Journal of Experimental Social Psychology 116 (2025) 104682

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

Ironic effects of prosocial gossip in driving inaccurate social perceptions☆

Samantha Grayson [a,*], Matthew Feinberg [b], Robb Willer [c], Jamil Zaki [d]

[a] Graduate School of Business, Columbia University, New York, NY, United States of America
[b] Rotman School of Management, University of Toronto, Toronto, Ontario, Canada
[c] Department of Sociology, Stanford University, Stanford, CA, United States of America
[d] Department of Psychology, Stanford University, Stanford, CA, United States of America

ARTICLE INFO

Keywords:
Cynicism
Gossip
Prosocial behavior

ABSTRACT

Gossip is often stereotyped as a frivolous social activity, but in fact can be a powerful tool for discouraging selfishness and cheating. In economic games, gossip induces people to act more cooperatively, presumably to avoid the cost of accruing a negative reputation. Might even this prosocial sort of gossip carry negative side effects? We propose that gossip might protect communities while simultaneously giving people the wrong idea about who's in them. Specifically, gossipers might disproportionately share information about cheaters in their midst, driving *cynical* perceptions among receivers of that gossip. To test these predictions, we first reanalyzed data from a prior study in which people played a public goods game and could gossip about their fellow players. These participants indeed produced negatively skewed gossip: writing much more frequently about cheaters than cooperators, even when most people in their public goods game groups acted generously. To examine the effect of this gossip on cynicism, we ran a new experiment in which a second generation of participants read these gossip notes, and then prepared to play their own public goods game. Gossip recipients inferred that the groups that produced these notes acted significantly more selfishly than they truly had–becoming both cynical and inaccurate based on gossip. However, this gossip did not affect second generation participants' forecasts of how their own group would behave, nor their own cooperative choices. Together, these findings suggest that gossip skews negative, and, therefore, encourages outside observers to draw more cynical conclusions about groups from which it comes.

"Great minds discuss ideas; average minds discuss events; small minds discuss people."

The quote above, typically attributed to Eleanor Roosevelt,[1] captures a stereotype that *gossip*—evaluative social information spread about absent individuals—is a vapid way to spend one's time. In some cases, the truth can be worse. Teens can use gossip to spread rumors and ostracize one another (Archer & Coyne, 2005); employees can use it to pit colleagues against each other in the workplace (De Gouveia et al., 2005). And yet, analysis of conversations in natural settings demonstrates that well over half concern social topics (Dunbar et al., 1997), and people spend almost an hour of their day on average engaging in conversations about absent third parties (Robbins & Karan, 2020). Perhaps most of us are cursed with small, mean minds.

Or perhaps this pessimism is misplaced. An alternative model holds that gossip evolved to serve prosocial functions. Dunbar (2004) describes it as akin to grooming among non-human primates, as a strategy for building social coalitions. And indeed, more recent evidence suggests that gossip during an economic game promotes a sense of connection between participants (Jolly & Chang, 2021).

In behavioral economics, gossip is often cast as a mechanism for protecting communities from cheaters. Consider the public goods game, in which individuals, typically anonymous to one another, choose how much money to contribute to a common fund. That fund is then doubled and distributed evenly between them, irrespective of each person's contribution. The optimal outcome for the group is for everyone to contribute as much as possible, doubling each person's money. But any one individual can gain more by "free-riding," contributing little or nothing but still receiving payouts from the common fund, taking

☆ This paper has been recommended for acceptance by Kim Peters.

* Corresponding author.
E-mail addresses: SGrayson28@gsb.columbia.edu (S. Grayson), matthew.feinberg@rotman.utoronto.ca (M. Feinberg), willer@stanford.edu (R. Willer), jzaki@stanford.edu (J. Zaki).

[1] This attribution is likely apocryphal; for more on the quote's lineage, see https://quoteinvestigator.com/2014/11/18/great-minds/

https://doi.org/10.1016/j.jesp.2024.104682
Received 9 March 2024; Received in revised form 16 August 2024; Accepted 23 September 2024
Available online 1 October 2024

advantage of others' contributions.

Over time, free-riding disincentivizes any participants from contributing, and in repeated versions of the public goods game, the common fund dwindles to near zero in later rounds of the game (Fehr & Gächter, 2000). However, this "tragedy of the commons" does not occur in all cases. When a group can punish free-riders, individuals contribute more and for longer in the game (Boyd & Richerson, 1992). Likewise, when individuals' actions in public goods games are known by others, they are more likely to contribute (Milinski et al., 2002), indicating that the cost of a bad reputation outweighs the benefits of free-riding.

Gossip is a major tool through which people create and manage each other's reputations, and this points to at least one way it can be used as a prosocial act. When someone cheats, steals, or lies, their victim may not be able to recoup what they've lost, but by spreading the word about harmful acts, can limit the damage cheaters can do in the future. Consistent with this view, Feinberg and colleagues found that people who witness antisocial behavior were most likely to gossip when they were high in prosocial values (Feinberg et al., 2012). Further, when participants in a public goods game could gossip about people's past behavior and select who they wanted to play with as the game continued, contributions rose instead of falling across rounds, eventually reaching over 400 % of the level made in a standard version of the public goods game (Feinberg et al., 2014).

These findings suggest that in at least some cases, people gossip to protect their community, and succeed in that goal (Pan et al., 2024). But might this prosocial process carry some antisocial side-effects? Here, we explore the idea that through social policing, gossipers might encourage cooperation within a group, but ironically encourage outside observers to draw more cynical conclusions about that group.

These predictions rest on two key ideas. The first is that gossipers might disproportionately share negative, as opposed to positive or neutral, social information. This trend is consistent across various domains. For instance, people share negative political information more readily than positive information on social media (Bellovary et al., 2021). Consumers are more likely to leave negative anonymous reviews of products than positive reviews (Deng et al., 2021). And, people are more likely to share information about an immoral act than a moral one (Hofmann et al., 2014). To the extent that people use gossip to police and protect communities, they might feel that the most effective way of accomplishing these goals is to "call out" negative actors, as opposed to praising cooperative ones.

Second, recipients of gossip might be disproportionately attuned to negative information. This is one example of a more general "negativity bias," through which human beings tend to focus on, remember, and weight negative over positive stimuli (e.g., Anderson, 1965; Baumeister et al., 2001; Peeters & Czapinski, 1990; Rozin & Royzman, 2001). Negativity bias applies to our social evaluations as well. People pay outsized attention to potentially harmful social actors, a phenomenon known as cheater detection (Cosmides et al., 2005; Van Lier et al., 2013). They also remember untrustworthy faces more easily than trustworthy ones (Rule et al., 2012), and when presented with information about a person who acts morally and immorally, tend to judge them based on their worst actions (Skowronski & Carlston, 1989).

Like gossip, cheater detection and negativity bias might protect individuals and communities from bad actors. But they might also skew observers' perception of how prevalent those bad actors are. People tend to underestimate others' prosociality and trustworthiness (Fetchenhauer & Dunning, 2009, 2010; Miller, 1999). Gossip might play a role in this process—spreading not merely information, but also cynicism.

Here, we explore this idea in two ways. First, through a re-analysis of data in Feinberg et al. (2014), we find that gossipers disproportionately share negative information about free-riders they play a public goods game with, as opposed to cooperative partners.

Second, we recruited a novel group of participants to receive gossip produced by the players in Feinberg et al.'s study. These new participants systematically underestimate how cooperative previous players had been, demonstrating that gossipers, while protecting their own community, give outsiders an unreasonably bleak view of it. However, we find that gossip-driven cynicism has its limits. Participants in our new study also played their own public goods game games. Negative gossip from previous players did not affect new players' forecasts for how their own partners would behave, nor their own contributions to the common fund.

1. Study 1: Gossip skews negative

In this study, we sought to understand whether gossip generated in a public goods game is negatively biased. When given the opportunity to pass on evaluative notes about other players' behaviors in the game, are people disproportionately focused on the worst contributors? We investigate this question in a public goods game setting in order to *quantify* bias. Players of the public goods game are inherently given a specific "menu" of topics to spread gossip about – selfish players who contribute a relatively low number of points, generous players who contribute a relatively high number of points, or middle ground players. This allows us to examine–given these clear examples of selfishness and generosity–which behavior people choose to gossip about. In this way, the setting used in Study 1 allowed for a more precise and quantifiable measurement of bias than naturalistic, everyday gossip.

1.1. Methods

1.1.1. Participants

Study 1 comprises a reanalysis of data collected by Feinberg et al. (2014). Two hundred sixteen participants (82 male, 134 female; mean age = 20.4 years) took part in that study.

The data and code for the analyses presented here are also available on the Open Science Framework at https://osf.io/4htpu/.

1.1.2. Procedure

Here, we focus on aspects of Feinberg et al. (2014) that are most relevant to the current reanalysis: the gossip notes participants in this study produced. More detailed explanation of the entire study can be found in the original article.

Participants in this study played a repeated measures public goods game with six rounds. In every round of the game, each person was randomly put in a group with three new people, where none of the four group members knew each other, and they played the game anonymously. They were given 10 "points" per round, and they were asked to decide how many of these points they wanted to keep for themself and how many they wanted to contribute to a group fund. For each group, every point contributed to the group fund was added together, doubled by the experimenter, and then equally split among all four of the group members.

After each group member made their contribution decision for a given round, their decisions were shared among the group members. Each member then had the opportunity to write a gossip note about one of the other players, based on the amount they contributed to the fund. If a player had a note written about them by one of their group mates, that note would travel with them to the next round of the game, where it would be shared with their new groupmates before contribution decisions were made (See Supplementary Materials for a visual illustration of the protocol).

In their original analysis of this data, Feinberg et al. (2014) observed that the threat of gossip promoted cooperation. Participants playing the public goods game with gossip contributed significantly more to the group fund than in a standard public goods game where gossip was not allowed, suggesting that the threat of social sanctions limited free-riding.

Here, we re-analyze the first round of the public goods game with gossip to understand how frequently negative evaluative information was shared. We only look at notes from the first round of the public

goods game, given that later rounds could be "contaminated" by gossip from earlier ones.

Out of the 216 participants, 183 shared a gossip note about one of their fellow group members in the first round of the game. A sensitivity analysis determined that this sample size ($N = 183$) provided 80 % power to detect an effect size of $w = 0.23$ or greater in a chi-square test with a 5 % false-positive rate. Feinberg et al. (2014) asked three independent coders who were blind to the study's hypotheses to rate the valence of each note as positive, negative, or neutral. All three coders made the same rating for 86 % of the notes. Any discrepancies observed between coders were resolved through discussion among them. These ratings were translated into a single score for each gossip note: positive (+1), negative (−1), or neutral (0).

1.2. Results

1.2.1. Gossip skews negative

Participants in Feinberg et al.'s first round shared 117 negative notes, 24 neutral notes, and 42 positive notes. A chi-square goodness of fit test revealed that this distribution would be highly unlikely to occur by chance, X^2 (2, $N = 183$) = 79.77, $p < 0.001$; Cohen's $w = 0.66$; see Fig. 1), suggesting that negative notes are indeed overrepresented in gossip from this cohort.

Simply producing negative notes does not mean that gossip is *unreasonably* skewed. If most players in a game cheat, it would be rational for gossipers to share mostly negative information. To further explore the relationship between cooperation and gossip, we began by looking at the extremes: free-riders, who contributed 0 points to the common fund, and "full contributors," who contributed all 10 of their points.

Of the 216 participants, only 21 (~10 %) free-rode in the first round of the game. Seventeen groups included 1 free-rider, and two groups included 2 each. Of the 19 groups that included at least one free-rider, 18 (95 %) produced gossip notes that *all* focused on those individuals. A *t*-test revealed that players in groups with free-riders were far more likely to gossip about the free-rider(s) than other, non-free-riding players in their group ($t(27) = -7.23$, $p < 0.001$, $d = -2.30$).

As a point of comparison, 56 players contributed all their funds to the common pool, making full contributors more plentiful than free-riders at 26 % of the sample. Thirty five of the 54 groups in the public goods game's first round included at least one full contributor. Of those 35 groups, 6 (17 %) produced notes that *all* focused on the full contributor (s). A t-test revealed that full contributors were the target of gossip *less* often than the average player in groups with full contributors ($t(53) = 2.57$, $p = 0.01$, $d = 0.64$).

This suggests that gossip is not merely negative, but negatively skewed: it is disproportionately focused on the worst players while eluding the more common cooperators. To further this exploration, we investigated whether people are more likely to gossip more about the lowest, middle, or highest contributor out of their other three groupmates, irrespective of whether their group included free-riders, full contributors, or neither. Participants shared 122 notes about the lowest contributor in their group, 4 notes about the middle contributor, and 51 notes about the highest contributor in their group. A chi-square goodness of fit test demonstrated that this distribution across categories would not be expected by chance, X^2 (2, $N = 177$) = 119.63, $p < 0.001$; Cohen's $w = 0.82$).

There was no clear relationship observed between a person's own contribution amount and the valence of the gossip they send along. A simple logistic regression revealed that for each additional point contributed to the group fund by an individual, on average, the odds of leaving a negative note increases by 3.7 % ($b = 0.037$, SE = 0.05, $z = 0.77$, $p = 0.443$, OR = 0.77). This tells us that whether they are a low, middle, or high contributor, the odds of leaving a negative note are not significantly different. People are eager to gossip negatively, despite how they acted in the game.

Together, these data suggest that in the context of a public goods game with the opportunity to gossip, participants, regardless of their behavior, are more likely to call out a free-riding or low contributing group member than praise a generous one or discuss a middle-of-the-pack contributor. This does not mean others' generosity holds no bearing on gossip. Correlation analyses revealed that groups with higher average contributions were more likely to produce positive gossip ($r(214) = 0.23$; $p < 0.001$) and less likely to produce negative gossip ($r(214) = -0.46$, $p < 0.001$; see Supplementary Materials for correlation matrix). That said, gossip also appears to be quite negatively skewed, highlighting harm while frequently ignoring positive action.

2. Study 2: Negative gossip spreads cynicism, but not mistrust or selfishness

In Study 1, we observed that gossip shared in the context of the public goods game was negatively biased. Players disproportionately shared information about low contributors, thus leaving out information about the actions of more prosocial players in the process. To assess the potential side effects of this negative skew, we pass these gossip notes generated by participants in Study 1 to a new set of participants in Study 2. This allows us to test whether exposure to this gossip informs a second generation's perceptions of the original targets of the gossip, their beliefs about future interactions, and their own actions. The experimental paradigm we employ also enables us to quantify how gossip skews people's judgments, in comparison to strictly factual information or no information about targets.

2.1. Methods

2.1.1. Participants

We recruited a sample of 772 participants (394 Female, Mean age = 40 years) for this study. Using pilot data we collected prior to running our full sample, we estimated the power to detect the effect of condition (gossip vs. numeric information) on one of our key dependent variables with an a priori power analysis, and calculated that a sample of 768 participants would give us >80 % power to detect the effect of interest.[2]

We recruited our sample through Prolific Academic Ltd. This sample was restricted to adult Prolific users in the United States with at least 50 previously completed studies, a minimum approval rate of 95 %, and English as a native language. Participants were excluded from analysis if they failed our attention check ($N = 39$). A sensitivity analysis determined that this sample size ($N = 733$) provided 80 % power to detect an effect size of $d = 0.25$ or greater in a two sample *t*-test with a 5 % false-positive rate.

All participants provided informed consent, and the study was approved by the Institutional Review Board at a large private university.

This study was preregistered on Open Science Framework at https://osf.io/gdcw5. The data and code for the analyses presented here are also available on the Open Science Framework at https://osf.io/gdcw5.

In this manuscript and supplementary materials, we report how we determined our sample size, all data exclusions, all manipulations, and all measures in the study.

2.1.2. Procedure

Participants in the study served as a "second generation" to participants from Feinberg et al. (2014): they received information from those "first generation" players, drew inferences about them, and then moved on to play their own public goods game.

Second generation participants were told about the public goods

[2] Our power analysis reported in our pre-registration was for 256 participants per condition (total $N = 768$). Prolific over-recruited because a few individuals returned their submission after completing the study. So, the raw data, prior to exclusion for failing the attention check, has 772 participants.

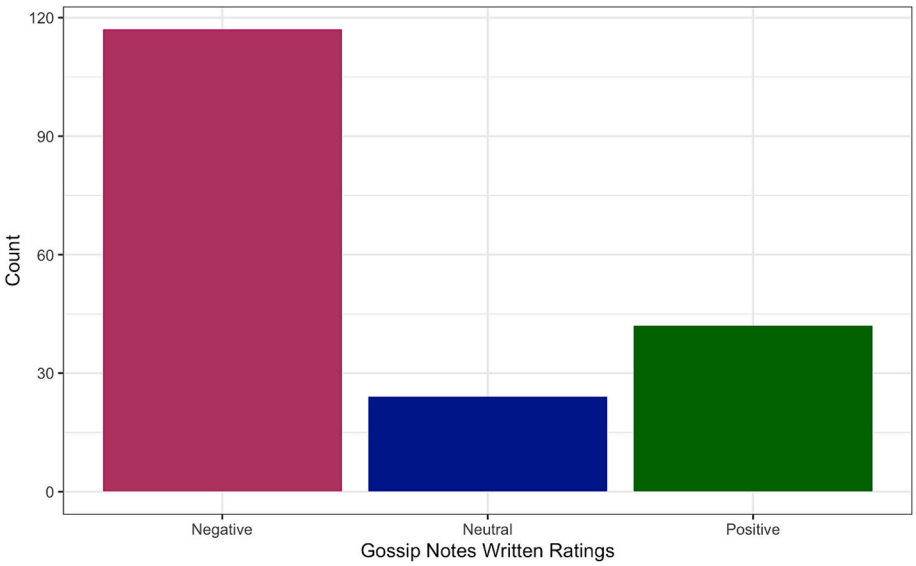

Fig. 1. ***Distribution of gossip notes by valence.*** Bars represent frequencies in each category of gossip valence.

game played by the first generation, and then completed a five-question comprehension check to ensure they understood the protocol of the original study. For example, they were asked how many points each first generation player was allotted at the beginning of the public goods game. Participants could not progress in the study without proving they understood the design and answering correctly.

Participants in the current study were then randomly assigned to one of three conditions: (1) gossip condition, (2) numeric condition, or (3) no information condition. (1) In the gossip condition, participants read four gossip notes provided by one group of players from the first generation public goods game study. (2) In the numeric condition, participants received a log of the amount each person in a first generation public goods game group contributed to their common fund. (3) In the control condition, participants did not receive any additional information about a first generation public goods game group.

In the gossip and numeric conditions, we used a yoked design – randomly assigning each four-person group collected in Feinberg et al. (2014) with two new four-person groups (8 participants) of our own. For those randomly assigned to the *gossip* condition, each group of generation 2 participants saw the gossip notes of the four players in the generation 1 group to which they were yoked.[3] For those randomly assigned to the *numeric* condition, each group of generation 2 participants instead saw the actual contribution amounts of the four players in the generation 1 group to which they were yoked.[4] As such, our players received unique, real, and naturalistic information about players in the previous study. In order to assure statistical power, we yoked four groups of generation 2 participants to each generation 1 group: two groups assigned to the gossip condition, and two groups assigned to the numerical condition. To have equal sample sizes across the three conditions, we assigned the same number of participants to the control condition, where they did not receive gossip or numeric information about a generation 1 group. For more information about the generation 1 groups, see *Stimuli* section below.

In all conditions, second generation participants were then asked how much, on average, they thought first generation participants contributed to the common fund. This measure serves as a metric of how the participants perceived the prosociality of the first generation public goods game players.

Participants then learned they would also be playing a round of the public goods game with three other online participants matched at random.[5] They were told that each player could independently and anonymously choose how many points they would like to share with the group fund, just like the players from the first generation of the game. However, their contribution amounts would not be made public to other players, and there would not be opportunities for notes to be written.

Before they began the game, players were asked to forecast how much the average member of their own public goods game group would contribute to the common fund. Then, they chose how many points to contribute to their group fund themselves. The experimenters multiplied the number of points in each group fund by two, and the total number of points in the group fund was equally split between the members of the group. The decisions made by each participant and the other subjects in their group determined their earnings. As such, their contribution is used as a measure of selfishness in actions. Bonuses were calculated and distributed via Prolific.

Finally, we collected demographic information, and other related measures, such as trust and cynicism. See Supplementary Materials for the complete list of questions used in the study.

2.1.3. Stimuli

The gossip notes and contribution amounts from real groups of four players were derived from Feinberg et al. (2014). We selected only

[3] Each generation 2 group in the gossip condition was assigned the gossip notes from a specific generation 1 group. Each generation 2 group was presented with four gossip notes written by members of their assigned generation 1 group about each other. For example, members of a generation 2 group assigned to the generation 1 group "IGVT" received the following four gossip notes: Player 1: "contributed 0"; Player 2: "Contributed 0 points. LAAAAAAAAAAAME."; Player 3: "They contributed 0, the least of the group"; Player 4: "Contributed 8, good player."

[4] Each generation 2 group in the numeric condition was assigned the contribution amounts from a specific generation 1 group. Each generation 2 group was presented with four numeric amounts documenting the contributions of the members of their assigned generation 1 group. For example, members of a generation 2 group assigned to the generation 1 group "IGVT" received the following four contribution amounts: Player 1: 7 points; Player 2: 8 points; Player 3: 5 points; Player 4: 0 points.

[5] In order to ensure the game operated quickly, the participants played the game asynchronously from their group. They knew that they were randomly assigned to a group of three other players where their contributions would be summed, doubled, and redistributed equally by the experimenters immediately following data collection. Otherwise, they were not provided with additional information about their fellow group-mates.

groups from Feinberg et al. in which all four participants left a note, to eliminate variance in the number of notes read by participants in our gossip condition.

Thirty two groups from Feinberg et al. fit this criterion. A *t*-test revealed no significant difference in contributions between members of groups that produced 4, versus less than 4, gossip notes ($t(185) = -0.28$, $p = 0.78$, $d = -0.04$). Although it's impossible to prove a null, Bayes factor analyses found that the data from these groups were 0.16 times as probable to occur if the underlying distributions were different, compared to under the null hypothesis that these groups did not differ.

Additionally, we probed possible differences in contribution variance between groups that produced 4, versus less than 4, gossip notes. In order to assess group variation, we calculated the difference score between the highest contribution and the lowest contribution in each generation 1 group to get a measure of its spread. A t-test revealed no significant difference in contribution variations between groups that produced 4, versus less than 4, gossip notes ($t(38) = -1.71$, $p = 0.10$, $d = -0.48$).

2.2. Results

2.2.1. (Inaccurate) cynicism about previous generations

First, we investigated the effect of gossip on generation 2 recipients' estimates of public goods game contributions among generation 1 gossipers. Pairwise *t*-tests revealed that participants exposed to gossip inferred that generation 1 participants had contributed significantly less to the common fund than participants who were exposed to numeric information about actual contribution amounts ($t(494) = 5.72$, $p < 0.001$, $d = 0.51$) and than control participants who received no information($t(483) = 5.81$, $p < 0.001$, $d = 0.53$; see Fig. 2).

Not only did the participants in the gossip condition draw negative conclusions about past players, we also found that they were inaccurate. We computed difference scores between the actual mean contribution of each generation 1 public goods game group on the one hand, and generation 2 participants' estimates of that group on the other hand. For example, if the members of a generation 1 group contributed 8 points on average, and a generation 2 participant reported that they believed the average contribution of that generation 1 group was 4 points, this would be a difference score of −4 points from the true generation 1 group mean. Because control participants in generation 2 did not receive information about a specific first generation group, we computed a difference score between their estimates and the true mean across all generation 1 participants. Participants who were exposed to gossip had significantly more inaccurate estimations than participants in the numeric condition ($t(494) = 6.74$, $p < 0.001$, $d = 0.61$) and than control participants ($t(483) = 6.76$, $p < 0.001$, $d = 0.61$; see Fig. 3). Interestingly, all groups significantly underestimated contribution amounts, consistent with the idea that people generally underperceive others' prosociality. But this tendency was intensified by the presence of gossip.

Did the particular nature of gossip generation 2 players received sway their inferences? For each group of four notes produced by a generation 1 public goods game group, we quantified the number of negative notes, producing a score from 0 (no negative notes shared in the group) to 4 (all notes written in the group were negative notes). A linear regression model with an interaction term revealed that note valence moderated the relationship between condition estimates of generation 1 groups' contributions. Participants in the gossip condition estimated that the first generation contributed significantly less for each additional negative gossip note ($b = -0.61$, SE = 0.17, $p < 0.001$, partial $\eta 2 = 0.026$, see Fig. 4).

Because the gossip data is skewed negative, we have very few observations in the 0 negative notes bin. We ran a robustness check to ensure that the data for 0 negative notes was not driving our results. We conducted the same linear regression model, excluding the observations in the 0 negative note bin, and found that our results hold ($b = -0.52$, SE = 0.19, $p < 0.01$, partial $\eta 2 = 0.017$; see Supplementary Materials for correlation matrices).

To further understand how the particular nature of gossip that generation 2 players received could have affected their inferences, we had two new coders, naive to our hypotheses, independently rate the generation 1 gossip notes for valence on a continuous scale from −4 (very negative) to +4 (very positive). The two coders had an average pairwise correlation of 0.9, and in instances of coder discrepancy, we took the average between their two scores (average inter-coder valence ratings).

We used these ratings to calculate an average group valence rating for each generation 1 group of notes. To calculate group average note rating, we summed the average inter-coder valence ratings of the four notes in each group and divided them by four. Using a linear regression model with an interaction term, we found that the group average note valence moderates the relationship between condition estimates of generation 1 groups' contributions. Participants in the gossip condition estimated that the first generation contributed less as the average valence of the gossip notes they were provided went down ($b = 0.69$, SE = 0.19, $p < 0.001$, partial $\eta 2 = 0.03$, see Fig. 5).

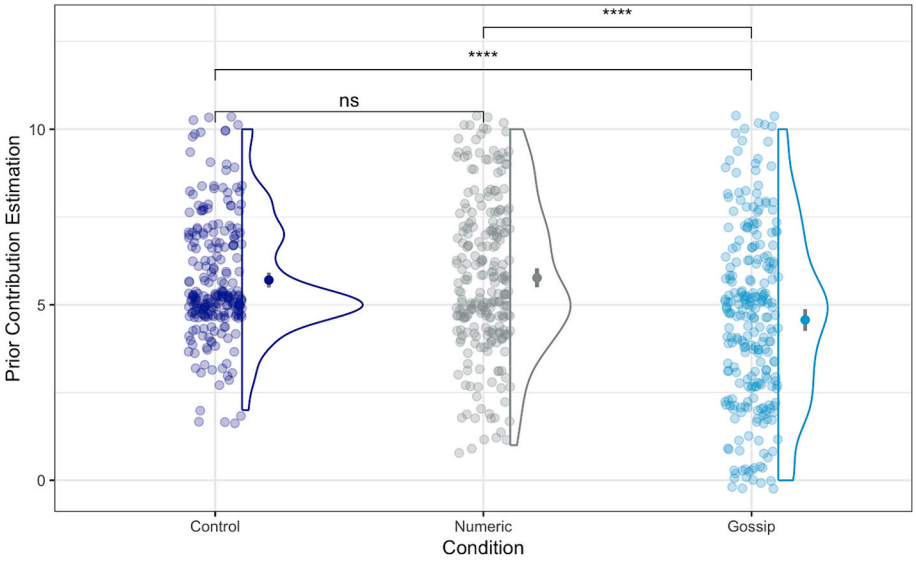

Fig. 2. ***Estimated prior group contributions by condition.*** Significance is indicated between all pairwise contrasts. Error bars represent 95 % confidence intervals. **** $p < 0.001$.

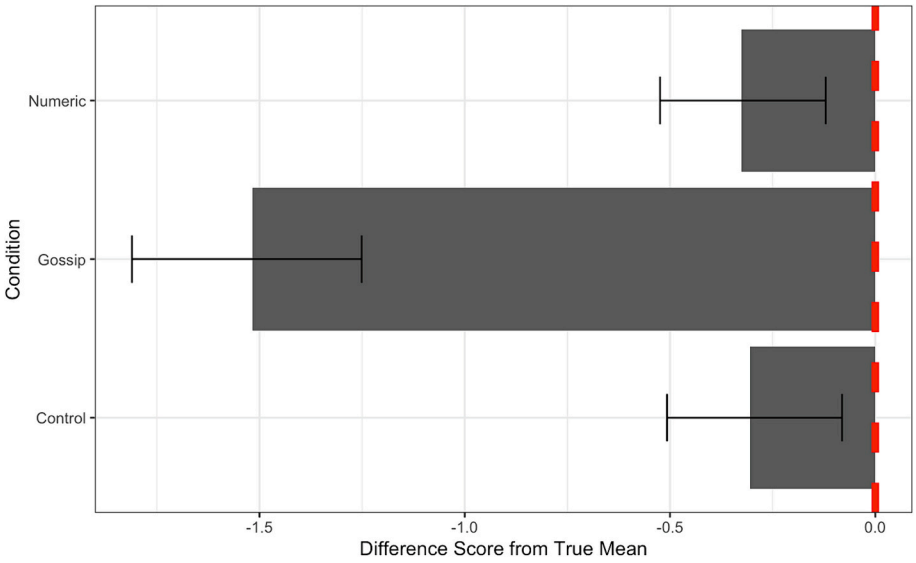

Fig. 3. ***Inaccuracy of estimated prior group contributions by condition.*** Bars represent the mean difference score in each condition. Error bars represent 95 % confidence intervals.

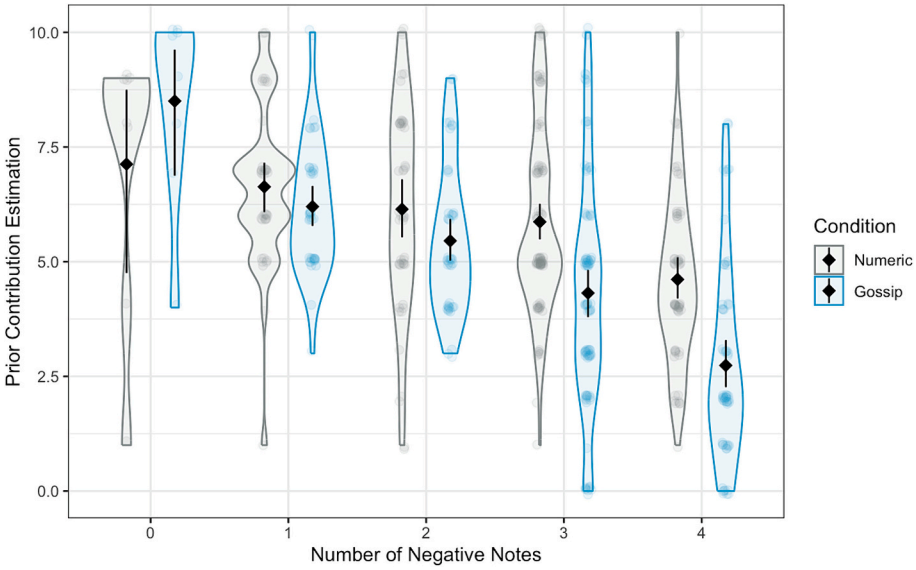

Fig. 4. ***Estimated prior group contributions by number of negative notes and condition.*** Diamonds represent mean prior group contribution estimations. Error bars represent 95 % confidence intervals.

2.2.2. Cynicism about future interactions

Next, we explored the effect of gossip from generation 1 on generation 2 participants' forecasts about a future public goods game–in particular, the average amount they expected players in their own public goods game groups to contribute. Pairwise *t*-tests revealed no differences between groups: participants who were exposed to gossip did not produce more pessimistic forecasts than participants in the numeric ($t(494) = 1.50$, $p = 0.133$, $d = 0.14$) or control conditions ($t(483) = 1.64$, $p = 0.10$, $d = 0.15$; see Fig. 6). These results deviated from our pre-registered hypothesis that gossip would not merely generate cynicism about the groups from which it comes, but also darken people's views of public goods game players in general.

2.3. Selfish actions

Lastly, we explored the effect of gossip on generation 2 participants' own contributions to the group fund in their own public goods game. Pairwise t-tests again revealed no differences across groups: participants who were exposed to gossip did not contribute less than participants in the numeric ($t(494) = 0.27$, $p = 0.79$, $d = 0.02$) or control conditions ($t(483) = -0.244$, $p = 0.81$, $d = -0.02$; see Fig. 7). These results deviated from our pre-registered hypothesis that gossip, by raising the expectation that others would free-ride, would inspire individuals to do the same themselves, in order to avoid being taken advantage of.

3. Discussion

Gossip is an ancient habit that remains misunderstood and too often maligned. Far from empty chit chat, it appears to serve multiple, crucial signaling functions–marking people's connections to each other, and also building and breaking reputations. Critically, it serves a role in collective action problems. When individuals face a social dilemma where they must choose between maximizing their own benefit or sacrificing for a larger group, gossip can serve as both carrot and stick:

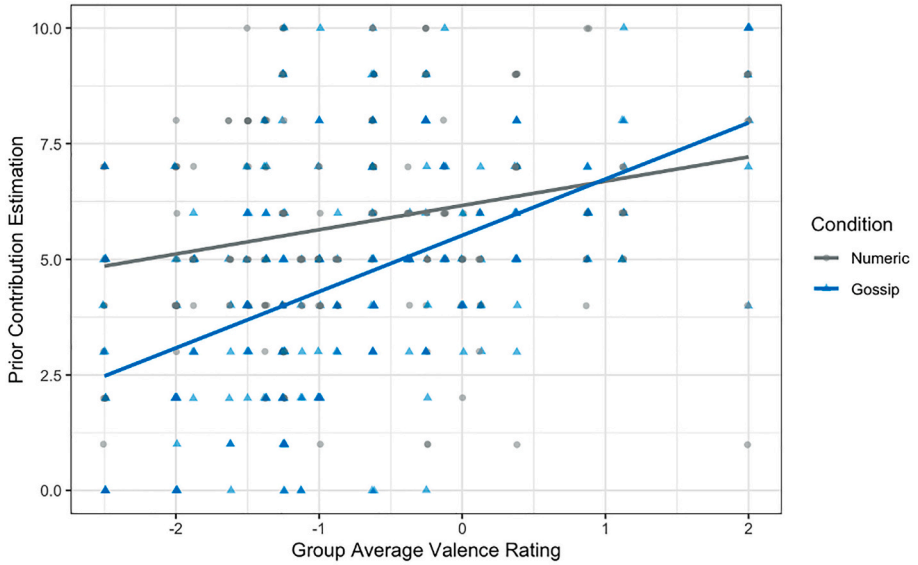

Fig. 5. ***Estimated prior group contributions by group average valence rating and condition.*** Lines represent the linear relationship between group average valence ratings and prior contribution estimations in the numeric and gossip conditions.

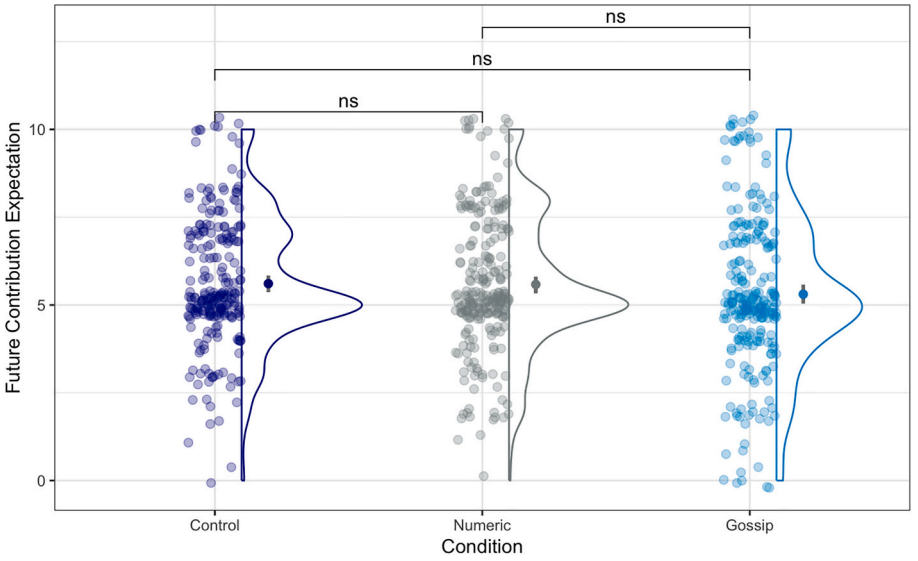

Fig. 6. ***Group contribution forecasts by condition.*** Significance is indicated between all pairwise contrasts. Error bars represent 95 % confidence intervals.

adding reputational benefits to prosocial choices and social costs to antisocial ones. In this work, we explored a hidden cost to this prosocial kind of gossip in driving cynical perceptions.

First, through archival reanalysis of prior work, we discovered that gossipers produce skewed information. In a public goods game setting, individuals who could pass along notes about other players were much more likely to write about free-riders than cooperative players. This was true even when cooperators outnumbered free-riders, suggesting that rather than sending notes that represented their group, gossipers focus on warning future players about the *worst* members of that group. This choice makes sense from an evolutionary and social perspective. If lawful citizens outnumber criminals, the most useful and actionable information one could provide is not about the majority, but about the small minority who do harm, so that others can avoid or sanction them.

However, we find that is not how recipients interpret gossip. Instead, they act as though gossip represents veridical information about the *entire* group from which it comes. In our study, second generation participants who received gossip from prior groups estimated that they contributed about 21 % less than second generation participants who received numerical information and about 20 % less than those who received no information. Their estimates were also 23 % less than the *actual* average contribution of first generation participants. In other words, gossip made people more cynical about the groups from which it came, and also more wrong about them.

We predicted that negative gossip would do even more: influencing not only people's inferences about the past, but generalizing to their expectations of the future and their actions. In particular, we hypothesized that people exposed to gossip from previous public goods game players who then were told they would take part in a public goods game of their own would (i) forecast that other players in their group would contribute relatively little and (ii) contribute relatively little themselves to hedge against other players' free-riding. Neither of these predictions was borne out.

This could signal a boundary condition for gossip. When a recipient hears about cheaters in another group, they might draw cynical inferences about *that group*, but not take their negative expectations any

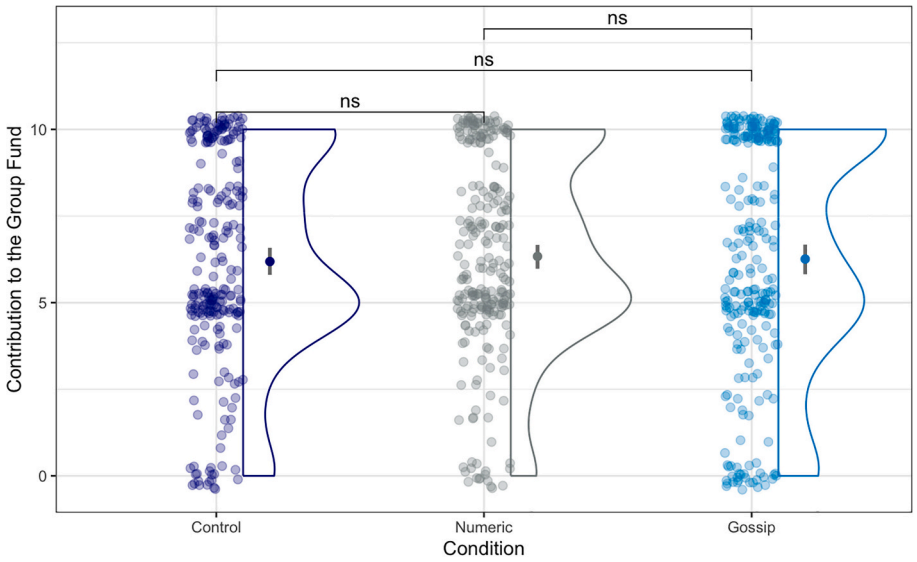

Fig. 7. ***Public goods game contributions by condition.*** Significance is indicated between all pairwise contrasts. Error bars represent 95 % confidence intervals.

further. This is consistent with the idea that people split the world into social groups, and feel motivated to protect the status of the groups to which they belong (Alicke & Sedikides, 2009; Turner et al., 1979). It's possible that generation 2 participants were happy to judge their predecessors based on gossip, but unwilling to judge their own groups or change their behavior based on this input.

There are other possible explanations for our null results. For example, it may be the case that different perceptions of the motives of generation 1 gossipers informs how generation 2 players respond. There is evidence that people trust gossipers less when they view the gossip as coming from a pro-self, as opposed to a prosocial motive (Testori et al., 2024). In our work, it is possible that varying perceptions of generation 1 motives led to different expectations and actions across our sample, thus wiping out an average effect. Future work could explore how motives interact with gossip skew to inform resulting beliefs and behaviors.

Our null results might also suggest that receiving negatively skewed gossip leads to more extreme behavior. For some, gossip may bring about a contrast effect, where receivers engage in downward social comparison to quarantine others' bad behavior. We observe evidence in favor of this hypothesis. The odds of participants exposed to gossip contributing all 10 points was 58.7 % higher than participants exposed to numeric information ($b = 0.46$, SE = 0.20, $z = 2.34$, $p = 0.02$, OR = 2.34) and 55.9 % higher than participants exposed to no information ($b = 0.44$, SE = 0.20, $z = 2.23$, $p = 0.03$, OR = 2.23). On the other side of the coin, for others, gossip might encourage assimilation, where gossip receivers behave selfishly. We also observe some evidence in favor of this hypothesis. The odds of participants exposed to gossip contributing zero points was 87 % higher than participants exposed to numeric information ($b = 0.63$, SE = 0.32, $z = 1.97$, $p = 0.049$, OR = 1.97) and 49.3 % higher than participants exposed to no information, although not significant ($b = 0.40$, SE = 0.30, $z = 1.32$, $p = 0.19$, OR = 1.32). Together, these findings suggest that exposure to negatively skewed gossip is polarizing. It can lead some participants to contrast and others to assimilate to the inferred behaviors of the past group of players.

Additionally, it is possible that our null effects represent the constrained context in which this study was conducted. Like much work on prosocial gossip, we relied on an anonymous task from behavioral economics, the public goods game. We combined this with a "note passing" task in which participants' gossip made its way to novel players. However, the public goods game is quite non-naturalistic. Much of social exchange is not monetary, and most gossip is not about people's financial contributions in an online game. This could have driven our effects on cynicism, because in the gossip condition, negative notes were the only source of information generation 2 participants had when drawing inferences. We feel this finding remains informative–especially because generation 1 participants generate gossip that is skewed and this produced incorrect inferences from outside observers–but future work should extend this multi-generational gossip approach beyond the context of economic games.

The constrained context of our work could also limit the generalizability of our findings. We chose to investigate gossip spread during a public goods game, specifically because this setting allows for people to observe obvious examples of cooperative and selfish behavior through contributions to the group. As such, we can examine which type of behavior people decide to gossip about when provided these clear cut instances of selfishness and generosity, and we can quantify the effects of that gossip on future generations of participants. However, in using this setting, we may limit the scope of our results. Daily life is far more varied, and obvious, quantifiable social dilemmas are likely relatively rare. This clarifies the value of our paradigm, in that we can examine gossip about canonical versions of selfishness and generosity. Yet, at the same time, it only allows us to investigate a slice of all daily gossip: evaluative information spread that occurs when agents engage in social dilemmas. Everyday gossip is much more broad and varied (Dores Cruz et al., 2021; Robbins & Karan, 2020). Future research could focus on gossip that occurs outside of a social dilemma to further understand how its valence and spread informs receivers perceptions, beliefs, and behaviors.

While the public goods game employed in Study 1 and Study 2 comes with this limitation, we specifically chose this paradigm because it also offers unique advantages. Beyond allowing for tight control of social incentives, our experimental public goods game design in Study 2 gives an idea of the substance that is communicated through gossip. By assigning participants to be exposed to (1) gossip about a group of players in a previous generation, (2) factual information about the behaviors of players in a previous generation, or (3) no additional information about players in a previous generation, we could specifically pit the effects of gossip spread against the effects of strictly factual information spread to see if prosocial gossip, meant to inform future interaction partners about a target's behavior, communicates something beyond fact.

Gossip is a means of informing people about others' actions, often seen as a vehicle for spreading accurate reputational information (e.g., Dunbar, 1993; Peters & Fonseca, 2020). Our results suggest that when it skews negatively, gossip instead paints a distorted picture of the targets. Participants who read gossip were more cynical and inaccurate in their

estimations of previous players' contributions compared to those who received factual logs of numeric contribution amounts. Our paradigm separates our findings from others in the literature who have reported that people update their evaluations of targets in response to gossip (Dores Cruz et al., 2021). We similarly show that people update their views, however, we also demonstrate that this updated view is more negative than if individuals had been presented with strictly factual information.

Our findings apply to contexts beyond the lab. For example, research shows that a tiny proportion of users on major online forums produce the majority of toxic content, much of it negative discussion about others (Kumar et al., 2023). These prolific gossipers could easily sway most people's perceptions beyond the forums they populate. Further, professional gossipers might do the same, even when their aims are prosocial. Many journalists, especially in investigative contexts, work to expose corruption, harm, and cheating of all sorts. As one journalist describes his industry's view in the *New York Times*, "society will get better when we show where it is going wrong" (Bornstein & Rosenberg, 2016). This is a moral mission, but one which incentivizes journalists to skew their gossip just like the players in our games did theirs. This skew has intensified across the 21st century, as news headlines have shifted to include more negative emotion words over time (Rozado et al., 2022). Readers appear to metabolize this negative gossip. For instance, political reporting increases audiences' cynicism about candidates and government (Cappella & Jamieson, 1997). And in the US, polls from 1990 to 2020 find that respondents believed violent crime was steadily increasing, even though FBI statistics demonstrate it actually *dropped* by 50 % over that span (Zaki, 2024). Although it's impossible to tell exactly where these perceptions come from, increased and increasingly sensationalistic representations of crime in the news coincide with Americans' cynical–and incorrect–beliefs.

Future work should probe how gossip, both personal and professional, guides people's perceptions and actions. Our work suggests that even when well-intended minds evaluate and spread information about people, they may make them seem morally smaller than they actually are.

Open practices

All data and code for both studies, as well as the pre-registration for Study 2 are available at https://osf.io/4htpu/.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

CRediT authorship contribution statement

Samantha Grayson: Writing – review & editing, Writing – original draft, Visualization, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Matthew Feinberg:** Supervision, Investigation, Data curation, Conceptualization. **Robb Willer:** Supervision, Conceptualization. **Jamil Zaki:** Writing – review & editing, Writing – original draft, Supervision, Resources, Methodology, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data/code is publicly available on OSF

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jesp.2024.104682.

References

Alicke, M. D., & Sedikides, C. (2009). Self-enhancement and self-protection: What they are and what they do. *European Review of Social Psychology, 20*(1), 1–48. https://doi.org/10.1080/10463280802613866

Anderson, N. H. (1965). Averaging versus adding as a stimulus-combination rule in impression formation. *Journal of Experimental Psychology, 70*(4), 394–400. https://doi.org/10.1037/h0022280

Archer, J., & Coyne, S. M. (2005). An integrated review of indirect, relational, and social aggression. *Personality and Social Psychology Review, 9*(3), 212–230. https://doi.org/10.1207/s15327957pspr0903_2

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology, 5*(4), 323–370. https://doi.org/10.1037/1089-2680.5.4.323

Bellovary, A. K., Young, N. A., & Goldenberg, A. (2021). Left- and right-leaning news organizations use negative emotional content and elicit user engagement similarly. *Affective Science, 2*(4), 391–396. https://doi.org/10.1007/s42761-021-00046-w

Bornstein, D., & Rosenberg, T. (2016, November 14). *Opinion | when reportage turns to cynicism.* The New York Times. https://www.nytimes.com/2016/11/15/opinion/when-reportage-turns-to-cynicism.html.

Boyd, R., & Richerson, P. J. (1992). Punishment allows the evolution of cooperation (or anything else) in sizable groups. *Ethology and Sociobiology, 13*(3), 171–195. https://doi.org/10.1016/0162-3095(92)90032-Y

Cappella, J. N., & Jamieson, K. H. (1997). *Spiral of cynicism: The press and the public good.* Oxford University Press.

Cosmides, L., Tooby, J., Fiddick, L., & Bryant, G. A. (2005). Detecting cheaters. *Trends in Cognitive Sciences, 9*(11), 505–506. https://doi.org/10.1016/j.tics.2005.09.005

De Gouveia, C. M., Van, V. L. J., & Crafford, A. (2005). Towards a typology of gossip in the workplace. *SA Journal of Human Resource Management, 3*(2), 56–68. https://doi.org/10.10520/EJC95806

Deng, L., Sun, W., Xu, D., & Ye, Q. (2021). Impact of anonymity on Consumers' online reviews. *Psychology & Marketing, 38*(12), 2259–2270. https://doi.org/10.1002/mar.21565

Dores Cruz, T. D., Thielmann, I., Columbus, S., Molho, C., Wu, J., Righetti, F., … Balliet, D. (2021). Gossip and reputation in everyday life. *Philosophical Transactions of the Royal Society, B: Biological Sciences, 376*(1838), 20200301. https://doi.org/10.1098/rstb.2020.0301

Dunbar, R. I., Marriott, A., & Duncan, N. D. (1997). Human conversational behavior. *Human Nature (Hawthorne, N.Y.), 8*(3), 231–246. https://doi.org/10.1007/BF02912493

Dunbar, R. I. M. (1993). Coevolution of neocortical size, group size and language in humans. *Behavioral and Brain Sciences, 16*(4), 681–694. https://doi.org/10.1017/S0140525X00032325

Dunbar, R. I. M. (2004). Gossip in evolutionary perspective. *Review of General Psychology, 8*(2), 100–110. https://doi.org/10.1037/1089-2680.8.2.100

Fehr, E., & Gächter, S. (2000). Cooperation and punishment in public goods experiments. *The American Economic Review, 90*(4), 980–994.

Feinberg, M., Willer, R., & Schultz, M. (2014). Gossip and ostracism promote cooperation in groups. *Psychological Science, 25*(3), 656–664. https://doi.org/10.1177/0956797613510184

Feinberg, M., Willer, R., Stellar, J., & Keltner, D. (2012). The virtues of gossip: Reputational information sharing as prosocial behavior. *Journal of Personality and Social Psychology, 102*(5), 1015–1030. https://doi.org/10.1037/a0026650

Fetchenhauer, D., & Dunning, D. (2009). Do people trust too much or too little? *Journal of Economic Psychology, 30*(3), 263–276. https://doi.org/10.1016/j.joep.2008.04.006

Fetchenhauer, D., & Dunning, D. (2010). Why so cynical?: Asymmetric feedback underlies misguided skepticism regarding the trustworthiness of others. *Psychological Science, 21*(2), 189–193. https://doi.org/10.1177/0956797609358586

Hofmann, W., Wisneski, D. C., Brandt, M. J., & Skitka, L. J. (2014). Morality in everyday life. *Science, 345*(6202), 1340–1343. https://doi.org/10.1126/science.1251560

Jolly, E., & Chang, L. J. (2021). Gossip drives vicarious learning and facilitates social connection. *Current Biology: CB, 31*(12), 2539–2549.e6. https://doi.org/10.1016/j.cub.2021.03.090

Kumar, D., Hancock, J., Thomas, K., & Durumeric, Z. (2023). Understanding the behaviors of toxic accounts on Reddit. In *Proceedings of the ACM Web Conference 2023* (pp. 2797–2807). https://doi.org/10.1145/3543507.3583522

Milinski, M., Semmann, D., & Krambeck, H.-J. (2002). Reputation helps solve the 'tragedy of the commons.'. *Nature, 415*(6870). https://doi.org/10.1038/415424a. Article 6870.

Miller, D. T. (1999). The norm of self-interest. *American Psychologist, 54*(12), 1053–1060. https://doi.org/10.1037/0003-066X.54.12.1053

Pan, X., Hsiao, V., Nau, D. S., & Gelfand, M. J. (2024). Explaining the evolution of gossip. *Proceedings of the National Academy of Sciences, 121*(9), Article e2214160121. https://doi.org/10.1073/pnas.2214160121

Peeters, G., & Czapinski, J. (1990). Positive-negative asymmetry in evaluations: The distinction between affective and informational negativity effects. *European Review of Social Psychology*. https://doi.org/10.1080/14792779108401856

Peters, K., & Fonseca, M. A. (2020). Truth, lies, and gossip. *Psychological Science, 31*(6), 702–714. https://doi.org/10.1177/0956797620916708

Robbins, M. L., & Karan, A. (2020). Who gossips and how in everyday life? *Social Psychological and Personality Science, 11*(2), 185–195. https://doi.org/10.1177/1948550619837000

Rozado, D., Hughes, R., & Halberstadt, J. (2022). Longitudinal analysis of sentiment and emotion in news media headlines using automated labelling with transformer language models. *PLoS One, 17*(10), Article e0276367. https://doi.org/10.1371/journal.pone.0276367

Rozin, P., & Royzman, E. B. (2001). Negativity Bias, negativity dominance, and contagion. *Personality and Social Psychology Review, 5*(4), 296–320. https://doi.org/10.1207/S15327957PSPR0504_2

Rule, N. O., Slepian, M. L., & Ambady, N. (2012). A memory advantage for untrustworthy faces. *Cognition, 125*(2), 207–218. https://doi.org/10.1016/j.cognition.2012.06.017

Skowronski, J. J., & Carlston, D. E. (1989). Negativity and extremity biases in impression formation: A review of explanations. *Psychological Bulletin, 105*(1), 131–142. https://doi.org/10.1037/0033-2909.105.1.131

Testori, M., Dores Cruz, T. D., & Beersma, B. (2024). Punishing or praising gossipers: How people interpret the motives driving negative gossip shapes its consequences. *Social and Personality Psychology Compass, 18*(2), Article e12924. https://doi.org/10.1111/spc3.12924

Turner, J. C., Brown, R. J., & Tajfel, H. (1979). Social comparison and group interest in ingroup favouritism. *European Journal of Social Psychology, 9*(2), 187–204. https://doi.org/10.1002/ejsp.2420090207

Van Lier, J., Revlin, R., & Neys, W. D. (2013). Detecting cheaters without thinking: Testing the automaticity of the cheater detection module. *PLoS One, 8*(1), Article e53827. https://doi.org/10.1371/journal.pone.0053827

Zaki, J. (2024). Hope for cynics. https://www.hachettebookgroup.com/titles/jamil-zaki/hope-for-cynics/9781538743065/?lens=grand-central-publishing.