Simulating Reputation Dynamics and their Manipulation

An Agent Based Model Built on Information Theory

Torsten Enßlin^{1,2,3}, Viktoria Kainz^{1,2}, and Céline Böhm⁴

¹ Max Planck Institute for Astrophysics, Karl-Schwarzschild-Str. 1, 85748 Garching, Germany

² Ludwig-Maximilians-Universität München, Geschwister-Scholl-Platz 1 80539 Munich, Germany

³ Excellence cluster ORIGINS, Boltzmannstr. 2 85748 Garching, Germany

⁴ School of Physics, The University of Sydney, Physics Road, Camperdown, NSW 2006, Australia

February 1, 2022

Abstract

Reputation is essential to human interactions and shapes group dynamics, however, it can be manipulated. The aim of this paper is to identify key aspects of reputation dynamics that eventually lead to a successful smearing campaign. For this purpose, we introduce an agent-based simulation which captures elements of communication dynamics that make agents more prone to believe lies, and to have an incorrect assessment of the honesty of others. The perceived honesty of agents being regarded here as their reputation, our simulations constitute a *reputation game*. Similarly to other works in the literature, the use of probability functions and Bayesian logic helps us to capture uncertainties in agents' beliefs and to model their associated cognitive limitations. However a major novelty of our work with respect to previous simulations on trust and reputation networks is how bounded rationality is dealt with. In our game, agents use information theory to compress all the information available to them in order to capture their complex environment in memories with finite capacity. The imperfect reasoning which results from such a compression process makes our agents vulnerable to deception. We find that this eventually leads to the emergence of communication and behavioral patterns which resemble reality, such as for example echo chambers, self-deception, deception symbiosis, and freezing of group opinions. As a result, the framework we propose could be used to develop methods to mitigate the impact of harmful communication strategies, and in the future safeguard the communication dynamics influenced by artificial intelligence systems in social media.

1 Introduction

Social media are becoming an increasingly important part of human communications. Through them a vast number of people, larger than ever, are directly communicating with each other. As a consequence, collective phenomena occur on a global scale; some being harmful to society or individuals. This includes the tendency for opinion polarization, radicalization, the emergence of filter bubbles, and hate speech. These phenomena are often enhanced by the presence of artificial intelligence (AI) systems in social media, which either influence the information flow pattern, e.g. by offering the participants different messages according to their profiles and thereby manipulate them,¹ or by participating directly in the exchange while pretending to be human [7–10]. The purpose of AI in these platforms might be to steer the debate into a certain direction or just to modify the general social atmosphere, for example

by spreading doubts about established facts, institutions, groups, or individuals [7, 11]. The intention could be for political purposes such as influencing an election, destabilizing a society, or misguiding public attention [12, 13]. Either way, such manipulations of public opinions often exploit the shortcomings and vulnerability of human minds and the way our brains process information. Furthermore, even without the presence of malicious agents, the complexity of social interactions on such a world-wide scale can easily lead to undesired socio-psychological effects, especially since social media is a new phenomenon and the human mind is not necessarily adapted to the complexity of opinion dynamics in such vast communication networks.

In this context, methods to detect and prevent unhealthy trends in the communication dynamics and malicious manipulations of the dynamics, are bound to become increasingly important. These methods require to understand what makes humans vulnerable to deception, what is the essence of successful manipulation strategies, and how these can be defeated. Especially methods to judge the trustworthiness of others and their statements play a central role in the interplay of manipulating and being manipulated. To

¹One objective of such AI systems is to keep the participants involved in the social media, and therefore to select messages that have a high chance of keeping their attention [1–3]. Extreme opinions can be particularly successful in that respect, potentially explaining why social media debates can become highly polarized [4–6].

address these questions from a theoretical perspective, we developed an agent based socio-psychological model.

This paper is organized as follows. In Sect. 2 we discuss agent based social models in general and in Sect. 3 the reputation game in particular. The results of reputation game simulations are presented in Sect. 4 and the therein emerging phenomena are discussed in Sect. 5. We conclude and give an outlook on future developments in Sect. 6.

2 Socio-Psychological Models

Agent based models [14, 15] have been used to construct predictive sociological simulations and study rumor dynamics [e.g. 16–19] and trust networks [20–23]. Many of the simulations that are modeling trust in a psychological, sociological and economical context are based on Bayesian probability theory [e.g. 24–34]. Bayesian probabilities are mathematical devices that enable to assign plausibility values to statements, and update these values following some basic principles whenever new information arrives [35, 36]. In this context it is no surprise that Bayesian probabilities provide an ideal language to describe uncertainties and doubts, which are central to human cognition. Consequently, *Bayesian models of cognition* [37] were developed and used in computational psychology [38–43]. These works assume that cognition, being an information processing activity, should follow a probabilistic reasoning [35, 36] and so they assign a probability value to every conceivable state of the world.² However, since there is an infinite number of such states and probability values, memorizing all the relevant information can easily exceed the capacity of all known information processing systems, including the human brain. Therefore, simplifications and mental shortcuts are necessary both in the real and virtual world. Here we show that such shortcuts could lead to a number of psychological effects, such as those discussed in Ref. [44], and therefore realistic psychological modeling should account for them [39, 45].

In a recent work [46], we have developed a socio-psychological agent based model, which emulates a reputation game and shows that psychological effects emerge from such limitations. Other works implement shortcomings of the human brain as memory loss over time [32, 47]. We instead assume that our agents have a finite memory

²For example, if a statement is supported by more evidence supporting it than against it, it would be assigned a probability value closer to one than to zero. This Bayesian interpretation of probabilities as indicators of truthfulness embraces and extends the perspective that regards probabilities as indicators for how frequent a statement turns out to be true. Thus, a probability of 0.6 can either mean that the associated statement is known to be true 60% of the time (probability as frequency indicator) or that the statement is either entirely true or false, however, it is not certain which way (probability as truthfulness indicator). Both of these roles probabilities play are essential here and are actually used in combination in this work. In particular, we need probabilities as truthfulness indicator on top of probabilities as frequency indicators in order to express the certainty of an agent about knowing how frequent some other agent lies.

capacity, and so information should be stored as effectively as possible at any time. This can be achieved by optimally compressing the memorized data at each memory update. For example, an agent’s brain in the reputation game maintains individual probabilities for each aspect of the world and constructs the joint probability over the world’s state as a direct product of those individual probabilities, but does not memorize the correlations between the different quantities to save memory.

The individual probabilities used by our agents are a function of two continuous parameters that are determined by the principle of *optimal belief approximation* [48]. These probabilities help each agent assess the honesty of others and assign a degree of credibility to the new information that they receive, which in turn helps them adjust their opinions about others. The novelty of this work lies in the fact that the rules for updating our agents’ opinions are based on information theory principles. This enables us to model complex opinion formation, and go beyond other works on reputation networks [49, 50]. For example, the way updates (or lack of) are implemented in the literature is such that if a message differs from the receiver’s opinion by an amount that exceeds a hard threshold, the information is rejected completely. On the contrary, if the discrepancy is lower than the pre-set threshold, the receiver’s beliefs are updated with the new information. In this paper, we do not require such a hard threshold. Indeed, with information theory, the rules of Bayesian logic generate a continuous transition between accepting and rejecting a new information and thus allow us to mimic human’s logic more accurately [51]. We will see that this smooth transition is critical as it leads to the emergence of real world phenomena, which other frameworks do not manage to reproduce as yet. In particular, we could demonstrate that some of the agent strategies that we propose here combined with some features of information theory successfully reproduce the key elements of a smearing campaign and other well-known behavioral phenomena. This indicates that we should be able to reverse engineering eventually and identify the strategy that an agent is using from the type of communication dynamics that unfolds. See the discussion of propaganda simulations in [46] for a good example of a dynamics that carries a visible signature of a manipulative strategy.

Our agents also possess a rudimentary theory of mind (ToM), which helps them to make more profitable decisions through the knowledge of their conversation partners’ opinions [52–54]. There exists already a rich literature of ToM-based psychological simulations. For example in PsychSim simulations [30, 55, 56], agents maintain a mental picture that contains beliefs of other agents’ mental states. They update their mental states when they receive new information, and use their belief states to anticipate other agents’ reactions to their actions. While Bayesian reasoning is used for these updates in PsychSim, the uncertainty of the world state is neither perceived, nor considered by agents to improve their strategy. This uncertainty is however captured

Figure 1: A conversation round of four agents. The initiator of each conversation chooses the conversation partner and the conversation topic (indicated next to the initiator’s communication arrow). The communication partner responds with a statement on the same topic.

in Bayesian ToM models (BToM), which take into account the probabilistic nature of belief states of agents [57, 58].

While BToM provides a more realistic description of the human mind, it also faces an important limitation. Indeed to reduce the calculation cost, BToM represents the different states of the world by a finite number of possibilities, and then assign probabilities values. This means that, in BToM, continuous variables need to be discretized in order to obtain a finite number of possibilities. Our approach differs in that we do not need to discretize the values of continuous variables. We are able to handle high-dimensional continuous probabilities using information compression techniques. The latter induces shortcomings, but our simulations show that they actually do capture some of the mistakes that victims of manipulators make and psychological vulnerabilities of the human minds. Therefore, as such, they are a key aspect of our modeling.

In the remainder of this paper, we shall summarize our simulations. Our model choices, the derivations of the mathematical framework, and discussions of the observed emergent phenomena can be found in great details in Ref. [46].

3 Reputation Game

The reputation game that we propose in this paper aims at capturing some of the relevant aspects of social communication dynamics. A central aspect of social communication is the credibility of the participants. Agents with the highest reputation are expected to tell the truth (even though they may not) while agents with the lowest reputation are meant to be liars and thus less trusted. To assess the trustworthiness of their network, each member of a social group thus needs to maintain a picture of the others’ honesty, which we identify here as agents’ reputation.

In the game, we introduce n agents $\mathcal{A} = \{\text{red, cyan, black, } \dots\}$, who are virtual entities with

certain human features exchanging messages among each others. The exchanges are organized in rounds consisting of conversations between initiating agents and their conversation partners. Each of the n agents initiates one conversation per round, see Fig. 1. In each conversation, the initiator picks a partner and a topic on which the conversation participants exchange statements. After this, they update their beliefs regarding the agent who is the topic of the conversation, themselves, and their conversation partner, see Fig. 2. The only topic of the messages in the game for now are the honesties of either themselves, their conversation partners, or a third agent in the game. Other topics are excluded in order to focus on the emerging trust network among the agents. The messages can be dishonest, and the frequency at which agents lie characterize them individually.

We define the honesty of an agent a as the frequency of their honest statements, $x_a \in [0, 1]$ and the set of all these values by the n -tuple $\underline{x} := (x_a)_{a \in \mathcal{A}}$. None of the agents’ real honesties are known by any of the agents at the start of the game. Hence each agent needs to infer the honesty of the other agents by communicating with them. The fact that agents do not know these values before starting the game makes agents more prone to deception and is closer to real life situations. Ideally, an agent a would maintain a probability function $\mathcal{P}(\underline{x}|I_a)$ over their n -dimensional world state, with I_a denoting an agent a ’s knowledge. However, this would require agent a to memorize a with n exponentially growing number of quantities, which easily gets out of hands if there are many agents and many communications. growing number of quantities, which easily gets out of hands if they are many agents and many communications. If, for example, ten honesty values were discriminated by an agent, that agent would need to store 10^n probability values to capture all possible states of the agent’s world. This is a with n exponentially growing number ($10^n = 10, 100, 1000, \dots$ for $n = 1, 2, 3, \dots$, respectively). To avoid this, an agent a only maintains individual belief states $\mathcal{P}(x_b|I_{ab})$ for the individual honesty x_b of each other agent b . Here I_{ab} expresses agent a ’s knowledge of agent b ’s honesty x_b . The joint belief state of agent a of all agents is then the direct product of individual probability densities

$$\mathcal{P}(\underline{x}|I_a) = \prod_{b \in \mathcal{A}} \mathcal{P}(x_b|I_{ab}). \quad (1)$$

The individual agent’s honesty probability densities are beta distributions,

$$\mathcal{P}(x_b|I_{ab}) := \frac{x_b^{\mu_{ab}}(1 - x_b)^{\lambda_{ab}}}{\mathcal{B}(\mu_{ab} + 1, \lambda_{ab} + 1)} = \text{Beta}(x_b|\mu_{ab}+1, \lambda_{ab}+1), \quad (2)$$

with $I_{ab} = (\mu_{ab}, \lambda_{ab})$ the set of parameters describing agent a ’s knowledge of agent b ’s honesty. More precisely the parameters μ_{ab} and λ_{ab} can be regarded as the number of honest and dishonest statements agent a believes agent b has made respectively.

Figure 2: Generating and receiving a message.

Studying how the perception of the other agents' honesties evolve is the central aim of the reputation game that is presented here, since this perception defines the trust network and the reputation of agents among each other. The immediate goal of our agents is to obtain a high reputation in order to be more trusted. This allows them to influence others.

Agents can benefit from conversations in two ways. First, they get information about other agents and about themselves. This helps them assess the honesty of agents in the game, and thus being informed about their social surrounding. Second, they have the opportunity to influence and even change others' opinions about their own reputation and that of the rest of the cohort.

Whenever agents receive a message, they try to determine whether they can trust the content. In real life humans can use several clues for this assessment so we try to mimic this in our simulations. For example, humans can detect lies when their conversation partner gives signs away, like blushing. So here we shall assume that agents can spot lies directly ten percent of the time. Like in real life, our agents can also judge the honesty of a message based on their prior knowledge. For example a statement might seem unbelievable, or they simply trust or distrust the message's sender. Finally, the messages that our agents receive might contain valuable new information on another agent. The exchange of information on an agent changes the opinions of the receiver and sender faster than simply interacting with the agent who is the topic of the conversation. However, when the information is accepted unfiltered, it can easily lead to a biased view of reality as well as the spread of rumors.

Thus, the recipients of a communication need to determine whether they should trust the message they received or not. If they trust the message fully, they will change their opinion about the agent who is the subject of the message if this is necessary and memorize that the sender has made an honest statement. If they absolutely distrust

Figure 3: Belief updates of an uncritical agent a receiving self-appraisal news from agent b . The initial belief state of a on b is $I_{ab} = (7, 5)$ (black solid line). This implies that a trusts b with $P(h|I_{ab}) = 0.57$ in general, independent of the message's content, as a is uncritical and therefore $P(h|d, I_{ab}) = P(h|I_{ab})$. The message is $J = (23, 1)$ (black dotted line) implying that b is very honest, $P(h|J) = 0.92$, if true. The Bayesian update results in a bimodal probability density (dashed orange line), which is a superposition of the cases of a regarding b 's message as entirely honest or dishonest, respectively. In contrast, the Bayesian update of a naive agent would only exhibit the peak at larger x_b , whereas a critical agent would mostly have the peak at lower x_b . The finally memorized compressed belief state of the uncritical agent (solid orange line) approximates the Bayesian probability density (dashed orange line).

the message, the receiving agent will ignore its content, and memorize that the speaker has lied. In most cases, however, a clear cut decision is not possible and only a probabilistic assessment of the message credibility can be made. In these cases, the correct Bayesian probability describing the knowledge of the situation is a superposition of two solutions: a credible message and a lie. The weights associated with these two cases are the probability that the communication is honest, $P(h|d, I_a)$, and dishonest, $P(\neg h|d, I_a) = 1 - P(h|d, I_a)$, respectively. Here, h means "honest message" and d is the data received in the communication, which includes the message content, the speaker's identity, and the observation whether or not the speaker has blushed. The resulting Bayesian probability density of such a superposition

$$\begin{aligned} \mathcal{P}(\underline{x}|d, I_a) &= \mathcal{P}(\underline{x}|h, d, I_a) P(h|d, I_a) + \\ &\quad \mathcal{P}(\underline{x}|\neg h, d, I_a) P(\neg h|d, I_a) \end{aligned} \quad (3)$$

is no longer of the form given by Eqs. 1 and 2. Thus, after receiving a message, the agents need to cast the correct Bayesian probability into the form of Eq. 2, and do so best in an information conserving manner. For this, our agents minimize the Kullback-Leibler (KL) divergence between the Bayesian correct probability density given by Eq. 3 and the

receiver	sender's reputation	blushing & confession	message surprise	sender's belief & intention
naive				
uncritical	✓	✓		
critical	✓	✓	✓	
smart	✓	✓	✓	✓

Table 1: Agents’ receiver strategies in terms of the used features for lie detection. Naive agents believe all messages.

parametric form given by Eqs. 1 and 2 used to memorize their beliefs,

$$\text{KL}_{\underline{x}}(I'_a, I''_a) := \mathcal{D}_{\text{KL}}(\mathcal{P}(\underline{x}|I'_a) || \mathcal{P}(\underline{x}|I''_a)) \quad (4)$$

$$:= \int d\underline{x} \mathcal{P}(\underline{x}|I'_a) \ln \left(\frac{\mathcal{P}(\underline{x}|I'_a)}{\mathcal{P}(\underline{x}|I''_a)} \right). \quad (5)$$

Here, $I'_a = (d, I_a)$ is the combination of agent a ’s knowledge that needs to be updated with the new data d that they receive, and $I''_a = (I''_{ab})_{b \in \mathcal{A}}$ the compressed information in the form of memorized parameters, with $I_{ab} = (\mu_{ab}, \lambda_{ab})$. The KL gets minimized with respect to the parameters of I''_a , which then constitute the new knowledge state. Such an information update, including the compression, is illustrated in Fig. 3.

The KL is the amount of information measured in nits (units of 1.44 bits) that an update of a probability function to some other probability function implies. Minimizing the KL ensures that only a minimal amount of information is lost in the compression. The compression conserves some integral quantities of the original probability distribution, namely the expectation values of $\ln(x_b)$ and $\ln(1 - x_b)$ of the original distribution on the honesty of agent b :

$$\langle \ln(x_b) \rangle_{(x_b|I'_a)} := \int d\underline{x} \mathcal{P}(\underline{x}|I'_a) \ln(x_b) \quad (6)$$

$$\langle \ln(1 - x_b) \rangle_{(x_b|I'_a)} := \int d\underline{x} \mathcal{P}(\underline{x}|I'_a) \ln(1 - x_b) \quad (7)$$

These moments characterize how surprising it is to find that agent b is honest or dishonest. In information theory, the surprise of an event is its negative log-probability. Therefore, $-\ln x_b$ measures how surprising an honest statement may seem, given an agent b ’s perceived honesty x_b . Information gets lost in the compression process, in particular the agents lose track of from which sources their beliefs stem. Thus, our agents cannot correct their beliefs on other agents even if they figure out that they trusted a deceptive source while acquiring these beliefs. This makes deceptive agents harder to spot and receiving agents more vulnerable to manipulations by deceptive agents and for longer.

Agents are equipped with both receiver and sender strategies. Receiver strategies can be naive, uncritical, critical, or smart. See Tab. 1 for an overview of these strategies and [46] for their mathematical implementations.

In order to determine whether a possible disagreement between a message and a belief (whether this belief stems from one’s own opinion, or the assumption of other’s opinion or intention) is relevant, agents maintain a characteristic “surprise” scale to measure how surprising an incoming message is. This characteristic scale adapts to the current social atmosphere agents find themselves in. Here social atmosphere refers to the diversity of opinions an agent is exposed to. An atmosphere is *frozen*, if the diversity of opinions is small, or *heated*, if opinions are very diverse.³

Ordinary agents employ a critical strategy when they receive a message that makes use of all of the above lie detection criteria except the ones that depend on the speaker’s opinion and intention. They accept messages only if these fit their own beliefs. Thus, in order to successfully spread lies, liars must send ordinary agents a message that is a small alteration of what they think is the victim’s belief. The direction in which the belief gets distorted depends on whether the subject of the message is a friend or an enemy of the sender. If an agent is the sender’s friend then the sender will depict them positively. If they are enemies then the sender will depict them negatively. Agents are considered as friends if they are remembered to have spoken positively about the sender the last time they made a statement on the sender.

Smart agents have similar receiver strategies as ordinary agents. However they also compare the message they receive to the belief and intention that they assume the sender has on the current topic. To do so, they need a rudimentary theory of mind, a mental representation of the belief state of their conversation partners.

A theory of mind is also needed to construct lies. Agents who are lying assume that their victims have a certain belief and modify it in a direction that suits them. The level of distortion is random, but is guided by their perception of their social network atmosphere and in particular the perceived surprise scale that marks the transition from honest to dishonest messages. In essence liars send messages that look to the receiver as honest as possible, but also contain misinformation on a level that has a good chance to go unnoticed by ordinary agents. *Smart* agents, however, are informed about the way lies are constructed and use this to detect them with higher confidence.

In addition to receiver strategies, agents possess also different sender strategies, see Tab. 2 for an overview and Ref. [46] for details. To initiate a conversation, agents need to choose a conversation partner, a conversation topic,

³In the simulation, simply the median surprise of the last ten messages an agent got is used as an estimator for the reference scale for lie detection and construction. This scale characterizes the perceived social atmosphere. The reasoning behind using this scale estimator is that on average, in the game’s setting, about half of the messages are honest and about half are lies and therefore the median of a sufficient large number of message surprises should mark the transition between typical honest statements and lies. Using a small number of surprises to determine the scale, like the ten we choose for our agents in the reputation game, allows for quick adaption to a changing social atmosphere.

sender	critical	smart	deceptive	strategic	anti-strat.	flattering	egocentric	aggressive	shameless
ordinary	✓								
clever	✓	✓	✓						
manipulative	✓	✓	✓		✓	✓			
dominant	✓	✓	✓	✓			✓		
destructive	✓	✓	✓	✓				✓	✓

Table 2: Summary of agents’ sender strategies.

and decide whether and how to lie, see Fig. 2. **Ordinary agents** make all these choices randomly, picking partners and topics, following a uniform distribution and lying at a frequency that characterizes them individually, and which is given by $1 - x_b$ for agent b . However, more efficient sender strategies exist. For example **clever agents** are *smart* and *deceptive*. By being *deceptive* all the time, they generate confusion and evade the risk to confess that they lie as slightly more honest agents occasionally do. By being *smart*, they become efficient liars thanks to better knowing the beliefs of other agents. **Manipulative agents** behave like clever agents, but are additionally *flattering* and *anti-strategic*. *Flattering* means that they make their conversation partner the topic of the conversation, whenever possible, and lie positively about them. *Anti-strategic* means that they prefer to converse with the agents that they perceive as having a low reputation. Manipulative agents befriend everyone, but especially frequent liars, to maximize their chances to be spoken about positively. **Dominant agents** are *clever*, *as well as egocentric* and *strategic*. Since *egocentric* agents preferentially pick themselves as conversation topic to praise themselves, and *strategic* agents prefer to address highly reputed others, the egocentric-strategic combination ensures that a dominant agent’s self-appraisal can efficiently spread in the game thanks to the high reputation of their victims. Finally, **destructive agents** are *clever*, *strategic*, *aggressive*, and *shameless*. *Aggressive* agents prefer to talk about their enemies to harm their reputation. They do not promote their friends nor themselves. When lying about the latter, they tell the receiver what they think the receiver believes on this topic. Being *shameless*, destructive agents also lie without blushing.

These sender strategies should to some extent mimic the communication patterns used by members of the dark triad. Of course, none of the strategies presented here is able to represent full personality types, but at least some aspects of them should be addressed. For example, the manipulative communication patterns typical of machiavellian are identifiable in our manipulative agent, the dominant and self-oriented appearance of narcissists is reflected in our dominant agent, and the antisocial behavior of many psychopaths is reflected in the destructive agent.

4 Results

We performed a number of simulations of the reputation game with three agents, agent black, cyan, and red. Agent black and cyan are always ordinary, but agent red has different strategies in the different setups. Agent black is very honest, with $x_{\text{black}} = 0.97$, agent cyan is mostly honest, with $x_{\text{cyan}} = 0.80$, and agent red is the least honest with $x_{\text{red}} \leq 0.14$. Each simulation setup corresponds to a different communication strategy used by red: ordinary, clever, manipulative, dominant, and destructive. For each of these setups we performed one hundred simulations with different random number sequences that are determining the agent’s decisions and analyzed the results. Fig. 4 shows the distributions of the agents’ relationships in terms of friendship and reputation. The vertical axis shows how reputed an agent a (e.g. black, cyan, or red) is in the eyes of another agent b on average within a simulation run while the horizontal axis indicates the fraction of the simulation run time during which agent a is regarded as a friend by b . The results of the hundred simulations are displayed as a density with a color corresponding to agent a . The contour line color corresponds to the agent b . For example, a red region with a cyan contour line, as shown in the top left panel of Fig. 4, displays how red is seen by cyan. In this example, red is regarded as being honest and being a friend by cyan since the colored area is in the top right corner of the plot. The true honesty of red is represented by a dashed red line on this plot as well. It is located well below the colored region, thus indicating that red’s true honesty is much lower than what cyan believes about red. This shows red’s influence on cyan, as it measures red’s capacity to make cyan believe something that is neither true nor supported by cyan’s direct observations of red.

Each other panel of Fig. 4 illustrates the impact of a different strategy by red. The top middle panel corresponds to the situation where red is an **ordinary agent**. Even though red is mostly regarded as a friend by the others, red has only a moderate reputation with black and cyan. However, it is larger than red’s intrinsic honesty which shows that red’s frequent lies served red’s purpose. In contrast, cyan and black reputations are lower than their respective honesties, showing that they are the victims of a smearing campaign from red. The reputation distribution of black, the most honest agent, is actually bimodal, with one peak closer to black’s intrinsic honesty and another one even below cyan’s honesty and reputation. Even though this is

Figure 4: Friendship and reputation relations $a \leftarrow b$ between agents a and b for all pairs of agents. The top left panel explains how to read the other panels. This panel shows a scenario in which agent a , who here is agent red, is regarded mostly as a friend as well as being perceived as honest by agent b , who here is agent cyan. Red’s true honesty is much lower and indicated by the dashed red line. Here and in the following panels, any contour is drawn in the color of agent b , who is having the respective view, and mark 10% of the peak value of the distribution shown in the color of the agent a , about whom the view is. The densities indicate in how many of 100 simulation b has a certain simulation-averaged view on a in terms of friendship and reputation. Other panels: The same for any of the views between the three agents black, cyan, and red for different scenarios. The scenarios differ by red being ordinary (top middle), clever (top right), manipulative (bottom left), dominant (bottom middle), and destructive (bottom right).

only about the honesty of other agents and not about facts or events this can be seen as the numerical illustration of the *Cassandra syndrome*, in which an honest agent ends up to be disbelieved by the others.

In the top right panel, agent red is a **clever agent**, meaning red is smart and fully deceptive. This makes red much more influential, reaching now higher reputation levels despite being completely dishonest. This is due to red’s lies, which are now based on a more accurate knowledge on the victim’s beliefs and are therefore more efficient. The fact that red never confesses to be a pathological liar also helps red to maintain a high reputation among others. This scenario should be compared to the remaining three ones, in which red is manipulative (bottom left), dominant (bottom middle), and destructive (bottom right). In those cases,

very different friendship and reputation relation patterns emerge.

The **manipulative agent** red manages to monopolize friendship. Both other agents regard red as a friend and do not consider the other agent as such. With the support of a less honest agent, namely cyan in this case, red manages to reach a higher reputation than that of cyan and black on average. Being manipulative is therefore a very successful strategy, clearly more effective than just being clever.

The **dominant agent** red less effectively befriends others, but more effectively reaches the highest reputation values due to incessant self-promotion. The dominant strategy turns out to be a high-risk high-gain strategy. In about half of the runs, the reputation dynamics freezes in with a relatively low level of reputation for red. In the other half,

Figure 5: Evolution of reputations (solid lines) and self-esteem (dashed lines) of the agents, as coded in the color of their names, in two simulations with identical setups but differing random numbers determining their decisions. In both red is using a dominant strategy while black and cyan are ordinary agents. Left: Simulation run, in which the opinions quickly freeze in, and actually reflect the real honesties of agents (square symbols on the right of the panels). Right: Simulation run, in which red succeeds to obtain a high reputation as well as a high self-esteem after a chaotic period with high variability of opinions. Note the occurrence of an extreme Cassandra syndrome for agent black during a period in between communication events 750 and 1250.

red reaches high reputations. Fig. 5 illustrates this with examples of the dynamics for these two regimes. There is an interesting phenomenon when red reaches a high reputation. Red’s reputation increases before red’s self-esteem does (self-esteem here is red’s opinion about their own honesty). This happens because red first has to convince the others of red’s high reputation, and only then the echo of these lies can raise red’s self-esteem. During this period, the KL divergence between red’s reputation and red’s self-esteem is large. The KL divergence measures the surprise (in terms of number of bits) that is associated with the amount of new information contained in a message. Hence a large KL divergence measures the amount of cognitive dissonance between a received message and the own belief. Being exposed to many surprising statements (with large KL) makes it harder for red to discriminate lies from honest statements. Our agents use the median of the last ten KL values to define an individual reference surprise scale which is key for them to detect lies. Since red uses the same scale to gauge the size of their own lies, these lies become extreme when red’s reputation strongly differs from red’s self-esteem. As a result, a *toxic social atmosphere* develops in the game with agents making critical statements being harshly attacked by red in particular. It is only when the positive feedback of the others raises red’s self-esteem that red’s reference scale relaxes. As a consequence, red’s statements become more moderate and the social atmosphere becomes less toxic. In the course of this evolution, red manages to manipulate the others to provide red with *narcissistic supply*, which in turn increases red’s self-esteem. In this simulation we see the occurrence of the Cassandra syndrome for agent black in Fig. 5.

The **destructive agent** red typically only reaches moderate reputation values compared to the manipulative and

dominant one, as can be seen in the bottom right panel of Fig. 4. However, the reputation of agent black and cyan are significantly lowered by the presence of the destructive agent red. For black, a frequent occurrence of the Cassandra syndrome is visible in form of a bimodal distribution of reputations, with a high reputation regime and a low reputation regime. The friendship distributions of cyan and black as perceived by red are also interesting as they are bimodal as well. An agent is either a friend or is red’s enemy during a simulation run. Destructive agents aim to decrease the reputation of their enemies, who in turn respond with statements that hurt red’s reputation. Owing to red’s focus on adversarial relationships, a friendly relationship is very unlikely to develop from an initially unfriendly one.

5 Emergent phenomena

The reputation game has a limited number of rules, that govern the actions of its players. Even though agents have many less means to orient themselves in their social environment and to influence it than humans have in real life, the reputation game shows a number of emergent phenomena (i.e. not implemented ab-initio) that resemble real world’s social and psychological effects. Some of these are not surprising, as being obvious consequences of the game’s setup. However some others are less obvious.

In particular, we observed communication patterns that resemble the Cassandra syndrome, echo chambers, group opinion freezing, self-deception, deception symbiosis, narcissistic supply, toxic social atmospheres etc. Some of these emergent phenomena were discussed in this letter, like the Cassandra syndrome, apparent in Fig. 4, and the narcissistic supply, shown in Fig. 5. For a more detailed overview

of all emergent phenomena that we saw appear in the simulations, see Ref. [46].

6 Conclusions and Outlook

Our simulations try to replicate the very complex reputation games real humans play [59, 60]. Although our agents’ minds are much less complex than that of humans, they mimic some basic functionalities that are common to human minds. For example, to protect themselves from deception agents have to assess the trustworthiness of the messages they receive. When the receiver is unsure about the sender’s honesty, this assessment requires to compare the message content to the agent’s prior knowledge and measure the resulting cognitive dissonance. This can be done by using the KL divergence, which is also convenient to measure the loss of information in the compression process that is implemented in this paper. In the real world, there may be critical factors other than cognitive dissonance (as is measured by our agents via the KL divergence) since humans pay particular attention to telltale signs [61]. However the honesty of a message will always need to be assessed and the amount of surprise compared to a characteristic scale in order to decide whether the new information is a lie. And this reference scale cannot be absolute; it has to reflect the environment in which the agents evolve. Therefore we have implemented a mechanism to learn and adapt this scale and found that including such a mechanism in the game allows others to manipulate this scale by bombarding an agent with a large quantity of confusing messages for example to undermine their ability to discriminate lies from honest statements in the long run. Propaganda and gaslighting strategies seem to be partly based on this effect.

Therefore, we believe that the reputation game presented here captures some of the dynamics between humans. If this is correct, a number of conclusions can be drawn that might be of interest to different communities. In the following we list these conclusions.

The **social scientist** might find interesting that simulating social dynamics with a minimal set of rules and parameters is able to reproduce known behavioral patterns of real humans and socio-psychological effects. Our framework could be used particularly as a testbed for checking hypothesis and see if certain rules do lead to the expected result when embedded in a communication dynamics. The simulation is a computational demonstration that despite being highly chaotic, the outcome of social dynamics is statistically predictable and might only depend on a very small number of parameters. In particular, we showed how the presence of a maliciously deceptive individual can drastically change the social atmosphere in their social surroundings. For the **cognition researcher**, a novel aspect of our model might relate to the ability and necessity for cognitive systems to compress information to memorize it. Indeed we showed that this makes agents more prone to manipulation. To **psychologists**, the mental states that we typically find

in agents following domination strategies might be of interest: There seem to be two states of low perceived cognitive dissonance, either when the agent has successfully built up a high reputation and a high self-esteem or when both these values are simultaneously low, and one of extreme high cognitive dissonance. The latter occurs when the agent’s reputation is high, but the self-esteem is still low. This happens when the “narcissistic supply” is too low to trump the low self-perception of the dominant agent. As a consequence, the dominant agent becomes very aggressive, leading often to toxic social atmospheres. For **policy makers** and **tech entrepreneurs** in social media, we have provided a numerical demonstration of the conditions under which a toxic social atmosphere can develop. Unexpectedly, this happens in our simulation when agents are exposed to too many statements that strongly conflict with their belief systems. Finally, our findings provide a numerical demonstration of how indoctrination via propaganda works and how its effect can be mitigated. This should be of interest to **everyone** who finds themselves in a toxic environment and tries to identify the mechanisms that lead to such a situation.

The motivation for the reputation game is to understand how to counteract the harmful communication strategies and help to mitigate their negative impact on society. We have shown that the agent based modeling together with the use of information theory techniques is a promising way to understand social phenomena qualitatively. The next step is to make our agents’ minds more realistic and to confront the outcome with real world data to improve the modeling. For example, we would like to include a more realistic treatment of friendship, and a better theory of mind, which captures more sophisticated strategies. In this respect, the concept of a Bayesian theory of mind, developed in [30, 34, 55–58], is very interesting. There, the agents anticipate possible futures that result from their action options and choose among those the most promising. Furthermore, in order to better reflect real world and social media situations the group size needs to be enlarged, as well as the language of the agents can be enriched, so that they can talk about several topics including the possibility to quote each other.

The behavior of such more complex agents can then be compared to communication patterns visible in social media, in order to test the model and adapt its parameters. For this purpose, a mapping between real world communication and the communication dynamics in our simulations needs to be established. This can be most conveniently done by training artificial neural networks. Analogous to existing methods of emotion recognition [62, 63], those networks could identify positive and negative references between social media participants and thereby provide large scale statistical data on social dynamics.

To complete these speculative thoughts about the further development of reputation game simulations, we can

imagine that once being sufficiently realistic, it can serve as a test bed for methods to detect and prevent abusive manipulations in social media. However, even before this, the reputation game can serve as a model system to test hypotheses about the inner workings of social dynamics and to get clues for yet unknown sociological effects that might exist in the real world.

Acknowledgement. We acknowledge helpful feedback on the manuscript by Prof. Dr. Sonja Utz.

References

- [1] *The Facebook Files*. <https://www.wsj.com/articles/the-facebook-files-11631713039>. Accessed: 2021-10-13. 2021.
- [2] *Frances Haugen: Facebook whistleblower reveals identity - BBC News*. <https://www.bbc.com/news/technology-58784615>. Accessed: 2021-10-13.
- [3] *Facebook Files: 5 things leaked documents reveal - BBC News*. <https://www.bbc.com/news/technology-58678332>. Accessed: 2021-10-13.
- [4] Pablo Barberá. “Social media, echo chambers, and political polarization.” In: *Social Media and Democracy: The State of the Field, Prospects for Reform* 34 (2020).
- [5] Ivan Dylko et al. “Impact of customizability technology on political polarization.” In: *Journal of Information Technology & Politics* 15.1 (2018), pp. 19–33.
- [6] Annamaria Silvana de Rosa et al. “Twitter as social media arena for polarised social representations about the (im) migration: The controversial discourse in the Italian and international political frame.” In: *Migration Studies* (2021).
- [7] Samantha Bradshaw and Philip Howard. “Troops, trolls and troublemakers: A global inventory of organized social media manipulation.” In: (2017).
- [8] Samantha Bradshaw and Philip N Howard. “Challenging truth and trust: A global inventory of organized social media manipulation.” In: *The Computational Propaganda Project* 1 (2018).
- [9] Liang Wu et al. “Misinformation in social media: definition, manipulation, and detection.” In: *ACM SIGKDD Explorations Newsletter* 21.2 (2019), pp. 80–90.
- [10] Emilio Ferrara et al. “Characterizing social media manipulation in the 2020 US presidential election.” In: *First Monday* (2020).
- [11] Chengcheng Shao et al. “The spread of fake news by social bots.” In: *arXiv preprint arXiv:1707.07592* 96 (2017), p. 104.
- [12] David MJ Lazer et al. “The science of fake news.” In: *Science* 359.6380 (2018), pp. 1094–1096.
- [13] Sinan Aral and Dean Eckles. “Protecting elections from social media manipulation.” In: *Science* 365.6456 (2019), pp. 858–861.
- [14] Dominik Klein and Johannes Marx. “Generalized Trust in the Mirror. An Agent-Based Model on the Dynamics of Trust.” In: *Historical Social Research / Historische Sozialforschung* 43.1 (163) (2018), pp. 234–258. ISSN: 01726404. URL: <http://www.jstor.org/stable/26395051>.
- [15] WARREN WEAVER. “SCIENCE AND COMPLEXITY.” In: *American Scientist* 36.4 (1948), pp. 536–544. ISSN: 00030996. URL: <http://www.jstor.org/stable/27826254>.
- [16] Guillaume Deffuant et al. “Mixing beliefs among interacting agents.” In: *Advances in Complex Systems* 3.01n04 (2000), pp. 87–98. DOI: 10.1142/S0219525900000078. URL: <http://www.lisc.clermont.cemagref.fr/imagesproject/finalreport/mixbel.pdf>.
- [17] Romualdo Pastor-Satorras et al. “Epidemic processes in complex networks.” In: *Reviews of Modern Physics* 87.3 (July 2015), pp. 925–979. DOI: 10.1103/RevModPhys.87.925. arXiv: 1408.2701 [physics.soc-ph].
- [18] Andrea E. F. Clementi et al. “Rational Fair Consensus in the GOSSIP Model.” In: *CoRR* abs/1705.09566 (2017). arXiv: 1705.09566. URL: <http://arxiv.org/abs/1705.09566>.
- [19] Daniel Geschke, Jan Lorenz, and Peter Holtz. “The triple-filter bubble: Using agent-based modelling to test a meta-theoretical framework for the emergence of filter bubbles and echo chambers.” In: *British Journal of Social Psychology* 58.1 (2019), pp. 129–149.
- [20] Gerardo Iniguez et al. “Effects of deception in social networks.” In: *Proceedings of the Royal Society B: Biological Sciences* 281.1790 (2014), p. 20141195.
- [21] Andrés Babino et al. “Maintaining trust when agents can engage in self-deception.” In: *Proceedings of the National Academy of Sciences* 115.35 (2018), pp. 8728–8733. ISSN: 0027-8424. DOI: 10.1073/pnas.1803438115. eprint: <https://www.pnas.org/content/115/35/8728.full.pdf>. URL: <https://www.pnas.org/content/115/35/8728>.
- [22] T.D. Huynh, N. R. Jennings, and N.R. Shadbolt. “An integrated trust and reputation model for open multi-agent systems.” In: *Autonomous Agents and Multi-Agent Systems* 13.2 (2006), pp. 119–154. URL: <https://eprints.soton.ac.uk/262593/>.
- [23] Y. Wang and J. Vassileva. “Trust and reputation model in peer-to-peer networks.” In: *Proceedings Third International Conference on Peer-to-Peer Computing (P2P2003)*. 2003, pp. 150–157. DOI: 10.1109/PTP.2003.1231515.

- [24] Jean Tirole. “A theory of collective reputations (with applications to the persistence of corruption and to firm quality).” In: *The Review of Economic Studies* 63.1 (1996), pp. 1–22.
- [25] Steven Tadelis. “What’s in a Name? Reputation as a Tradeable Asset.” In: *American Economic Review* 89.3 (1999), pp. 548–563.
- [26] Rosaria Conte and Mario Paolucci. *Reputation in artificial societies: Social beliefs for social order*. Vol. 6. Springer Science & Business Media, 2002.
- [27] Lik Mui et al. *Ratings in Distributed Systems: A Bayesian Approach*. 2001.
- [28] Lik Mui, Mojdeh Mohtashemi, and Ari Halberstadt. “Notions of reputation in multi-agents systems: a review.” In: *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1*. 2002, pp. 280–287.
- [29] Lik Mui. “Computational models of trust and reputation: Agents, evolutionary games, and social networks.” PhD thesis. Massachusetts Institute of Technology, 2002.
- [30] David V Pynadath and Stacy C Marsella. “Psych-Sim: Modeling theory of mind with decision-theoretic agents.” In: *IJCAI*. Vol. 5. 2005, pp. 1181–1186.
- [31] Wako Yoshida, Ray J Dolan, and Karl J Friston. “Game theory of mind.” In: *PLoS computational biology* 4.12 (2008), e1000254.
- [32] Hien Trang Nguyen, Weiliang Zhao, and Jian Yang. “A Trust and Reputation Model Based on Bayesian Network for Web Services.” In: *2010 IEEE International Conference on Web Services*. 2010, pp. 251–258. DOI: 10.1109/ICWS.2010.36.
- [33] Manuel Förster, Ana Mauleon, and Vincent J Vanetelbosch. “Trust and manipulation in social networks.” In: *Network Science* 4.2 (2016), pp. 216–243.
- [34] Jin Joo Lee, Fei Sha, and Cynthia Breazeal. “A Bayesian theory of mind approach to nonverbal communication.” In: *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE. 2019, pp. 487–496.
- [35] R. T. Cox. “Probability, Frequency and Reasonable Expectation.” In: *American Journal of Physics* 14 (Jan. 1946), pp. 1–13. DOI: 10.1119/1.1990764.
- [36] E.T. Jaynes et al. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003. ISBN: 9780521592710. URL: <https://books.google.de/books?id=tTN4HuUNXjgC>.
- [37] Thomas L Griffiths, Charles Kemp, and Joshua B Tenenbaum. “Bayesian models of cognition.” In: (2008).
- [38] Alison Gopnik and Joshua B. Tenenbaum. “Bayesian networks, Bayesian learning and cognitive development.” In: *Developmental Science* 10.3 (2007), pp. 281–287. DOI: <https://doi.org/10.1111/j.1467-7687.2007.00584.x>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-7687.2007.00584.x>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-7687.2007.00584.x>.
- [39] Thomas R Shultz. “The Bayesian revolution approaches psychological development.” In: *Developmental Science* 10.3 (2007), pp. 357–364.
- [40] *The Cambridge Handbook of Computational Psychology*. Cambridge Handbooks in Psychology. Cambridge University Press, 2008. DOI: 10.1017/CB09780511816772.
- [41] Jerome R Busemeyer et al. *The Oxford handbook of computational and mathematical psychology*. Oxford University Press, 2015.
- [42] Alison Gopnik and Elizabeth Bonawitz. “Bayesian models of child development.” In: *WIREs Cognitive Science* 6.2 (2015), pp. 75–86. DOI: <https://doi.org/10.1002/wcs.1330>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/wcs.1330>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/wcs.1330>.
- [43] David M Condon and René Möttus. *A role for information theory in personality modeling, assessment, and judgment*. Jan. 2021. DOI: 10.31234/osf.io/me8j4. URL: psyarxiv.com/me8j4.
- [44] Amos Tversky and Daniel Kahneman. “Judgment under Uncertainty: Heuristics and Biases.” In: *Science* 27 (Sept. 1974). URL: <http://links.jstor.org/sici?sici=0036-8075%2819740927%293%3A185%3A4157%3C1124%3AJUUHAB%3E2.0.CO%3B2-M>.
- [45] Aaron Sloman. “Putting the pieces together again.” In: *Cambridge Handbook on Computational Psychology* (2008), pp. 684–709.
- [46] Torsten Enßlin, Viktoria Kainz, and Céline Böhm. “Theoretical Modeling of Communication Dynamics.” In: *arXiv e-prints*, arXiv:2106.05414 (June 2021), arXiv:2106.05414. arXiv: 2106 . 05414 [physics.soc-ph].
- [47] Poong Oh and Soojong Kim. “An Evolutionary Model of the Emergence of Meanings.” In: *Communication Methods and Measures* 0.0 (2020), pp. 1–18. DOI: 10.1080/19312458.2020.1768519. eprint: <https://doi.org/10.1080/19312458.2020.1768519>. URL: <https://doi.org/10.1080/19312458.2020.1768519>.
- [48] Reimar Leike and Torsten Enßlin. “Optimal Belief Approximation.” In: *Entropy* 19.8 (Aug. 2017), p. 402. DOI: 10.3390/e19080402. arXiv: 1610.09018 [math.ST].

- [49] Sarah Shugars. “Good Decisions or Bad Outcomes? A Model for Group Deliberation on Value-Laden Topics.” In: *Communication Methods and Measures* 0.0 (2020), pp. 1–19. DOI: 10.1080/19312458.2020.1768521. eprint: <https://doi.org/10.1080/19312458.2020.1768521>. URL: <https://doi.org/10.1080/19312458.2020.1768521>.
- [50] Dongwon Lim et al. “Opinion Formation in the Digital Divide.” In: *Journal of Artificial Societies and Social Simulation* 17.1 (2014), pp. 1–13. URL: <https://ideas.repec.org/a/jas/jasssj/2013-11-3.html>.
- [51] Muzafer Sherif and Carl I Hovland. “Social judgment: Assimilation and contrast effects in communication and attitude change.” In: (1961).
- [52] Wako Yoshida, Ray J. Dolan, and Karl J. Friston. “Game Theory of Mind.” In: *PLOS Computational Biology* 4.12 (Dec. 2008), pp. 1–14. DOI: 10.1371/journal.pcbi.1000254. URL: <https://doi.org/10.1371/journal.pcbi.1000254>.
- [53] David V. Pynadath and Stacy C. Marsella. “PsychSim: Modeling Theory of Mind with Decision-Theoretic Agents.” In: *in Proceedings of the International Joint Conference on Artificial Intelligence*. Morgan Kaufman Publishers Inc, 2005, pp. 1181–1186.
- [54] Harmen De Weerd, Rineke Verbrugge, and Bart Verheij. *Theory of mind in the Mod game: An agent-based model of strategic reasoning*.
- [55] Stacy C Marsella, David V Pynadath, and Stephen J Read. “PsychSim: Agent-based modeling of social interactions and influence.” In: *Proceedings of the international conference on cognitive modeling*. Vol. 36. 2004, pp. 243–248.
- [56] Jonathan Y Ito, David V Pynadath, and Stacy C Marsella. “A decision-theoretic approach to evaluating posterior probabilities of mental models.” In: *AAAI-07 workshop on plan, activity, and intent recognition*. 2007.
- [57] Chris Lawrence Baker. “Bayesian theory of mind: Modeling human reasoning about beliefs, desires, goals, and social relations.” PhD thesis. Massachusetts Institute of Technology, 2012.
- [58] Chris L Baker and Joshua B Tenenbaum. “Modeling human plan recognition using Bayesian theory of mind.” In: *Plan, activity, and intent recognition: Theory and practice* 7 (2014), pp. 177–204.
- [59] David Waller and Rupert Younger. *The reputation game: The art of changing how people see you*. Simon and Schuster, 2017.
- [60] Grahame R Dowling. *Winning the reputation game: Creating stakeholder value and competitive advantage*. MIT Press, 2016.
- [61] John J Skowronski and Donal E Carlston. “Social judgment and social memory: The role of cue diagnosticity in negativity, positivity, and extremity biases.” In: *Journal of personality and social psychology* 52.4 (1987), p. 689.
- [62] Erdenebileg Batbaatar, Meijing Li, and Keun Ho Ryu. “Semantic-Emotion Neural Network for Emotion Recognition From Text.” In: *IEEE Access* 7 (2019), pp. 111866–111878. DOI: 10.1109/ACCESS.2019.2934529.
- [63] Peng Wu et al. “Social media opinion summarization using emotion cognition and convolutional neural networks.” In: *International Journal of Information Management* 51 (2020), p. 101978. ISSN: 0268-4012. DOI: <https://doi.org/10.1016/j.ijinfomgt.2019.07.004>. URL: <https://www.sciencedirect.com/science/article/pii/S0268401218313690>.