

The Majority Illusion in Social Networks

Kristina Lerman, Xiaoran Yan, and Xin-Zeng Wu

USC Information Sciences Institute, 4676 Admiralty Way, Marina del Rey, CA 90292

(Dated: June 10, 2015)

Social behaviors are often contagious, spreading through a population as individuals imitate the decisions and choices of others. A variety of global phenomena, from innovation adoption to the emergence of social norms and political movements, arise as a result of people following a simple local rule, such as copy what others are doing. However, individuals often lack global knowledge of the behaviors of others and must estimate them from the observations of their friends' behaviors. In some cases, the structure of the underlying social network can dramatically skew an individual's local observations, making a behavior appear far more common locally than it is globally. We trace the origins of this phenomenon, which we call "the majority illusion," to the friendship paradox in social networks. As a result of this paradox, a behavior that is globally rare may be systematically overrepresented in the local neighborhoods of many people, i.e., among their friends. Thus, the "majority illusion" may facilitate the spread of social contagions in networks and also explain why systematic biases in social perceptions, for example, of risky behavior, arise. Using synthetic and real-world networks, we explore how the "majority illusion" depends on network structure and develop a statistical model to calculate its magnitude in a network.

I. INTRODUCTION

An individual's behavior often depends on the actions of others [8, 9, 17, 29, 30, 33]. This phenomenon is manifested daily in the decisions people make to adopt a new technology [28, 31] or idea [7, 33], listen to music [29], engage in risky behavior [4], or join a social movement [17, 30]. As a result, a variety of behaviors are said to be 'contagious', because they spread through the population as people observe others adopting the behavior and then adopt it themselves. In some cases, behaviors will spread from a small number of initial adopters to a large portion of the population, resulting in fads, hit songs, successful political campaigns, epidemics, and social norms. Researchers have examined the conditions under which such global outbreaks occur, especially in a networked setting, where individuals interact with a subset of the population, i.e., their network neighbors or friends. Studies have linked the onset of global outbreaks to the topology of underlying network [9, 32], including the presence of highly connected individuals [21, 25] and small clusters of connected people [8, 33]. However, network structure can affect the emergence of global outbreaks in a subtler way. As we show in this paper, the configuration of initial adopters on a network can systematically skew the observations people make of their friends' behavior. This can make some behavior appear much more popular than it is, thus creating conditions for its spread.

Networks often have counter-intuitive properties. One of the better known of these is the

friendship paradox, which states that, on average, most people have fewer friends than their friends have [14]. Despite its almost nonsensical nature, the friendship paradox has been used to design efficient strategies for vaccination [12], social intervention [22], and early detection of contagious outbreaks [11, 16]. In a nutshell, rather than monitoring random people to catch a contagious outbreak in its early stages, the friendship paradox suggests monitoring their random friends, because they are more likely to be better connected and not only to get sick earlier, but also to infect more people once sick.

Recently, friendship paradox was generalized for attributes other than degree, i.e., number of network neighbors. For example, your co-authors are cited more often than you [13], and the people you follow on Twitter post more frequently than you do [19]. In fact, any attribute that is correlated with degree will produce a paradox [13, 23]. Thus, if heavy drinkers also happen to be more popular, then people examining their friends' drinking behavior will conclude that, on average, their friends drink more than they do. This may help explain why adolescents systematically overestimate their friends' alcohol consumption and drug use [3, 6].

In this paper, we describe a novel variation of the friendship paradox that is essential for understanding contagious behaviors. The paradox applies to networks in which nodes have attributes, in the simplest case a binary attribute, such as "has red hair" vs "does not have red hair" or "purchased an iPhone" vs "did not purchase an iPhone". We refer to nodes that have this attribute as "active", and the rest are "in-

active.” We show that under some conditions, a large fraction of nodes will observe most of their neighbors in the active state, even when it is globally rare. For this reason, we call the paradox the “majority illusion.”

As a simple illustration of the “majority illusion” paradox, consider the two networks in Figure 1. The networks are identical, except for which of the few nodes are colored. Imagine that colored nodes are active and the rest of the nodes are inactive. Despite this apparently small difference, the two networks are profoundly different: in the first network, every inactive node will examine its neighbors to observe that “at least half of my neighbors are active,” while in the second network no node will make this observation. Thus, even though only three of the 14 nodes are active, it appears to all inactive nodes in the first network that most of their neighbors are active.

The “majority illusion” can dramatically impact social contagions. Researchers use the threshold model to describe the spread of social contagions in networks [10, 17, 32]. At each time step in this model, an inactive individual observes the current states of its k neighbors, and becomes active if more than ϕk of the neighbors are active; otherwise, it remains inactive. The fraction $0 \leq \phi \leq 1$ is the activation *threshold*. It represents the amount of social proof an individual requires before switching to the active state [17]. Threshold of $\phi = 0.5$ means that to become active, an individual has to have a majority neighbors in the active state. Though the two networks in Figure 1 have the same topology, when the threshold is $\phi = 0.5$, all nodes will eventually become active in the network on the left, but not in the network on the right. This is because the “majority illusion” alters local neighborhoods of the nodes, distorting their perceptions of the prevalence of the active state.

This paper describes and analyzes the “majority illusion” paradox. We measure the strength of the paradox as the fraction of network nodes with a majority active neighbors. Using synthetic and real-world networks, we study how network structure and configuration of active nodes contributes to the paradox. We demonstrate empirically, as well as through theoretical analysis, that the paradox is stronger in networks in which the better-connected nodes are active, and also in networks with a heterogeneous degree distribution. Network structure also amplifies the paradox via degree correlations. The paradox is strongest in networks where low degree nodes have the tendency to

connect to high degree nodes. Activating the high degree nodes in such networks biases the local observations of many nodes, which in turn impacts collective phenomena emerging in networks, including social contagions. Our statistical model quantifies the strength of this effect.

II. RESULTS

A network’s structure is partly specified by its degree distribution $p(k)$, which gives the probability that a randomly chosen node in an undirected network has k neighbors (i.e., degree k). This quantity also affects the probability that a randomly chosen edge is connected to a node of degree k , otherwise known as neighbor degree distribution $q(k)$. Since high degree nodes have more edges, they will be over-represented in the neighbor degree distribution by a factor proportional to their degree; hence, $q(k) = kp(k)/\langle k \rangle$, where $\langle k \rangle$ is the average node degree.

Networks often have structure beyond that specified by their degree distribution: for example, nodes may preferentially link to others with a similar (or very different) degree. Such degree correlation is captured by the joint degree distribution $e(k, k')$, the probability to find nodes of degrees k and k' at either end of a randomly chosen edge in an undirected network [27]. This quantity obeys normalization conditions $\sum_{k,k'} e(k, k') = 1$ and $\sum_{k'} e(k, k') = q(k)$. Globally, degree correlation in an undirected network is quantified by the assortativity coefficient, which is simply the Pearson correlation between degrees of connected nodes:

$$\begin{aligned} r &= \frac{1}{\sigma_q^2} \sum_{k,k'} kk' [e(k, k') - q(k)q(k')] \\ &= \frac{1}{\sigma_q^2} \left[\left(\sum_{k,k'} kk' e(k, k') \right) - \langle k \rangle_q^2 \right]. \end{aligned}$$

Here, $\sigma_q^2 = \sum_k k^2 q(k) - [\sum_k k q(k)]^2$. In assortative networks ($r > 0$), nodes have a tendency link to similar nodes, e.g., high-degree nodes to other high-degree nodes. In disassortative networks ($r < 0$), on the other hand, they prefer to link to dissimilar nodes. A star composed of a central hub and nodes linked only to the hub is an example of a disassortative network.

We can use Newman’s edge rewiring procedure [27] to change a network’s assortativity without changing its degree distribution $p(k)$. The rewiring procedure randomly chooses two pairs of connected nodes and swaps their edges

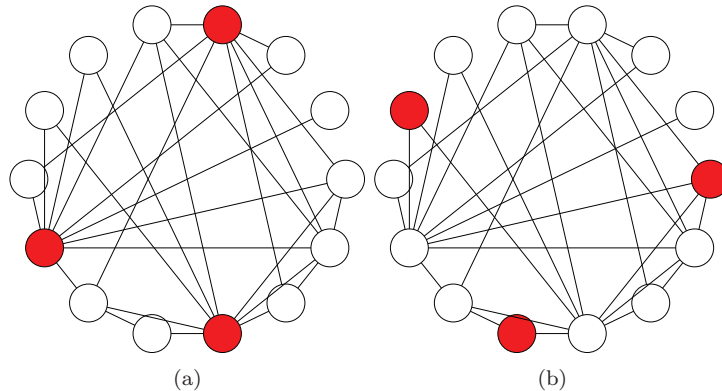


FIG. 1. An illustration of the “majority illusion” paradox. The two networks are identical, except for which three nodes are colored. These are the “active” nodes and the rest are “inactive.” In the network on the left, all “inactive” nodes observe that at least half of their neighbors are “active,” while in the network on the right, no “inactive” node makes this observation.

if doing so changes their degree correlation. This can be repeated until desired assortativity is achieved.

The configuration of attributes in a network is specified by the joint probability distribution $P(x, k)$, the probability that node of degree k has an attribute x . In this work, we consider binary attributes only, and refer to nodes with $x = 1$ as active and those with $x = 0$ as inactive. The joint distribution can be used to compute ρ_{kx} , the correlation between node degrees and attributes:

$$\begin{aligned} \rho_{kx} &\equiv \frac{1}{\sigma_x \sigma_k} \sum_{x,k} xk [P(x, k) - P(x)p(k)] \quad (1) \\ &= \frac{1}{\sigma_x \sigma_k} \sum_k k [P(x = 1, k) - P(x = 1)p(k)] \\ &= \frac{P(x = 1)}{\sigma_x \sigma_k} [\langle k \rangle_{x=1} - \langle k \rangle]. \end{aligned}$$

In the equations above, σ_k and σ_x are the standard deviations of the degree and attribute distributions respectively, and $\langle k \rangle_{x=1}$ is the average degree of active nodes.

Randomly activating nodes creates a configuration with ρ_{kx} close to zero. We can change it by swapping attribute values among the nodes. For example, to increase ρ_{kx} , we randomly choose nodes v_1 with $x = 1$ and v_0 with $x = 0$ and swap their attributes if the degree of v_0 is greater than the degree of v_1 . We can continue swapping attributes until desired ρ_{kx} is achieved (or it no longer changes).

A. “Majority Illusion” in Synthetic and Real-world Networks

Synthetic networks allow us to systematically study how network structure affects the strength of the “majority illusion” paradox. First, we looked at scale-free networks. We generated networks with $N = 10,000$ nodes and degree distribution of the form $p(k) \sim k^{-\alpha}$. Such networks are used to model the heterogeneous structure of many real-world networks, which contain a few high degree hubs and many low degree nodes. To create a scale-free network, we first sampled a degree sequence from a distribution with exponent α , where exponent α took three different values (2.1, 2.4, and 3.1), and then used the configuration model to create an undirected network with that degree sequence (see Appendix Sec. S1). We activated $P(x = 1) = 0.05$ of nodes and used edge rewiring and attribute swapping procedures describe above to change the network’s degree assortativity r and degree–attribute correlation ρ_{kx} .

Figure 2 shows the fraction of nodes with a majority active neighbors in these scale-free networks as a function of the degree–attribute correlation ρ_{kx} and for different values of degree assortativity r . The fraction of nodes experiencing the “majority illusion” can be quite large. For $\alpha = 2.1$, 60%–80% of the nodes may observe a majority active neighbors, even though only 5% of the nodes are, in fact, active. The “majority illusion” is exacerbated by three factors: it becomes stronger as the degree–attribute correlation increases, and as the network becomes more disassortative (i.e., r decreases) and heavier-

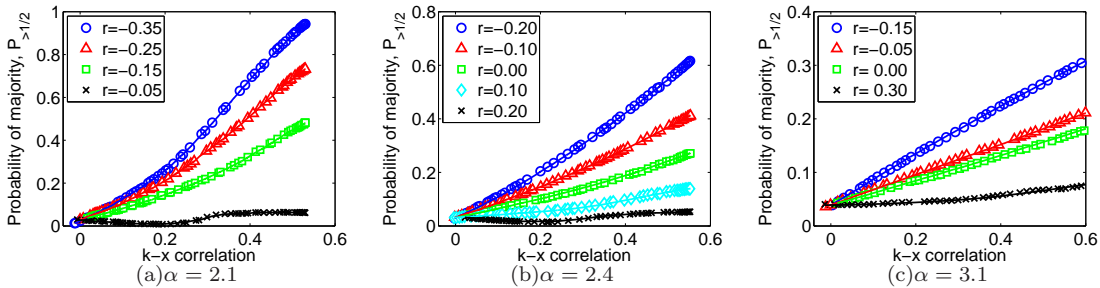


FIG. 2. Magnitude of the “majority illusion” in scale-free networks as a function of degree–attribute correlation ρ_{kx} and for different values of degree assortativity r . Each network has 10,000 nodes and degree distribution of the form $p(k) \sim k^{-\alpha}$. The fraction of active nodes in all cases is 5%. The lines represent calculations using the statistical model of Equation 4.

tailed (i.e., α becomes smaller). However, even when $\alpha = 3.1$, under some conditions a substantial fraction of nodes will experience the paradox. The lines in the figure show theoretical estimates of the paradox using Equation 4, as described in the next subsection.

“Majority illusion” can also be observed in networks with a Poisson degree distribution. We used the Erdős-Rényi model to generate networks with $N = 10,000$ and average degrees $\langle k \rangle = 5.2$ and $\langle k \rangle = 2.5$ (see Appendix Sec. S1). We randomly activated 5%, 10%, and 20% of the nodes, and used edge rewiring and attribute swapping to change r and ρ_{kx} in these networks. Figure 3 shows the fraction of nodes in the paradox regime. Though much reduced compared to scale-free networks, we still observe some amount of the paradox, especially in networks with a greater fraction of active nodes.

We also examined whether “majority illusion” can be manifested in real-world networks. We looked at three typical social and communications networks: the co-authorship network of high energy physicists (HepTh) [24], social media follower graph (Digg) [20], and the network representing links between political blogs (blogs) [2]. All three networks are undirected. To make the Digg follower graph undirected, we kept only the mutual follow links, and further reduced the graph by extracting the largest connected component. Except for political blogs ($r = -0.22$), the networks were assortative with $r = 0.23$ and $r = 0.12$ for HepTh and Digg respectively. These data are described in Appendix Sec. S1.

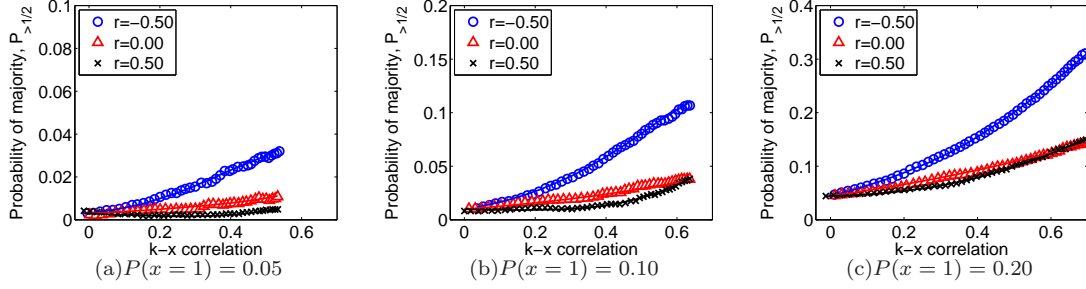
Figure 4 shows the fraction of nodes experiencing the “majority illusion” for different fractions of active nodes $P(x = 1) = 0.05, 0.1, 0.2$ and 0.3 . As degree–attribute correlation ρ_{kx} in-

creases, a substantial fraction of nodes experience the paradox in almost all networks. The effect is largest in the disassortative political blogs network, where for high enough correlation, as many as 60%–70% of nodes will have a majority active neighbors, even when only 20% of the nodes are active. The effect also exists in the Digg network of mutual followers, and to a lesser degree in the HepTh co-authorship network. Although positive assortativity reduces the magnitude of the effect, compared with synthetic networks, local perceptions of nodes in real-world networks can also be skewed. If the attribute represents an opinion, under some conditions, even a minority opinion can appear to be extremely popular locally.

B. Modeling “Majority Illusion” in Networks

Having demonstrated empirically some of the relationships between “majority illusion” and network structure, we next develop a model that includes network properties in the calculation of paradox strength. Like the friendship paradox, the “majority illusion” is rooted in differences between degrees of nodes and their neighbors [14, 18]. These differences result in nodes perceiving that, not only are their neighbors better connected [14], on average, but that they also have more of some attribute than they themselves have [19]. The latter paradox, which is referred to as the generalized friendship paradox, is enhanced by correlations between node degrees and attribute values ρ_{kx} [13]. In binary attribute networks, where nodes can be either active or inactive, a configuration in which higher degree nodes tend to be active causes the

Erdős-Rényi network with $N = 10,000$ and $\langle k \rangle = 5.2$



Erdős-Rényi network with $N = 10,000$ and $\langle k \rangle = 2.5$

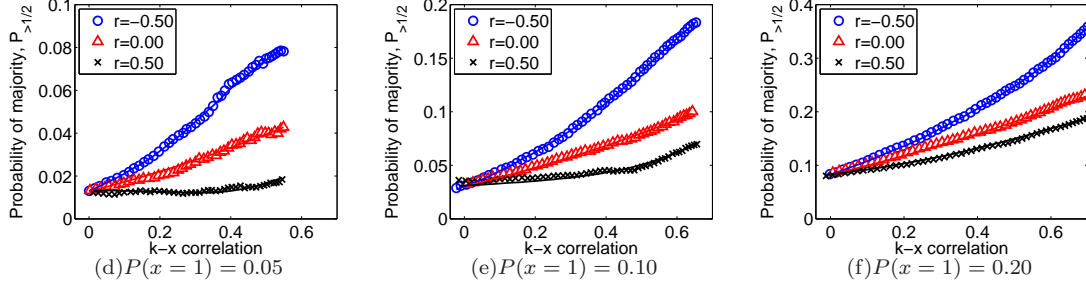


FIG. 3. Magnitude of the “majority illusion” in Erdős-Rényi-type networks as a function of degree–attribute correlation ρ_{kx} and for different values of degree assortativity r . Each network has 10,000 nodes with $\langle k \rangle = 5.2$ (top row) or $\langle k \rangle = 2.5$ (bottom row), and different fractions of active nodes. The lines represent calculations using the statistical model of Equation 4.

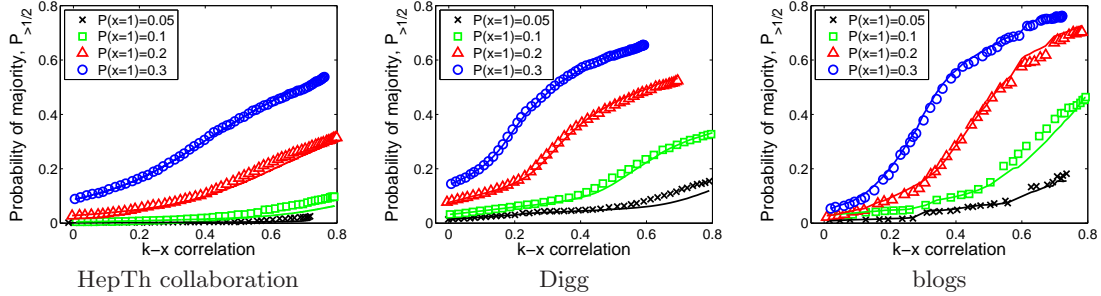


FIG. 4. Magnitude of the “majority illusion” in real-world networks as a function of degree–attribute correlation ρ_{kx} for different fraction of active nodes $P(x = 1)$. The lines represent calculations using the statistical model of Equation 4.

remaining nodes to observe that their neighbors are more active than they are (see Appendix Sec. S1 A).

While heterogeneous degree distribution and degree–attribute correlations give rise to friendship paradoxes even in random networks, other elements of network structure, such as correlation between degrees of connected nodes, may also affect observations nodes make of their neighbors. To understand why, consider the probability that a node has an active neighbor,

conditioned on that node having a degree k :

$$\begin{aligned}
 P(x' = 1|k) &= \sum_{k'} P(x' = 1|k')P(k'|k) \quad (2) \\
 &= \sum_{k'} P(x' = 1|k') \frac{e(k, k')}{q(k)}.
 \end{aligned}$$

In the equation above, $e(k, k')$ is the joint degree distribution. Globally, the probability that any

node has an active neighbor is

$$\begin{aligned}
P(x' = 1) &= \sum_k P(x' = 1|k)p(k) \\
&= \sum_k \sum_{k'} P(x' = 1|k') \frac{e(k, k')}{q(k)} p(k) \\
&= \sum_k \sum_{k'} \frac{P(x' = 1, k')}{p(k')} e(k, k') \frac{\langle k \rangle}{k} \\
&= \sum_{k'} \frac{P(x' = 1, k')}{q(k')} \sum_k \frac{k'}{k} e(k, k').
\end{aligned}$$

Given two different networks with the same degree distribution $p(k)$ and the same configuration of active nodes, the probability that a node in each network observes an active neighbor $P(x' = 1)$ is a function of $\sum_{k, k'} (k'/k) e(k, k')$. Since assortativity r is a function of $\sum_{k, k'} k k' e(k, k')$, we can see that the two expressions weigh the $e(k, k')$ term in opposite ways. This suggests that the probability of having an active neighbor increases as assortativity decreases and vice versa. Thus, we expect stronger paradoxes in disassortative networks.

To quantify the “majority illusion” paradox, we calculate the probability that a node of degree k has more than a fraction ϕ of active neighbors, i.e., neighbors with attribute value $x' = 1$:

$$\begin{aligned}
P_{>\phi}(k) &= \sum_{n>\phi k}^k \binom{k}{n} \times \\
&P(x' = 1|k)^n (1 - P(x' = 1|k))^{(k-n)}.
\end{aligned} \tag{3}$$

Here $P(x' = 1|k)$ is the conditional probability of having an active neighbor, given a node with degree k , and is specified by Eq. 2. Although the threshold ϕ in Eq. 3 could be any fraction, in this paper we focus on $\phi = \frac{1}{2}$, which represents a straight majority. Thus, the fraction of all nodes most of whose neighbors are active is

$$\begin{aligned}
P_{>\frac{1}{2}} &= \sum_k p(k) \sum_{n>\frac{k}{2}}^k \binom{k}{n} \times \\
&P(x' = 1|k)^n (1 - P(x' = 1|k))^{(k-n)}.
\end{aligned} \tag{4}$$

Using Equation 4, we can calculate the strength of the “majority illusion” paradox for any network whose degree sequence, joint degree distribution $e(k, k')$, and conditional attribute distribution $P(x|k)$ are known. The solid lines in Figures 2–4 report these calculations for each network. We used the empirically determined joint probability distribution $P(x, k)$ to calculate $P(x = 1|k)$ in the equation above, as well

as ρ_{kx} for the figure. For some “well-behaved” degree distributions, $P(x = 1|k)$ can be determined analytically (rather than empirically) by approximating the joint distribution $P(x, k)$ as a bivariate normal distribution (see Appendix Sec. S1 B). However, this “gaussian” approximation breaks down as degree distributions become heavier tailed. Overall, our statistical model does a good job explaining most of the empirical observations. Although the global degree assortativity r is an important contributor to the “majority illusion,” a more detailed view of the structure using joint degree distribution $e(k, k')$ is necessary to accurately estimate the magnitude of the paradox. As demonstrated in Appendix Sec. S1 C, two networks with the same $p(k)$ and r can display different amounts of the paradox.

III. DISCUSSION

Local prevalence of some attribute among a node’s network neighbors can be very different from its global prevalence, creating an illusion that the attribute is far more common than it actually is. In a social network, this illusion may cause people to reach wrong conclusions about how common a behavior is, leading them to accept as a norm a behavior that is globally rare. This may explain how global outbreaks can be triggered by very few initial adopters, and why people overestimate how much their friends engage in risky behaviors, such as alcohol and drug use.

We quantified this paradox, which we call the “majority illusion”, and studied its dependence on network structure and attribute configuration. As in the friendship paradox [13, 14, 19, 23], “majority illusion” can ultimately be traced to the power of high degree nodes to skew the observations of many others. This is because such nodes are overrepresented in the local neighborhoods of other nodes. This, by itself is not surprising, given that high degree nodes are expected to have more influence and are often targeted by influence maximization algorithms [21]. However, the ability of high degree nodes to bias the perceptions of others depends on other aspects of network structure. Specifically, we showed that the paradox is much stronger in disassortative networks, where high degree nodes tend to link to low degree nodes. In other words, given the same degree distribution, the high degree nodes in a disassortative network will have greater power to skew the observations of oth-

ers than those in an assortative network. This suggests that some network structures are more susceptible than others to influence manipulation and the spread of external shocks [32]. Furthermore, small changes in network topology, assortativity and degree–attribute correlation may further exacerbate the paradox even when there are no actual changes in the distribution of the attribute. This may explain the apparently sudden shifts in public attitudes witnessed during the Arab Spring and on the question of gay marriage.

The “majority illusion” is an example of class size bias effect. When sampling data to estimate average class or event size, more popular classes and events will be overrepresented in the sam-

ple, biasing estimates of their average size [15]. Thus, the average class size that students experience at college is larger than the college’s average class size. Similarly, people experience highways, restaurants, and events to be more crowded than they normally are. In networks, sampling bias affects estimates of network structure, including its degree distribution [1, 18]. Our work suggests that network bias also affects an individual’s local perceptions and the collective social phenomena that emerge.

Acknowledgements

Authors are grateful to Nathan Hodas and Farshad Kooti for their inputs into this work.

-
- [1] Dimitris Achlioptas, Aaron Clauset, David Kempe, and Christopher Moore. On the Bias of Traceroute Sampling; or, Power-law Degree Distributions in Regular Graphs. In *Proc. 37th ACM Symposium on Theory of Computing (STOC)*, 2005.
 - [2] Lada A Adamic and Natalie Glance. The political blogosphere and the 2004 us election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*, pages 36–43. ACM, 2005.
 - [3] J. S. Baer, A. Stacy, and M. Larimer. Biases in the perception of drinking norms among college students. *Journal of studies on alcohol*, 52(6):580–586, November 1991.
 - [4] Jonathan M. Bearak. Casual Contraception in Casual Sex: Life-Cycle Change in Undergraduates’ Sexual Behavior in Hookups. *Social Forces*, 93(2):sou091–513, October 2014.
 - [5] Edward A Bender and E Rodney Canfield. The asymptotic number of labeled graphs with given degree sequences. *Journal of Combinatorial Theory, Series A*, 24(3):296–307, 1978.
 - [6] Alan D Berkowitz. An overview of the social norms approach. *Changing the culture of college drinking: A socially situated health communication campaign*, pages 193–214, 2005.
 - [7] Luís Bettencourt, Ariel Cintron-Arias, David I Kaiser, and Carlos Castillo-Chavez. The power of a good idea: Quantitative modeling of the spread of ideas from epidemiological models. *Physica A: Statistical Mechanics and its Applications*, 364:513–536, 2006.
 - [8] Damon Centola. The Spread of Behavior in an Online Social Network Experiment. *Science*, 329(5996):1194–1197, September 2010.
 - [9] Damon Centola and Andrea Baronchelli. The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proceedings of the National Academy of Sciences*, 112(7):1989–1994, February 2015.
 - [10] Damon Centola, Víctor M. Eguíluz, and Michael W. Macy. Cascade dynamics of complex propagation. *Physica A: Statistical Mechanics and its Applications*, 374(1):449–456, January 2007.
 - [11] Nicholas A. Christakis and James H. Fowler. Social Network Sensors for Early Detection of Contagious Outbreaks. *PLoS ONE*, 5(9):e12948+, September 2010.
 - [12] Reuven Cohen, Shlomo Havlin, and Daniel ben Avraham. Efficient Immunization Strategies for Computer Networks and Populations. *Phys. Rev. Lett.*, 91:247901, Dec 2003.
 - [13] Young-Ho Eom and Hang-Hyun Jo. Generalized friendship paradox in complex networks: The case of scientific collaboration. *Scientific Reports*, 4, April 2014.
 - [14] Scott L. Feld. Why Your Friends Have More Friends Than You Do. *American Journal of Sociology*, 96(6):1464–1477, May 1991.
 - [15] Scott L. Feld and Bernard Grofman. Variation in Class Size, the Class Size Paradox, and Some Consequences for Students. *Research in Higher Education*, 6(3), 1977.
 - [16] Manuel Garcia-Herranz, Esteban Moro, Manuel Cebrian, Nicholas A. Christakis, and James H. Fowler. Using Friends as Sensors to Detect Global-Scale Contagious Outbreaks. *PLoS ONE*, 9(4):e92413+, April 2014.
 - [17] Mark Granovetter. Threshold Models of Collective Behavior. *American Journal of Sociology*, 83(6):1420–1443, 1978.
 - [18] Sidharth Gupta, Xiaoran Yan, and Kristina Lerman. Structural Properties of Ego Networks. In *International Conference on Social Computing, Behavioral Modeling and Prediction*, 2015.

- [19] Nathan Hodas, Farshad Kooti, and Kristina Lerman. Friendship Paradox Redux: Your Friends Are More Interesting Than You. In *Proc. 7th Int. AAAI Conf. on Weblogs And Social Media*, 2013.
- [20] Tad Hogg and Kristina Lerman. Social dynamics of digg. *EPJ Data Science*, 1(5), June 2012.
- [21] David Kempe, Jon Kleinberg, and Eva Tardos. Maximizing the spread of influence through a social network. In *KDD '03: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, New York, NY, USA, 2003. ACM Press.
- [22] David A. Kim, Alison R. Hwang, Derek Stafford, D. Alex Hughes, A. James O’Malley, James H. Fowler, and Nicholas A. Christakis. Social network targeting to maximise population behaviour change: a cluster randomised controlled trial. *The Lancet*, May 2015.
- [23] Farshad Kooti, Nathan O. Hodas, and Kristina Lerman. Network Weirdness: Exploring the Origins of Network Paradoxes. In *International Conference on Weblogs and Social Media (ICWSM)*, March 2014.
- [24] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1):2, 2007.
- [25] J. O. Lloyd-Smith, S. J. Schreiber, P. E. Kopp, and W. M. Getz. Superspreading and the effect of individual variation on disease emergence. *Nature*, 438(7066):355–359, 2005.
- [26] Michael Molloy and Bruce Reed. A critical point for random graphs with a given degree sequence. *Random structures & algorithms*, 6(2-3):161–180, 1995.
- [27] M. E. J. Newman. Assortative Mixing in Networks. *Phys. Rev. Lett.*, 89:208701, Oct 2002.
- [28] Everett M. Rogers. *Diffusion of Innovations, 5th Edition*. Free Press, 5th edition, August 2003.
- [29] M. J. Salganik, P. S. Dodds, and D. J. Watts. Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science*, 311, 2006.
- [30] Thomas C. Schelling. Hockey Helmets, Concealed Weapons, and Daylight Saving: A Study of Binary Choices with Externalities. *The Journal of Conflict Resolution*, 17(3), 1973.
- [31] Thomas W. Valente. *Network Models of the Diffusion of Innovations (Quantitative Methods in Communication Subseries)*. Hampton Press (NJ), 1995.
- [32] Duncan J. Watts. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9):5766–5771, April 2002.
- [33] H. Peyton Young. The dynamics of social innovation. *Proceedings of the National Academy of*

Sciences, 108(Supplement 4):21285–21291, December 2011.

APPENDIX

S1. DATA

We use the configuration model [5, 26], as implemented by the SNAP library (<https://snap.stanford.edu/data/>), to create a scale-free network with a specified degree sequence. We generated a degree sequence from a power law of the form $p(k) \sim k^{-\alpha}$. Here, p_k is fraction of nodes that have k half-edges. The configuration model proceeded by linking a pair of randomly chosen half-edges to form an edge. The linking procedure was repeated until all half-edges have been used up or there were no more ways to form an edge.

To create Erdős-Rényi-type networks, we started with $N = 10,000$ nodes and linked pairs at random with some fixed probability. These probabilities were fixed to produce average degree similar to the average degree of the scale-free networks.

The statistics of real-world networks we studied, including the collaboration network of high energy physicist (HepTh),¹ Digg follower graph (<http://www.isi.edu/~lerman/downloads/digg2009.html>), and a network of political blogs (<http://www-personal.umich.edu/~mejnetdata/>) are summarized below.

network	nodes	edges	$\langle k \rangle$	assortativity
HepTh	9,877	25,998	5.3	0.2283
Digg	25,454	175,892	13.8	0.1160
Political blogs	1,490	19,090	25.6	-0.2212

A. Friendship Paradox

Node degree distribution $p(k)$ gives the probability that a randomly chosen node in an undirected network has k neighbors or edges. Neighbor degree distribution $q(k)$ gives the probability that a randomly chosen edge in an undirected network is connected to a node of degree k . It is easy to demonstrate that average neighbor degree $\langle k \rangle_q$ is larger than average node degree $\langle k \rangle$. The difference between these quantities is

$$\langle k \rangle_q - \langle k \rangle = \sum_k \frac{k^2 p(k)}{\langle k \rangle} - \langle k \rangle = \frac{\langle k^2 \rangle - \langle k \rangle^2}{\langle k \rangle} = \frac{\sigma_k^2}{\langle k \rangle},$$

where σ_k is the standard deviation of the degree distribution $p(k)$. Since $\sigma_k \geq 0$, $\langle k \rangle_q - \langle k \rangle \geq$

0. This confirms that the friendship paradox, which says that average neighbor degree is larger than node’s own degree, has its origins in the heterogeneous degree distribution [14], and is more pronounced in networks with larger degree heterogeneity σ_k .

Heterogeneous degree distribution also contributes to nodes perceiving that their neighbors have more of some attribute than they themselves have — what is referred to as the generalized friendship paradox [13]. Let’s consider again a network where nodes have a binary attribute x . For convenience, we will refer to nodes with the attribute value $x = 1$ as active, and those with $x = 0$ as inactive. The probability that a random node is active is $P(x = 1) = \sum_k P(x = 1|k)p(k)$. The probability that a random neighbor is active is $Q(x = 1) = \sum_k P(x = 1|k)q(k)$. Using Bayes’ rule, this can be rewritten as

$$\begin{aligned} Q(x = 1) &= \sum_k \frac{P(x = 1, k)}{p(k)} \frac{k p(k)}{\langle k \rangle} \times \\ &= \sum_k \frac{P(x = 1, k)}{P(x = 1)} \frac{k P(x = 1)}{\langle k \rangle} \\ &= \frac{P(x = 1)}{\langle k \rangle} \sum_k k P(k|x = 1) \\ &= P(x = 1) \frac{\langle k \rangle_{x=1}}{\langle k \rangle}, \end{aligned}$$

where $\langle k \rangle_{x=1}$ is the average degree of active nodes. This quantity and the average degree $\langle k \rangle$ are related via the correlation coefficient $\rho_{kx} = \frac{P(x=1)}{\sigma_x \sigma_k} [\langle k \rangle_{x=1} - \langle k \rangle]$ (Eq. 1). Hence, the strength of the generalized friendship paradox is

$$Q(x = 1) - P(x = 1) = \rho_{kx} \frac{\sigma_x \sigma_k}{\langle k \rangle},$$

which is positive when node degree and attribute are positively correlated ($\rho_{kx} > 0$) and increases with this correlation [13].

B. Gaussian Approximation

The conditional probability $P(x' = 1|k')$ required to calculate the strength of the “majority illusion” using Eq. 4 can be specified analytically for some “well-behaved” degree distributions, such as scale-free distributions of the form $p(k) \sim k^{-\alpha}$ with $\alpha > 3$ or the Poisson distributions of the Erdős-Rényi random graphs in near-zero assortativity. In these cases, probability $P(x' = 1|k')$ can be acquired by approxi-

mating the joint distribution $P(x', k')$ as a multivariate normal distribution, and we have

$$\langle P(x'|k') \rangle = \langle P(x') \rangle + \rho_{kx} \frac{\sigma_x}{\sigma_k} (k' - \langle k \rangle),$$

resulting in

$$P(x' = 1|k') = \langle x \rangle + \rho_{kx} \frac{\sigma_x}{\sigma_k} (k' - \langle k \rangle).$$

Figure S1 reports the “majority illusion” in the same synthetic scale-free networks as Fig. 2, but with theoretical lines (dashed lines) calculated using the Gaussian approximation for estimating $P(x' = 1|k')$. The Gaussian approximation fits results quite well for the network with degree distribution exponent $\alpha = 3.1$. However, theoretical estimate deviates significantly from data in a network with a heavier-tailed degree distribution with exponent $\alpha = 2.1$. The approximation also deviate from the actual result when the network has very positive or negative assortativity. When this approximation is not applicable, we need to determine the joint distribution $P(x', k')$ empirically to calculate both ρ_{kx} and $P(x' = 1|k')$ for a specific joint distribution.

C. Influence of Network Structure

A network is not fully specified by its degree sequence and degree assortativity r . In fact, many different structures are possible with an identical degree sequence and r . These structural difference may affect the “majority illusion.” Here, we report some comparison of selected degree sequences in scale-free synthetic networks.

We generated scale-free networks with some degree sequence and assortativity. We then used edge rewiring to change network’s structure while keeping the assortativity r (and degree sequence) the same. Existing structural constraints restrict the range of assortativity values that a given degree sequence may attain. Thus there are fewer choices for extreme values of assortativity. Figure S2 reports the fraction of nodes in the paradox regime in these networks. Identical symbols are for the same value of assortativity. Results don’t change much in cases where structural constraints prevent varying the structure while keeping assortativity fixed. On the other hand, the fraction of nodes in the paradox regime can vary somewhat in the mid-assortativity range.

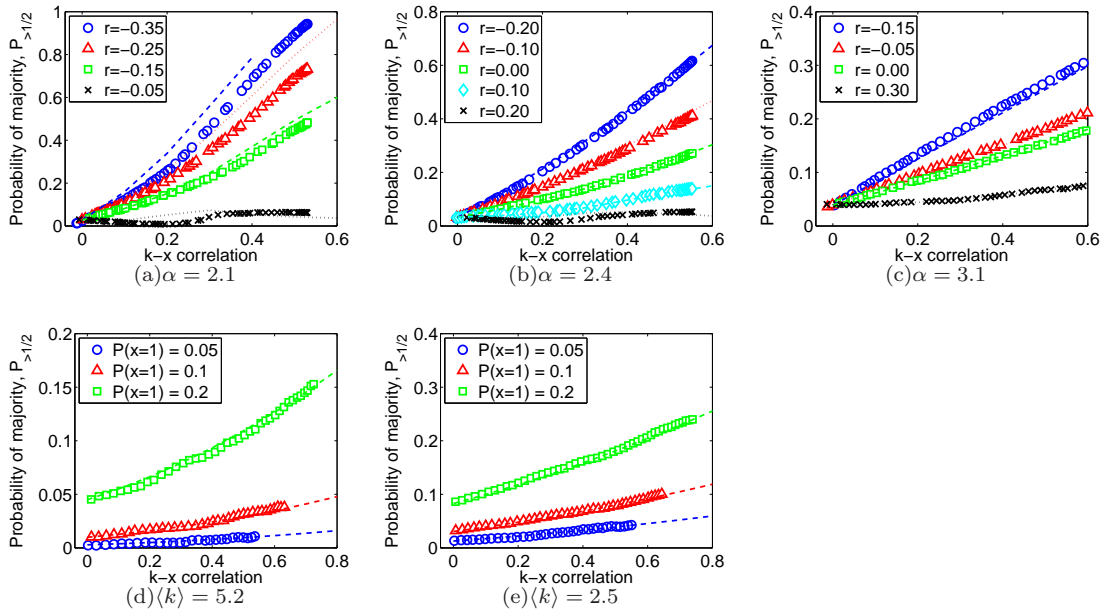


FIG. S1. Strength of the “majority illusion” with Gaussian approximation. Symbols show the empirically determined fraction of nodes in the paradox regime (same as in Fig. 2 and Fig. 3), while dashed lines show theoretical estimates using the Gaussian approximation.

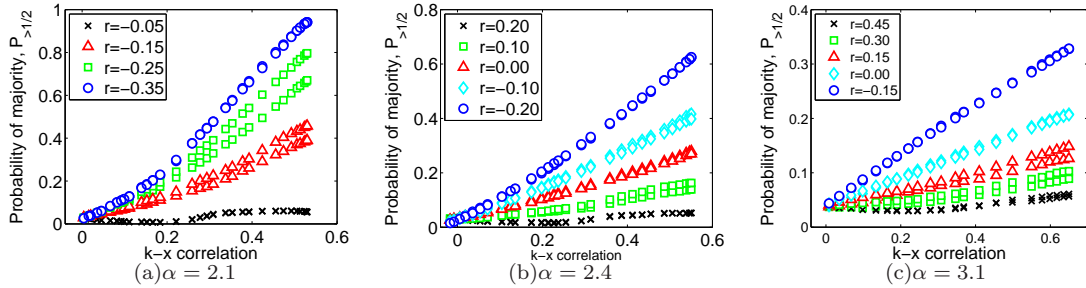


FIG. S2. Strength of the majority illusion in synthetic scale-free networks showing that networks with the same degree sequence and assortativity can manifest differing levels of the paradox. Identical symbols in the same plot are for networks with the same assortativity.