

Dynamics of deceptive interactions in social networks

Rafael A. Barrio^{1,2}, Tziipe Govezensky³, Robin Dunbar^{4,2}, Gerardo Iníguez ^{*5,2}, and
Kimmo Kaski^{2,4,6,7}

¹Instituto de Física, Universidad Nacional Autónoma de México, 01000 México D.F., Mexico

²Department of Computer Science, Aalto University School of Science, FI-00076 AALTO, Finland

³Instituto de Investigaciones Biomédicas, Universidad Nacional Autónoma de México, 04510 Mexico D.F.,
Mexico

⁴Department of Experimental Psychology, University of Oxford, OX1 3UD, United Kingdom

⁵Centro de Investigación y Docencia Económicas, Consejo Nacional de Ciencia y Tecnología, 01210 México
D.F., Mexico

⁶CABDyN Complexity Centre, Saïd Business School, University of Oxford, OX1 1HP, United Kingdom

⁷Center for Complex Network Research (CCNR), Department of Physics, Northeastern University, Boston MA
02115, USA

September 4, 2015

Abstract

In this paper we examine the role of lies in human social relations by implementing some salient characteristics of [deceptive](#) interactions into an opinion formation model, so as to describe the dynamical behaviour of a social network more realistically. In this model we take into account such basic properties of social networks as the dynamics of the intensity of interactions, the influence of public opinion, and the fact that in every human interaction it might be convenient to deceive or withhold information depending on the instantaneous situation of each individual in the network. We find that lies shape the topology of social networks, especially the formation of tightly linked, small communities with loose connections between them. We also find that agents with a larger proportion of deceptive interactions are the ones that connect communities of different opinion, and in this sense they have substantial centrality [in the network](#). We then discuss the consequences of these results for the social behaviour of humans and predict the changes that could arise due to a varying tolerance for lies in society.

keywords: social networks, anthropology, self-organised systems, mathematical modelling

1 Background

Deception, withholding information, making misleading statements, or blunt lies, are attitudes that most societies abhor, and parents, mentors and educators invest a great deal of effort in teaching that such behaviour is wrong and damages society [1, 2, 3, 4]. Yet it is also true that deception and lies are present

*gerardo.iniguez@cide.edu

in practically all human interactions and societies [5, 6, 7, 8]. This being so, we must conclude that there is a fundamental reason that prevents the social world from being totally honest.

Broadly speaking, trust-based exchange relationships play an important role in the emergence of cooperation and complex structure in many social, economic, and biological systems [9, 10, 11]. In human societies trust promotes people’s willingness to engage in reciprocity [12], while deception is likely to destroy the stability of such relationships by only favouring particular individuals [13, 14]. However, most research has been focused on how to detect and police deception [15, 16], rather than on the mechanisms regulating the appearance of lies and their implications for the structure of social networks.

Previously we have studied deception and its societal consequences by means of an agent-based opinion formation model [17] where the state of an agent i is described with two time-dependent variables, i.e. its true (but not public) opinion x_i and its public opinion y_i , in principle different from x_i . Their difference $d = |x_i - y_i|$ quantifies the lies that are told by agent i to its neighbours, which are linked by weighted links A_{ij} representing social interactions. Agents and links constitute a highly structured social network where opinion formation takes place. Both state variables evolve with a characteristic time scale dt , while link weights change on a different time scale D . In addition, the network structure co-evolves with the opinion dynamics via a network rewiring process with its own slower time scale, such that the weakest links are cut and the same number of new links are randomly created to conserve the average degree of the network. In the model, deception is defined as a partially truthful exchange of information between agents (that is, a linear combination of x_i and y_j) with the overall proportion of honesty in the system regulated by a single parameter. Thus lies may be considered as *pro-* or *anti-social* interactions if the information passed from agent i to agent j is proportional to y_j or $-y_j$, respectively. The selection of pro- or anti-social deception mimics the agent’s intention to be as similar or different as possible from its neighbour [18]. In this context, pro-social lies are those that benefit the recipient rather than the perpetrator, for example by continuing to reinforce the dyadic relationship between them. Common examples might be ‘liking’ something on someone’s social media page even though one does not really, or asserting that something is fine when in fact it is not.

This quite simple model already gives some hints about what the fundamental utility for lying might be. We discovered that, although anti-social lies destroy the connectivity of the social network, a certain frequency of pro-social deception actually enhances the centrality of liars (who serve as links between small communities of honest people). However, in this model individuals are assumed to pursue a fixed strategy: they are always honest individuals, pro-social liars or anti-social liars. In more realistic scenarios, of course, there are enormous fluctuations away from this simple fixed strategy set and individuals vary their behaviour between the three strategies according to circumstances, even though they may naturally tend towards one strategy most of the time. An important step in exploring the dynamics of social deception, then, is to develop a model that incorporates a significant amount of strategic flexibility at the individual level. Apart from adding more realism to the model, this has the important consequence of allowing individuals and populations to evolve towards a natural equilibrium, as individuals adjust their own behaviour in accordance with the cost and benefit regimes they encounter [13, 19, 20].

The fundamental question in modelling deception is: why do people lie? In human deception, the key issue must be related to the benefits and costs when deciding what information to pass on in an interaction with another person. From this point of view, lying is a decision-making problem with an optimal solution dependent on the gains and risks of lying [21, 22]. It is therefore important to include in the model some way of deciding the best possible answer in every instantaneous and directed dyadic interaction. In this paper we propose a more realistic model for the dynamics of deception including

these features. First we describe the model in general, including the dynamics of link weights and the decision-making process. Then we discuss the results of our numerical simulations and make concluding remarks.

2 Methods

Like in our earlier study [23], the basic dynamical equation for the opinion of an agent can be written as,

$$\frac{\partial x_i}{\partial t} = f_s(i)|x_i| + \alpha_i f_l(i), \quad (1)$$

where the state variable x_i is bounded by $[-1,1]$ and represents the instantaneous opinion of agent i , such that -1 corresponds to total disagreement and +1 to total agreement with a given topic of discussion. The first term on the right-hand side describes an exchange of information between a pair of agents through discussion, i.e. the interaction is short range. The second term stands for the influence of the overall opinion in the network on agent i , and hence the interaction is long range. Both terms evolve with a time scale dt called ‘transaction time’. The parameter α_i is a random bounded variable that represents the attitude of agent i to the overall opinion $f_l(i)$, being near -1 if the agent is inclined to go against the crowd and near +1 otherwise.

In accord with our earlier model of deceptive interactions between agents [17], we define a second state variable y_i corresponding to other agents’ public perception of the true but private x_i , from which y_i may differ in value if agent i is not totally honest. The difference $d = |x_i - y_i|$ stands for the amount of dishonesty or the size of the lie. Hence the overall opinion $f_l(i)$ should be formed with the publicly available information (through social meetings, rumours, and news in the media) represented here by the time-dependent variable y_i ,

$$f_l(i) = \sum_{\ell=2}^{\ell_{max}} \frac{1}{\ell} \sum_{j \in m_\ell(i)} y_j(t), \quad (2)$$

where the second sum is over the set $m_\ell(i)$ of all agents j separated from agent i by a shortest-path length $\ell = 2, \dots, \ell_{max}$. We assume that the influence of an agent decays with the distance ℓ , i.e. the smallest number of links needed to reach j from i in the network. Without loss of generality we also consider a $1/\ell$ -dependence.

In Eq. (1) the short-range term is the direct interaction between agents with $\ell = 1$,

$$f_s(i) = \sum_{j \in m_1(i)} w_{ij}(t), \quad (3)$$

where $w_{ij}(t)$ is the instantaneous information that agent j passes to i (see Eq. 5). Observe that in general the matrix \mathbf{w} is not symmetric; that is, the information that agent i gives to j , $w_{ji} \neq w_{ij}$. Therefore, the sum of the elements of a row in \mathbf{w} gives $f_s(i)$, while the sum of the elements of each column in \mathbf{w} is proportional to the average apparent opinion the agent holds in the network,

$$y_i(t) = \frac{1}{k_i} \sum_{j \in m_1(i)} w_{ji}(t), \quad (4)$$

where $k_i = |m_1(i)|$ is the degree of agent i . Explicitly, the public opinion y_i is the average of the instantaneous information w_{ji} received by all neighbours j , and is thus bounded between -1 and +1.

Finally, we define the instantaneous information w_{ij} as,

$$w_{ij}(t) = \phi_0(j, i, t), \quad (5)$$

where the optimal opinion ϕ_0 that agent j shares with agent i (i.e. between truth and pro- or anti-social lies) is the result of an individual decision-making process, as explained in section 2.2.

The nature of direct transactions is illustrated in Fig 1. For example, the terms $w_{ij}|x_i|$ in Eq. (1) imply that if w_{ij} has the same sign as x_i , agent i will reinforce its position and get closer to the extreme opinion $\text{sign}(x_i)$. Next we introduce the dynamical processes involved in our deception model, as described in the sections below.

2.1 Dynamics of link weights and rewiring scheme

In social networks individuals are connected by weighted links that vary over time in the course of their dyadic interactions and decision making. We assume that ‘bad’ decisions (not necessarily due to lies) are punished by weakening the link weight A_{ij} between agents i and j . This can be incorporated into the model by introducing a simple dynamics for link weights,

$$\frac{\partial A_{ij}}{\partial t} = D T_{ij}(t), \quad (6)$$

where D sets the time scale of change and T_{ij} is a function of the four site variables associated with a link, namely (x_i, y_i) and (x_j, y_j) . Since A_{ij} depends on two agents, we choose the following symmetric form,

$$T_{ij}(t) = \{ |[x_i(t) + y_j(t)] + [y_i(t) + x_j(t)] - 1 \} - P_{ij}(t), \quad (7)$$

where the first square bracket represents similarity between agents according to the information agent i has at its disposal, the second bracket is the corresponding term for agent j , and $P_{ij}(t)$ is the instant punishment for lying. Observe that the term in $\{ \dots \}$ varies between 3 and -1, such that links with $T_{ij} < 0$ are at risk of being cut, as A_{ij} approaches zero. The matrix \mathbf{T} should be symmetric under exchange between i and j , in which case the punishment the society imposes on liars reads as follows,

$$P_{ij}(t) = (1 - e) (|w_{ji}(t) - x_i(t)| + |w_{ij}(t) - x_j(t)|), \quad (8)$$

where e is a parameter that measures the tolerance of society against lies, being 0 if it is intolerant, and 1 if it does not punish liars. Thus, the punishment P_{ij} is proportional to the difference between the true opinion of an agent and the instantaneous information it shares with its neighbour.

In a real social network, its topology may coevolve with the dyadic interactions between individuals [24]. Thus we introduce a rewiring scheme with dynamics dependent on link weights. We proceed by cutting links with negative weight ($A_{ij} \leq 0 \rightarrow A_{ij} \equiv 0$) and immediately creating a new link (with initial weight 1) to substitute the cut link, in order to maintain the average degree of the network constant. This sets the time scale for network rewiring larger than Δt and variable, unlike in our former model in which rewiring is performed at fixed intervals [17, 23].

The creation of new links is performed as follows. First we identify the two agents involved in the cut link and choose the one with the least number of neighbours (i.e. the most isolated agent); then we look at the second neighbours of this individual and create a link with the second neighbour (friend of a [possibly lost] friend) that has the lowest degree. This bias for favouring agents with only a few links

assumes that such agents are more keen on making new friends. If there are no second neighbours, then we create a link with one of the agents with the lowest degree in the whole network. As a further remark, we note that the instantaneous information w_{ij} is not necessarily the same for everyone all the time (see Eq. 5), the net effect of which is that the rewiring time is variable and controlled for each link by the slope T_{ij} .

2.2 Decision-making process

In the Background we state that a key issue for human deceptive interactions is related to the benefit and cost of lying, which an individual needs to evaluate in order to pass optimal information to others. This means that in each transaction, acting agent j has to make a decision whether to lie or not **to neighbour i , by** finding the extremal values of a utility function R that includes all gains and costs of being deceitful or truthful,

$$R(\phi) = H - L(\phi) = G_H - C_H - G_L(\phi) + C_L(\phi), \quad (9)$$

where ϕ is the opinion agent j decides to share **with i ,** either the truth ($\phi = x_j$) or a lie ($\phi \neq x_j$). Note that the gain G_H and the cost C_H of being honest do not depend on ϕ , while the gain G_L and the cost C_L of being dishonest depend on the particular **opinion ϕ** that agent j wishes to share. Then, the optimal opinion $\phi_0(j, i, t) = w_{ij}(t)$ is a stationary point of R **(either a maximum or minimum)** in the permissible interval $[-1, 1]$, implicitly defined by,

$$\left. \frac{\partial R}{\partial \phi} \right|_{\phi_0} = 0. \quad (10)$$

Under **these** conditions, the decision-making process for agent j is as follows. When interacting with neighbour i , agent j finds the optimal opinion ϕ_0 by solving Eq. (10). If $R(\phi_0) > 0$, then agent j ignores ϕ_0 and shares its true opinion **(i.e. $w_{ij} = x_j$)**, since being truthful is a ‘better’ decision than not being truthful. Otherwise, agent j shares the optimal opinion ϕ_0 . Note that in general ϕ_0 stands for a lie, except for the case when $\phi_0 = x_j$. This particular case could be interpreted as a situation where an agent (that has initially decided to lie) finds that the optimal decision is to be honest.

For the decision-making process to be complete, we need to find concrete expressions for the gains and costs in Eq. (9), based on the available sociological knowledge about interactions between individuals. The gain for being honest is considered to be ‘prestige’ or ‘reputation’ [8], which in our context is measured by the degree k_j . This is based on a previously studied sociological assumption [25, 26], namely that the more connected you are, the more prestige you have, **which means that you are considered trustworthy**. Therefore, we write the gain as,

$$G_H = \frac{k_j - \min(\{k\})}{\max(\{k\}) - \min(\{k\})}, \quad (11)$$

where we have normalised the degree to compare agents within and between communities.

The risk associated with being honest is proportional to the apparent disparity of opinion, as this distance increases antagonism between agents. In other words, people tend to use small, ‘white’ lies to protect social links rather than put them at risk, since the difference in opinion corresponding to complete honesty may create tension in the relationship [27]. Then we write,

$$C_H = \frac{|x_j - y_i|}{2}, \quad (12)$$

which is normalised to make the gain and cost terms comparable.

If the main aim of an agent's deception is to avoid rejection by strengthening its social links, then everyday experience suggests that the gain due to lying has two components. First, the liar benefits by not 'losing face', that is, by minimising the distance $\phi - y_j$ between its lie and public opinion so that the lie ϕ is not discovered easily. Second, the agent j gains by mimicking the response w_{ji} that agent i is giving back, i.e. by pretending to be more similar to its peers than it is in reality. In this case we write,

$$G_L(\phi) = \frac{1}{2} \left[1 - \frac{|\phi - y_j|}{2} \right] + \frac{1}{2} \left[1 - \frac{|\phi - w_{ji}|}{2} \right]. \quad (13)$$

The risk of lying is also two-fold: agent j could pass information that is similar to its true opinion (x_j) and risk a large mismatch of opinions. The bigger this difference, the higher the penalty (or cost) that the liar will incur from being found out [28]. Simultaneously, the agent could try to mimic an agreement with the public opinion of agent i , thereby risking a mismatch if agent i is deceptive: the bigger the difference between the lie and public opinion, the bigger the cost the liar bears from being found out. This being so, the risk is the product of the two possibilities,

$$C_L(\phi) = \frac{\beta}{4} |\phi - x_j| \times |\phi - y_i|. \quad (14)$$

We have normalised Eqs. (11)-(13) such that all of them vary between zero and one. The coefficient β in Eq. (14) is a quantity that controls the relative weight of the cost of lying, which could depend on other social and cultural properties. We have examined the behaviour of the utility function R and determined that $\beta = 4$ balances the gains and costs between lying and being honest (see Supplementary Information).

To summarise, the dynamics of our model is highly non-linear: the elements of the adjacency matrix \mathbf{A} depend on the vector \mathbf{y} (the other agents' perception of an agent's true opinion), which in turn is calculated every interaction using \mathbf{w} . The matrix \mathbf{w} , in turn, is the instantaneous flow of information through each link, resulting from an agent's decision about the optimal information to pass on (ϕ_0). Our new approach of casting transactions between agents as an optimised decision-making process constitutes a major difference from our earlier model [17]. The benefit is that we now avoid predefining individuals as either pro- or anti-social liars. Nevertheless, we can still classify the lie ϕ_0 in a binary way by comparing the distances $|\phi_0 - y_i|$ and $|\phi_0 + y_i|$, being pro-social if the former is smaller and anti-social otherwise. Then, the threshold that classifies ϕ_0 as a pro- or anti-social lie is 0, the midpoint between $\pm y_i$. We emphasise that $\phi_0 = \phi_0(j, i, t)$ is a function of j , i , and t , obtained by finding a stationary point of $H - L$ as given in Eq. (9). Allowing deception to vary in this way is more realistic than our previous approach of having fixed, predefined phenotypes that do not vary in their behaviour: everyday experience tells us that people do not say the same thing to everybody all the time.

3 Results

Using the model described above, we perform extensive numerical simulations for networks of different sizes ($N = 100, 200$ and 500), starting with random initial conditions for the state variables x_i , y_i and attitude parameters α_i , and by following a simple Euler integration scheme for a sufficient number of time steps to obtain extreme opinions ($x_i = \pm 1$) for all agents. In some cases, we find a small number of agents remaining locked in an undecided state $|x_i| < 1$. This number depends on the values of the parameters e (tolerance of society against lies) and D (time scale for the growth of link weights), the

only two parameters whose variation we consider. We can follow the time history of the process and monitor the state variables and also the amount of instantaneous lying for all agents in the system. We may also distinguish anti-social lies from pro-social ones by monitoring the optimal opinion $\phi_0(j, i, t)$. If this quantity is nearer to y_i , then we consider it as a pro-social lie, and if it is nearer to $-y_i$ we take it as an anti-social lie [17]. **As simulation results are qualitatively unaffected by network size, from now on we only consider $N = 100$.**

In Fig. 2 we show typical realisations of the dynamics while keeping $D = 3$ for the two extreme values of the parameter e . Observe that honest agents with similar opinion form small clusters (i.e. **the network is assortative with respect to x_i [29]**), but there is also a large number of liars that serve as weak links between these tightly bonded communities and can dwell within very small communities. The effect of increasing social tolerance ($e = 1$) is small, but surprisingly the relative number of liars is smaller when there is no punishment for lying. This result is in qualitative agreement with empirical observations made in schools with and without punishment [30], where the authors report that “a punitive environment not only fosters increased dishonesty but also children’s abilities to lie in order to conceal their transgressions”.

In Fig. 3 we show the proportion of pro- and anti-social lies in the instantaneous response of each agent to all of its neighbours, for the case $e = 1$ of Fig. 2. Observe that many agents lie pro-socially all the time and to all their neighbours. In contrast, there are very few anti-social lies and they do not persist, but instead disappear completely for some agents while for others they become intermittent. If we reduce the social tolerance for lying, anti-social behaviour disappears completely. Notice also that, despite using ‘ideal’ conditions for the appearance of big lies ($e = 1$), there are always some agents that behave in a totally honest manner all the time.

To analyse these results further, we find it convenient to quantify separately various groups of agents. We focus our attention on those agents who are totally honest throughout the time line of the model, those who tell only pro-social lies, those who tell anti- or pro-social lies indiscriminately, and those who only lie anti-socially. Notice that for this kind of analysis to succeed, we need many realisations to obtain an average value for each case. Also, we need to look at probability distributions rather than well-defined categories, since the freedom to decide produces strategy changes in all agents. The model output suggests that agents who only lie anti-socially are very few in number, as can be seen from Fig. 4 where we show the probability distribution of the proportion of anti-social lies for the case of zero tolerance ($e = 0$) and no punishment ($e = 1$). Note that social tolerance to lying has very little effect on the appearance of anti-social lies, and that most of the agents turn out to tell very few lies.

In Fig. 5 we show the probability distribution of the proportion of lies per dyadic interaction, r , for agents who lie indiscriminately (Anti-Pro case) and for those who tell only pro-social lies (Pro case), for the two extreme values of the social parameter e . Explicitly, r is the fraction of the total number of interactions that are lies. These results suggest that nearly 50% of the agents lie pro-socially a small amount of time ($< 10\%$ of the total time). However, there are always a few agents who lie more frequently: 20% of agents lie all the time, regardless of the level of social tolerance. This result implies that it is disadvantageous to lie all the time. Fig. 5 also suggests that the Pro- and Anti-Pro strategies are qualitatively quite similar, in the sense that many agents lie sporadically (small r) and a few agents ($\sim 20\%$) lie most of the time. Obviously, the relative numbers depend on the social tolerance parameter as well. An interesting observation here is that the lack of punishment for lies makes very small difference between the appearance of anti- or pro-social lies.

Ref. [5] reports a statistical study of the number and nature of lies told by a group of 144 individuals, the results of which are summarised in Table 2 therein for comparison with the results of our model.

For instance, the percentage of honest people (those who never tell lies) is 1.3% for individuals recruited from a college population, and 8.6% for individuals from a local community. Our results show that 2.7% of agents are honest if there is punishment, and 3.5% if there is not. Furthermore, the mean number of lies per day measured was roughly 2, and the mean number of social interactions 6, of which only 61% were dyadic interactions [5]. This means that 50% of the dyadic interactions were lies. The area under the curve in Fig. 5 (without distinguishing between pro- and anti-social lies) gives about 53%, thus roughly agreeing with the experimental findings. In addition, we predict that the number of lies per social interaction (obtained by calculating the mean value for the amount of dishonesty or the size of the lie, d) is 0.38, in close agreement with the value 0.31 ± 0.11 reported in Table 2 of the experimental study [5].

We now investigate the social advantages of lying. This is done by examining network measures such as the weighted clustering coefficient (WCC) and betweenness centrality (BC). In a weighted network or graph, WCC is the geometric average of link weights in a subgraph, where weights are normalised by the maximum weight in the network [31, 32]. BC is the sum of the fraction of all-pairs shortest paths that pass through a node [33]. With these measures we see that liars serve as bridges between communities (Fig. 2); hence they sacrifice their WCC (belonging to highly connected clusters) in order to improve their BC (easy communication with all members of the network).

It is possible for a deceptive agent to increase its clustering coefficient provided it tells small lies ($d \leq 0.1$), irrespective of whether these are anti- or pro-social, even in the face of social punishment. In Fig. 6, we show WCC averaged over 300 runs of the model for $N = 100$ agents. The conclusion is that from the perspective of clustering, there is no benefit to lying unless the lie is small. We also see that a society with total tolerance to lying does not provide liars with much advantage. However, when there is punishment, the agents who lie pro- and anti-socially have an advantage over totally honest agents, provided the number of lies is small. This can be seen in Fig. 7, where we show the WCC probability distribution for selected values of the proportion of lies per dyadic interaction r .

In conclusion, the real advantage to being a liar is that BC increases for pro-social liars, provided they tell small lies. This could be interpreted as a mechanism people use to fit into a social network better. In Fig 8 we show the BC median taken over 300 runs of the model as a function of the size of lies, and for different groups of agents. We present the median, instead of the average, because for the form of distribution functions we have here, the median is a more robust quantity. In Fig 8(a) we show the case of zero tolerance ($e = 0$), where only pro-social liars have an advantage over honest agents, provided their lies are very small. In Fig 8(b) we show the same for tolerance $e = 1$.

4 Discussion and conclusions

Our model for the dynamics of lying in a social network of humans incorporates the relevant fact that the individual act of lying corresponds to a flexible, personal and instantaneous decision. Hence we have mapped this action to a decision-making problem for individuals in society so that they can adjust their behaviour to the situations they face. In contrast to our earlier model [17], where agents have fixed behavioural strategies, the present model is more realistic as the information an agent passes on (either as truth, or as a pro- or anti-social lie) is a function of the circumstances the agent encounters. In effect, we assume that agents learn and adjust their behaviour in the light of experience. In this respect, the present model lies at the opposite extreme from our previous model in that it does not assume that agents have inherited psychological predispositions to behave in a particular way. In all likelihood, of

course, the real world probably lies somewhere between these two extremes. The fact that the findings from the two models are in broad agreement is therefore comforting, in that it suggests that irrespective of where reality actually lies our findings will be robust.

The model studied here does not have a network rewiring time scale that is proportional to the fundamental transaction time scale dt . Nevertheless, the rewiring time scale can still be tuned by using one of two parameters: D (the time scale for the growth of link weights) or e (the tolerance of society against lies). In addition, we see that as the tolerance parameter e increases, society is more tolerant of lies and the time at which bonds are deleted increases, thus making the process slower. Furthermore, in Fig. 2 we see that communities are much better defined when $e = 0$; as a result intolerance to lies and a potential for high punishment shortens the mean life of liars, segregating the network into communities with strong internal links.

In all our simulations we find that the number of anti-social lies diminishes, while pro-social lies persist to considerable numbers throughout the dynamical progress of the system. Here we see that the social tolerance parameter e has little effect on the proportion of anti-social lies, although it regulates the total number of lies. Most of the agents lie sporadically and only very few seem to lie all the time. This indicates that ‘true’ liars are very rare in society, [although](#) they are nonetheless very special since they have large [BC](#). We also find that liars who tell small lies ($d < 0.1$) have larger [WCC](#). In addition, we observe that the dynamics favours the formation of cliques of purely honest agents, and that liars are usually found to be in the perimeter of cliques and connected by weak links.

We also show that in general being honest pays off, but in some circumstances liars acquire an advantage over honest agents. For instance, agents who occasionally tell small lies have larger [WCC](#) and [BC](#) than honest agents (see Fig.7 for $k < 0.2$). Moreover, an agent who tells a fair number of medium-sized lies ($d = 1$) could attain a larger [BC](#) than when it chooses to be honest.

In summary, it is interesting to note that for small lies, all liars are better off than honest agents. Even more interesting is the fact that there is a maximal advantage for people who tell sizeable anti-social lies. In short, anti-social lies yield considerable benefits for liars in appropriate circumstances. We know that anti-social lies normally destroy the social network when they are widely distributed throughout society [17]. However, our findings suggest that, in certain specific circumstances, they could have the opposite effect and make the network more robust. This implies that we need to identify the conditions [under which such a situation arises](#), by examining the local circumstances of those agents who present this peculiar property. Paradoxically, it might then be possible to increase the information flow in the network by adding appropriate motifs that allow agents to have both high [BC](#) and [WCC](#).

Competing interests: We have no competing interests.

Authors’ contributions: All authors conceived, designed and coordinated the study. RAB and GI developed and analysed the model. RAB and TG carried out the numerical and statistical analyses. All authors helped draft the manuscript and gave final approval for publication.

Funding: RAB acknowledges support from Conacyt project No. 179616. RD’s research is funded by a European Research Council Advanced grant. GI and KK acknowledge support from EU’s FP7 FET Open STREP Project ICTeCollective No. 238597, and GI from the Academy of Finland.

References

- [1] M. Lachmann, S. Számádó, and C. T. Bergstrom. Cost and conflict in animal signals and human language. *Proc. Natl. Acad. Sci. USA*, 98:13189–13194, 2001. DOI: 10.1073/pnas.231216498.

- [2] C. S. Wang, A. D. Galinsky, and J. K. Murnighan. Bad drives psychological reactions, but good propels behavior responses to honesty and deception. *Psychol. Sci.*, 20:634–644, 2009. DOI: 10.1111/j.1467-9280.2009.02344.x.
- [3] F. Xu, X. Bao, G. Fu, V. Talwar, and K. Lee. Lying and truth-telling in children: From concept to action. *Child Dev.*, 81:581–596, 2010. DOI: 10.1111/j.1467-8624.2009.01417.x.
- [4] M. Popliger, V. Talwar, and A. Crossman. Predictors of children’s prosocial lie-telling: Motivation, socialization variables, and moral understanding. *J. Exp. Child Psychol.*, 110:373–392, 2011. DOI: 10.1016/j.jecp.2011.05.003.
- [5] B. M. DePaulo, D. A. Kashy, S. E. Kirkendol, M. M. Wyer, and J. A. Epstein. Lying in everyday life. *J. Pers. Soc. Psychol.*, 70:979–995, 1996. DOI: 10.1037/0022-3514.70.5.979.
- [6] J. T. Hancock, J. Thom-Santelli, and T. Ritchie. Deception and design: The impact of communication technology on lying behavior. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 129–134. ACM, New York, 2004. DOI: 10.1145/985692.985709.
- [7] K. B. Serota, T. R. Levine, and F. J. Boster. The prevalence of lying in america: Three studies of self-reported lies. *Hum. Commun. Res.*, 36:2–25, 2010. DOI: 10.1111/j.1468-2958.2009.01366.x.
- [8] M. S. Panasiti, E. F. Pavone, A. Merla, and S. M. Aglioti. Situational and dispositional determinants of intentional deceiving. *PloS ONE*, 6:e19465, 2011. DOI: 10.1371/journal.pone.0019465.
- [9] M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity. *Nature*, 437:1291–1298, 2005. DOI: 10.1038/nature04131.
- [10] G. Roberts and J. S. Renwick. The development of cooperative relationships: An experiment. *Proc. R. Soc. B*, 270:2279, 2003. DOI: 10.1098/rspb.2003.2491.
- [11] W.-S. Kim. Effects of a trust mechanism on complex adaptive supply networks: An agent-based social simulation study. *J. Artif. Soc. Soc. Simul.*, 12:4, 2009.
- [12] S. V. Paunonen. You are honest, therefore I like you and find you attractive. *J. Res. Pers.*, 40:237–249, 2006. DOI: 10.1016/j.jrp.2004.12.003.
- [13] W. A. Searcy and S. Nowicki. *The Evolution of Animal Communication: Reliability and Deception in Signaling Systems*. Princeton University Press, Princeton, 2005.
- [14] A. Sutcliffe and D. Wang. Computational modelling of trust and social relationships. *J. Artif. Soc. Soc. Simul.*, 15:3, 2012.
- [15] D. Nyberg. *The Varnished Truth: Truth Telling and Deceiving in Ordinary Life*. University of Chicago Press, Chicago, 1993.
- [16] K. H. Craik. *Reputation: A Network Interpretation*. Oxford University Press, USA, 2008.
- [17] G. Iñiguez, T. Govezensky, R. Dunbar, K. Kaski, and R. A. Barrio. Effects of deception in social networks. *Proc. R. Soc. B*, 281:20141195, 2014. DOI: 10.1098/rspb.2014.1195.
- [18] F. Gino and L. Pierce. Dishonesty in the name of equity. *Psychol. Sci.*, 20:1153–1160, 2009. DOI: 10.1111/j.1467-9280.2009.02421.x.

- [19] S. Számadó. Cheating as a mixed strategy in a simple model of aggressive communication. *Anim. Behav.*, 59:221–230, 2000. DOI: 10.1006/anbe.1999.1293.
- [20] J. T. Rowell, S. P. Ellner, and H. K. Reeve. Why animals lie: How dishonesty and belief can coexist in a signaling system. *Am. Nat.*, 168:E180–E204, 2006. DOI: 10.1086/508809.
- [21] U. Gneezy. Deception: The role of consequences. *Am. Econ. Rev.*, 95:384–394, 2005.
- [22] S. Mitri, D. Floreano, and L. Keller. The evolution of information suppression in communicating robots with conflicting interests. *Proc. Natl. Acad. Sci. USA*, 106:15786–15790, 2009. DOI: 10.1073/pnas.0903152106.
- [23] G. Iñiguez, J. Kertész, K. K. Kaski, and R. A. Barrio. Opinion and community formation in coevolving networks. *Phys. Rev. E*, 80:066119, 2009. DOI: 10.1103/PhysRevE.80.066119.
- [24] T. Gross and H. Sayama. *Adaptive Networks*. Springer-Verlag, Berlin, 2009.
- [25] J. Henrich and F. J. Gil-White. The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evol. Hum. Behav.*, 22:165–196, 2001. DOI: 10.1016/S1090-5138(00)00071-4.
- [26] T. Dávid-Barrett and R. I. M. Dunbar. Social elites can emerge naturally when interaction in networks is restricted. *Behav. Ecol.*, 25:58–68, 2014. DOI: 10.1093/beheco/art085.
- [27] R. I. M. Dunbar and A. Machin. Sex differences in relationship conflict and reconciliation. *J. Evol. Psych.*, 12:109–133, 2014. DOI: 10.1556/JEP-D-13-00024.
- [28] J. T. Hancock, C. Toma, and N. Ellison. The truth about lying in online dating profiles. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 449–452. ACM, New York, 2007. DOI: 10.1145/1240624.1240697.
- [29] M. E. J. Newman. The structure and function of complex networks. *SIAM Rev.*, 45:167–256, 2003. DOI: 10.1137/S003614450342480.
- [30] Victoria Talwar and Kang Lee. A punitive environment fosters children’s dishonesty: A natural experiment. *Child Dev.*, 82:1751–1758, 2011. DOI: 10.1111/j.1467-8624.2011.01663.x.
- [31] J.-P. Onnela, J. Saramäki, J. Kertész, and K. Kaski. Intensity and coherence of motifs in weighted complex networks. *Phys. Rev. E*, 71:065103(R), 2005. DOI: 10.1103/PhysRevE.71.065103.
- [32] J. Saramäki, M. Kivelä, J.-P. Onnela, K. Kaski, and J. Kertész. Generalizations of the clustering coefficient to weighted complex networks. *Phys. Rev. E*, 75:027105, 2007. DOI: 10.1103/PhysRevE.75.027105.
- [33] U. Brandes. On variants of shortest-path betweenness centrality and their generic computation. *Soc. Networks*, 30:136–145, 2008. DOI: 10.1016/j.socnet.2007.11.001.

Short title for page headings: Dynamics of deception in social networks




Figure 1: (Online version in colour.) Diagram of the transaction dynamics of the model. Agent i perceives the opinion w_{ij} from neighbour j and changes its own true opinion according to Eq. (1).






Figure 2: Examples of networks with $N = 100$ agents after $t_f = 600$ time iterations of the dynamics, for $e = 0, 1$ and $D = 3$. The color code for the agents is: honest and $x_i = 1$ (white), honest and $x_i = -1$ (black), and liars (light blue or dark green) for $x_i = 1$ or $x_i = -1$, respectively. The blue/red circles are undecided agents with $0 < x_i < 1$ and $-1 < x_i < 0$, respectively. **Observe that the network is assortative with respect to the opinion x_i .** The width of the links indicates their **weight**.



prosocial lies



antisocial lies

Figure 3: Time history of the proportion of pro- and anti-social lies in the instantaneous response of each agent to all of its neighbours, for $e = 1$. Dark/light dots correspond to high/low proportions, respectively. Observe that pro-social lies are abundant and persistent. On the contrary, anti-social lies are very few and intermittent.




Figure 4: (Online version in colour.) Probability distribution of the proportion of anti-social lies, obtained from 300 runs of networks with $N = 100$ agents. The two extreme cases of social tolerance to lies ($e = 0, 1$) are shown.




Figure 5: (Online version in colour.) Probability distribution of the proportion of lies per dyadic interaction, obtained from 300 runs of networks with $N = 100$ agents. Dotted lines include pro-social lies only, while continuous lines indicate both anti- and pro-social lies. Orange/blue (light/dark gray) lines correspond to $e = 1, 0$, respectively. The normalisation factor is the total time span of each dynamics.




Figure 6: (Online version in colour.) Average weighted clustering coefficient (WCC) as a function of the size of the lie d for: (a) zero tolerance ($e = 0$), and (b) no punishment ($e = 1$). Triangles correspond to pro-social lies only, and dots to both anti- and pro-social lies. The WCC for totally honest networks is indicated by a large dot at $d = 0$.




Figure 7: (Online version in colour.) WCC probability distribution for several values of r in a society with zero tolerance ($e = 0$), for the: (a) Anti-Pro and (b) Pro cases. Notice the advantage over honest people that agents who lie indiscriminately gain, in contrast to those who only lie pro-socially.




Figure 8: (Online version in colour.) Median of betweenness centrality (BC), obtained from 300 runs in networks with $N = 100$ agents. Symbols are as in Fig. 6. Notice the advantage over honest individuals for agents who lie indiscriminately, in contrast to those who only lie pro-socially, for most lie sizes d .