






<https://doi.org/10.1038/s42005-021-00579-3>

OPEN

The fragility of opinion formation in a complex world

Matúš Medo ^{1,2,3✉}, Manuel S. Mariani ^{1,4} & Linyuan Lü ^{1,5,6}

How does the complexity of the world around us affect the reliability of our opinions? Motivated by this question, we quantitatively study an opinion formation mechanism whereby an uninformed observer gradually forms opinions about a world composed of subjects interrelated by a signed network of mutual trust and distrust. We show numerically and analytically that the observer's resulting opinions are highly inconsistent (they tend to be independent of the observer's initial opinions) and unstable (they exhibit wide stochastic variations). Opinion inconsistency and instability increase with the world's complexity, intended as the number of subjects and their interactions. This increase can be prevented by suitably expanding the observer's initial amount of information. Our findings imply that an individual who initially trusts a few credible information sources may end up trusting the deceptive ones even if only a small number of trust relations exist between the credible and deceptive sources.

¹Institute of Fundamental and Frontier Sciences, University of Electronic Science and Technology of China, Chengdu, People's Republic of China.

²Department of Radiation Oncology, Inselspital, University Hospital of Bern, and University of Bern, Bern, Switzerland. ³Department of Physics, University of Fribourg, Fribourg, Switzerland. ⁴URPP Social Networks, Universität Zürich, Zürich, Switzerland. ⁵Yangtze Delta Region Institute (Huzhou), University of Electronic Science and Technology of China, Huzhou, People's Republic of China. ⁶Beijing Computational Science Research Center, Beijing, People's Republic of China. ✉email: matus.medo@unifr.ch

Identifying potential mechanisms behind the formation of opinions in society is vital to understand how polarization emerges in society¹, how misinformation spreads and can be prevented², and how science can be effectively communicated to the public³. Despite recent advances in how opinions propagate in social networks¹, how artificial agents promote low-credibility content in social media^{4,5}, and how rapidly misinformation spreads compared to reliable content^{2,6}, misinformation still thrives in our society. This is well-exemplified by the recent growth of anti-vaccination views and the related existence of anti-vaccination clusters in online social networks⁷. The popularity of unreliable opinions—which is especially dangerous during global emergencies, such as the recent COVID-19 pandemic⁸—calls for a deeper investigation of the possible drivers behind the process whereby individuals form opinions in a society.

Models of opinion formation in social systems have been studied in statistical physics^{9–11}, sociology^{12–16}, and game theory¹⁷. They focus on opinion propagation on a social network of influence and allow us to model phenomena such as the emergence of consensus in the population, for example. Most existing models, however, neglect the potential interconnectedness among topics, which is instead a key aspect of modern social and information environments^{18,19}. These models view opinions of an individual on different topics as results of independent realizations of opinion formation processes, and they neglect direct connections among opinions on distinct topics.

Departing from existing models, we consider a situation where an individual observer forms opinions about a set of subjects interrelated by positive and negative links. The subjects on which the opinions are formed can represent governments, politicians, news media, or other individuals. Connections between subjects can take various forms²⁰. One credible news source can, for example, tend to report in agreement with another credible source (a positive link) whereas a misinformation source can tend to disagree with both credible information sources (negative links). Such systems with signed relations are described by Heider's social balance theory^{21–23} as documented using data on armed conflicts among countries^{24,25} and large-scale social media^{26–28}, for example. Of particular interest is the tendency for the subjects to form two opposing camps^{24,29,30}, although these macroscopic structures are generally imperfect. Two countries, for example, can belong to the same alliance whose members generally have positive relations, yet their mutual relation can be negative due to historic or economic reasons. In science, there are multiple Nobel prize recipients who endorsed unscientific theories³¹.

We study how an observer equipped with prior information about some subjects can use noisy signed relations between subjects divided into two camps to form opinions on the remaining subjects. Forming a reliable opinion in a complex situation

requires effortful reasoning. However, psychological research indicates that humans tend to be rather driven by simple heuristics when forming opinions about complex topics, sometimes reaching opinions that violate basic logic rules^{32,33}. The limitations of our cognition have important consequences. For example, the susceptibility to partisan fake news was recently found to be driven more by “lazy reasoning” than by partisan bias³⁴. For this reason, we focus here on local rules (heuristics) of opinion formation using a signed network of relations.

We find that even a small fraction of misleading links in the relation network among the subjects (e.g., a link of mutual trust between a scientific and low-credibility information source) leads to the resulting opinions that are both inconsistent with the observer's seed opinions and vary significantly between model realizations. We determine analytically the relation between average opinion consistency and the world complexity, represented by the number of subjects, which demonstrates that opinion inconsistency grows as the world complexity increases. This increase can be prevented by suitably increasing the observer's initial number of independent opinions. Although opinion consistency depends on network topology and can be improved by considering a more sophisticated local opinion formation mechanism, our main conclusions are robust to variations of the network topology and the opinion formation mechanism.

Our findings point that even in the absence of social influence, opinion formation in a world composed of many interrelated subjects is inherently fragile. Since subjects may represent co-existing scientific or low-credibility information sources, our model presents a contributing mechanism for how misinformation sources may gain their audience. Our work paves the way to studying strategies to increase the reliability of opinion formation.

Results

Opinion formation model. We consider an individual observer who gradually develops opinions on a world composed of N interrelated subjects (Fig. 1a). The number of subjects represents the complexity of the world. Each opinion is for simplicity assumed to take one of three possible states: no opinion, a positive opinion (trust), or a negative opinion (distrust). The observer's opinions can be formally represented by an N -dimensional opinion vector \mathbf{o} whose element o_i represents the opinion on subject i ; $o_i \in \{-1, 0, 1\}$ corresponds to a negative opinion, no opinion, and a positive opinion, respectively. The subjects form a signed undirected network of relations. These relations are represented by a symmetric $N \times N$ relation matrix whose element R_{ij} represents the trust relation between subjects i and j ; $R_{ij} \in \{-1, 0, 1\}$ corresponds to a negative relation, no relation, and a positive relation, respectively. We emphasize the main difference between this setting and traditional opinion formation models based on propagation on networks of social




Fig. 1 The opinion formation model. **a** An observer faces a set of subjects that are interconnected by mutual relations of trust (solid green lines) or distrust (dotted red lines). Starting from a small set of seed opinions (here one positive seed opinion marked with + in a black circle) and a world of unknown subjects (gray circles), the observer gradually forms opinions on all subjects. **b** The formed opinion is determined as a product of the opinion on a chosen source subject (one of the seed subjects or any other subject on which an opinion has been already made) and the sign of the relation between the source subject and the target subject. A positive opinion is formed when the source opinion and the relation are both positive or both negative; a negative opinion is formed otherwise.

influence^{9,10,13,14}: In existing models, simulating the opinion formation on N subjects would require running N independent realizations of the opinion formation process, which would miss the interconnectedness among subjects; by contrast, in the proposed approach, the interconnectedness among subjects is naturally encoded in the relation matrix \mathbf{R} .

The observer's opinion formation starts from an initial condition where the observer has an initial opinion on a "seed" subset of subjects, \mathcal{S} (seed opinions). The observer then gradually forms an opinion on each of the remaining subjects via sequential opinion formation, until opinions on all subjects are formed. Once formed, the opinions are not updated (we relax this assumption later). In one step, a target subject i is chosen at random from the pool of subjects with no opinion ($o_i = 0$). The observer then attempts to form an opinion on i . From all subjects j with an opinion ($o_j \neq 0$) that are adjacent to i ($R_{ij} \neq 0$), we choose source subject k at random. The opinion o_i is then set to $o_k R_{ki}$ (see Fig. 1b). As a result, a positive opinion on i is formed if either: (1) the observer has a positive opinion on k and the relation between k and i is positive ("the friends of my friends are also my friends") or (2) the observer has a negative opinion on k and the relation between k and i is negative (formalizing the ancient proverb "the enemies of my enemies are my friends"). A negative opinion on i is formed otherwise. Note that this mechanism produces a balanced triad consisting of the observer and subjects i and k (in Heider's original sense of heterogeneous triads that can include both individuals as well as objects^{21,35}). The observer then continues with a next subject until opinions on all subjects have been formed. This opinion formation process—which we refer to as the random neighbor rule as it forms opinions using neighboring subjects chosen at random—is purposely simple as it intends to imitate an observer with limited cognitive resources (see ref. ³⁴ for a recent account on susceptibility to fake news driven by "lack of reasoning"). We study a more thorough process (majority rule) below.

The opinion formation outcome is not deterministic (except for the special case when all paths in the subject network are balanced; see Sec. S1 in the Supplementary Information (SI)) as it is influenced by: (1) the order in which subjects are chosen as targets and (2) the choice of the source subject for each target subject. For a given relation network and a set of seed opinions, various resulting opinions are thus possible. Individual realizations of the opinion formation process correspond to a population of independent individuals or, alternatively, various possible "fates" of a single individual. To characterize statistical properties of the resulting opinions, we study outcomes of multiple model realizations. For simulations on synthetic relation networks, we additionally average over various network realizations to remove possible effects of a particular network realization on the results.

Opinion formation simulations on synthetic networks. We now study the opinion formation model on a specific relation network where the subjects form two camps. This scenario is relevant to various real situations^{24,29,30}: The two camps can represent two opposing political parties (such as democrats and republicans), standard news outlets and false news outlets, or scientists and conspiracy theorists, for example. In synthetic networks, each camp consists of $N/2$ subjects. Every subject is connected by signed links with z random subjects, thus creating a regular random network of trust with node degree z . If subjects from the same camp are linked, the sign of their relation is $+1$ with probability $1 - \beta$ and -1 otherwise. Similarly, if subjects from different camps are linked, the sign of their relation is -1 with probability $1 - \beta$ and $+1$ otherwise. Parameter $\beta \in [0, 0.5]$ thus

plays the role of structural noise. As β grows, the negative relations become more common within each camp and positive relations become more common across the camps. When $\beta = 0.5$, the two camps become indistinguishable by definition. The network's level of structural balance^{22,36} is the ratio of the number of balanced triads to all triads in the network. In our case,

$$B = (1 - \beta)^3 + 3(1 - \beta)\beta^2, \quad (1)$$

which corresponds to either all links of a triad (the first term) or one link of a triad (the second term) respecting the two-camp structure, producing a balanced triad as a result. B grows monotonously with β . The equation can be inverted, yielding $\beta = (1 + \sqrt[3]{1 - 2B})/2$, which can be used to write our results in terms of B instead of β .

We assume that the observer has initially a positive opinion on N_S seed subjects from camp 1, and we examine whether the observer ends up with a positive opinion on other subjects from camp 1 and a negative opinion on subjects from camp 2, or not. If the two camps represent scientists and conspiracy theorists, for example, the corresponding practical question is whether an observer who initially trusts a scientist would end up predominantly trusting scientists or conspiracy theorists. Without noise ($\beta = 0$), the opinion formation leads to a definite outcome: A positive opinion on all subjects from camp 1 and a negative opinion on all subjects from camp 2. In such a case, we say that the opinions are perfectly consistent with the underlying two-camp structure of the relation network among the subjects. Opinion consistency of a resulting opinion vector, \mathbf{o} , can be measured as

$$C(\mathbf{o}, \mathbf{T}) = \frac{1}{N - N_S} \sum_{j \notin \mathcal{S}} o_j T_j \quad (2)$$

where \mathcal{S} is the set of seed subjects and \mathbf{T} is the vector representing the ground-truth structure of the relation network (in our case, $T_j = 1$ for j from camp 1 and $T_j = -1$ for j from camp 2). If the observer's opinions are chosen at random, the resulting consistency is zero on average. A zero or small consistency value thus indicates that the observer's opinions are independent of the seed opinion and thus inconsistent with the two-camp structure of the relation network. Negative consistency is also possible: The observer starts with a positive opinion on subjects from camp 1 but ends with more positive opinions in camp 2 than in camp 1.

Knowing that opinion consistency is one in the absence of noise, how does it change as the noise parameter β grows? Numerical simulations for a set of 100 subjects and one seed opinion show (Fig. 2a) that opinion consistency decreases rapidly with β . Indeed, if the relationship between consistency and noise was linear, we would have expected $C_o(\beta) = 1 - 2\beta$ which starts at one when $\beta = 0$ and reaches zero when $\beta = 0.5$ as the two camps then cannot be distinguished by definition. By contrast, we observe a substantially faster decay of the mean consistency $\mu_C(\beta)$. In addition, the consistency values vary strongly between model realizations. For $\beta = 0.02$, for example, mean consistency is only 0.80 and there are model realizations with consistency below 0.54 and above 0.97 (the 10th and 90th percentile, respectively, of the obtained consistency values for $z = 4$). This means that even when the noise is small, some sets of formed opinions are in a dramatic disagreement with the observer's seed opinion. To appreciate the level of noise in real data, Moore²⁴ reported that 80% of triads among middle East countries are balanced. Equation (1) shows that such a level of structural balance is achieved at $\beta \approx 0.08$ in our two-camp networks. In Fig. 2b, mean opinion consistency at $\beta = 0.08$ is as low as 0.42 (for $z = 10$). These results confirm our initial hypothesis that a realistic level of

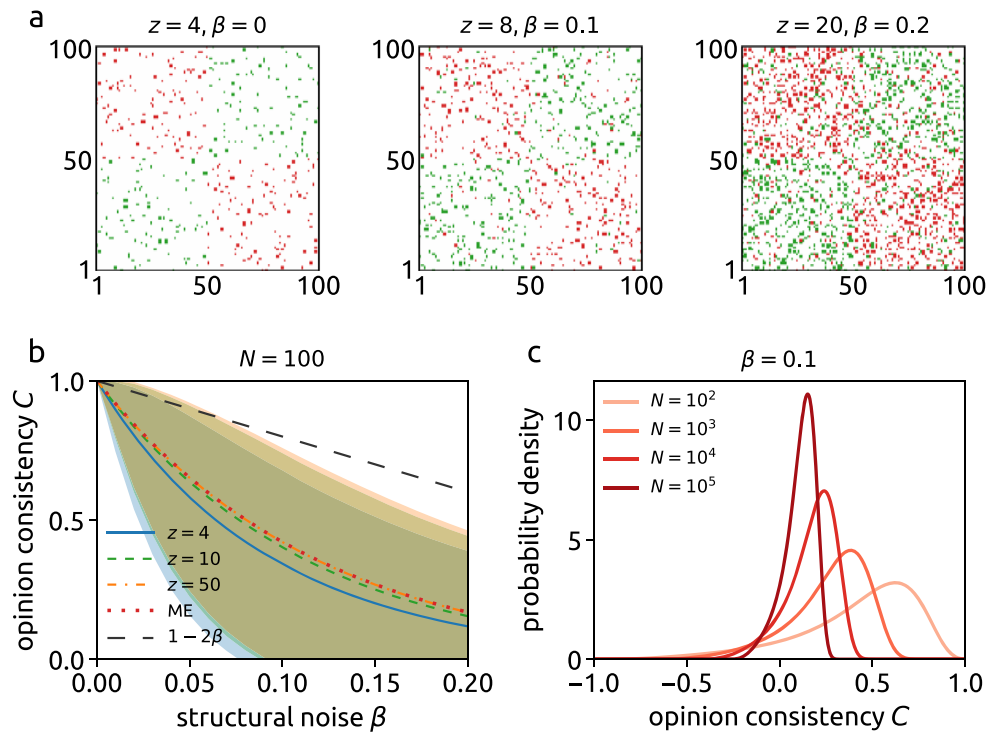


Fig. 2 Opinion formation on random relation networks with a two-camp structure. **a** Examples of adjacency matrices of regular random networks with a two-camp structure for 100 subjects and various values of node degree, z , and structural noise, β . Subjects 1–50 are from camp 1 and subjects 51–100 are from camp 2. The green, red, and white points represent positive, negative, and absent relations, respectively. **b** Resulting opinion consistency for $N = 100$ subjects, one seed opinion, and various values of z . The lines show mean values and the shaded regions show the 10th–90th percentile ranges (both computed from 1000 model realizations for each of 1000 realizations of the relation network). The dotted line labeled ME shows the mean consistency computed by numerically solving the master equation given by Eq. (3). The dark dashed line shows the expected consistency if its relationship with the level of noise was linear [$C_O(\beta) = 1 - 2\beta$]. These results demonstrate that consistency decreases quickly with structural noise. **c** The consistency distributions obtained using Eq. (3) for $\beta = 0.1$, one seed subject, and a growing number of subjects, N . As N increases, the distributions shrink and their peaks shift toward zero.

noise leads to the adoption of a large fraction of opinions that do not align with the observer’s initial opinion.

Master equation for opinion consistency and its solution. The opinion formation with the two-camp structure can be studied analytically under the assumption of homogeneous mixing³⁷. It is advantageous to study the problem in terms of the number of formed opinions, n , and the number of consistent opinions, c (that is, the opinions that are consistent with the seed opinions and the two-camp structure). By rewriting the sum $\sum_{j \neq i} o_j T_j$ in Eq. (2) as $2c - N + N_S$, we obtain opinion consistency as $C = (2c - N + N_S)/(N - N_S)$.

When the observer forms a new opinion, n increases by one and c either increases by one (if the new opinion is consistent) or remains constant. We introduce the probability distribution of c when n opinions have been formed, $P(c; n)$, for which the master equation (see “Methods” for the derivation) has the form

$$P(c; n) = P(c - 1; n - 1) \frac{c(1 - 2\beta) + \beta(n + 1) - 1}{n - 1} + P(c; n - 1) \left[1 - \beta - \frac{c(1 - 2\beta)}{n - 1} \right] \quad (3)$$

The initial condition $P(N_S; N_S) = 1$ represents that all N_S seed opinions are consistent. Equation (3) can be solved numerically and the obtained solution $P(c; n)$ can be used to compute the corresponding mean opinion consistency. The numerical solution agrees well with the model simulations (Fig. 2b), in particular when the relation network is not sparse ($z \gtrsim 10$).

Equation (3) allows us to analytically study the dependence between opinion consistency and the world complexity, represented

by the number of subjects, N . A surprising finding is that as the number of subjects increases, the distribution of C obtained by solving Eq. (3), $P(C)$, does not approach a well-defined limit distribution, but instead steadily shifts toward $C = 0$ and becomes narrower in the process (Fig. 2c). We study this behavior by computing the mean opinion consistency, $\mu_C(N)$, and the standard deviation of consistency, $\sigma_C(N)$ (see Sec. S2 in the SI for details and additional analytical results).

Multiplying Eq. (3) with c and summing it over $c = N_S, \dots, N$ yields the recurrence equation

$$\langle c(n) \rangle = \frac{n - 2\beta}{n - 1} \langle c(n - 1) \rangle + \beta \quad (4)$$

with the initial condition $\langle c(N_S) \rangle = N_S$ (the seed opinions are assumed to be correct). This recurrence equation can be solved in general, yielding

$$\langle c(N) \rangle = \frac{1}{2} \left[N + \frac{\Gamma(N + 1 - 2\beta)\Gamma(N_S + 1)}{\Gamma(N_S + 1 - 2\beta)\Gamma(N)} \right]. \quad (5)$$

For $N_S = 1$, the corresponding mean consistency is

$$\mu_C(N) = \left[\frac{\Gamma(N + 1 - 2\beta)}{\Gamma(2 - 2\beta)\Gamma(N)} - 1 \right] / (N - 1) \quad (6)$$

which simplifies to $\mu_C(N) = 1$ when $\beta = 0$, as expected. The leading contribution of Eq. (6) is

$$\mu_C(N) = N^{-2\beta} / \Gamma(2 - 2\beta) + O(1/N). \quad (7)$$

This shows that when $\beta > 0$, the mean opinion consistency vanishes in the limit $N \rightarrow \infty$. The leading-term contribution to

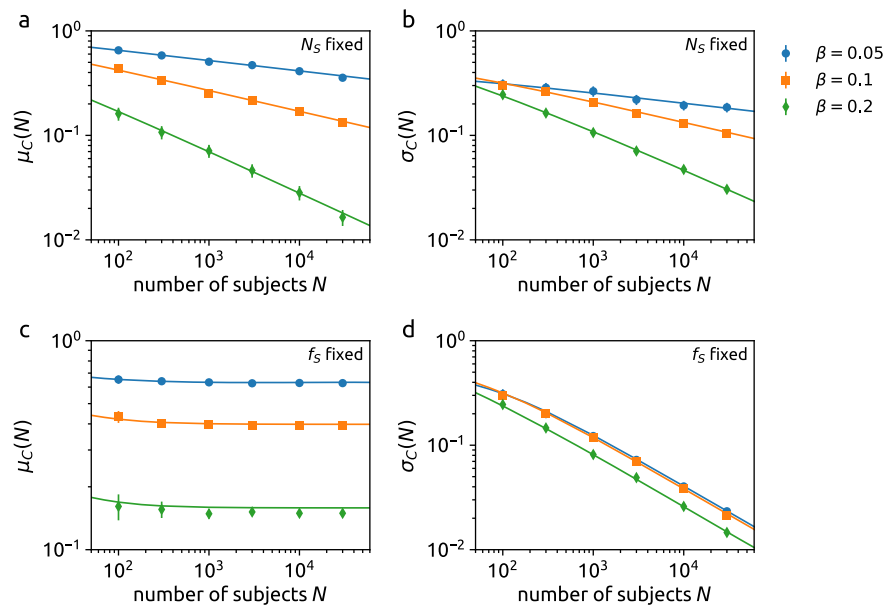


Fig. 3 Scaling of opinion consistency in synthetic worlds. Mean opinion consistency, μ_C , and the standard deviation of consistency, σ_C , as functions of the number of subjects, N , for regular random networks with node degree $z = 50$ and various levels of structural noise, β . The symbols and error bars show the mean and three times the standard error of the mean, respectively, obtained by running the model on 1000 network realizations. The solid lines show $\mu_C(N)$ and $\sigma_C(N)$ obtained by solving the master equation (Eq. (6) and Eq. (S6) in the SI). **a, b** Results for a single seed opinion. In this scenario, both μ_C and σ_C converge to zero as N grows: There is a tension between opinion consistency and the world's complexity. **c, d** Results for a fixed fraction of seed opinions, f_S , when the number of seed opinions is $N_S = f_S N$; here $f_S = 0.01$. In this scenario, the tension between consistency and complexity is prevented: μ_C converges to $f_S^{2\beta}$ (**c**) and σ_C converges to zero (**d**).

$\sigma_C(N)$ is also proportional to $N^{-2\beta}$ when $\beta \leq 1/4$. When $\beta > 1/4$, the leading term becomes proportional to $N^{-1/2}$. These analytic results agree with numerical simulations of the model (Fig. 3a, b and Fig. S1 in the SI).

The behavior demonstrated by Figs. 2b and 3, and supported by the analytic solution above, has important consequences. It shows that as the world complexity increases, the formed opinions become on average less consistent with the seed opinions and the two-camp structure of the subject network. Crucially, the opinion consistency is zero in the limit of an infinite number of subjects for any positive level of noise, β , in the subject relation network: In the limit of an infinite-complexity world, even a tiny amount of noise is enough to nullify opinion consistency.

The convergence of opinion consistency to zero as $N \rightarrow \infty$ can be avoided if the number of seed opinions grows linearly with N so that the fraction of seed opinions remains constant. Assuming that $N_S = f_S N$, Eq. (5) can be used to show that the mean consistency approaches to

$$\mu_C = f_S^{2\beta} \quad (8)$$

in the limit $N \rightarrow \infty$ and the standard deviation of consistency vanishes as $1/\sqrt{N}$ (see Sec. S2.2 in the SI). This scaling relation determines the necessary proportion of seed opinions, f_S , needed to achieve a desired opinion consistency, μ_C , for given β . These results are confirmed by numerical simulations shown in Fig. 3c, d and Fig. S2 in the SI. Despite having a positive limit value, opinion consistency still decreases quickly with noise in the relation network when f_S is small.

While our findings hold qualitatively when a different topology of the relation network is used, the mean opinion consistency values are heavily affected by the network topology (see Fig. 4). We run the opinion formation model on a growing preferential attachment network, a configuration model (CN) network with a power-law degree distribution, and Watts–Strogatz networks with

various values of the rewiring probability, p_r (see Sec. S3 in the SI for details on the network construction and additional simulation results). We find that networks with broad degree distributions lead to higher opinion consistency which decays with N slower (see Figs. S3 and S4 in the SI) than in the previously studied random networks. By contrast, Watts–Strogatz networks yield lower opinion consistency which further decreases as the networks become more regular through lowering the rewiring probability, p_r .

Opinion formation using the majority rule. The results described above hold for the opinion formation model where a random neighbor of a target subject is chosen as the reference. We chose this model to study the consequences of a cognitively easy opinion formation model. At this stage, one might object that the observed sensitivity of opinion consistency to noise might be because each formed opinion directly relies on only one previously formed opinion, and it might disappear if the observer incorporates the information from more neighbors before forming an opinion. To rule out this potential argument, we investigate a model where all neighbors of a target subject are considered before forming the opinion. Denote the numbers of neighbors leading to the adoption of a positive and a negative opinion (determined as in Fig. 1b) as n_p and n_N , respectively. If $n_p > n_N$, the observer forms a positive opinion. If $n_N > n_p$, the observer forms a negative opinion. If $n_p = n_N$, a random opinion is formed. We refer to this as the majority opinion formation rule. It is more demanding than the original random neighbor rule based on choosing a random neighbor as it assumes that the observer carefully collects all evidence for forming an opinion on a target subject. The majority rule is nevertheless still a local rule as it only considers direct neighbors of a target subject.

Using the majority rule, a scaling analogous to Fig. 3 can be observed (Fig. 5a, b) when the number of seed opinions is fixed. The important difference is that the scaling exponent now depends on both β and z whereas a higher mean degree, z ,




Fig. 4 Opinion consistency for various topologies of synthetic relation networks. The dependence of mean opinion consistency, μ_C , on structural noise, β , for one seed opinion and different numbers of subjects: **a** $N = 100$ and **b** $N = 1000$. The results are averaged over 1000 model realizations on each of 1000 network realizations, mean degree $z \approx 4$ for each network topology. The error bars are too small to be shown. CN and p_r stand for the configuration model and the rewiring probability in the Watts-Strogatz model, respectively.




Fig. 5 Scaling of opinion consistency achieved by the majority rule in synthetic worlds. **a, b** The dependencies of μ_C and σ_C on the number of subjects, N , are affected by both structural noise, β , and mean degree, z (the solid, dashed, and dotted lines are used for z equal 50, 10, and 4, respectively). One seed opinion is used ($N_S = 1$). The symbols and error bars show the mean and three times the standard error of the mean, respectively, obtained by running the model on 1000 realizations of regular random networks. **c** The dependence of mean opinion consistency, μ_C , on structural noise, β , for a fixed fraction of seed opinions, $f_S = 0.01$, $N = 10,000$, and $N_S = f_S N$. The results are averaged over model realizations on 1000 regular random networks, the error bars are too small to be shown. The dotted line shows the analytical result $\mu_C = f_S^{2\beta}$ for the random neighbor rule when the same fixed fraction of seed opinions, $f_S = 0.01$, is used.

generally leads to $\mu_C(N)$ and $\sigma_C(N)$ decaying slower with N . Except for the smallest used noise and the highest used degree ($\beta = 0.05$ and $z = 50$), all slopes of linear fits between $\ln \mu_C(N)$ and $\ln N$ are significantly positive. Since the majority rule does not lend itself to analytical computation, whether the limit of $\mu_C(N)$ is indeed positive when β is sufficiently small and z is sufficiently high remains an open question. When the fraction of seed opinions, f_S , is fixed, $\mu_C(N)$ converges to a positive value and $\sigma_C(N)$ vanishes as N grows (Fig. S5 in the SI). The scaling of the majority rule and the random neighbor rule is thus qualitatively the same when f_S is fixed. Figure 5c shows mean opinion consistency as a function of structural noise for fixed f_S when the number of subjects is large. We see that when the network density is low ($z = 4$), the majority rule achieves results that are comparable with those of the random neighbor rule. When z increases, the majority rule leads to significantly more consistent opinions than the random neighbor rule. It has to be noted, though, that when z is large, the cost for the observer to collect and analyze all information for opinion formation is large too.

The effect of opinion updating on opinion consistency. To focus on the opinion formation mechanism alone, the original model assumes that the formed opinions cannot be changed anymore. In real life, however, we occasionally re-evaluate our positions. This re-evaluation can be included in the model by allowing the opinions to be updated as it is, in fact, common in the opinion formation literature^{17,38,39}.

To incorporate an updating mechanism while preserving the spirit of the model, we assume that in the first phase, all opinions are formed using the original model. In the second phase, target subjects for opinion updating are repeatedly chosen from all subjects at random. Opinions are updated in the same way as they

have been formed in the first phase, i.e., using the random neighbor rule or the majority rule. Simulation results (see Fig. S6 in the SI) show that the effect of opinion updating markedly differs between the two opinion formation rules. When the random neighbor rule is used, opinion consistency substantially decreases with the number of updated opinions. The reason for this decrease lies in the opinion formation rule which relies on one randomly chosen neighbor of the target subject. Noise in the relation network thus continues to accumulate through updating which thus lowers opinion consistency. By contrast, opinion consistency substantially increases with the number of updated opinions when the majority rule is used. Thanks to its thorough (albeit local) consideration of opinions on neighbors of the target subject, the majority rule is capable of identifying and correcting imbalanced triads in the relation network, and thus increase opinion consistency.

Opinion formation simulations on real networks. The trust consistency metric requires information on the ground truth structure of the relation network (such as the assignment of subjects to one of the two camps in the case of a two-camp structure). Before analyzing empirical data, we aim to introduce a proxy for opinion consistency that does not require such information which is typically not available for real data. To this end, we introduce opinion stability, S , which measures the extent to which elements of the opinion vector are the same in independent realizations of the opinion formation model (see “Methods” for the definition). If an opinion on a given subject always ends up positive (or always negative), it is a sign of a robust opinion and it contributes positively to opinion stability. Small opinion stability indicates that the opinion formation outcomes are highly volatile and, in turn, they do not comply with the division of subjects in