groups of "audiences" (Experiments 1 and 2) who were asked to report their

inferences about narrators' goals (reputational and informational goal inference) as well as their actual evaluations of the narrators based on the contents of their narratives (moral character judgment). We first examined how accurately audiences detect narrators' actual reputational goals. Given the meaning-disambiguating function of goal recognition, narrators' reputational goals should be readily recognizable by audiences, and audiences' moral character judgments should align with what they inferred to be narrators' reputational goals. For example, an audience who thinks a narrator is trying to appear morally good should be more likely to view them as morally good. However, if audiences are vigilant toward reputational goals, their goal detection may cause judgments about moral character to systematically deviate from inferred reputational goals. For example, if audiences generally assume that narrators will self-enhance, they may always judge narrators to be worse than they are trying to appear. Importantly, if differences between inferred reputational goals and moral character judgments are driven by vigilance toward the narrator's trustworthiness, they should be accompanied by inferences of low informational goals.

Methods and Materials

Overview of Experiments

To explore audience inferences about narrators' goals, we first obtained narratives varying in reputational goals from a previous study of moral narratives (Colombatto et al., under review) (**Fig. 1, left**). These "Narrators" (n=164) were first instructed to write about morally ambiguous events from their lives, with no particular instruction with respect to reputational goals ('Original' narratives). Three days later, they were invited to take part in a follow-up survey where they were incentivized to write about the same event again, two times – once trying to appear like a morally bad person ('Worst Person' narratives), and once trying to appear like a morally good person ('Best Person' narratives). This design resulted in sets of narratives

describing the same events but with different reputational goals, allowing us to test a direct link

between narrators' reputational goals and audience inferences of those goals.

Next, we ran two experiments testing perceptions of these narratives in naive

"Audiences" (Experiment 1: n=296 and Experiment 2: n=296). In Experiment 1, we conducted

an initial test of goal detection in moral narratives: after reading about the narrative elicitation

procedure, audience participants were shown all three narratives from each of several narrators

and were asked to match each of the three narrative with its reputational goal ('Original', 'Best',
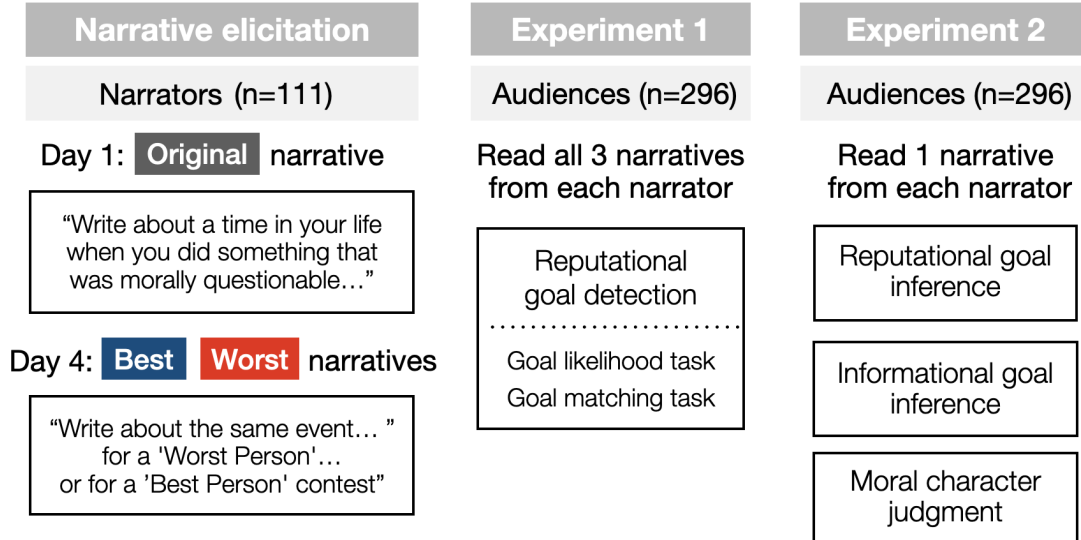
'Worst') (**Fig. 1, middle**).

Figure 1. Overview of experimental procedures. <u>Left</u>: Narrative elicitation in narrators from a

previous study (n=111). Narrators were asked to write about the same event three times, first

with no particular instruction with respect to reputational goals ('Original'), and then with

positive and negative reputational goals ('Best' and 'Worst' person). <u>Middle</u>: In Experiment 1,

audiences (n=296) read all three narratives from a given narrator and matched each with a

reputational goal. <u>Right</u>: In Experiment 2, audiences (n=296) read one narrative from a given

narrator and made inferences about the narrators' reputational goals, informational goals, and

moral characters.

In Experiment 2, we asked whether reputational goals could be detected spontaneously: audience participants read only one narrative from each narrator and were asked to rate the narrator's reputational goals, informational goals, and moral character (**Fig. 1, right**). These measures allowed us to test how audiences detect goals in moral narratives, and how these impressions may impact more global judgments of the authors' characters.

Narrative Elicitation

To investigate audience inferences about moral narratives written with different reputational goals, we used narrative collected as part of a previous study (Colombatto et al., under review). These narratives were collected in a two-part study conducted over two days, four days apart. For this and all subsequent studies, all procedures were approved by the Yale University Institutional Review Board, and all participants provided informed consent and were compensated for their time. Note that all examples have been paraphrased to protect participant anonymity.

On Day 1, 340 participants ('Narrators'; 153 female, 182 male; four nonbinary; one unspecified; mean age=31.77) were recruited via Prolific (www.prolific.co). Narrators were instructed to *"write about a time in [their]life when [they] did something that was morally questionable: some people might think what [they]did was morally wrong, but others might think it wasn't so bad."* They were specifically asked to write about *"something that truly happened.",* avoiding illegal activities and omitting identifiable information. Next, narrators were asked to provide a short title for their narrative, and were asked for permission for us to share their de-identified stories. Of the 340 participants, 53 were excluded for failure to follow the prompt, and the remaining 287 were invited to participate in the second part.

On Day 4, narrators were invited to take part in a follow-up experiment; out of 287 participants who were invited, 225 (101 female, 120 male; three nonbinary; one unspecified; mean age=32.05) completed this second survey. They were asked to write about the same event

again, and that these new stories would be entered into two contests – a "Worst Person" and a

"Best Person" contest. Each participant wrote entries for both contests, one at a time, in a

randomized order. Participants were told that judges would then later read their entries along

with stories written by other participants, and select the story that made "the author seem like the

[worst/best] person". The winner of each contest would be awarded a $50 bonus. Narrators were

thus asked to write the story from the beginning (as judges would not have access to their

original stories), but now "trying to convince the judges that [they] were a very [immoral/moral]

person." At the end of this two-part experiment, we excluded those who did not follow the

prompt (e.g., recounting a different event on the second day), and ended up with sets of three

valid narratives from 164 narrators.

Next, we prepared these narratives to be shown to audiences. First, we excluded those

from narrators who did not follow the prompt (e.g. recounting different events in their three

stories; N=61), or who did not grant us permission to share their three stories (N=32). From the

remaining 132 narratives, we next removed those that were unfit to present to audiences due to

excessive length, grammatical errors, de-anonymized materials, or leading statements – leaving a

total of N=76 narratives for Experiment 1. Because in Experiment 2 audiences read only one

narrative from each narrator rather than all three, we had more flexibility in narrative selection,

and we only excluded those that explicitly mentioned the prompt – leaving a total of N=111

narratives for Experiment 2.

Experiment 1: Detection of reputational goals

Participants

A nationally representative sample of US participants was recruited from Prolific. We

recruited 300 participants, with this pre-registered sample size chosen as the minimum to achieve

a representative sample on Prolific. Of 325 participants who completed the study, we excluded

according to our pre-registered criteria those who reported having encountered problems (N=5), failed to answer our debriefing questions sensibly (N=6), or selected the wrong option in a comprehension check (N=18; see Experimental procedures for full text). This left a valid sample of N=296 participants (138 female, 140 male; two nonbinary; one gender neutral; 15 unspecified; mean age=44.21).

Experimental procedures

Audiences were first given all details about how narrators were prompted. They were informed that narrators were other participants on Prolific who were asked to write about a true event in from their own lives where they did something morally questionable, and that they wrote the same story three times (first, as they remembered the events; second, while trying to appear like a very morally good person to win a 'Best Person' contest for a bonus, and third, while trying to appear like a very morally bad person to win a 'Worst Person' contest) (**Fig. 2**).

To ensure that participants understood the structure of narrators' task, they completed a comprehension check where they had to select what the narrators were asked to do ("Write about a time in their lives where they did something nice for others."; "Write about the same story three times: one where they appear to be good, one where they appear to be bad, one where they just tell us what happened."; or "Write about three different stories: one where they did something good, one where they did something bad, and one where they did something neither good nor bad."). Audiences were then shown the three narratives from the same narrator (labeled "Story 1", "Story 2", "Story 3") in randomized order and on separate screens (**Fig. 2, left**). Participants were allowed to move on to the next narrative only after 10 seconds to ensure they fully read each story.

After reading all three narratives, participants were asked to indicate how likely that story was to be the 'Original', 'Best Person', and 'Worst Person' story, on a scale of 'Not at all likely' to 'Very likely' (likelihood ratings) (**Fig. 2, right**). Audience participants made these ratings for

each narrative on separate screens, with all three versions shown at the top of the page as a reminder. Audiences were also asked to make a forced choice by dragging the story labels into 'Original', 'Best Person', and 'Worst Person' boxes (matching task). Again, all three stories were displayed on the page, allowing participants to read them again for comparison. Finally, audiences made judgments about the narrators' likability, trustworthiness, and moral character (not analyzed in this paper; see Colombatto et al., under review). Each audience participant completed this task for sets of narratives from 5 randomly selected narrators, followed by some demographic questions.
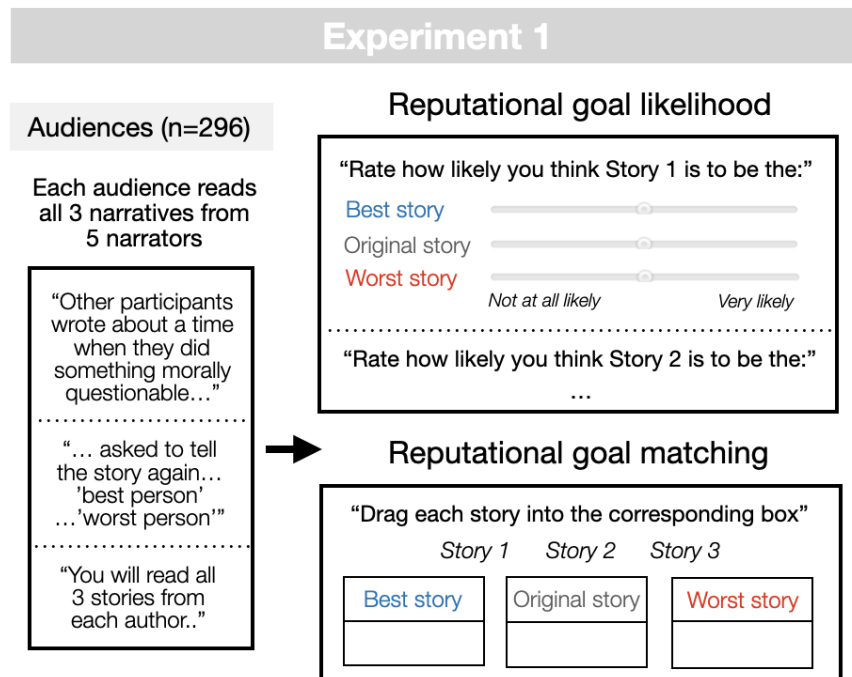

Figure 2. Procedures for Experiment 1. <u>Left</u>: Audiences were informed about the instructions narrators received. <u>Right</u>: Likelihood ratings and goal matching tasks used to probe reputational goal detection.

Experiment 2: Goal detection and reputational judgments

Participants

A nationally representative sample of US participants was recruited from Prolific. Of 300 participants who completed the study, we excluded according to our pre-registered criteria failed to answer our debriefing questions sensibly (N=4), and no participants reported having encountered problems. This left a valid sample of N=296 participants (136 female, 139 male; three nonbinary; 18 unspecified; mean age=45.20)

Experimental procedures

In contrast to Experiment 1, audience members in Experiment 2 read only one narrative from each narrator (**Fig. 3**). Each audience member read 12 narratives (4 'Original', 4 'Best', 4 'Worst', each written by a different narrator and shown in a randomized order). Because in this experiment we aimed at probing spontaneous goal inference, we no longer informed audience participants about the reputational goal manipulation, but rather simply informed them that they would read narratives from other participants: *"In this study, you are going to read some stories about moral situations written by other participants on Prolific. The participants [...] could write about any event they wanted to as long as it was something that truly happened to them."*

After reading each story for at least 10 seconds, audience participants answered three questions probing reputational goal inference ("How much was the author trying to appear like a morally bad vs. good person?", with slider labels "Want to appear very morally bad" , "Neutral/neither" , and "Want to appear very morally good"); informational goal inference ("How much did the author care about communicating what happened accurately?", with slider labels "Not at all" and "Very much"), and moral character judgment ("Based on what they did, how morally bad vs. good do you think the author is?", with slider labels "Very morally bad", "Neutral/neither", and "Very morally good") (**Fig. 3**).
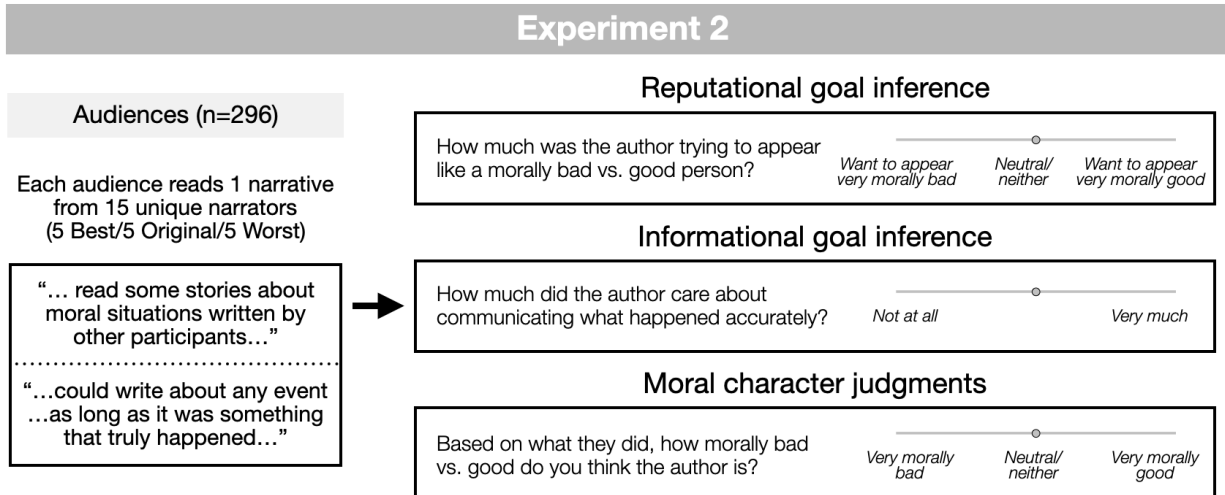
Figure 3. Procedures for Experiment 2. Audiences only read one narrative from each narrator (N=111). <u>Left</u>: Audiences were only told that narrators wrote about a true event involving some moral situation. <u>Right</u>: Questions probing audience inferences about the narrator's reputational goal, informational goal, and moral character.

All data and analysis scripts are available at

https://osf.io/x3t5e/?view_only=90cfcc9768924193944df28a6c90015d. Note that per our IRB

protocol, the original narratives cannot be shared publicly.

Results

Audiences accurately detect narrators' reputational goals

Using three tasks (likelihood, matching, inference) across two experiments, we examined how accurately audiences detect narrators' reputational goals. In Experiment 1, audiences were told that each narrator wrote three narratives with different reputational goals. After reading all three narratives from the same narrator, audiences rated the <u>likelihood</u> of each narrative corresponding to the three reputational goals (Best/Original/Worst) (**Fig. 4, left**). They additionally <u>matched</u> each of the narratives to one of the three goals (**Fig. 4, middle**). In

Experiment 2, audiences were not told about the reputational goal manipulation and only saw one narrative from each narrator. Audiences then indicated their spontaneous <u>inferences</u> about whether the narrator wanted to appear very morally bad or good (**Fig. 4, right**). Goal detection accuracy was high across all tasks.
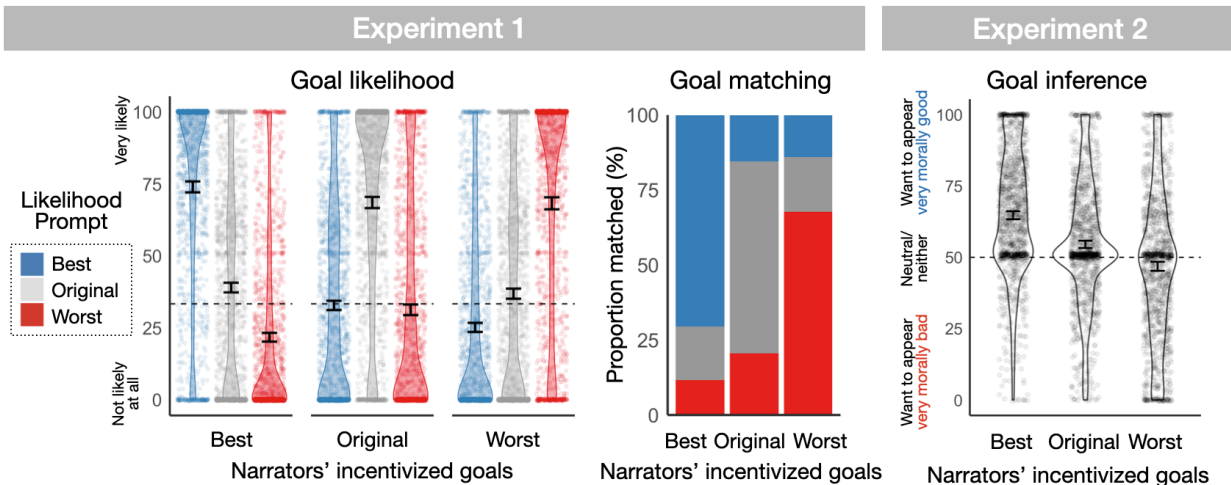

Figure 4. Audience inferences about narrators' reputational goals across Experiments 1 and 2. Results are presented according to narrators' actual goals on the x-axis (i.e., whether their narrative was generated in response to Best/Worst Person contest, or was the Original narrative). <u>Left</u>: likelihood ratings ("Rate how likely you think the Story is the [Best/Original/Worst]"). <u>Middle</u>: proportion of each narrative type matched as Best/Original/Worst. <u>Right</u>: goal inference ratings from Experiment 2. Error bars are mean +/- 95% confidence intervals subtracting out the shared variance.

Following our pre-registered plan, likelihood ratings (**Fig. 4, left**) were analyzed using a linear mixed effect model with the narrator's incentivized reputational goal ('Original', 'Best', 'Worst') and audience rating prompts ('Original', 'Best', 'Worst') as interacting fixed effects (dummy coded with Original as reference), as well as audience participants as random effects. (Our pre-registered analysis included narrative nested within narrators as random effects as well,

but this resulted in a singular fit, and so we report a model with simplified random effects

structure.) As would be expected if audiences consistently discriminated between narratives

written with different narrator goals, there was an interaction between audience rating prompts

and narrator's incentivized goals (F(4, 13004)=1076.49, p<.001). There was a significant main

effect of audience prompt (F(2, 13004)=60.03, p<.001), but not of narrators' goals

(F(2,13004)=2.18, p=.113).

Post-hoc tests (Bonferroni corrected) confirmed that audiences were highly accurate at

reputational goal detection: for each rating prompt, likelihood ratings were highest for the

prompts that matched the narrator's incentivized goal. Best Person narratives received higher

Best likelihood ratings compared to Original (B=34.95, SE=1.22, t(13004)=28.63, p<.001) and

Worst ratings (B=52.21, SE=1.22, t(13004)=42.76, p<.001). Similarly, Original narratives

received higher Original likelihood ratings compared to Best (B=35.75, SE=1.22,

t(13004)=29.28, p<.001) and Worst ratings (B=37.28, SE=1.22, t(13004)=30.54, p<.001).

Likewise, Worst Person narratives received higher Worst likelihood ratings compared to Original

(B=31.40, SE=1.22, t(13004)=25.72, p<.001) and Best ratings (B=43.11, SE=1.22,

t(13004)=35.30, p<.001).

Best and Worst Person narratives were unlikely to be confused as the other – Best

narratives received higher Original compared to Worst likelihood ratings (B=17.26, SE=1.22,

t(13004)=14.14, p<.001), and Worst narratives received higher Original compared to Best

likelihood ratings (B=11.71 SE=1.22, t(13004)=9.59, p<.001). Further, Original narratives

received similar Best and Worst likelihood ratings (B=1.53, SE=1.22, t(13004)=1.25, p=0.211).

Matched goals (Experiment 1) were analyzed using a multinomial logistic regression,

with the narrator's incentivized reputational goal as a fixed effect, and narrators and audiences as

random effects (**Fig. 4, middle**). As with the likelihood ratings, audiences were very accurate at

detecting narrators' reputational goals. There was a significant effect of narrators' goals on the

audience's matching choices (chi^2(4) = 2185.2, p<.001). Narratives written with Best Person goals were much more likely to be matched by audiences as Best than as Original (log odds, B=2.40, SE=0.12, t(6)=20.66, p<.001) or Worst (B=2.89, SE=0.12, t(6)=23.58, p<.001), Original narratives were more likely to be matched as Original than Best (B=2.28, SE=0.11, t(6)=20.27, p<.001) or Worst (B=1.93, SE=0.11, t(6)=17.68, p<.001), and Worst Person narratives were more likely to be matched as Worst than Best (B=2.55, SE=0.12, t(6)=21.86, p<.001) or Original (B=2.24, SE=0.11, t(6)=19.82, p<.001).

Similar to the likelihood ratings, Best Person narratives were more likely to be matched as Original than Worst (log odds, B=0.49, SE=0.11, t(6)=4.30, p=0.005), and Worst Person narratives were more likely to be matched as Original than Best (B=0.31, SE=0.11, t(6)=2.80, p=.031). However, Original narratives were more likely to be Worst than Best (B=0.35, SE=0.11, t(6)=3.30, p=0.016).

Audience inferences about reputational goals in Experiment 2 were accurate as well, even though by design audiences could not compare the three narratives from each narrator (**Fig. 4, right**). Goal inference ratings were analyzed using a linear mixed effects model with narrative type as a fixed effect and audiences and narratives nested within narrators as random effects. There was a main effect of narrator goal (F(2, 206)=81.73, p<.001) where narrators of Best Person narratives were inferred as wanting to appear morally good more than narrators of Original (B=10.08, SE=1.39, t(210)=7.26, p<.001) and Worst Person narratives (B=17.73, SE=1.39, t(210)=12.74, p<.001), and narrators of Original Person narratives were inferred as wanting to appear morally good more than narrators of Worst Person narratives (B=7.65, SE=1.39, t(210)=5.51, p<.001).

Note that unlike in Experiment 1, audiences in Experiment 2 did not know that narrators were given any goals at all. This allowed us to explore audiences' spontaneous expectations about narrators' goals. Goal inference ratings showed a positive bias (all reputational goal ratings

> 50: t(3549)=13.44, p<.001; M=55.37, CI[54.59 56.16]) (**Fig. 4, right**). Reputational goal

ratings for both Best and Original narratives tended toward "want to appear morally good" (for

Best > 50: t(1183)=23.30, p<.001; M=64.73, CI[63.49 65.97]; for Original > 50: t(1182)=7.51,

p<.001; M=54.53, CI[53.35 55.71]). Since audiences in Experiment 2 did not know that narrators

had been assigned specific reputational goals, this bias towards positive reputational goals

suggests that audiences may assume narrators are more likely to be driven by positive

reputational goals.

Audiences discount more when positive reputational goals are inferred

To summarize so far, we found across two experiments that audiences accurately detect

narrators' reputational goals, and audiences expect narrators to have positive reputational goals.

In the following analyses, we investigate each of these findings further, asking whether the

effects reflect strategies taken by narrators to gain audience trust, and whether such strategies are

successful.

In Experiment 2, audiences were asked to judge the narrator's moral character.

Comparing audiences' inferred reputational goals and actual character judgments allows us to

examine whether audiences' impressions of narrators matched the impressions audiences thought

narrators desired. Indeed, inferred reputational goal ratings and character judgments were highly

correlated (**Fig. 5, left**, r=0.51, t(3548)=35.61, p<.001). These results are consistent with the

explanation that audiences' judgments are likely to align with a narrator's reputational goal when
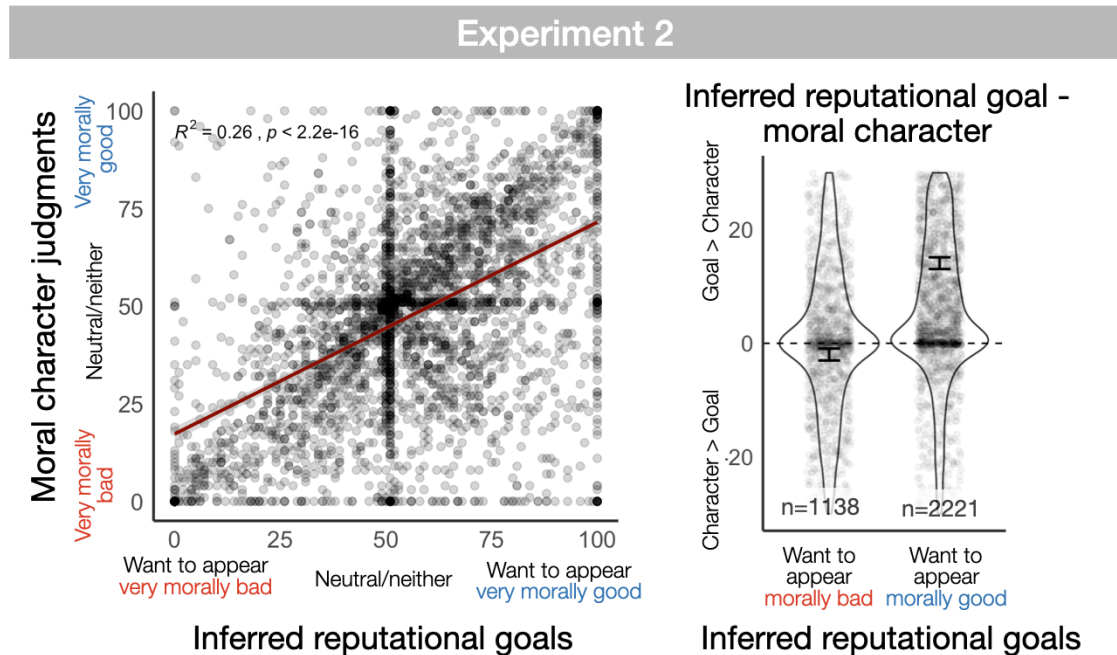
they accurately infer it.

Figure 5. Comparison between audience inferences of narrators' reputational goals vs. their judgments of narrators' character. <u>Left</u>: Correlation between inferred reputational goal and character judgments. <u>Right</u>: the amount of adjustment (deviation of character judgments from inferred goals) split by whether the inferred goal was positive (want to appear morally good) or negative (want to appear morally bad). Error bars are mean +/- 95% confidence intervals confidence intervals subtracting out the shared variance.

Character judgments, however, were not perfectly aligned with inferred reputational goals. Given our previous result that audiences expect narrators' reputational goals to be more positive than negative, we tested whether they also expect narrators to always self-*enhance* (i.e., try to appear better than they actually are). Note that inferred positive and negative reputational goals refer to inferences about whether narrators are trying to appear morally good or bad, not whether they are trying to appear morally *better* or *worse*. Audiences could in principle infer either a positive or negative reputational goal and infer that the narrator is truthfully representing their own moral character, trying to appear better, or trying to appear worse. If audiences

generally expect self-enhancement, they should downwardly adjust their character judgments (i.e., judge the narrator to be morally worse than they are trying to appear) regardless of the inferred reputational goal (i.e., positive or negative).

We examined the difference between audience's inferences of narrators' reputational goals and their character judgments, split by the direction of the inferred goal (positive: ratings > 50, corresponding to responses toward "want to appear very morally good" or negative: ratings < 50, corresponding to responses toward "want to appear very morally bad") (**Fig. 5, right**). For positive inferred goals, audiences judged narrators' actual moral characters to be worse than they are trying to present ((goal - actual) for positive: M=14.1, SD=0.52, CI[13.07 15.13], t(2220)=26.88, p<.001). In contrast, when the inferred goal was negative, audiences judged the narrator to be slightly *better* than they are trying to appear ((goal - actual) for negative: M=-3.69, SD=0.53, CI[-3.02 -0.92], t(1328)=-3.69, p<.001). This suggests that audiences expect narrators' goals to be *exaggerations* in either positive or negative directions (i.e., trying to appear better than they are if they have a positive goal and trying to appear worse than they are if they have a negative goal).

Alternatively, this result could reflect a floor and ceiling effect or different uses of scales for the inferred reputational goal and moral character judgment questions. However, the discrepancy between inferred reputational goals and character judgments was also much larger when positive goals were inferred (t(2740)=21.2, p<.001). This suggests that audiences disproportionately discount positive reputational goals as self-enhancements. This result is consistent with audiences being more likely to expect positive, self-*enhancing* goals than negative, self-deprecating goals. Even when negative goals are detected, audiences think most narrators are unlikely to want to appear *much worse* than they actually are.

Achieving positive reputational goals requires audience inference of high informational goals

We next tested the possibility that audiences' inferences of narrators' informational goals impacts their trust in narrators. In our main pre-registered analysis, we looked at the interactive effects of inferred informational and reputational goals on actual character judgments using another linear mixed effects model (Character ~ Reputational goal * Informational goal + same random effects as before) (**Fig. 6**). While inferred informational goals alone did not have an effect on character judgments (B=0.03, SE=0.03, CI[-0.01 0.09], t(3388)=1.18, p=0.2), there was a significant interactive effect of informational and reputational goals (B=0.003, SE=0.0049, CI[0.00 0.00], t(3359)=6.48, p<.001) and a main effect of reputational goal (B=0.16, SE=0.038, CI[0.09 0.24], t(3432)=4.29, p<.001). Comparing a model without the interaction also confirmed that the full model is better (χ^2(1)=41.75, p<.001).

This interaction effect shows that if a narrator is inferred to have a negative reputational goal, whether they are inferred as trying to be informative does not matter much for character judgment – audiences will judge them as similarly morally bad. In contrast, if a narrator is inferred to have a positive reputational goal, then appearing informative is crucial: the presence of both goals are necessary for narrators to be judged as morally good. An informative narrator who is trying to appear good is judged to be (almost) as good as they want to appear, but an uninformative narrator with a positive goal is judged to be morally bad. Importantly, the interaction between informational and reputational goals also suggests that the deviation of audiences' moral character judgments from inferred reputational goals are driven by suspicion that narrators are not being truthful.
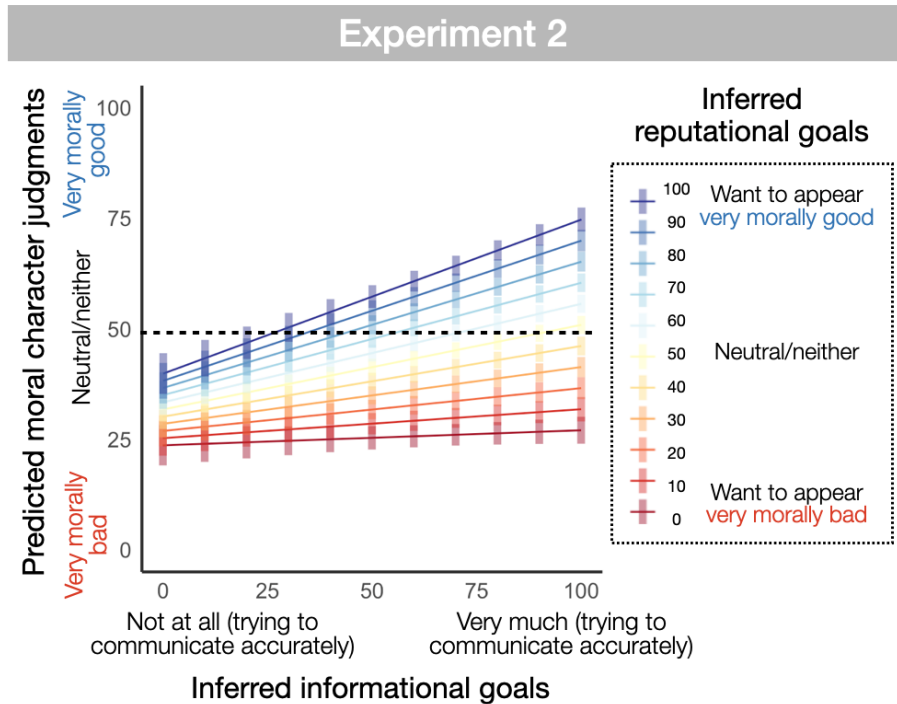
Figure 6. Interactive effect of inferred informational and reputational goals on moral impressions. Model outputs for the interactive effects of inferred informational and reputational goals on moral character judgments. Y-axis shows predicted actual reputation ratings, x-axis shows inferred informational goal ratings, and colors indicate levels of inferred reputational goal ratings.

Narrators try harder to appear informative when reputational goals are positive

To summarize the results so far: first, narrators' reputational goals (and positive ones in particular) are easily detectable. One benefit of making these goals detectable for audiences is that their beliefs about narrators' actual reputations align closely with what they think narrators want them to believe. However, we also find evidence that audiences do not completely accept narrators' reputational goals: they expect narrators to generally want to appear better than they are. Importantly, inferred reputational goals interact with inferred informational goals. Being seen as wanting to be informative is particularly important for trust when the audience perceives a positive reputational goal.

A key implication of this last result is that narrators *should* try harder to appear informative if they have a positive reputational goal. To test this hypothesis, we returned to analyzing inferred informational goals as a function of what reputational goals narrators were assigned to have (Best/Worst/Original). Indeed, when we look at inferences about informational goals ("did the narrator care about communicating what happened accurately") by narrative type, audiences rated Best narratives as having the highest informational goals (LMEM predicting goal ratings, effect of narrative type: $F(2,194.12)=11.09$, $p<.001$; contrasts of Best vs. Original: $B=4.26$, $SE=1.01$, $t(194.38)=3.95$, $p<.001$, Best vs. Worst: $B=4.52$, $SE=1.08$, $t(193.19)=3.49$, $p<.001$; **Fig. 7, left**).

As this was not predicted, we conducted further exploratory analyses to understand why Best narratives were perceived as most informative. One difference we observed across narratives is that narrators tended to write more for the Best narratives, followed by Worst, then Original (word count for Best: $M=164.43$ words, $SD=88.84$; Original: $M=96.6$, $SD=43.1$; Worst: $M=135.6$, $SD=81.68$; Wilcoxon signed-rank tests, two-tailed for Best vs. Worst: $z=4.41$, $p<.001$; Best vs. Original: $z=7.72$, $p<.001$; Original vs. Worst: $z=5.50$, $p<.001$) (**Fig. 7, middle**). Word length was in turn positively correlated with inferred informational goals ($r=0.37$, $p<.001$) (**Fig. 7, right**). These results suggest that narrators might have very accurate understandings of audiences' likely inferences, and put more effort into writing the Best narratives so that they will be viewed as having high informational goals and be accepted.
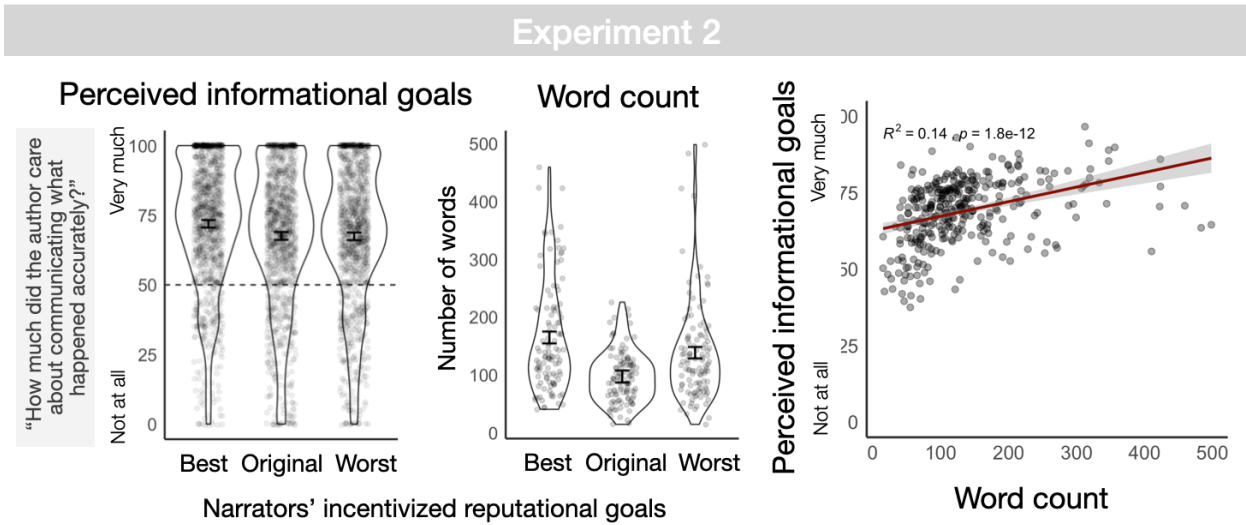
Figure 7. Inferences about narrators' informational goals. <u>Left</u>: audience's inferences of narrators' informational goals, arranged by the narrators' instructed reputational goals. <u>Middle</u>: word length for each type of instructed reputational goal. <u>Right</u>: correlation between word length and inferred informational goals. Error bars are mean +/- 95% confidence intervals subtracting out the shared variance.

Discussion

People frequently give testimony about their own moral character by telling. Because narrators often want to appear morally better than they are, audiences need to exercise epistemic vigilance (Sperber et al., 2010). Often, as audiences, we have to make inferences about narrators' goals (e.g., informational and reputational goals) based on the content of their testimony alone (e.g., Kim & Crockett, 2022). How good are audiences at detecting narrators' positive and negative reputational goals? How much do audiences accept or discount what they infer narrators as wanting to communicate? And how do inferred informational and reputational goals affect audiences' acceptance of moral narratives?

Across two experiments, we find that audiences detect narrators' incentivized reputational goals with high accuracy. When audiences did not know that the narrators were

given particular reputational goals (Experiment 2), they tended to perceive narrators as having positive reputational goals. While inferred reputational goals and moral character judgments in audiences were highly correlated, there were systematic deviations such that moral character was judged to be *worse* than the narrator wanted to appear when they were inferred as wanting to appear morally good, and morally *better* than the narrator desired when they were inferred as wanting to appear morally bad. Further, there was an asymmetry such that the deviation was greater for positive reputational goals.

Importantly, we also found an interactive effect of inferred informational and reputational goals on moral character judgments: the deviation between inferred reputational goals and moral character were greatest when low informational goals were detected, only for positive reputational goals. This result suggests that audiences adjusted from what they inferred narrators as wanting them to believe out of suspicion toward the narrator's trustworthiness. Finally, we found indirect evidence that narrators might have written their narratives with accurate predictions about audiences' likely inferences in mind: narrators wrote longer narratives when they had positive reputational goals, which resulted in audience inferences of higher informational goals. In other words, narrators may have tried harder to appear informative when their reputational goals were positive because they knew audiences would be more vigilant.

The asymmetrical discounting between positive and negative reputational goals is consistent with the idea that audiences may have prior beliefs that narrators are more likely to want to appear better than they actually are (as opposed to appearing worse), especially if they are trying to appear morally good (rather than bad). While negative self-presentation can occur due to low self-esteem or related issues like depression, people generally want to appear good to others (Koszegi et al., 2019; Vonk, 1998; Steimetz, Sezer & Sedikides, 2017; Tversky & Marsh, 2004; Bourdage, Roulin & Tarraf, 2018). Therefore, audiences may think it unlikely that a narrator with a negative reputational goal is deceptively presenting themselves as morally bad

when they believe themselves to be good. In contrast, for a narrator who only describes positive information about themselves, there is more room for the truth to be a lot worse than is being presented.

At the same time, audiences did not automatically doubt all positive reputational goals. The discounting we observed with positive reputational goals was small and hinged on additional inferences of lower informational goals. This may be explained by audiences taking a 'stance of trust', where they remain skeptical and wary without actively disbelieving a testimony (Sperber et al., 2010; Sperber & Mercier, 2013; Levine, Mikhail & Leslie, 2018). Whether audiences are justified in maintaining this amount of skepticism should be tested by examining how much narrators deceive – for example, by looking more precisely at whether audience expectations of narrators' rates of self-enhancement match how much narrators spontaneously self-enhance. In a sense, audiences were "correct" in the current study to not suspect narrators at high rates – our narrators were instructed to only write about things that actually happened. The different narratives they produced did not contain lies, but rather, reflected changes in framing, where narrators changed which details were included, emphasized, and worded.

Another possible explanation for why we did not see bigger deviations between inferred reputational goals and moral character judgments is that even when audiences distrusted a narrator, they had trouble rejecting the narrative they had entertained. Such "belief perseverance" or failure to course correct in the face of known misinformation has been frequently documented in memory (e.g., Anderson et al., 1983; Marsh et al., 2003; Green & Donahue, 2011). In Green & Donahue (2011), for example, audiences read a story and then were told the story was false. Upon finding out that the author of the story intended deception, audiences judged them to be immoral, and discounted some parts of the story, but continued to believe many of the facts revealed to be false. One explanation offered by Green & Donahue is that even when deception is revealed, it can be difficult to construct an alternative or know exactly how much to correct as

well as knowing precisely which details to discount (e.g., if you realize someone is trying to appear like a morally good person, do you decide that they are now a "bad" person, or just a slightly less good person).

Intriguingly, there are proposals that the format of narratives may make discounting particularly difficult for audiences. Narratives are thought to "transport" and immerse readers into a narrative world, making them less aware of their real surroundings (Green & Brock, 2000). Fraser (2021) goes beyond transportation to argue that the structure of narratives creates a deeper dependence than other forms of communication, shaping what audiences infer and pay attention to. Such perspectival dependence could lead to difficulty by audiences to form a belief that is different from the one being presented to them, even if they are able to recognize the narrators' intentions and exercise vigilance toward them. A related but contrary view is that vigilance and discounting increases with a narrators' ability to mask their own goals (Reboul, 2017; Bietti, Tilston & Bangerton, 2018). According to this account, the more audiences are made to feel like they arrived at a belief themselves (rather than through the narrators' machinations), the more audiences are likely to align to the narrators' intended belief. In this study, we found that audiences' beliefs about moral character aligned with narrators' inferred reputational goals in spite of accurate goal recognition. An open question is whether audiences become even less vigilant when narrators' goals are harder to detect (e.g., when narratives are particularly engaging and effective at transporting the audience).

Together, these investigations of goal inference in moral narratives demonstrate how audiences' inferences about narrators' goals can both support and hinder the acceptance of the narrators' intended meanings. Accurate goal recognition increases the chance that audiences will make judgments intended by narrators, but it can also lead to vigilance about accuracy, depending on their expectations about particular goals. More broadly, by combining insights from narrative theory, moral psychology, impression management, and linguistics, we show how

comprehension and trust can be studied in complex, naturalistic forms of communication, such as narratives.

Author contributions

J.K.: Conceptualization, Methodology, Formal analysis, Visualization, Writing - Original draft; C.C: Conceptualization, Methodology, Formal analysis, Visualization, Writing - Review & Editing; M.J.C: Conceptualization, Supervision, Writing - Review & Editing, Funding acquisition.

References

Anderson, C. A. (1983). Abstract and concrete data in the perseverance of social theories: When weak data lead to unshakeable beliefs. *Journal of experimental social psychology*, 19(2), 93-108.

Aronowitz, S., & Lombrozo, T. (2020). Experiential explanation. *Topics in Cognitive Science*, 12(4), 1321-1336.

Banerjee, R., Heyman, G. D., & Lee, K. (2020). The development of moral self-presentation. *The Oxford handbook of moral development: An interdisciplinary perspective*, 92-109.

Bénabou, R., Falk, A., & Tirole, J. (2018). Narratives, imperatives, and moral reasoning (No. w24798). *National Bureau of Economic Research.*

Baumeister, R. F., Zhang, L., & Vohs, K. D. (2004). Gossip as cultural learning. Review of general psychology, 8(2), 111-121.

Bietti, L. M., Tilston, O., & Bangerter, A. (2019). Storytelling as adaptive collective sensemaking. *Topics in cognitive science,* 11(4), 710-732.

Bourdage, J. S., Roulin, N., & Tarraf, R. (2018). "I (might be) just that good": Honest and deceptive impression management in employment interviews. *Personnel Psychology*, 71(4), 597-632.

Boyd, R., & Richerson, P. J. (2005). The origin and evolution of cultures. *Oxford University Press.*

Brown, P., & Levinson, S. C. (1978). Universals in language usage: Politeness phenomena. In *Questions and politeness: Strategies in social interaction* (pp. 56-311). Cambridge University Press.

Camp, E. (2018). Insinuation, common ground. *New work on speech acts,* 40.

Colombatto et al. (under revision). Motivated abstraction of personal transgressions in moral narratives.

Crant, J. M. (1996). Doing more harm than good: When is impression management likely to evoke a negative response?. *Journal of Applied Social Psychology*, *26*(16), 1454-1471.

Currie, G. (2010). Narratives and narrators: A philosophy of stories. *Oxford University Press*, USA.

Dahlstrom, M. F. (2021). The narrative truth about scientific misinformation. *Proceedings of the National Academy of Sciences*, *118*(15), e1914085117.

Davis, W. A.. (2016). Implicature. *Irregular Negatives, Implicatures, and Idioms*, 51-84.

Dunbar, R. I. (2004). Gossip in evolutionary perspective. *Review of general psychology*, 8(2), 100-110.

Fraser, R. (2021). Narrative testimony. *Philosophical Studies*, 178, 4025-4052.

Genette, G. (1980). Narrative Discourse: An Essay in Method. 1972. Trans. Jane E. Lewin. Ithaca: Cornell UP.

Green, M. C., & Brock, T. C. (2000). The role of transportation in the persuasiveness of public narratives. *Journal of personality and social psychology*, 79(5), 701.

Green, M. C., & Donahue, J. K. (2011). Persistence of belief change in the face of deception: The effect of factual stories revealed to be false. *Media Psychology*, 14(3), 312-331.

Grice, H. P. (1975). Logic and conversation. *In Speech acts* (pp. 41-58). Brill.

Harris, P. L., & Koenig, M. A. (2006). Trust in testimony: How children learn about science and religion. *Child development*, 77(3), 505-524.

Harris, P. L., Koenig, M. A., Corriveau, K. H., & Jaswal, V. K. (2018). Cognitive foundations of learning from testimony. *Annual Review of Psychology*, 69, 251-273.

Haugh, M. (2009). Intention (ality) and the conceptualization of communication in pragmatics. *Australian Journal of Linguistics*, *29*(1), 91-113.

Hills, A. (2013). Moral testimony. *Philosophy Compass*, 8(6), 552-559.

Horn, L. R. (2010). Multiple negation in English and other languages. *The expression of negation*, 111-148.

Kim, J., & Crockett, M.J (2022). "Narrating the" what" and" why" of our moral actions." In *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 44, no. 44.

Lafrenière, M. A. K., Sedikides, C., Van Tongeren, D. R., & Davis, J. (2016). On the perceived intentionality of self-enhancement. *The Journal of Social Psychology,* 156(1), 28-42.

Levine, S., Mikhail, J., & Leslie, A. M. (2018). Presumed innocent? How tacit assumptions of intentional structure shape moral judgment. *Journal of Experimental Psychology: General*, 147(11), 1728.

Li, L., & Tomasello, M. (2021). On the moral functions of language. *Social Cognition*, 39(1), 99-116.

Mahr, J. B., & Csibra, G. (2020). Witnessing, Remembering, and Testifying: Why the Past Is Special for Human Beings. *Perspectives on Psychological Science*, 15(2), 428.

Marsh, E. J., & Tversky, B. (2004). Spinning the stories of our lives. *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition*, 18(5), 491-503.

Marsh, E. J., Meade, M. L., & Roediger III, H. L. (2003). Learning facts from fiction. *Journal of Memory and Language*, 49(4), 519-536.

Mascaro, O., & Sperber, D. (2009). The moral, epistemic, and mindreading components of
    children's vigilance towards deception. *Cognition*, 112(3), 367-380.

Mazzarella, D., Reinecke, R., Noveck, I., & Mercier, H. (2018). Saying, presupposing and
    implicating: How pragmatics modulates commitment. *Journal of Pragmatics*, 133, 15-27.

Meibauer, J. (2014). Lying at the semantics-pragmatics interface (Vol. 14). Walter de Gruyter
    GmbH & Co KG.

Mercier, H. (2020). Not born yesterday: The science of who we trust and what we believe.
    Princeton University Press.

Mercier, H., & Sperber, D. (2011). Why do humans reason? Arguments for an argumentative
    theory. *Behavioral and brain sciences*, 34(2), 57-74.

Pasupathi, M., & Wainryb, C. (2010). Developing moral agency through narrative. *Human
    Development*, *53*(2), 55-80.

Reboul, A. (2017). Is implicit communication a way to escape epistemic vigilance. *Pragmatics
    at its Interfaces*, 17, 91.

Schlenker, B. R. (2003). Self-presentation. *Handbook of self and identity,* 2, 542-570.

Sedikides, C., Hoorens, V., & Dufner, M. (2015). Self-enhancing self-presentation: Interpersonal,
    relational, and organizational implications.

Sperber, D. (2001). An evolutionary perspective on testimony and argumentation. *Philosophical
    topics*, *29*(1/2), 401-413.

Sperber, D., & Wilson, D. (1986). Relevance: Communication and cognition (Vol. 142).
    Cambridge, MA: Harvard University Press.

Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., & Wilson, D. (2010).
    Epistemic vigilance. *Mind & language*, 25(4), 359-393.

Steinmetz, J., Sezer, O., & Sedikides, C. (2017). Impression mismanagement: People as inept
    self‑presenters. *Social and Personality Psychology Compass*, 11(6), e12321.

Vullioud, C., Clément, F., Scott-Phillips, T., & Mercier, H. (2017). Confidence as an expression

    of commitment: Why misplaced expressions of confidence backfire. E*volution and*

    Human Behavior, 38(1), 9-17.

Wiessner, P. W. (2014). Embers of society: Firelight talk among the Ju/'hoansi Bushmen.

    Proceedings of the National Academy of Sciences, *111*(39), 14027-14035.

Yerkovich, S. (1977). Gossiping as a way of speaking. *Journal of Communication*.

Yoon, E. J., Tessler, M. H., Goodman, N. D., & Frank, M. C. (2020). Polite speech emerges from

    competing social goals. *Open Mind*, 4, 71-87.

Yus, F. (1999). Misunderstandings and explicit/implicit communication. *Pragmatics. Quarterly*

    Publication of the International Pragmatics Association (IPrA), *9*(4), 487-517.