

What matters most in supporting cooperation, the gossip content or the gossiper's intention?

Simulating motive interpretation in gossip dynamics

Abstract

Gossip provides individuals a great volume of information, which allows them to make informed decisions and better adapt to the environment around them. Like all pieces of information, however, if not correctly interpreted, gossip can lead to harmful consequences for individuals. Indeed, computational models have portrayed a complex picture on how gossip impacts cooperation, identifying several limitations of the mechanism. Recent theoretical models and empirical studies have shown how interpreting the information received through gossip is a key component to understand how gossip influences individuals and groups. Thus, we built an agent-based model where we examine two reaction mechanisms based on different reputation systems, in which agents first interpret the motive behind gossip and then react on the basis of this interpretation. While the first mechanism relies on an encompassing reputation system in which all pieces of information are used to inform future decisions with other group members, the second mechanism comprises a two-layer reputation system, in which agents' actions are separate from agents' reliability as gossipers. Our results support previous empirical findings asserting gossip as a cheap and effective way to support cooperation, and offer a solution for gossip driven by negative motives: as long as gossip receivers ignore the information provided by gossipers they deem unreliable, but don't punish them by refraining from cooperative interactions with them, cooperation can be sustained.

Keywords: Gossip, Reputation, Cooperation, Agent-Based Model, Multi-layer reputation, Motives

Introduction

Gossip - defined as a sender communicating to a receiver about a target who is absent or unaware of the content (Terence D. Dores Cruz, Nieper, Testori, Martinescu, & Beersma, 2021) - is a naturally occurring, widespread behaviour. Gossip has been argued to promote and sustain cooperation in groups (see Wu, Balliet, and Van Lange (2015, 2016) for experimental works, Giardini & Wittek 2019 for theoretical work, and Meng, Zhang, and Sun (2018); Smith (2014) for computational models). Indeed, research has shown that gossip helps the spread of information about group members, and, as such, provides the information needed to avoid and/or punish defectors in future interactions (Dunbar, 2004; Giardini, Paolucci, Villatoro, & Conte, 2014; Wu et al., 2016), which supports in-group cooperation (Terence D Dores Cruz, Beersma, Dijkstra, & Bechtoldt, 2019; Giardini, Balliet, Power, Számadó, & Takács, 2022; Peters & Fonseca, 2020).

The importance of gossip in supporting cooperation seems well-established, but computational models are consistently showing that the complete picture is way more complex. For instance, Giardini and Vilone (2016) revealed in their study that when noise is introduced in the transmission of gossip, cooperators fail to correctly identify defectors, reducing the positive contribution of gossip to cooperation (Nakamaru & Kawata, 2004; Roberts, 2008). Similarly, recent agent-based models further showed the limitations of gossip as a means to promote cooperation. According to Testori and colleagues, gossip needs to be truthful to promote slightly higher levels of cooperation compared to first-hand reputation (Testori, Hemelrijk, & Beersma, 2022). Moreover, in order to sustain cooperation, gossip should contain multiple pieces of information, making the message more complex and time consuming. That is, gossip should not only convey private assessments of others, but must also include elements of perspective taking and information about tolerance thresholds (Righi & Takács, 2022).

Thus, computational models portray a many-sided picture on how gossip impacts cooperation: while some models praise gossip as a powerful mechanism to sustain high levels of cooperation (Meng et al., 2018; Smith, 2014), other models show how its positive role is contingent on the applications of restrictive conditions (Giardini & Vilone, 2016; Righi & Takács, 2022; Testori et al., 2022), questioning whether gossip can indeed be considered as an efficient and effective mechanism to promote cooperation. Being a complex human behaviour, we should not consider gossip as a completely beneficial or harmful action, especially because its consequences for someone's reputation can be difficult to predict or even to assess in the short run. A person's inferences about another's reputation may be built on the basis of many potential factors (Ames & Fiske, 2013), and different actors may react differently to the same information. Evidence from punishment games shows how important motive attribution is in the choice of response: intentional harm is judged and punished more severely than unintentional harm (Ames & Fiske, 2013; Darley & Huff, 1990). Experimental work shows that even

in abstract situations such as a dictator game, the evaluation of an action depends on the choices available to an actor (Bardsley, 2008; List, 2007), and this is even more likely outside a controlled laboratory experiment.

In a recent theoretical model, Lee and Barnes (2020) argued that how gossip receivers interpret gossip is a key aspect in determining the consequences that work-place gossip has on cooperation. Drawing from attribution theory (Heider & Weiner, 2002; Kelley, 1967), they suggested that the characteristics of gossip (e.g., valence, work-relatedness, and credibility) as well as the context in which gossip occurs shape how receivers interpret gossip, which in turns influences gossip consequences. They identified three main motives that receivers can infer: self-interested motives, which focus on benefitting the sender; relational motives, which focus on creating interpersonal connections between the sender and receiver of gossip; and pro-social motives, which focus on the wellbeing of the group. They argued that when receivers interpret gossip as driven by pro-social or relational motives, they see the sender as more trustworthy and consider the gossip as more reliable than when they interpret gossip as driven by self-interested motives.

Given the critical role of the interpretation of motives to disentangle the impact of gossip on cooperation, we built an agent-based model in which receivers' reaction to gossip depends on their interpretation of the motives behind it (Beersma, Van Kleef, & Dijkstra, 2019; Lee & Barnes, 2021; Testori, Dores Cruz, & Beersma, In preparation). Our model tests whether introducing reaction mechanisms that depend on the receivers' interpretation of the senders' motives influences group cooperation. Does gossip support cooperation when agents' choices depend not on the content of the information received or the reliability of the source, but on the attribution of different motives to the gossiper? Differences in the interpretation of gossip motives has been shown to lead to different reactions to gossip in one-shot interactions in the lab (Testori et al., In preparation), and motive interpretations could also have a major impact on gossip and reputation dynamics in the long run, fostering or hampering cooperation. In our model, agents first interpret the motive behind gossip. In turn, this interpretation drives their reaction toward both the sender and the target of the gossip. Thus, to bring clarity on the effects of gossip on cooperation, we argue that a key aspect has been missing in previous models investigating gossip consequences. That is, gossip has always been taken at face value: receivers do not discern between the motives of gossipers, but simply use the information received from any gossiper to inform their future actions. Here we compare the face value mechanism (complete acceptance, baseline) to two mechanisms that agents can adopt when reacting to gossip, which reflect two different reputation systems.

The first mechanism (single reputation) relies on empirical findings showing that individuals tend to respond negatively to gossipers that they perceive as motivated by selfish motives (Berman & Silver, 2022; Farley, 2011; Peters & Kashima, 2015; Reeder, Kumar, Hesson-McInnis, & Trafimow, 2002); while they tend to respond positively to gossipers that they perceive as motivated by pro-social motives (Beersma & Van Kleef, 2012; Miguel A Fonseca & Peters, 2021; Peters & Kashima, 2015;

Testori et al., In preparation). In this case, agents update their reputation of the target and the sender of gossip as future interaction partners, based on their interpretation of the gossip motive (see Methods for more details). That is, if receivers interpret the gossip as motivated by pro-social motives, they will be more likely to cooperate with the sender and less likely to cooperate with the target in the future, and vice versa if gossip is interpreted as driven by selfish motives.

The second mechanism (double reputation) is rooted in the idea that individuals might hold several distinctive reputations of others, that differ depending on the task at hand. That is, reputation is often context specific (Garfield et al., 2021; Raub & Weesie, 1990; Takács et al., 2021): while it is hard to establish whether someone is good or bad overall, we can establish whether someone is good or bad when performing a specific action (e.g., one could have a positive reputation of their colleague as project manager but not as tennis partner). As seen in the formalisation of closely related concepts (see Hardin (1993); Lo Iacono and Testori (2021); Luhmann (2018) for similar debates on trust), we argue that reputation can be portrayed as a three-part relation: an agent i holds an opinion about an agent j with respect to a certain action a . Based on such literature, we built a mechanism in which agents hold two separate reputations: one as interaction partner (which reflects the likelihood, in the eyes of agent i , that agent j will cooperate in the next interaction; action a is the likelihood of cooperation in the next interaction), and one as gossiper (which indicates how reliable agent i expects the gossip shared by agent j to be; action a is the reliability of the gossip shared). With regards to the latter reputation, the reputation as gossiper, if agents interpret gossip as driven by pro-social motives, this implies that it is likely to be reliable, and they increase the reputation of the sender as a gossiper. On the other hand, if agents interpret gossip as selfishly motivated, this implies that gossip is not reliable, and they decrease the sender's reputation as a gossiper. Note that in these mechanisms, one's reputation as a gossiper and one's reputation as an interaction partner, are independent, which implies that if receivers interpret gossip as motivated by pro-social or pro-self motives, this will not impact the extent to which they will cooperate with the gossiper. Besides updating the reputation of the gossip sender, receivers increase or decrease the reputation of the target as an interaction partner proportionally to the reputation the gossip sender has as a gossiper. That is, if the sender has a positive reputation as a gossiper, the weight assigned to this sender's gossip will be higher compared to the weight assigned to a sender with a negative reputation as gossiper.

A potential advantage of separating the two reputation systems is that it allows groups to prevent significant drawbacks due to misinterpretation of gossip motives. That is, in a single reputation system, if gossip motives are correctly interpreted and the sender's reputation as interaction partner is updated, then selfish gossipers should be ostracized regardless of their cooperative actions, which could have a perverse effect at the system level. This spillover effect could be prevented by a double reputation system in which the sender's reputation as gossiper is updated upon receiving gossip, but not their reputation as an interaction partner. We expect that when reputations as interaction partner and gossiper

are separate, it is possible to discriminate more accurately between agents who engage in uncooperative actions (defectors) and agents who gossip for selfish reasons (unreliable gossipers).

This model contributes to the current interdisciplinary debate on the complex effects of gossip on cooperation (Giardini, Vilone, Sánchez, & Antonioni, 2021) in several ways, by providing 1) a more realistic model of gossip behaviour based on information processing theory (Lee & Barnes, 2021); 2) a better understanding of the gossip triad and how micro-level processes (motive interpretation) give rise to complex group-level dynamics and contribute to sustaining or hampering cooperation levels (Beersma et al., 2019; Lee & Barnes, 2021); 3) a comparison between different reaction mechanisms and reputation systems that could provide insightful findings on how to deal with gossip in groups.

Methods

We built an agent-based model with the aim of exploring how the receivers' reactions to gossip, based on their interpretation of the senders' motives, influence cooperation. Cooperation is calculated in the model as the number of cooperative decisions made by the agents in each round, and we report the fraction of agents cooperating in the population over rounds (i.e., the density of such cooperation over rounds). Our model is based on the work of Testori et al. (2022). The code is available on OSF (https://osf.io/keuc3/?view_only=5348fc2a1dcb44c798c9fe92e70c4e15)

Model description

The model can be divided into four parts: (1) initialization, (2) interaction, (3) gossip, (4) strategy update.

Let us consider a population of $N \in \{1, 2, \dots, n\}$ agents. Every agent has a gossip motive (see below for specifications) and an initial likelihood to cooperate $c_{i,0}$ drawn from a normal distribution $\mathcal{N}(0.5, 0.05)$. Each round, agents are randomly matched in pairs and interact with one another, and each agent can either cooperate (C) or defect (D). Their likelihood to cooperate depends on the reputation of their interaction partner $R_{i,j} \in \mathbb{R}$ (the reputation that agent i holds of agent j as an interaction partner). After each interaction, agents update their reputation $R_{i,j}$ upwards/downwards by a value ω_d if agent j cooperated/defected.

Once all agents interacted and updated their reputation of their interaction partners, agents share x gossip statements (default: $x=2$, see SI for $x=\{1, 5\}$) about their previous interaction partners. From the perspective of the gossip receiver, each gossip statement affects the reputation of the target as interaction partner and can also affect the reputation of the sender, as either an interaction partner or as a gossiper. The sender's motive determines the content of gossip (see below for further explanation),

and agents can use different mechanism to react to gossip (see below for further explanation). The simulation ends after 10^4 rounds.

Gossip motives and gossip content

People can have different motives to engage in gossip (Beersma & Van Kleef, 2012; Beersma et al., 2019; Terence D Dores Cruz, Balliet, et al., 2019), such as to protect the group, to negatively influence the gossip target, to vent one's emotions, to gather information, and to enjoy oneself. The content of gossip is likely to depend on these motives (Miguel A. Fonseca & Peters, 2018; Lee & Barnes, 2021; Peters & Fonseca, 2020).

Individuals motivated to protect group members should be more likely to send truthful information about the target (Miguel A. Fonseca & Peters, 2018; Lee & Barnes, 2021). By providing reliable information about the target, gossipers can protect receivers from norm-violators while sustaining cooperative behaviours among group members (Feinberg, Willer, & Schultz, 2014; Peters, Jetten, Radova, & Austin, 2017). Further, if gossip is motivated by a desire to gather and share useful information about the target, gossipers are expected to share their honest view of the target. Similarly, if people gossip to vent their emotions, they are more likely to share a true event that happened to them, so that they can let go of negative feelings linked to such experience and find comfort in the receiver's empathetic response (Terence D Dores Cruz, Beersma, et al., 2019; Martinescu, Janssen, & Nijstad, 2019).

In contrast, if individuals gossip to discredit the target, they are more likely to share negative information no matter whether this information is truthful or not, given that their goal is to shed an unfavourable light on the target (Beersma & Van Kleef, 2012; McAndrew, Bell, & Garcia, 2007; Peters & Fonseca, 2020). Lastly, no clear pattern can be delineated for the social enjoyment motive: people may share both positive and negative, truthful, and false information about the target when they aim to enjoy themselves, without any strategic reason. Indeed, research shows that both positive and negative gossip can foster bonds between senders and receivers (Terence D Dores Cruz et al., 2021; Peters et al., 2017).

Following this reasoning and recent developments in the gossip literature, it possible to divide gossip motives more broadly into motives that seem to align with largely pro-social concerns to benefit others, and motives that seem to align largely with pro-self concerns to benefit oneself (Beersma et al., 2019; Hess & Hagen, 2019; McAndrew et al., 2007; Testori et al., In preparation). Whereas pro-social motives are associated with sending truthful information, pro-self motives are associated with sending false information when this is beneficial for the gossipers. To capture this, we modelled two gossiping rules based on gossip motives: *always-true* and *always-negative*:

- (1) *Always-true*: Agents always gossip truthfully, sharing whether the target cooperated or defected in the previous interaction.
- (2) *Always-negative*: Agents always report that the target defected, no matter whether the target cooperated or defected.

These two rules have been operationalised in such a way as to simplify the complexity of human behaviour when gossiping: it is unlikely that someone would *always* share truthful or negative information (i.e., gossip motives likely vary per instance of gossip). Nevertheless, this operationalisation allows us to examine the impact of these two empirically grounded behaviours without introducing internal noise.

Reaction to gossip

In this model, we explored three different mechanisms of reaction to gossip (see Table 1 for overview).

Complete acceptance (Baseline): Receivers take the gossip at face value. Agents either increase or decrease their reputation of the target as an interaction partner by a value $\omega_{g,t}$, depending on whether the target cooperated or defected, without updating the reputation of the sender. This mechanism is the most common way to operationalise gossip, i.e., as an information exchange about a third agent, without any further psychological specification or reputational consequences for gossipers.

In the next two mechanisms, receivers first interpret the senders' motive to gossip. Receivers can correctly interpret the gossipers' motive (always-true and always-negative) with a certain probability $\rho \in \{0,0.25,0.5,0.75,1\}$.

Single reputation: After interpreting the senders' motive, receivers increase/decrease their reputation of the sender as an interaction partner by a value $\omega_{g,s}$ depending on whether the gossip is interpreted as always-true/always-negative. Receivers also increase/decrease their reputation of the target as interaction partner by a value $\omega_{g,t}$ if the target cooperated/defected. That is, gossip impacts not only the likelihood with which the receiver cooperates with the target but also the likelihood that the receiver cooperates with the sender.

Double reputation: Here, agents hold two separate reputations; one as interaction partner ($R_{i,j}$, which reflects the likelihood, in the eyes of agent i , that agent j will cooperate in the next interaction) and one as gossipers ($G_{i,j} \in [0,1]$, which indicates how reliable agent i expects the gossip shared by agent j to be). After interpreting the senders' motive, receivers increase or decrease their

reputation of the sender as a gossiper ($G_{i,j}$) by a value $\omega_{g,s}$ if the gossip is interpreted as driven by always-true or always-negative motives. The senders' reputation as gossipers does not affect the likelihood that the receiver will cooperate with the sender in future interactions but affects how the receiver uses the gossip to inform future interactions with the target: receivers increase or decrease their reputation of the target as interaction partner by a value $g_{i,j} * \omega_{g,t}$ if the target cooperated or defected. In other words, gossip impacts the likelihood that the receiver cooperates with the target in a manner proportional to the sender's reputation as gossiper, but it has no effect on the likelihood that the receiver cooperates with the sender.

Table 1. Overview of the reaction mechanisms

	<i>Complete acceptance</i>	<i>Single reputation</i>	<i>Double reputation</i>
Is the gossip motive interpreted?	No	Yes	Yes
Reputation systems updated following gossip	Target's reputation as interaction partner	Target's and sender's reputation as interaction partners	Target's reputation as interaction partner and sender's reputation as gossiper
Reaction rule towards the <u>target</u>	Increase/decrease reputation as interaction partner by a fixed value if target cooperated/defected	Increase/decrease reputation as interaction partner by a fixed value if target cooperated/defected	Increase/decrease reputation as interaction partner if target cooperated/defected. The value of the update depends on the reputation of the sender as gossiper
Reaction rule towards the <u>sender</u>	-	Increase/decrease reputation as interaction partner by a fixed value if receiver interprets the gossip motive as <i>always-true/always-negative</i>	Increase/decrease reputation as gossiper by a fixed value if receiver interprets the gossip motive as <i>always-true/always-negative</i>

Results

We ran our simulation with different values of the model parameters. The results hereafter use the default values of such parameters (see Table 2) and are averaged over 100 independent realisations. For robustness, the entire range has been tested and results are reported in the SI. To understand how gossip influences cooperation, we examined the percentage of cooperation (i.e., the fraction of agents that choose to cooperate in each interaction) over 10^4 rounds.

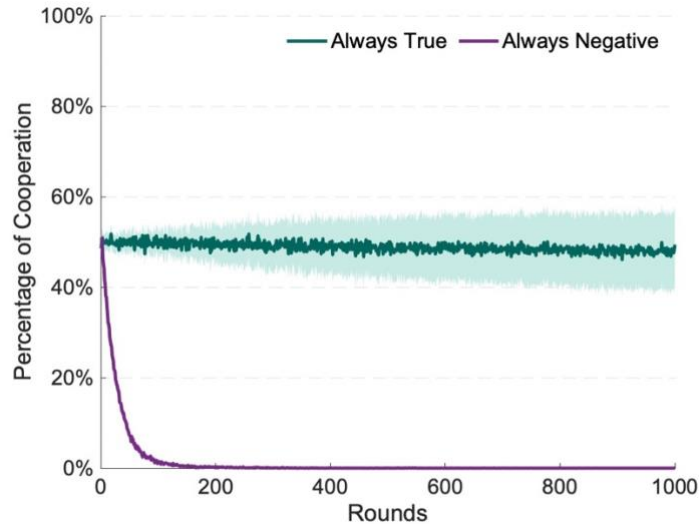
Table 2. Summary of model parameters

Variable	Description	Default	Range
N	Group size	50	{50,100,250}
x	Number of gossip statements exchanged per interaction	2	{1,2,5}
ω_d	Size of the update following direct interaction	0.5	{0.1,0.3,0.5}
$\omega_{g,s}$	Size of the update following gossip for senders	0.5	{0.1,0.3,0.5}
$\omega_{g,t}$	Size of the update following gossip for targets	0.5	{0.1,0.3,0.5}
ρ	Probability of correctly interpreting gossip motives	-	{0,0.25,0.5,0.75,1}

Complete acceptance

First, we report the results for our baseline mechanism. Here, gossip content is shaped by the senders' motives (always-true, always-negative), there is no interpretation of such motives by the receivers (complete acceptance), and receivers only update their reputation of the target as interaction partners. Figure 1 shows the percentage of cooperation throughout rounds depending on the gossip shared by agents (either always-true or always-negative).

Fig 1. Percentage of cooperation over rounds for *always-true* and *always-negative* gossip motives. Reaction mechanism: *complete acceptance*. The shadows indicate the 95% confidence intervals, calculated over 100 independent runs.




Always-negative gossip depletes cooperation in groups. Always-true gossip sustains the group's initial cooperation, by reducing the percentage of exploited cooperation (cases in which one agent cooperated and the other defected), and optimising the initial cooperation of the group. Indeed, the group is initialised with cooperation levels drawn from a normal distribution with mean 0.5.

Figure 2 shows the percentage of interactions reported in Figure 1 that resulted in cooperation being exploited (one agent cooperated and the other defected), and in cooperation being reciprocated (both agents cooperated). In the first 100 rounds, half of the cooperative actions are “wasted” as agents do not know with whom to cooperate and with whom to defect. As information circulates, the exploitation of such cooperative actions decreases when agents share always-true gossip (as seen in Fig.2a) while successful interactions (in which both agents cooperate) increase (see Fig.2b). That is, groups maintain levels of cooperation similar to those they were initialised with (in line with Testori, Hemelrijk & Beersma, 2022), but use such cooperation in a more efficient way.

On the other hand, when always-negative gossip is shared, agents decrease their likelihood to cooperate with other group members, which leads all agents to defect with one another, making it impossible for cooperation thrive.

Fig 2. Percentage of interactions in which (2a) one agent cooperated and the other defected (CD), and (2b) both agents cooperated over rounds (CC) for *always-true* and *always-negative* gossip. Reaction mechanism: *complete acceptance*.

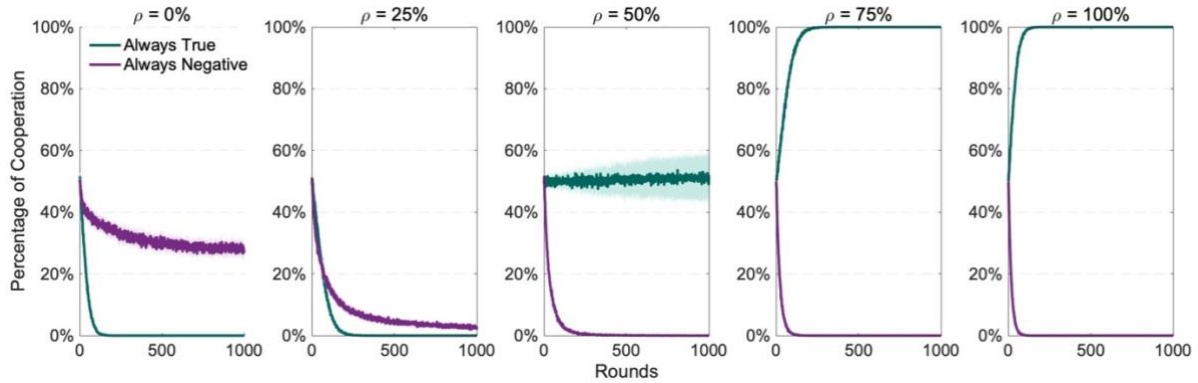


Single reputation

Receivers here interpret the motive behind the gossip received and base their reaction on such interpretation. If gossip is interpreted as always-true, then receivers increase their likelihood to cooperate with the sender (by increasing the sender’s reputation as interaction partner). On the other hand, if gossip is interpreted as driven by negative motives (always-negative), then receivers decrease their likelihood to cooperate with the gossip sender (by decreasing the sender’s reputation as interaction partner).

Fig 3. Percentage of cooperation over rounds for *always-true* and *always-negative* gossip motives. Reaction mechanism: *single reputation*. The panels show the results for different probabilities of

correctly interpreting the gossip motives ($\rho = 0\%$ means that gossip motives are always wrongly interpreted; $\rho = 100\%$ means gossip motives are always correctly interpreted). The shadows indicate the 95% confidence intervals, calculated over 100 independent runs.




If always-true gossip is mostly interpreted as such ($\rho > 50\%$), then it promotes and sustains high levels of cooperation, well above the initial group cooperation. In this case, the initial likelihood to cooperate rises until reaching 100% of cooperation. If always-true gossip is mostly incorrectly interpreted as always-negative ($\rho < 50\%$), then it depletes cooperation in groups (see Fig 3).

Always-negative gossip depletes the initial cooperation of the group, leading to no cooperation in all cases in which some correct interpretation occurs ($\rho > 0\%$). If always-negative gossip is always interpreted as truthful ($\rho = 0\%$), then cooperation stabilises around 30%. This level of cooperation is due to the positive response that gossipers receive when they are evaluated as honest gossipers (sharing always-true motivated gossip). That is, since receivers interpret gossip as always-true, they reward the gossipers by increasing the likelihood to cooperate with them in the future, which allows cooperation to stabilise around 30%.

Double reputation

Finally, we tested what happens to cooperation when gossipers interpret the gossip motive and, use this information to update the reputation of the senders as gossipers. This double reputation system is expected to support cooperation because it discriminates between malicious gossipers and defectors.

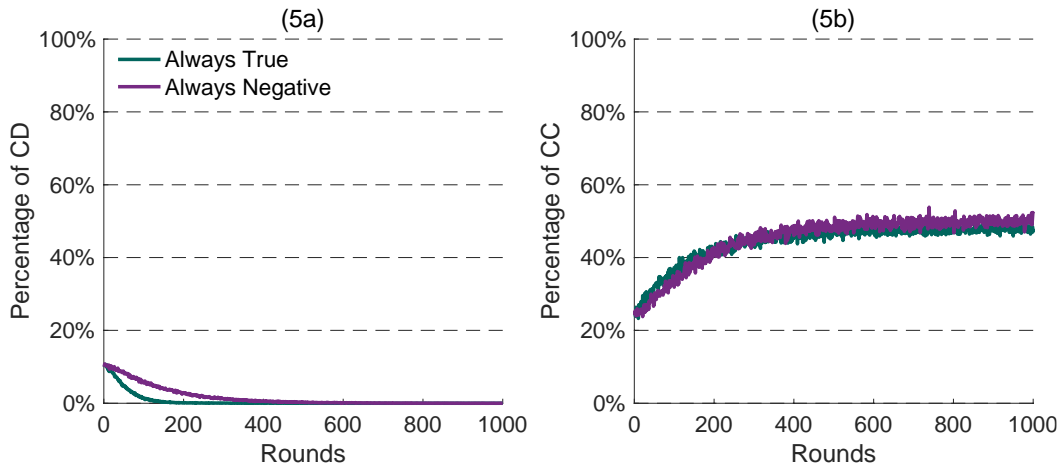
Fig 4. Cooperation density over rounds for *always-true* and *always-negative* gossip motives. Reaction mechanism: *double reputation*. The panels show the results for different probabilities of correctly interpreting the gossip motives ($\rho = 0\%$ means that gossip motives are always wrongly interpreted; $\rho = 100\%$ means that gossip motives are always correctly interpreted). The shadows indicate the 95% confidence intervals, calculated over 100 independent runs.



Always-true gossip sustains the initial cooperation of the group, regardless of how it is interpreted (similarly to the complete acceptance mechanism). If always-true gossip is mostly correctly interpreted ($\rho > 50\%$), then the content of gossip is used to inform future interaction with the targets. However, given that the gossip does contain true information and agents are initialized with a 50% chance to cooperate, the gossip reports positive (i.e., the target cooperated) and negative (i.e., the target defected) information with similar frequency. Positive and negative information, hence, cancels out, hampering cooperative cycles to emerge and cooperation to thrive. On the other hand, if always-true gossip is mostly incorrectly interpreted ($\rho < 50\%$), then the information received through gossip has little to no weight, and the agents base their actions mainly on their direct interactions (see Fig 4).

Noticeably, when always-negative gossip is correctly interpreted ($\rho = 100\%$), then senders quickly lose their credibility as gossipers and their gossip is disregarded ($G_{i,j} \rightarrow 0, \forall i, j$). Agents therefore base their actions solely on their direct interactions, and groups can sustain their initial cooperation. Over time, agents learn with whom to cooperate/defect, thus the percentage of interactions resulting in one agent cooperating and the other defecting decreases, while it increases the percentage of interactions resulting in both agents cooperating (see Fig 5a, 5b). However, if always-negative gossip is misinterpreted ($\rho < 100\%$), then cooperation is depleted. Indeed, when gossipers are not identified as sharing self-interested gossip (always-negative), receivers use the gossip to inform future interactions with the targets. Given that gossip always contains negative information about the target (the target defected), a negative loop of defection is established, and cooperation drops.

Fig 5. Percentage of interactions in which (5a) one agent cooperated and the other defected (CD), and (5b) both agents cooperated over rounds (CC) for *always-true* and *always-negative* gossip. Reaction mechanism: *double reputation*, $\rho = 100\%$.



Comparison between reactions mechanisms – How to foster higher levels of cooperation at the group level?

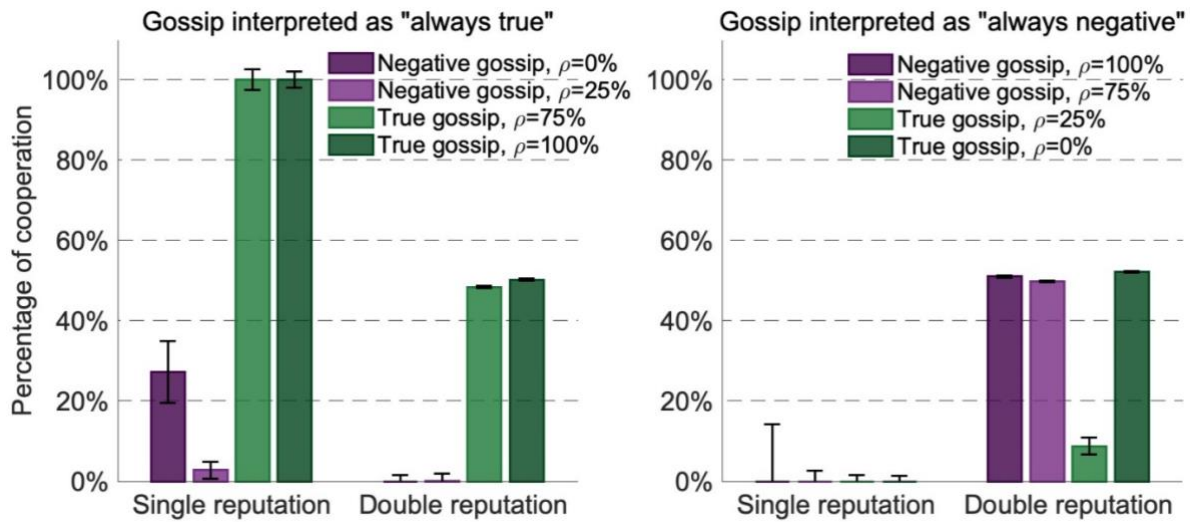
Our results show that cooperation is strongly influenced by how receivers interpret the sender’s motive. Here, we explored which reaction mechanism led to the highest cooperation levels for the population, on the basis of receivers’ interpretation.

Figure 6a compares the two reaction mechanisms (single reputation – left side – and double reputation – right side –) when gossip was interpreted as always-true. The green bars represent gossip that was driven by negative motives, but was incorrectly interpreted by the receiver. The purple bars represent gossip that was driven by positive (always-true) motives and was correctly interpreted by the receiver. Results show that, if gossip is interpreted as always-true, increasing the likelihood to cooperate with the sender (by increasing the sender’s reputation as interaction partner) and using the information received about the target to inform future actions leads to higher cooperation at the group level (see Fig 6a). This holds regardless of the original motivation with which gossip was shared (comparison between green and purple bars across the left and right side of the plot).

Figure 6b compares the cooperation density for the two reaction mechanisms (single reputation – left side – and double reputation – right side –) when gossip was interpreted as always-negative. The green bars represent gossip that was initially shared with negative motives and was correctly interpreted by the receiver. The purple bars represent gossip that was initially shared with honest (always-true) motives but was incorrectly interpreted by the receiver. If gossip is interpreted as always-negative, updating the sender’s reputation as gossiper and weighting the information based on the senders’ reputation is the most effective mechanism to achieve higher group cooperation. In other words, if agents interpret the gossip as negative, they should decrease the senders’ reputation as gossipers, but not decrease their cooperation towards them. By decreasing the senders’ reputation as gossipers, receivers are also less affected by the gossip content because the targets’ reputation as interaction partner changes proportionally to the reputation of the gossipers, and this avoids the emergence of loops

of defections. This mechanism allows agents not to be influenced by what they interpret as unreliable gossip, without hindering future interactions with the senders and the target of gossip.

Fig6. Level of cooperation in the last round when gossip is interpreted as (a) *always-true* (b) *always-negative*. Comparison between the two reactions: (left columns) single reputation; (right columns) double reputation. Green bars always refer to gossip initially shared with always-true motives. Purple bars always refer to gossip initially shared with always-negative motives.



Conclusions

Gossip comprises a large proportion of daily life communication (Robbins & Karan, 2020). Gossip provides information about peers' actions and behaviours, and as such it permits individuals to make more accurate predictions about others' behaviours, which can then be used to inform future interactions (Giardini & Wittek, 2019).

Like all pieces of information, however, if not correctly interpreted, gossip can lead to harmful consequences for individuals (Testori et al., 2022). In line with this argumentation, recent theoretical models and empirical studies have shown how interpreting the information received through gossip is a key component to understand how gossip influences individuals and groups (Lee & Barnes, 2021; Testori et al., In preparation). From previous research we know that the content shared through gossip can be influenced by the motivation with which gossip was sent (Beersma & Van Kleef, 2012; Terence D Does Cruz, Balliet, et al., 2019). If senders have malicious intentions when gossiping, it is likely that gossip will contain unreliable information. On the other hand, if senders gossip to protect others, gossip content is likely to provide accurate and reliable information about the target (Lee & Barnes, 2021; Wu et al., 2021). Hence, to efficiently use gossip to inform future decisions, it is crucial that receivers interpret the motivations behind the gossip they receive and adapt their reaction according to such interpretation.

The current study investigated how different reaction mechanisms based on the receivers' interpretation of gossip motives affect group cooperation. To this purpose, we built two reaction mechanisms based on different reputation systems and compared them to a baseline condition in which gossip was accepted at face value. In the new reaction mechanisms, agents first interpret the motive behind gossip and then react on the basis of this interpretation. While the first mechanism relies on an encompassing reputation system in which all pieces of information are used to inform future decisions with other group members, the second mechanism comprises a two-layer reputation system, in which agents' actions are separate from their reliability as gossipers. In the latter case, gossip information is only used if the source is believed to be reliable, and the reliability of the source in turn depends on the receiver's interpretation of the sender's motives.

Our results showed that when the single reputation mechanism is used, gossip promotes high levels of cooperation if receivers correctly interpret senders' pro-social motivations, but not if pro-social motivations are wrongly interpreted as selfish ones (wrong interpretation of more than 50% of the cases). On the other hand, if senders have selfish motivations, the highest cooperation level is achieved when receivers misinterpret the senders' motivation as pro-social. That is, if agents interpret the gossip as driven by pro-social motives, cooperation is more likely to increase because of the positive reputation of senders: when receivers interpret the senders as moved by pro-social motives, they are more likely to cooperate with them in future interactions, thus increasing the overall group cooperation.

When agents react to gossip using the double reputation mechanism, this sustains the initial level of cooperation of the group if the gossip was shared with pro-social intentions, regardless of whether they correctly or incorrectly interpret the senders' motive. On the other hand, if senders share pro-selfishly motivated gossip, the two-layer reputation mechanism prevents cooperation from declining only if agents correctly identify the motives behind gossip.

It is important to note some limitations of our model. First, our model does not elaborate on *how* agents interpret gossip motives. That is, populations are set to have a fixed probability with which they correctly interpret the motivation behind gossip. However, in real life people might use different cues to assess the motives behind gossip, such as body language, context, and the relationship between the gossip actors. Future models could be developed to assess which combination of these factors lead to the most accurate motive interpretation and how this affects cooperation.

Second, in our model agents do not change their motivation behind gossip. In real life, however, individuals might have different reasons to gossip, depending on who the target or the receiver of the gossip are (Giardini & Wittek, 2019). Thus, future models could investigate how different network structures impact agents' motivations to gossip and how they affect group cooperation.

Nevertheless, our results shed light on motive interpretation as a central factor for understanding whether and how gossip affects group cooperation, regardless of the agents' motivation to gossip. Empirical studies show how difficult it is to correctly interpret the motivations behind gossip and how individuals are prone to misinterpret others' intentions (Testori et al., In preparation; Walmsley &

O'Madagain, 2020). To simplify such a complicated process, our model identifies which reaction mechanism yields the highest cooperation levels, based on the receivers' interpretation of the motives behind gossip and not on the original motive of the sender. Comparing two reaction mechanisms leads to the conclusion that if gossip is interpreted as driven by selfish motives, agents should use the double-reputation system to ensure high levels of cooperation, while they should adopt the single-reputation system if they interpret gossip as driven by pro-social motives. That is, if agents interpret the gossip as negatively motivated, they should decrease the senders' reputation as gossipers, but not decrease their cooperation towards them; and if agents interpret the gossip as pro-socially motivated, they should increase their likelihood to cooperate with the senders in future interactions. Our results support the numerous empirical findings asserting gossip as a cheap and effective way to support cooperation (Peters & Fonseca, 2020; Wu et al., 2015) and offer a solution for gossip driven by negative motives (Testori et al., 2022): as long as gossip receivers ignore the information provided by gossipers they deem unreliable, but don't punish them by refraining from cooperative interactions with them, cooperation can be sustained.

References

- Ames, D. L., & Fiske, S. T. (2013). Intentional harms are worse, even when they're not. *Psychological Science, 24*(9), 1755-1762.
- Bardsley, N. (2008). Dictator game giving: altruism or artefact? *Experimental Economics, 11*(2), 122-133.
- Beersma, B., & Van Kleef, G. A. (2012). Why people gossip: An empirical analysis of social motives, antecedents, and consequences. *Journal of Applied Social Psychology, 42*(11), 2640-2670.
- Beersma, B., Van Kleef, G. A., & Dijkstra, M. T. (2019). Antecedents and consequences of gossip in work groups. *The Oxford handbook of gossip and reputation, 417-434*.
- Berman, J. Z., & Silver, I. (2022). Prosocial behavior and reputation: When does doing good lead to looking good? *Current opinion in psychology, 43*, 102-107.
- Darley, J. M., & Huff, C. W. (1990). Heightened damage assessment as a result of the intentionality of the damage-causing act. *British Journal of Social Psychology, 29*(2), 181-188.
- Dores Cruz, T. D., Balliet, D., Sleebos, E., Beersma, B., Van Kleef, G. A., & Gallucci, M. (2019). Getting a grip on the grapevine: extension and factor structure of the motives to gossip questionnaire. *Frontiers in Psychology, 10*, 1190.
- Dores Cruz, T. D., Beersma, B., Dijkstra, M. T. M., & Bechtoldt, M. N. (2019). The Bright and Dark Side of Gossip for Cooperation in Groups. *Front Psychol, 10*, 1374. doi:10.3389/fpsyg.2019.01374
- Dores Cruz, T. D., Nieper, A. S., Testori, M., Martinescu, E., & Beersma, B. (2021). An Integrative Definition and Framework to Study Gossip. *Group & Organization Management, 46*(2), 252-285. doi:10.1177/1059601121992887
- Dores Cruz, T. D., Thielmann, I., Columbus, S., Molho, C., Wu, J., Righetti, F., . . . Beersma, B. (2021). Gossip and reputation in everyday life. *Philosophical Transactions of the Royal Society B, 376*(1838), 20200301.
- Dunbar, R. I. (2004). Gossip in evolutionary perspective. *Review of General Psychology, 8*(2), 100-110.
- Farley, S. D. (2011). Is gossip power? The inverse relationships between gossip, power, and likability. *European Journal of Social Psychology, 41*(5), 574-579.
- Feinberg, M., Willer, R., & Schultz, M. (2014). Gossip and ostracism promote cooperation in groups. *Psychological Science, 25*(3), 656-664.
- Fonseca, M. A., & Peters, K. (2018). Will any gossip do? Gossip does not need to be perfectly accurate to promote trust. *Games and Economic Behavior, 107*, 253-281. doi:10.1016/j.geb.2017.09.015
- Fonseca, M. A., & Peters, K. (2021). Is it costly to deceive? People are adept at detecting gossipers' lies but may not reward honesty. *Philosophical Transactions of the Royal Society B, 376*(1838), 20200304.
- Garfield, Z. H., Schacht, R., Post, E. R., Ingram, D., Uehling, A., & Macfarlan, S. J. (2021). The content and structure of reputation domains across human societies: a view from the evolutionary social sciences. *Philosophical Transactions of the Royal Society B, 376*(1838), 20200296.
- Giardini, F., Balliet, D., Power, E. A., Számadó, S., & Takács, K. (2022). Four Puzzles of Reputation-Based Cooperation. *Human Nature, 33*(1), 43-61.

- Giardini, F., Paolucci, M., Villatoro, D., & Conte, R. (2014). Punishment and gossip: sustaining cooperation in a public goods game. In *Advances in social simulation* (pp. 107-118): Springer.
- Giardini, F., & Vilone, D. (2016). Evolution of gossip-based indirect reciprocity on a bipartite network. *Sci Rep*, *6*, 37931. doi:10.1038/srep37931
- Giardini, F., Vilone, D., Sánchez, A., & Antonioni, A. (2021). Gossip and competitive altruism support.
- Giardini, F., & Wittek, R. (2019). *The Oxford handbook of gossip and reputation*: Oxford University Press.
- Hardin, R. (1993). The street-level epistemology of trust. *Politics & society*, *21*(4), 505-529.
- Heider, F., & Weiner, B. (2002). Attribution theory. *The Motivation Handbook*, 31.
- Hess, N. H., & Hagen, E. H. (2019). Gossip, reputation, and friendship in within-group competition. *The Oxford handbook of gossip and reputation*, 275-302.
- Kelley, H. H. (1967). *Attribution theory in social psychology*. Paper presented at the Nebraska symposium on motivation.
- Lee, S. H., & Barnes, C. M. (2021). An attributional process model of workplace gossip. *Journal of Applied Psychology*, *106*(2), 300.
- List, J. A. (2007). On the interpretation of giving in dictator games. *Journal of Political economy*, *115*(3), 482-493.
- Lo Iacono, S., & Testori, M. (2021). Who are We Trusting? A Category-Based Formalisation of Trust. *Trust Matters: Cross-Disciplinary Essays*, 35.
- Luhmann, N. (2018). *Trust and power*: John Wiley & Sons.
- Martinescu, E., Janssen, O., & Nijstad, B. A. (2019). Gossip as a resource: How and why power relationships shape gossip behavior. *Organizational Behavior and Human Decision Processes*, *153*, 89-102.
- McAndrew, F. T., Bell, E. K., & Garcia, C. M. (2007). Who do we tell and whom do we tell on? Gossip as a strategy for status enhancement 1. *Journal of Applied Social Psychology*, *37*(7), 1562-1577.
- Meng, X., Zhang, Y., & Sun, C. (2018). *Impact of multi-valued gossip computational methods on cooperation*. Paper presented at the 2018 37th Chinese Control Conference (CCC).
- Nakamaru, M., & Kawata, M. (2004). Evolution of rumours that discriminate lying defectors. *Evolutionary Ecology Research*, *6*(2), 261-283.
- Peters, K., & Fonseca, M. A. (2020). Truth, Lies, and Gossip. *Psychol Sci*, *31*(6), 702-714. doi:10.1177/0956797620916708
- Peters, K., Jetten, J., Radova, D., & Austin, K. (2017). Gossiping about deviance: Evidence that deviance spurs the gossip that builds bonds. *Psychological Science*, *28*(11), 1610-1619.
- Peters, K., & Kashima, Y. (2015). Bad habit or social good? How perceptions of gossip morality are related to gossip content. *European Journal of Social Psychology*, *45*(6), 784-798.
- Raub, W., & Weesie, J. (1990). Reputation and efficiency in social interactions: An example of network effects. *American Journal of Sociology*, *96*(3), 626-654.
- Reeder, G. D., Kumar, S., Hesson-McInnis, M. S., & Trafimow, D. (2002). Inferences about the morality of an aggressor: the role of perceived motive. *Journal of Personality and Social Psychology*, *83*(4), 789.
- Righi, S., & Takács, K. (2022). Gossip: Perspective Taking to Establish Cooperation. *Dynamic Games and Applications*, 1-15.

- Robbins, M. L., & Karan, A. (2020). Who gossips and how in everyday life? *Social Psychological and Personality Science*, *11*(2), 185-195.
- Roberts, G. (2008). Evolution of direct and indirect reciprocity. *Proceedings of the Royal Society B: Biological Sciences*, *275*(1631), 173-179.
- Smith, E. R. (2014). Evil acts and malicious gossip: A multiagent model of the effects of gossip in socially distributed person perception. *Personality and social psychology review*, *18*(4), 311-325.
- Takács, K., Gross, J., Testori, M., Letina, S., Kenny, A. R., Power, E. A., & Wittek, R. P. (2021). Networks of reliable reputations and cooperation: a review. *Philosophical Transactions of the Royal Society B*, *376*(1838), 20200297.
- Testori, M., Hemelrijk, C. K., & Beersma, B. (2022). Gossip promotes cooperation only when it is pro-socially motivated. *Scientific Reports*, *12*(1), 1-12.
- Walmsley, J., & O'Madagain, C. (2020). The worst-motive fallacy: A negativity bias in motive attribution. *Psychological Science*, *31*(11), 1430-1438.
- Wu, J., Balliet, D., & Van Lange, P. A. (2015). When does gossip promote generosity? Indirect reciprocity under the shadow of the future. *Social Psychological and Personality Science*, *6*(8), 923-930.
- Wu, J., Balliet, D., & Van Lange, P. A. (2016). Gossip versus punishment: The efficiency of reputation to promote and maintain cooperation. *Scientific Reports*, *6*(1), 1-8.
- Wu, J., Számadó, S., Barclay, P., Beersma, B., Dores Cruz, T. D., Iacono, S. L., . . . Tiokhin, L. (2021). Honesty and dishonesty in gossip strategies: a fitness interdependence analysis. *Philosophical Transactions of the Royal Society B*, *376*(1838), 20200300.