

# The Susceptibility Paradox in Online Social Influence

Luca Luceri\*, Jinyi Ye\*, Julie Jiang, Emilio Ferrara

Information Sciences Institute, University of Southern California  
 Thomas Lord Department of Computer Science, University of Southern California  
 lluceri@isi.edu, jinyiy@usc.edu, yioujian@usc.edu, emiliofe@usc.edu

## Abstract

Understanding susceptibility to online influence is crucial for mitigating the spread of misinformation and protecting vulnerable audiences. This paper investigates susceptibility to influence within social networks, focusing on the differential effects of influence-driven versus spontaneous behaviors on user content adoption. Our analysis reveals that influence-driven adoption exhibits high homophily, indicating that individuals prone to influence often connect with similarly susceptible peers, thereby reinforcing peer influence dynamics, whereas spontaneous adoption shows significant but lower homophily. Additionally, we extend the *Generalized Friendship Paradox* to influence-driven behaviors, demonstrating that users' friends are generally more susceptible to influence than the users themselves, *de facto* establishing the notion of *Susceptibility Paradox* in online social influence. This pattern does not hold for spontaneous behaviors, where friends exhibit fewer spontaneous adoptions. We find that susceptibility to influence can be predicted using friends' susceptibility alone, while predicting spontaneous adoption requires additional features, such as user metadata. These findings highlight the complex interplay between user engagement and characteristics in spontaneous content adoption. Our results provide new insights into social influence mechanisms and offer implications for designing more effective moderation strategies to protect vulnerable audiences.

## Introduction

On social media, social influence plays a significant role in disseminating information and shaping beliefs, attitudes, and behaviors. Politically, platforms like Twitter and Facebook are arenas for spreading biased information and fake news, which can potentially disrupt the integrity of voting events (González-Bailón et al. 2023; Stella, Ferrara, and De Domenico 2018). Economically, brands can sway consumers' purchasing decisions through targeted advertising and influencer partnerships (Aral and Walker 2012; Bakshy et al. 2012a). Socio-culturally, social media is pivotal in accelerating social movements and activism (Den Hond and De Bakker 2007; Mundt, Ross, and Burnett 2018).

Research has shown that online social influence can be a powerful tool for persuading behavioral change, especially

when messages are tailored to specific audiences (Oyibo and Vassileva 2019). Identifying groups highly susceptible to such influences is, therefore, essential. While many studies have examined various factors that contribute to susceptibility to influence—such as personality traits (Oyibo and Vassileva 2019), political leanings (Calvillo, Rutchick, and Garcia 2021), information credibility (Traberg and van der Linden 2022), and exposure frequency (Hassan and Barber 2021)—they often overlook the relational aspect and networked nature of influence dynamics. Consequently, the interconnectedness and potential peer influence within networks are not fully accounted for understanding and predicting susceptibility. For instance, how an individual's likelihood of being influenced relates to the susceptibility levels of their social connections remains underexplored.

## Contributions of This Work

In this paper, we present an empirical framework designed to understand, model, and predict users' susceptibility to influence on social media platforms. This framework takes into account the networked nature of social influence, with a focus on homophilous influence dynamics, and investigates how the *Generalized Friendship Paradox* manifests in influence-driven versus spontaneous behaviors, thereby establishing the concept of *Susceptibility Paradox*. In detail, we address the following Research Questions (RQs):

- RQ1: *Is susceptibility to influence homophilous within a social network?* We aim to investigate whether a user's susceptibility is correlated with the susceptibility of their friends.
- RQ2: *Does the Generalized Friendship Paradox extend to user susceptibility to influence?* Our objective is to explore whether a user's friends are typically more susceptible to influence than the user themselves.
- RQ3: *How effectively can a user's susceptibility be predicted by their friends' susceptibility?* We seek to measure to what extent a user's level of susceptibility can be predicted using information from their social network.

By leveraging two large-scale Twitter datasets covering political and public health discussions, our study reveals that susceptibility to influence is homophilous within social networks. We demonstrate that influence-driven behaviors exhibit greater homophily than spontaneous behaviors. Addi-

\*These authors contributed equally.

tionally, we find that users' friends are generally more susceptible to influence and less prone to spontaneous actions than the users themselves, supporting that the *Generalized Friendship Paradox* holds for influence-driven behavior but not for spontaneous action in friendship networks. These insights inform predictive models, which can predict users' rate of influence-driven adoptions by leveraging only their friends' susceptibility to influence. In contrast, spontaneous behaviors are more difficult to predict using solely friends' influence signals, suggesting a complex interplay of user engagement and preferences in spontaneous actions.

Our findings are consistent across the two large-scale Twitter datasets and provide crucial insights into diverse behaviors, vulnerabilities, and influence dynamics on social media platforms. This knowledge can inform network interventions designed to protect vulnerable audiences and mitigate risks associated with influence operations, propaganda, and misinformation campaigns.

## Related Work

### Susceptibility to Social Influence

Numerous studies have investigated the factors that make individuals susceptible to influence, categorizing these factors as either personal or external. From a *personal* standpoint, demographics such as age and gender play significant roles, with younger individuals and men often being more susceptible to influence when adopting new products (Aral and Walker 2012). Personality traits like neuroticism, openness, and conscientiousness are also recognized as consistent predictors of susceptibility to influence strategies (Oyibo and Vassileva 2019). In the political sphere, individuals with right-leaning ideologies are more susceptible to misinformation and deception (Calvillo, Rutchick, and Garcia 2021; Luceri et al. 2019; Pennycook and Rand 2019), particularly those who exhibit lower trust in scientific authority or a tendency towards conspiracy thinking (Saling et al. 2021).

*External* factors also play a role in susceptibility. Individuals are more likely to accept claims from sources they perceive as credible (Traberg and van der Linden 2022), and this tendency is further amplified when the persuasive message aligns with their partisanship or ideological views (Guess, Nyhan, and Reifler 2018; Moravec, Minas, and Dennis 2018). Additionally, the "illusory truth" effect occurs when repeated exposure to specific information makes a claim appear more believable, regardless of its veracity, thereby increasing its acceptance (Hassan and Barber 2021).

**Influence and Network Dynamics.** A user network may be an under-explored external factor that affects an individual's susceptibility to social influence. Individuals with similar characteristics often cluster within networks, leading to homophilous patterns of content exposure and adoption behaviors (McPherson, Smith-Lovin, and Cook 2001; Aral and Walker 2012). Furthermore, the structure of these networks often enhances the visibility of information among similar users, potentially fostering echo chambers (Cinelli et al. 2021). These effects are intensified by algorithmic recommendations that tailor content exposure to match the prevailing interests within one's network (Guess, Nyhan, and Rei-

fler 2020). Such biased exposure can contribute to the "majority illusion," where individuals mistakenly believe that the views or behaviors prevalent in their immediate social circle are representative of the general population (Lerman, Yan, and Wu 2016). A structural cause of this phenomenon is known as the "friendship paradox," which we will explore in the following section.

### The Generalized Friendship Paradox

Network structure can skew local observations. The friendship paradox states that, on average, your friends have more friends than you do (Feld 1991). This occurs because high-degree nodes (individuals with many friends) are more likely to be sampled when considering the friends of a random individual. The friendship paradox is more pronounced in networks with heavy-tailed degree distributions, where a small number of nodes have a large number of connections (Lerman, Yan, and Wu 2016).

The *Generalized Friendship Paradox* (GFP) extends the original concept to node attributes beyond the number of friends. It posits that the characteristics of the neighbors of a node in any network will, on average, be greater than those of the node itself (Eom and Jo 2014). For example, Hodas, Kooti, and Lerman (2013) examined the GFP in the context of online social networks, showing that users' friends are generally more active and receive more viral content on platforms like Twitter. Additionally, Higham (2019) found that individuals' friends are more central in social networks.

Research also suggests that the GFP could bias local perceptions and individuals' understanding of societal norms (Alipourfard et al. 2020), potentially amplifying or diminishing the spread of information and behaviors. Christakis and Fowler (2013) highlight that the adoption of health-related habits, such as smoking, drinking, and exercising, is influenced by perceptions of peers' actions. Similarly, an individual's likelihood of adopting information may be influenced by exposure to friends' sharing behaviors (such as the frequency and type of content they share), and might increase when content is popular within an ego network. Although this foundational work provides valuable insights into influence dynamics on social media, the relationship between a user's susceptibility to influence and that of their friends remains underexplored. In this work, our aim is to fill this gap by investigating how network structural properties—specifically homophily and the GFP—shape patterns of susceptibility to influence within social networks.

## Methodology

### Framework for Susceptibility Modeling

Susceptibility to online influence is commonly modeled in probabilistic terms, indicating the likelihood that a user's behavior and/or beliefs may be influenced by online social interactions (Zhou and Zafarani 2019; Romero, Meeder, and Kleinberg 2011; Goyal, Bonchi, and Lakshmanan 2010; Luceri, Braun, and Giordano 2018). The existing literature consistently highlights two primary components of susceptibility to influence: *Exposure* to content and the potential subsequent *Adoption* of that content (Aral and Walker 2012;

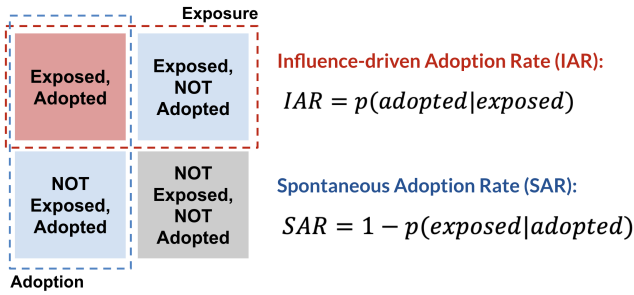


Figure 1: An illustration of our *Susceptibility Framework*.

Hoang and Lim 2016). We adhere to this formulation to construct our *Susceptibility Framework*, which aims to estimate users’ susceptibility to influence based on the dynamics between exposure and adoption.

**Exposure.** On social media, exposure to a piece of content typically refers to a user encountering or viewing the content. Measuring exposure from observational social media data can be challenging, as detailed user behavioral data, such as scrolling and clicking, remain largely inaccessible to academics. Additionally, it is increasingly difficult to get access to the structure of the user follower network and, therefore, to whose content each user is exposed to from those whom they follow. The situation is further complicated by *algorithmic feed curation*, i.e., recommendation algorithms displaying potentially relevant content to a user that does not originate from the users they are following (Guess, Nyhan, and Reifler 2020). Despite these challenges, several studies have approximated exposure by analyzing user interactions (Ferrara and Yang 2015; Rao, Morstatter, and Lerman 2022; Sasahara et al. 2021; Ye et al. 2024). The core assumption is that any form of engagement—whether retweeting, quoting, or replying—indicates a user’s potential exposure to another user’s content. These interactions imply that the user has not only viewed but also interacted with the content, thereby increasing the likelihood of subsequent exposures to that user’s content. The rationale is to leverage past interactions as a proxy to predict future exposure.

Therefore, we operationalize exposure as follows: If a user  $u_t$  interacts with another user  $u_s$  through a retweet, quote, or reply, we consider  $u_t$  exposed to all tweets posted by  $u_s$  after their first observed interaction.

**Adoption.** In this study, we define adoption as the action of sharing a URL, whether through an original tweet, a retweet, a quoted tweet, or a reply. We focus on URLs because they are commonly embedded in tweets and correspond to specific pieces of information, including multimodal content, such as images and videos, allowing for straightforward tracking of their diffusion and adoption within a social network. Sharing a URL, whether through original content or interactions, represents a form of information adoption, reflecting a user’s active decision to disseminate specific content. However, it is crucial to note that sharing or adopting a URL does not necessarily imply trust or endorsement of its content. Although retweeting is of-

ten seen as an endorsement (Boyd, Golder, and Lotan 2010; Metaxas et al. 2015), other sharing activities can serve to question or critique content (Hemsley et al. 2018).

**Susceptibility Metrics.** Building on our conceptualization of *Exposure* and *Adoption*, we develop a framework to estimate user *susceptibility* to influence (see Figure 1). Specifically, by analyzing user activity, we evaluated two crucial dimensions: potential exposure to content and potential adoption of content. It is important to acknowledge that our data is ultimately an incomplete rendition of the real world: There may be instances where users have shared or been exposed to content not included in our dataset.

Henceforth, we classify the relationship between a user and a piece of content into four distinct scenarios:

- (i) *Exposed and adopted*: the user viewed and then shared the content;
- (ii) *Exposed but not adopted*: the user viewed the content but did not share it;
- (iii) *Not exposed but adopted*: the user shared the content without any prior exposure;
- (iv) *Neither exposed nor adopted*: the user neither viewed nor ever shared the content.

Based on this framework, and in line with the definitions provided by (Aral and Walker 2012), we categorize adoption into two main types: *influence-driven adoption*, where the content is shared post-exposure, and *spontaneous adoption*, where the content is shared without any prior exposure. Recognizing these distinctions is crucial, particularly since adoptions can occur independently of direct influence from others. Motivated by this categorization, we introduce two susceptibility metrics: the *Influence-Driven Adoption Rate (IAR)* and the *Spontaneous Adoption Rate (SAR)*.

**Influence-driven Adoption Rate (IAR).** This metric calculates the ratio of URLs that a user adopted after being exposed to them to the total number of URLs the user was exposed to, yielding a value between 0 and 1. Mathematically, it is expressed as:

$$IAR = p(\text{adopted} | \text{exposed}) = \frac{|E_{u_t} \cap A_{u_t}|}{|E_{u_t}|},$$

where  $E_{u_t}$  denotes the set of all tweets containing URLs to which the user  $u_t$  was exposed, and  $A_{u_t}$  denotes the set of all tweets containing URLs that the user  $u_t$  actually adopted. Following the methodology in (Bakshy et al. 2012b; Ye et al. 2024), we implement a buffer period of two months for  $A_{u_t}$ . This period counts adoptions only if they occur after the first two months of our dataset, ensuring there is enough time to observe potential exposures before any subsequent adoptions.

**Spontaneous Adoption Rate (SAR).** This metric measures the proportion of URLs that a user adopts without prior exposure, with values ranging from 0 to 1. It is mathematically expressed as:

$$SAR = 1 - p(\text{exposed} | \text{adopted}) = 1 - \frac{|E_{u_t} \cap A_{u_t}|}{|A_{u_t}|},$$

where  $E_{u_t}$  and  $A_{u_t}$  are defined similarly to the IAR metric.

	ELECTION		COVID	
All tweets w/ URL	64M	76M		
Total users	4.5M	10.8M		
Target users	0.7M	1.0M		
Original tweets	7.2M	18.8M		
Retweets	27.5M	38.4M		
Quoted tweets	18.5M	10.9M		
Reply tweets	10.8M	7.9M		
$G(V, E)$	$ V $	$ E $	$ V $	$ E $
Interaction	49,794	117,384	73,960	163,974
Retweet	25,077	49,679	53,161	90,800
Mention	45,289	72,554	56,060	81,969

Table 1: Number of users and tweets, along with the number of nodes ( $|V|$ ) and edges ( $|E|$ ) in the friendship network  $G(V, E)$ .

The distributions of the IAR and SAR metrics for both datasets are shown in the *Appendix* (Figure 8). The strong negative Spearman correlation between IAR and SAR (ELECTION dataset:  $\rho = -0.855$ ,  $p < 0.001$ ; COVID dataset:  $\rho = -0.857$ ,  $p < 0.001$ ) indicates that users who frequently adopt content after being exposed to it (i.e., high IAR) tend to adopt content without prior exposure less frequently (i.e., low SAR), and vice versa. This observation supports the validity of our susceptibility framework and particularly these two metrics, in capturing distinct behaviors on social media—whether users are more influenced by external sources (high IAR) or more prone to independent (i.e., spontaneous) content sharing (high SAR). It is important to note that a high susceptibility level is typically associated with a high IAR and a low SAR. We maintain these two distinct metrics to more accurately reflect the nuances of influence dynamics and user behaviors.

### Data Collection and Curation

To ensure the robustness of our findings across different contexts, we conduct our experiments on two datasets from Twitter (now X), referred to as the ELECTION dataset and the COVID dataset. These datasets focus on two major events: the 2020 US Presidential Election and the COVID-19 global pandemic, covering a comprehensive range of topics related to political and public health discussions.

The ELECTION dataset (Chen, Deb, and Ferrara 2021) was collected beginning January 1st, 2020, and the COVID dataset (Chen, Lerman, and Ferrara 2020) beginning January 21st, 2020. For this research, we utilize data from the 12-month period of the same year, ending on December 31, 2020. Both datasets were collected via Twitter’s streaming API, which gathered a sample of all tweets using a list of manually curated, topic-specific keywords, accounts, and hashtags. The tweets were fetched in real-time, ensuring data completeness and avoiding well-known re-hydration issues (Pfeffer et al. 2023).

Since our methodology relies on tracking users’ exposed to and adoption of information via URLs embedded in shared tweets, in this study, we only include tweets that con-

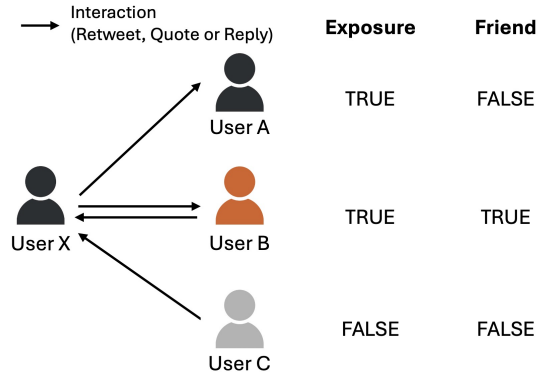


Figure 2: Differentiation between exposure and friendship. For User X, the unidirectional interaction with User A is modeled as *exposure*, while the bidirectional interaction with User B is modeled as *friendship*.

tain at least one URL. Detailed statistics on the number of tweets with URLs, the number of engaged users, and the distribution of tweet categories are provided in Table 1.

**Selecting “Target Users”.** We further refine our dataset to isolate a specific subset of *target users*. In line with the methodology of Nikolov, Flammini, and Menczer (2021), we establish a minimum inclusion criterion: we study only users who shared at least 10 URLs of any kind during the overall observation period (12 months). This threshold ensures that we focus on users with sufficient data to reliably assess their susceptibility. Applying this criterion, our data narrows down to 698,808 users for the ELECTION dataset and 1,047,566 users for the COVID dataset.

**Identifying “Friends” and Constructing Friendship Networks.** Friendship networks are social structures formed by interconnected individuals who are linked through mutual recognition, shared activities, interactions, or sentiments (Feld 1981; Zeggelink, Stokman, and Van De Bunt 1996). We define two target users as “friends” if they engage in active reciprocal interactions. This means both users need to actively engage with the other via retweets, quotes, or replies. Although the *followers* network could provide a more robust definition of friendship, such data is challenging to obtain using the Twitter API (Martha, Zhao, and Xu 2013). In lieu of the *followers* data, we use reciprocal interactions (retweet, quote, and reply) as a proxy of friendship, consistent with previous studies (Jiang, Ren, and Ferrara 2023; Rao, Morstatter, and Lerman 2022; Ye et al. 2024). This filtering results in 49,794 and 73,960 users from the ELECTION and COVID datasets, respectively. It is important to note that exposure to a user’s tweet, as defined by unidirectional interactions, does not necessarily indicate friendship with that user. However, the reverse is true: if two users have mutual interactions and are considered friends, they are exposed to each other’s tweets. Figure 2 illustrates the distinction between exposure and friendship.

We construct three distinct friendship networks based on the types of interactions: *Interaction*, *Retweet*, and *Men-*

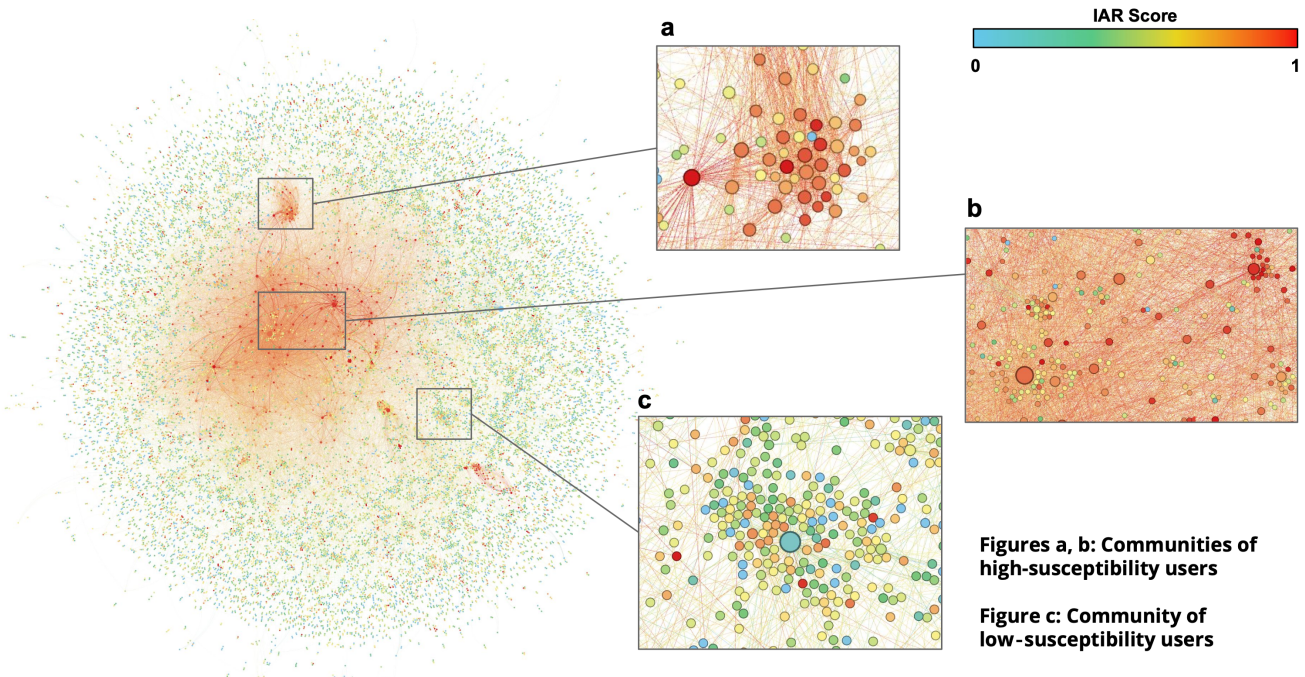


Figure 3: Homophily of influence-driven adoption rates (IAR) in the *Retweet* network of the ELECTION dataset. Nodes are colored by IAR score (blue = low susceptibility, red = high susceptibility), scaled in size by node degree, and connected by edges representing mutual retweet interactions. Panels (a–c) display communities in which connected users share similar IAR scores, illustrating homophily.

*tion* networks. The *Interaction* network includes all forms of interaction between users, such as retweets, quotes, and replies. In the *Retweet* network, connections capture retweets only, whereas the *Mention* network includes interactions through replies or quotes. We differentiate the *Retweet* and *Mention* networks due to their varying levels of social implications and endorsement, a distinction further explored in the “Adoption” section. In Table 1, we present statistical descriptors of each network in the two datasets.

## Results

### RQ1: Susceptibility and Homophily

We first examine whether susceptibility to influence exhibits homophily in a friendship network. To this end, we observe whether a user’s susceptibility metrics are related to their friends. Following Jiang et al. (2023), we measure homophily by calculating the Pearson correlation between each user’s susceptibility score and the weighted average susceptibility scores of all their friends. This method enables us to consider the frequency of user interactions and accounts for cases where users may have varying levels of influence across different connections. A coefficient of  $-1$  suggests that users preferentially connect with others who have opposing susceptibility scores, while a coefficient of  $1$  indicates a preference for connections between users with similar susceptibility scores.

The results presented in Figure 4 show that susceptibil-

ity is homophilous with respect to both the IAR and SAR scores, as indicated by the statistically significant and positive weighted average correlations. The results are consistent across all types of networks in both datasets. To further confirm the reliability of these findings, we compared our results with those from two null models. The first null model maintains the original network’s degree distribution by randomly selecting and swapping the endpoints of edge pairs. The second null model preserves only the number of edges by randomly reassigning neighbors to nodes at one end of each edge. Both null models result in weighted average correlation measures that are not significant, therefore confirming that our results on the real data cannot be attributed to statistical flukes. The interested reader can refer to the *Appendix* for further details.

Our analysis reveals that there is indeed homophily in user susceptibility. Moreover, this homophily, which refers to the tendency for individuals to form connections with others who are similar to them, appears to be more prevalent in influence-driven adoption (IAR) than in spontaneous adoption (SAR). Figure 3 visualizes the distribution of IAR scores in the *Retweet* network of the ELECTION dataset. This suggests that within social networks, individuals connected by social ties tend to exhibit similar inclinations toward adopting the behaviors of their friends. This is likely due to the stronger impact of peer influence within close social circles, where interactions are more frequent and intense (De Klepper et al. 2010). As a result, a reinforcing loop

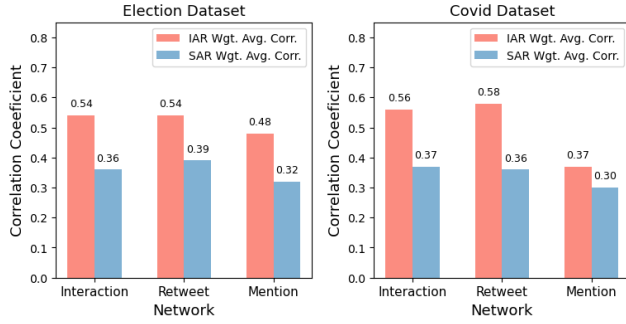


Figure 4: The susceptibility metrics (IAR and SAR) exhibit homophily within different friendship networks, as evidenced by the weighted average correlation (Wgt. Avg. Corr.) between users and their friends. All correlation coefficients are significant ( $p < 0.001$ ).

emerges: individuals in tight-knit groups are more likely to adopt similar behaviors or attitudes due to the influence of their peers.

**Summary.** Susceptibility to influence is homophilous. In particular, influence-driven behavior (as measured by IAR) appears to be more homophilous than spontaneous-driven activity (as measured by SAR).

## RQ2: The Generalized Friendship Paradox of Susceptibility to Influence

When examining whether the *Generalized Friendship Paradox* (GFP) applies to susceptibility to influence, we seek to determine whether individuals’ friends are generally more susceptible to influence than they are. Consistent with Eom and Jo (2014), in this study, we formulate the GFP at both the *individual* and *network* levels. At the individual-level, a node exhibits GFP if its susceptibility metric is lower than the average of its friends. At the network-level, the GFP holds if the average susceptibility metric of all nodes is less than the average susceptibility metric of their respective friends.

**Individual-level GFP.** The individual-level GFP holds for one user  $i$  if the following condition is satisfied:

$$s_i < \frac{\sum_{j \in \mathcal{N}_i} s_j}{k_i}, \quad (1)$$

where  $s_i$  is the susceptibility metric of user  $i$ ,  $\mathcal{N}_i$  represents the set of  $i$ ’s friends, and  $k_i$  is the degree of  $i$ . Essentially, it calculates whether a user’s susceptibility is less than the average of their friends. Based on this formulation, we define the paradox holding probability  $p(k, s)$  as the proportion of nodes with degree  $k$  and susceptibility metric  $s$  that satisfies the condition in Eq. (1).

**Network-level GFP.** At the network-level, the GFP examines the phenomenon across the entire network. It is determined by comparing the average susceptibility metric of every node  $\langle s \rangle$  with the average susceptibility metric of all

their friends  $\langle s \rangle_{nn}$ , computed as follows:

$$\langle s \rangle_{nn} = \frac{\sum_{i=1}^N k_i s_i}{\sum_{i=1}^N k_i}, \quad (2)$$

where a node  $i$  with degree  $k_i$  is counted as a neighbor  $k_i$  times. If the average susceptibility metric of a node’s friends exceeds the node’s own average susceptibility metric ( $\langle s \rangle < \langle s \rangle_{nn}$ ), then the GFP holds at the network-level.

**GFP: Comparing SAR vs. IAR.** Figure 5 depicts the paradox holding probability  $p(k, s)$  as a function of degree  $k$  and susceptibility metric  $s$  using the *Interaction* network in the COVID dataset. Similar findings from the ELECTION dataset are presented in the *Appendix*. This three-dimensional representation reveals that, with a constant degree  $k$ ,  $p(k, s)$  decreases as  $s$  (IAR or SAR) increases, aligning with findings from Eom and Jo (2014). This implies that users with very low susceptibility  $s$  tend to connect with those having higher  $s$  values, resulting in a  $p(k, s)$  close to 1. Conversely, users with very high susceptibility values exhibit a  $p(k, s)$  approaching 0. Users across this spectrum exhibit a consistent trend, wherein  $p(k, s)$  remains close to 1 when  $s$  and  $k$  are strongly, positively correlated. This effect is particularly evident for IAR, indicating a specific pattern where the paradox holds.

Table 2 presents the empirical results of the individual- and network-level GFP across three friendship networks constructed from the ELECTION dataset. Similar results are observed in the COVID dataset. For each susceptibility metric  $s$  (IAR or SAR), we calculate the average paradox holding probability at the individual-level  $P$ , integrating  $p(k, s)$  over  $k$  and  $s$ , following Eom and Jo (2014), and the Spearman correlation coefficient between node degree  $k$  and susceptibility metric  $s$ , denoted as  $\rho_{ks}$ . For the network-level GFP, we compare the node average  $\langle s \rangle$  and their friends’ average  $\langle s \rangle_{nn}$ . We also employ two baseline models for comparison based on the null models described above.

For IAR, the value of  $P$  exceeds 0.5 in every friendship network, indicating that the GFP holds at the individual-level. Furthermore, the condition  $\langle s \rangle < \langle s \rangle_{nn}$  is consistently satisfied across networks, indicating that the GFP also holds at the network-level. This implies that, on average, individuals’ friends are more likely to adopt content they are exposed to than the individuals themselves. In contrast, for SAR,  $P$  consistently falls below 0.5, and  $\langle s \rangle$  is always higher than  $\langle s \rangle_{nn}$ . These observations suggest that the GFP does not hold for spontaneous adoption. In simpler terms, individuals are less inclined to adopt unseen content compared to their friends.

Moreover, it is worth noting that while the GFP holds for IAR, it is less pronounced compared to the two null models. The discrepancy between individual IAR and that of their friends is smaller in the observed network than in a randomly constructed network. For example, at the individual-level, in the *Interaction* network of the ELECTION dataset, a random baseline model predicts that in 79.8% of cases, an individual’s friends, on average, would exhibit higher IAR than the individual. However, the actual observed rate in the real-world network is only 54.1%. At the network-level, the

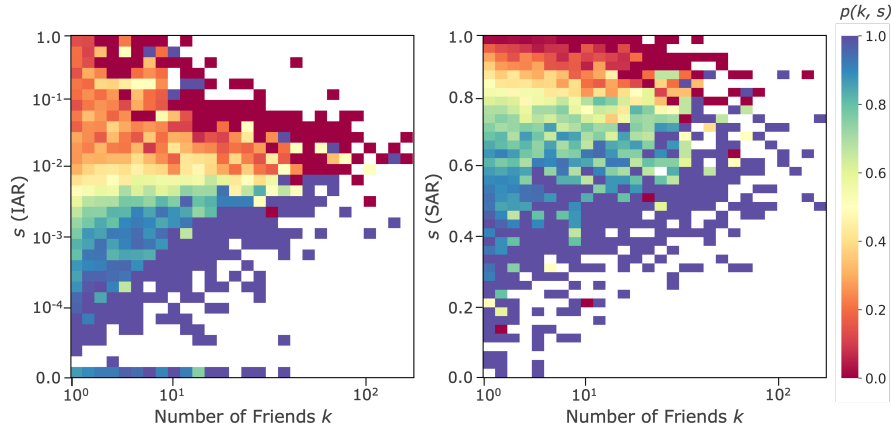


Figure 5: Paradox holding probability  $p(k, s)$  at a varying degree  $k$  within the *Interaction* network of the COVID dataset.

Network	Metric $s$	$\rho_{ks}$	Individual-level GFP			$\langle s \rangle$		Network-level GFP		
			Real-Network $P$	Baseline1 $P$	Baseline2 $P$			Real-Network $\langle s \rangle_{nn}$	Baseline1 $\langle s \rangle_{nn}$	Baseline2 $\langle s \rangle_{nn}$
Interaction	IAR	0.179	0.541	0.672	0.798	0.007	<	0.008	0.019	0.028
	SAR	-0.166	0.409	0.326	0.249	0.880	>	0.868	0.810	0.751
Retweet	IAR	0.298	0.501	0.662	0.824	0.013	<	0.014	0.022	0.040
	SAR	-0.294	0.425	0.355	0.250	0.821	>	0.809	0.722	0.670
Mention	IAR	0.111	0.544	0.623	0.697	0.005	<	0.006	0.009	0.015
	SAR	-0.099	0.383	0.354	0.325	0.896	>	0.880	0.852	0.834

Table 2: Empirical results of the *Generalized Friendship Paradox* (GFP) in the three friendship networks of the ELECTION dataset. For each susceptibility metric  $s$  (both IAR and SAR), we calculate the Spearman correlation coefficient  $\rho_{ks}$  between degree  $k$  and  $s$  ( $p < 0.001$  for all  $\rho_{ks}$ ), average paradox holding probability  $P$ , and the average susceptibility  $\langle s \rangle$  of nodes and their friends  $\langle s \rangle_{nn}$ .

differences between  $\langle s \rangle$  and  $\langle s \rangle_{nn}$  are also smaller in the actual network. This discrepancy between random baselines and actual networks is similarly observed for SAR.

These observations corroborate the study by Lerman, Yan, and Wu (2016), which demonstrates that the “majority illusion,” a phenomenon resulting from the friendship paradox, is more intense within networks characterized by heterogeneous degree distributions and disassortative structures. Our observations of real-world social media networks, which exhibit greater homophily compared to random baselines, indeed reveal a reduced paradox holding probability. This implies that the structural attributes of a network can significantly impact the manifestation of the GFP.

**Summary.** Our findings indicate that the *Generalized Friendship Paradox* holds for influence-driven behavior but not for spontaneous actions on social networks. This suggests that users’ friends are generally more susceptible to influence and less prone to spontaneous sharing activities than the users themselves.

### RQ3: Predicting Susceptibility to Online Influence

As our findings in RQ1 and RQ2 provide valuable insights into the relationship between an individual’s susceptibility

and that of their friends, we explore the potential to predict a user’s susceptibility based on the susceptibility levels of their friends. Furthermore, we evaluate whether incorporating additional features, such as user and friend metadata along with network properties, can enhance prediction accuracy. To accomplish this, we employ both simple linear regression and random forest models.

**Linear Regression Model.** First, we assess the effectiveness of predicting a user’s IAR (resp. SAR) based on the average IAR (resp. SAR) of their friends. The linear regression model is expressed as:

$$s_{user} = \beta_0 + \beta_1 \times \bar{s}_{friend} + \epsilon, \quad (3)$$

where  $s_{user}$  and  $\bar{s}_{friend}$  refer to a user IAR (resp. SAR) and their friends’ average IAR (resp. SAR).

Table 3 shows the results of the simple linear regression analysis. From Table 3, we note that both IAR and SAR models yield positive regression coefficients ( $\beta_1$ ), suggesting a positive relationship between a user’s susceptibility metrics and those of their friends. Notably, the  $R^2$  values for the IAR models are consistently higher than those of the SAR models, indicating that influence-driven adoptions are more accurately explained by friends’ susceptibility compared to spontaneous adoptions.

Dataset	$s_{user}$	$\bar{s}_{friend}$	$\beta_0$	$\beta_1$	$R^2$
ELECTION	IAR	IAR	0.002	0.622	0.363
ELECTION	SAR	SAR	0.550	0.377	0.137
COVID	IAR	IAR	0.005	0.533	0.283
COVID	SAR	SAR	0.537	0.392	0.134

Table 3: Summary of the linear regression results for predicting user susceptibility metrics based on the friends’ average susceptibility metrics. All regression coefficients reported are significant ( $p < 0.001$ ).

Additionally, we also attempt to predict user IAR using friends’ SAR and vice versa; however, these models produce  $R^2$  values all below 0.05 in both datasets, highlighting that although strongly negative-correlated, IAR and SAR do not adequately predict each other.

**Random Forest Model.** To assess whether a more advanced model equipped with a broader range of features can enhance predictive accuracy, we employ an interpretable ensembling machine learning algorithm. This approach also aims to enhance our understanding of the factors that impact users’ susceptibility to social influence. Specifically, we utilize a random forest regression model that incorporates not only susceptibility metrics but also network properties and user metadata. This feature set includes friends count, followers count, favorites count (the total number of tweets a user has liked), and statuses count (the total number of tweets a user has posted), along with network properties from the *Interaction* network such as degree centrality, eigenvector centrality, and clustering coefficient.

We divide users into training and testing sets and use the variance inflation factor (VIF) to assess multicollinearity among the features. In both the ELECTION and COVID datasets, all VIF values are below 5, indicating no significant multicollinearity issues (James et al. 2013). We optimize the random forest regressor using a randomized search over the hyperparameter space to efficiently predict user IAR and SAR scores. The randomized search is configured to sample 100 different parameter settings using 5-fold cross-validation. After selecting the optimal parameters (see the *Appendix* for details), we assess each model’s performance on the test set using the  $R^2$  metric.

Table 4 provides an overview of the results. The random forest regressor shows promising performance, particularly for the IAR metric. It slightly outperforms the linear regressor, as indicated by the  $R^2$  difference labeled as  $\Delta R^2$  in Table 4. This difference is more pronounced for the SAR metric, while it remains quite small (less than 0.05) for the IAR metric, suggesting that IAR can be effectively predicted using solely friends’ IAR, as previously demonstrated in the linear regression model. An analysis of feature importance further supports this intuition.

Figure 6 illustrates the feature importance (top 5 features) for the random forest regressor in predicting IAR and SAR within the COVID dataset. Notably, *friends’ IAR* emerges as the most significant feature for predicting IAR, with other features proving substantially less relevant. This

<i>Susceptibility Metric</i>	ELECTION		COVID	
	IAR	SAR	IAR	SAR
Random Forest $R^2$	0.406	0.389	0.332	0.312
Linear Regression $R^2$	0.363	0.137	0.283	0.134
$\Delta R^2$	0.043	0.252	0.049	0.178

Table 4: Summary of the performance of the random forest and linear regressors.  $R^2$  denotes the coefficient of determination and  $\Delta R^2$  represents the  $R^2$  difference between random forest and linear regression models.

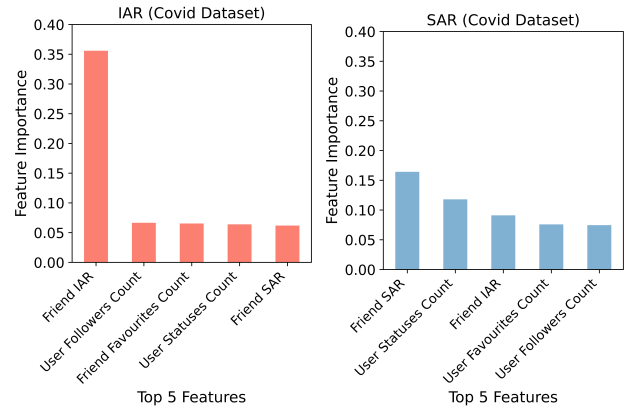



Figure 6: The top 5 most important features from the random forest model in the COVID dataset for predicting IAR and SAR.

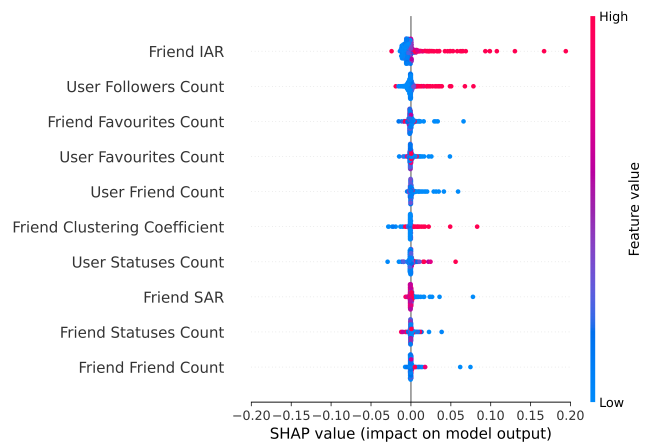
finding confirms our previous intuition and remains consistent in the ELECTION dataset. For SAR predictions, while *friends’ SAR* is the most critical feature, additional factors such as *statuses count*, *friends’ IAR*, and *favorites count* also contribute. This pattern holds in the ELECTION dataset, where *statuses count* is the most important feature.

**Output Explainability.** To further elucidate the results from the feature importance analysis and quantify the specific contribution of each feature to individual predictions, we compute the SHAP (SHapley Additive exPlanations) values (Lundberg and Lee 2017) of the dependent variables predicted in the test set. SHAP values provide a nuanced measure of the impact of each feature on individual instances. Each feature is assigned a SHAP value per instance, which can be aggregated to offer a comprehensive view of the direction and magnitude of its impact on the model’s predictions. Figure 7 shows a SHAP summary plot for the SAR and IAR metrics of the COVID dataset (additional results are detailed in the *Appendix*). On the y-axis, features are listed in order of impact, while the x-axis displays the SHAP values, indicating the change in log odds caused by each feature. Points represent individual data instances, colored red (resp. blue) for high (resp. low) feature values.

As shown by Figure 7, the most impactful features affecting the SAR predictions in the COVID dataset are *friends’ SAR*, *user statuses count*, and *user clustering coefficient*. The features *friends’ SAR* and *user statuses count* have signif-



(a) SHAP value distributions for SAR predictions.



(b) SHAP value distributions for IAR predictions.

Figure 7: SHAP value distributions for SAR (a) and IAR (b) predictions in the COVID dataset. Each distribution illustrates the direction and magnitude of the impact of individual features on the model’s output, with features ordered by importance from top to bottom.

icant opposing impacts on the predictions, indicating that these features can notably increase or decrease the SAR value. Specifically, high values of *friends’ SAR* tend to increase the SAR score, while higher values of *user statuses count* tend to decrease it. This relationship between *user statuses count* and SAR suggests that a high volume of shared tweets correlates with a small rate of spontaneous adoptions. Interestingly, the *user clustering coefficient* also emerges as a significant factor, where higher clustering coefficients typically decrease SAR, possibly indicating that users with more cohesive social networks may experience fewer spontaneous adoptions. This finding is consistent in the ELECTION dataset (see the *Appendix*), where user statuses count, friends’ SAR, and user favorites count are the most impactful features.

For IAR prediction (see Figure 7), the most impactful feature is *friends’ IAR*, with high values of *friends’ IAR* significantly increasing the IAR score, supporting our previous results related to homophily and GFP in IAR.

**Summary.** Users’ level of influence-driven adoption can be effectively predicted by solely leveraging friends’ IAR. In contrast, the predictability of a user’s SAR is significantly enhanced by incorporating additional features, such as counts of tweets and likes. Explainable models reveal that SAR is affected by a wider array of factors, suggesting a more complex interplay of user engagement and characteristics in the spontaneous adoption of online content.

### Sensitivity Analysis

To address potential sampling biases related to user network construction, we conducted a sensitivity analysis by testing different thresholds for identifying target users. Specifically, we applied thresholds of 5 and 20 shared URLs, in addition to the previously used threshold of 10 URLs, to select the set of target users under scrutiny. This approach allowed

us to assess whether varying levels of user activity affected our results. The results, detailed in the *Appendix*, remained consistent across the different thresholds. Both the IAR and SAR metrics exhibit homophily. Additionally, the *Generalized Friendship Paradox* (GFP) holds for influence-driven behaviors (IAR) but not for spontaneous actions (SAR), further validating our findings and conclusions.

## Discussion and Conclusions

This study, based on the analysis of millions of tweets, explores the networked dynamics of user susceptibility to online influence from the dual perspective of influence-driven and spontaneous adoption. We present four key insights drawn from our research.

**Finding 1: Susceptibility is homophilous.** We find that susceptibility is homophilous in social networks, as evidenced by the positive correlations between users’ adoption rates (either influence-driven or spontaneous) and those of their friends. This indicates that influenceable individuals often connect with similarly susceptible peers. Spontaneous adoption shows significant but lower homophily.

**Finding 2a: The Susceptibility Paradox holds for influence-driven behavior.** We observe the *Generalized Friendship Paradox* (GFP) for user susceptibility to influence, termed the *Susceptibility Paradox*, across two dimensions: First, individuals in a social network tend to be connected to others who are more likely to adopt content due to prior exposure, as reflected in a higher IAR. Second, users’ friends engage in less spontaneous content adoption with respect to the users themselves, as indicated by their lower SAR. Collectively, these observations suggest that “your friends are generally more susceptible to social influence than you.” Additionally, we observe that IAR and SAR are negatively correlated, suggesting that users with

higher influence-driven adoption are less likely to share content spontaneously.

**Finding 2b: The Susceptibility Paradox is less pronounced when homophily is present.** While we observe that the GFP holds for IAR but not for SAR, both IAR and SAR demonstrate homophily. These homophilous patterns indicate that individuals are prone to connect with others who exhibit similar rates of influence-driven and spontaneous adoption. This reveals two countervailing forces: while an individual’s friends generally have a higher IAR but a lower SAR compared to the individual, their overall susceptibility scores tend to remain similar. Comparing the individual- and network-level GFP in homophilous real-world networks against non-homophilous baseline networks, our findings suggest that the GFP is less pronounced in homophilous networks, where high-IAR (or SAR) nodes tend to connect with other high-IAR (or SAR) nodes. This corroborates the findings of Lerman, Yan, and Wu (2016), which indicate that the “majority illusion” is more pronounced in heterogeneous, disassortative networks.

**Finding 3: IAR can be predicted by friends’ IAR and is more predictable than SAR.** Our analysis indicates that IAR can be predicted solely based on friends’ IAR, whereas SAR is influenced by a broader range of factors. This distinction adds depth to our understanding of susceptibility to influence. On the one hand, our results suggest that influence-driven behaviors (IAR) tend to follow more predictable patterns. This likely stems from the strong relationship between influence-driven adoption and topological properties, such as homophily (RQ1) and GFP (RQ2), observed in this study. Also, research has shown that targeted marketing and influence campaigns can be effectively strategized based on previous adoption patterns of a consumer’s network (Hill, Provost, and Volinsky 2006). On the other hand, spontaneous behaviors (SAR) are influenced by a myriad of internal and external, often stochastic factors, making them less predictable.

**Implications.** Our research offers practical insights for social media platforms, researchers, policymakers, and marketers. The observed homophily in susceptibility suggests that targeted interventions in fields like public health and marketing can effectively target influenceable audiences. By identifying these cohesive, vulnerable groups, intervention strategies can be tailored to mitigate the spread of harmful content and encourage positive behaviors more efficiently (Centola 2011). Further, the observation that the *Generalized Friendship Paradox* applies to influence-driven adoption could inspire efficient sampling algorithms to identify highly susceptible populations (Alipourfard et al. 2020).

Predicting users’ IAR and SAR could be valuable in two key scenarios. First, for new users joining an existing social network, their IAR and SAR can be approximated based on interactions with more longevous users, leveraging metrics and attributes of their friends. This allows for an early assessment of a new user’s susceptibility to influence. Second, for less active users or those who rarely share URLs—falling below the threshold required to compute their IAR or

SAR—we can infer their susceptibility leveraging the susceptibility metrics of their friends. Overall, this enables social media providers to account for the susceptibility of users with limited activity. Since IAR is more predictable, and individuals with high IAR are more influenced by their social connections, platforms can strategically target vulnerable audiences by nudging them with accurate, fact-checked content to help prevent the endorsement and spread of misinformation (Pennycook et al. 2021).

Future research will extend our susceptibility framework across various media and content modalities. Since the framework is based solely on the dynamics between exposure and adoption, it is platform-agnostic, allowing us to analyze susceptibility on different social media networks. This could involve studying not only URLs in general, but also disentangling specific content types, such as videos, images, and memes. Additionally, our research will explore application scenarios where predicting user susceptibility is critical, such as modeling the impact of coordinated influence operations (Luceri et al. 2024) and exposure to AI-generated disinformation (Augenstein et al. 2024) on content adoption. Validating our model in these domains will further strengthen its real-world applicability and relevance.

**Limitations.** We acknowledge several limitations in our research. First, our study draws on datasets specifically related to the 2020 US election and the COVID-19 pandemic, capturing only a subset of users over an extended but limited period. Consequently, generalizing our results to broader user networks, different contexts, or other social media platforms necessitates further validation. Second, we infer exposure based on observable user interactions, which might not comprehensively capture exposures within Twitter and across different media. Nonetheless, this method for assessing exposure was tested and validated in prior research (Ferrara and Yang 2015; Rao, Morstatter, and Lerman 2023; Ye et al. 2024). Third, our network construction method excludes less active users (those sharing fewer than 5, 10, or 20 URLs), which may limit the generalizability of the observed homophily and friendship paradox. Additionally, since we use URL sharing as the adoption criterion, there may be users who rarely share URLs, further restricting the size of the target users under scrutiny. However, our sensitivity analysis shows consistent findings (see related section in the *Appendix*). Varying the target user thresholds (sharing 5, 10, or 20 URLs) did not affect the significance of our results, indicating the robustness of our findings despite potential selection bias. Fourth, although friends’ IAR explains 36% and 28% of the variation of user IAR in ELECTION and COVID datasets, respectively, we recognize that social influence is a complex phenomenon and may be influenced by various unobserved personal, societal, and cultural factors. Despite these limitations, the consistency of our results across multiple intertwined analyses on distinct datasets supports the reliability and robustness of our findings.

## Acknowledgments

Work supported in part by DARPA (contract #HR001121C0169) and the NSF (award #2331722).

## References

- Alipoufard, N.; Nettasinghe, B.; Abeliuk, A.; Krishnamurthy, V.; and Lerman, K. 2020. Friendship paradox biases perceptions in directed networks. *Nature Communications*, 11(1): 707.
- Aral, S.; and Walker, D. 2012. Identifying influential and susceptible members of social networks. *Science*, 337(6092): 337–341.
- Augenstein, I.; Baldwin, T.; Cha, M.; Chakraborty, T.; Ciampaglia, G. L.; Corney, D.; DiResta, R.; Ferrara, E.; Hale, S.; Halevy, A.; et al. 2024. Factuality challenges in the era of large language models and opportunities for fact-checking. *Nature Machine Intelligence*, 1–12.
- Bakshy, E.; Eckles, D.; Yan, R.; and Rosenn, I. 2012a. Social influence in social advertising: evidence from field experiments. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 146–161.
- Bakshy, E.; Rosenn, I.; Marlow, C.; and Adamic, L. 2012b. The role of social networks in information diffusion. In *WWW 2012*, 519–528.
- Boyd, D.; Golder, S.; and Lotan, G. 2010. Tweet, tweet, retweet: Conversational aspects of retweeting on Twitter. In *HICSS 2010*, 1–10. IEEE.
- Calvillo, D. P.; Rutchick, A. M.; and Garcia, R. J. 2021. Individual differences in belief in fake news about election fraud after the 2020 US election. *Behavioral Sciences*, 11(12): 175.
- Centola, D. 2011. An experimental study of homophily in the adoption of health behavior. *Science*, 334(6060): 1269–1272.
- Chen, E.; Deb, A.; and Ferrara, E. 2021. #Election2020: The first public Twitter dataset on the 2020 US presidential election. *Journal of Computational Social Science*, 5: 1–18.
- Chen, E.; Lerman, K.; and Ferrara, E. 2020. Tracking social media discourse about the COVID-19 pandemic: Development of a public coronavirus Twitter data set. *JMIR Public Health and Surveillance*, 6(2): e19273.
- Christakis, N. A.; and Fowler, J. H. 2013. Social contagion theory: Examining dynamic social networks and human behavior. *Statistics in Medicine*, 32(4): 556–577.
- Cinelli, M.; De Francisci Morales, G.; Galeazzi, A.; Quattrociocchi, W.; and Starnini, M. 2021. The echo chamber effect on social media. *PNAS*, 118(9): e2023301118.
- De Klepper, M.; Sleenbos, E.; Van de Bunt, G.; and Agneessens, F. 2010. Similarity in friendship networks: Selection or influence? The effect of constraining contexts and non-visible individual attributes. *Social Networks*, 32(1): 82–90.
- Den Hond, F.; and De Bakker, F. G. 2007. Ideologically motivated activism: How activist groups influence corporate social change activities. *Academy of Management Review*, 32(3): 901–924.
- Eom, Y.-H.; and Jo, H.-H. 2014. Generalized friendship paradox in complex networks: The case of scientific collaboration. *Scientific Reports*, 4(1): 1–6.
- Feld, S. L. 1981. The focused organization of social ties. *American Journal of Sociology*, 86(5): 1015–1035.
- Feld, S. L. 1991. Why your friends have more friends than you do. *American Journal of Sociology*, 96(6): 1464–1477.
- Ferrara, E.; and Yang, Z. 2015. Measuring emotional contagion in social media. *PLoS ONE*, 10(11): e0142390.
- González-Bailón, S.; Lazer, D.; Barberá, P.; Zhang, M.; Allcott, H.; Brown, T.; Crespo-Tenorio, A.; Freelon, D.; Gentzkow, M.; Guess, A. M.; et al. 2023. Asymmetric ideological segregation in exposure to political news on Facebook. *Science*, 381(6656): 392–398.
- Goyal, A.; Bonchi, F.; and Lakshmanan, L. V. 2010. Learning influence probabilities in social networks. In *WSDM 2010*, 241–250.
- Guess, A.; Nyhan, B.; and Reifler, J. 2018. Selective exposure to misinformation: Evidence from the consumption of fake news during the 2016 US presidential campaign.
- Guess, A. M.; Nyhan, B.; and Reifler, J. 2020. Exposure to untrustworthy websites in the 2016 US election. *Nature Human Behaviour*, 4(5): 472–480.
- Hassan, A.; and Barber, S. J. 2021. The effects of repetition frequency on the illusory truth effect. *Cognitive Research: Principles and Implications*, 6(1): 38.
- Hemsley, J.; Stromer-Galley, J.; Semaan, B.; and Tanupabrungsun, S. 2018. Tweeting to the target: Candidates’ use of strategic messages and @ mentions on Twitter. *Journal of Information Technology & Politics*, 15(1): 3–18.
- Higham, D. J. 2019. Centrality-friendship paradoxes: When our friends are more important than us. *Journal of Complex Networks*, 7(4): 515–528.
- Hill, S.; Provost, F.; and Volinsky, C. 2006. Network-based marketing: Identifying likely adopters via consumer networks. *Statistical Science*, 21(2): 256–276.
- Hoang, T.-A.; and Lim, E.-P. 2016. Tracking virality and susceptibility in social media. In *CIKM 2016*, 1059–1068.
- Hodas, N.; Kooti, F.; and Lerman, K. 2013. Friendship paradox redux: Your friends are more interesting than you. In *ICWSM 2013*, 225–233.
- James, G.; Witten, D.; Hastie, T.; Tibshirani, R.; et al. 2013. *An introduction to statistical learning*, volume 112. Springer.
- Jiang, J.; Luceri, L.; Walther, J. B.; and Ferrara, E. 2023. Social approval and network homophily as motivators of online toxicity. *arXiv preprint arXiv:2310.07779*.
- Jiang, J.; Ren, X.; and Ferrara, E. 2023. Retweet-BERT: Political leaning detection using language features and information diffusion on social networks. In *ICWSM 2023*, 459–469.
- Lerman, K.; Yan, X.; and Wu, X.-Z. 2016. The “majority illusion” in social networks. *PLoS One*, 11(2): e0147617.
- Luceri, L.; Braun, T.; and Giordano, S. 2018. Social influence (deep) learning for human behavior prediction. In *CompleNet 2018*, 261–269. Springer.

- Luceri, L.; Deb, A.; Badawy, A.; and Ferrara, E. 2019. Red bots do it better: Comparative analysis of social bot partisan behavior. In *Companion Proceedings of the 2019 World Wide Web Conference*, 1007–1012.
- Luceri, L.; Pantè, V.; Burghardt, K.; and Ferrara, E. 2024. Unmasking the web of deceit: Uncovering coordinated activity to expose information operations on twitter. In *WWW 2024*, 2530–2541.
- Lundberg, S. M.; and Lee, S.-I. 2017. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30.
- Martha, V.; Zhao, W.; and Xu, X. 2013. A study on Twitter user-follower network: A network based analysis. In *ASONAM*, 1405–1409.
- McPherson, M.; Smith-Lovin, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1): 415–444.
- Metaxas, P.; Mustafaraj, E.; Wong, K.; Zeng, L.; O’Keefe, M.; and Finn, S. 2015. What do retweets indicate? Results from user survey and meta-review of research. In *ICWSM 2015*, 658–661.
- Moravec, P.; Minas, R.; and Dennis, A. R. 2018. Fake news on social media: People believe what they want to believe when it makes no sense at all. *Kelley School of Business Research Paper*, 18–87.
- Mundt, M.; Ross, K.; and Burnett, C. M. 2018. Scaling social movements through social media: The case of Black Lives Matter. *Social Media + Society*, 4(4): 2056305118807911.
- Nikolov, D.; Flammini, A.; and Menczer, F. 2021. Right and left, partisanship predicts (asymmetric) vulnerability to misinformation. *The Harvard Kennedy School Misinformation Review*, 1(7).
- Oyibo, K.; and Vassileva, J. 2019. The relationship between personality traits and susceptibility to social influence. *Computers in Human Behavior*, 98: 174–188.
- Pennycook, G.; Epstein, Z.; Mosleh, M.; Arechar, A. A.; Eckles, D.; and Rand, D. G. 2021. Shifting attention to accuracy can reduce misinformation online. *Nature*, 592(7855): 590–595.
- Pennycook, G.; and Rand, D. G. 2019. Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, 188: 39–50.
- Pfeffer, J.; Matter, D.; Jaidka, K.; Varol, O.; Mashhadi, A.; Lasser, J.; Assenmacher, D.; Wu, S.; Yang, D.; Brantner, C.; et al. 2023. Just another day on twitter: A complete 24 hours of twitter data. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 17, 1073–1081.
- Rao, A.; Morstatter, F.; and Lerman, K. 2022. Partisan asymmetries in exposure to misinformation. *Scientific Reports*, 12(1): 15671.
- Rao, A.; Morstatter, F.; and Lerman, K. 2023. Retweets amplify the echo chamber effect. In *ASONAM 2023*, 30–37.
- Romero, D. M.; Meeder, B.; and Kleinberg, J. 2011. Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on Twitter. In *WWW 2011*, 695–704.
- Saling, L. L.; Mallal, D.; Scholer, F.; Skelton, R.; and Spina, D. 2021. No one is immune to misinformation: An investigation of misinformation sharing by subscribers to a fact-checking newsletter. *PLoS ONE*, 16(8): e0255702.
- Sasahara, K.; Chen, W.; Peng, H.; Ciampaglia, G. L.; Flammini, A.; and Menczer, F. 2021. Social influence and unfollowing accelerate the emergence of echo chambers. *Journal of Computational Social Science*, 4(1): 381–402.
- Stella, M.; Ferrara, E.; and De Domenico, M. 2018. Bots increase exposure to negative and inflammatory content in online social systems. *PNAS*, 115(49): 12435–12440.
- Traberg, C. S.; and van der Linden, S. 2022. Birds of a feather are persuaded together: Perceived source credibility mediates the effect of political bias on misinformation susceptibility. *Personality and Individual Differences*, 185: 111269.
- Ye, J.; Luceri, L.; Jiang, J.; and Ferrara, E. 2024. Susceptibility to unreliable information sources: Swift adoption with minimal exposure. In *WWW 2024*, 4674–4685.
- Zeggelink, E. P.; Stokman, F. N.; and Van De Bunt, G. G. 1996. The emergence of groups in the evolution of friendship networks. *Journal of Mathematical Sociology*, 21(1-2): 29–55.
- Zhou, X.; and Zafarani, R. 2019. Network-based fake news detection: A pattern-driven approach. *ACM SIGKDD Explorations Newsletter*, 21(2): 48–60.

## Paper Checklist

### 1. General questions:

- (a) **Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures?**
  - Yes, the study adheres to ethical guidelines, focusing on aggregate data analysis without violating individual privacy.
- (b) **Do your main claims in the abstract and introduction accurately reflect the paper’s contributions and scope?**
  - Yes, the main claims regarding the Susceptibility Paradox in influence on social media are supported by the findings.
- (c) **Do you clarify how the proposed methodological approach is appropriate for the claims made?**
  - Yes, the methodological approach, including statistical analysis and machine learning models, is detailed and appropriate for the research questions.
- (d) **Do you clarify what are possible artifacts in the data used, given population-specific distributions?**
  - Yes, the paper discusses potential limitations related to data collection and the representativeness of the sample populations.
- (e) **Did you describe the limitations of your work?**
  - Yes, the manuscript includes a section on limitations, acknowledging the constraints of the data and methods used.
- (f) **Did you discuss any potential negative societal impacts of your work?**
  - Yes, the potential societal impacts, particularly in terms of misuse for targeted manipulation or misinformation, are discussed.
- (g) **Did you discuss any potential misuse of your work?**
  - Yes, the manuscript addresses potential misuse, especially regarding the exploitation of susceptibility metrics for unethical influence campaigns.
- (h) **Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings?**
  - Yes, measures such as data anonymization and ethical guidelines for data usage are mentioned.
- (i) **Have you read the ethics review guidelines and ensured that your paper conforms to them?**
  - Yes, the authors confirm adherence to ethics review guidelines throughout the study.

### 2. Hypotheses testing:

- (a) **Did you clearly state the assumptions underlying all theoretical results?**
  - Yes, the assumptions underlying the statistical models and hypotheses are clearly stated.

- (b) **Have you provided justifications for all theoretical results?**
  - Yes, the theoretical results are justified with empirical data and statistical analysis.
- (c) **Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results?**
  - Yes, the paper discusses alternative explanations and competing theories regarding social influence and susceptibility.
- (d) **Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study?**
  - Yes, alternative mechanisms and explanations are explored to ensure the robustness of the findings.
- (e) **Did you address potential biases or limitations in your theoretical framework?**
  - Yes, the potential biases and limitations in the theoretical framework are acknowledged and discussed.
- (f) **Have you related your theoretical results to the existing literature in social science?**
  - Yes, the results are related to existing social science literature on social influence and susceptibility.
- (g) **Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain?**
  - Yes, the implications for policy, practice, and further research are thoroughly discussed.

### 3. Theoretical proofs:

- (a) **Did you state the full set of assumptions of all theoretical results?**
  - Yes, all assumptions for the theoretical results are stated clearly.
- (b) **Did you include complete proofs of all theoretical results?**
  - Yes, mathematical formulations are included where applicable.

### 4. Machine learning experiments:

- (a) **Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)?**
  - Yes, details on code and data will be provided to ensure reproducibility.
- (b) **Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)?**
  - Yes, the training details, including data splits and hyperparameters, are specified.
- (c) **Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)?**
  - Yes, error bars and statistical significance are reported where applicable.
- (d) **Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)?**

- NA
  - (e) **Do you justify how the proposed evaluation is sufficient and appropriate to the claims made?**
    - Yes, the evaluation methods are justified as appropriate for the research claims.
  - (f) **Do you discuss what is “the cost” of misclassification and fault (in)tolerance?**
    - Yes, the costs of misclassification and prediction inaccuracy are discussed in the context of the study.
5. Code & Data Assets:
- (a) **If your work uses existing assets, did you cite the creators?**
    - Yes, existing assets and their creators are properly cited.
  - (b) **Did you mention the license of the assets?**
    - Yes, the licenses of the assets used are mentioned and followed.
  - (c) **Did you include any new assets in the supplemental material or as a URL?**
    - Yes, upon acceptance, we will release our code and data, and document it in a repository linked via URL.
  - (d) **Did you discuss whether and how consent was obtained from people whose data you’re using/curating?**
    - NA
  - (e) **Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content?**
    - Yes, the associated manuscripts address concerns about personally identifiable information.
  - (f) **If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see FORCE11 (2020))?**
    - NA
  - (g) **If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see Gebru et al. (2021))?**
    - NA

### Ethical Statement

Throughout our research process, we have adhered to stringent ethical standards. We have ensured that no personally identifiable information was utilized, and all analyses were conducted on aggregated data. Additionally, we have thoroughly considered the societal impacts of our research. We acknowledge the potential misuse of our findings by malicious entities to target vulnerable individuals susceptible to manipulation. However, we firmly believe that the significance of our work outweighs these risks. Our research has the potential to inform the development of online systems that prioritize safety and responsibility for all individuals, thereby contributing positively to the integrity of digital environments.

## Appendix

### Frequency Distributions of IAR and SAR

Figure 8 shows the frequency distributions of the IAR and SAR metrics in ELECTION and COVID datasets.




Figure 8: Frequency distributions of the IAR and SAR metrics in ELECTION and COVID datasets.

### Explanation of the Null Models

**Baseline 1: Edge Swapping** This null model is designed to test network robustness while preserving the original network’s degree distribution and total number of edges. The model operates by implementing an edge swapping mechanism: (1) Pairs of edges are randomly selected within the network. Care is taken to ensure that these selections do not result in multiple edges between the same pair of nodes (unless the original network allows for multiple edges) or self-loops; (2) The endpoints of these selected edge pairs are swapped—this means that if edges  $(u, v)$  and  $(x, y)$  are chosen, they might be swapped to  $(u, y)$  and  $(x, v)$ ; (3) This process is repeated multiple times to ensure thorough randomization while keeping the number of connections (degree) each node has intact.

**Baseline 2: Random Neighbor Assignment** This model tests the robustness of network properties against random reassignment of edges: (1) Each node in the network retains its degree but has its neighbors randomly reassigned. This means that while each node will still have the same number of connections, the actual connections (neighbors) will be randomly chosen; (2) The total number of edges and the overall degrees are preserved, similar to Baseline 1. However, unlike edge swapping, this method completely disrupts

any existing neighbor relationships. The objective here is to create a completely randomized structure where only the degree sequence of the original network is preserved.

### Paradox Holding Probability $p(k, s)$ at Varying Degrees $k$

Figure 9 depicts the paradox holding probability  $p(k, s)$  as a function of degree  $k$  and susceptibility metric  $s$  using the *Interaction* network in the ELECTION dataset.

### Random Forest Model Parameters

Parameters tested included the number of trees (“n\_estimators”; 100 to 1000 in increments of 50), features per split (“max\_features”; scaled linearly from 1 to the total number of features), maximum tree depth (“max\_depth”; 10 to 100 and unlimited), and minimum samples for splits and leaves (“min\_samples\_split” and “min\_samples\_leaf”), using values of 2, 5, 10, and 1, 2, 4, respectively. Table 5 shows a summary of the random forest model parameters.

Table 5: Summary of the random forest model parameters.

Susceptibility Metric	ELECTION		COVID	
	IAR	SAR	IAR	SAR
n_estimators	750	800	750	800
max_features	4	4	4	4
max_depth	70	80	70	80
min_samples_split	5	10	5	10
min_samples_leaf	2	2	2	2

### SHAP Values of IAR and SAR Prediction

Figure 10 shows the SHAP value distributions for SAR predictions in the ELECTION dataset. Figure 11 shows the SHAP value distributions for IAR predictions in the ELECTION dataset.

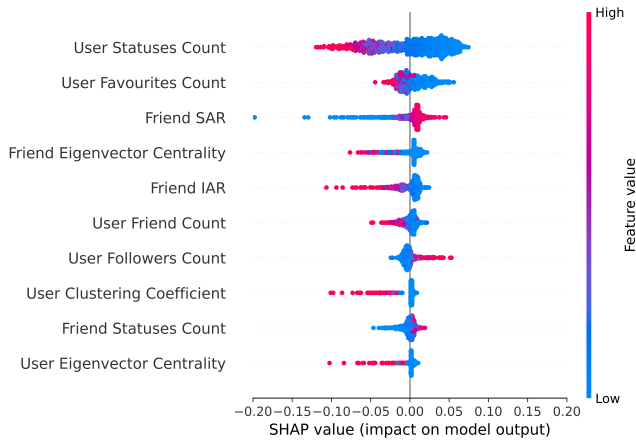


Figure 10: SHAP value distributions for SAR predictions in the ELECTION dataset. Each distribution illustrates the direction and magnitude of the impact of individual features on the model’s output, with features ordered by importance from top to bottom.

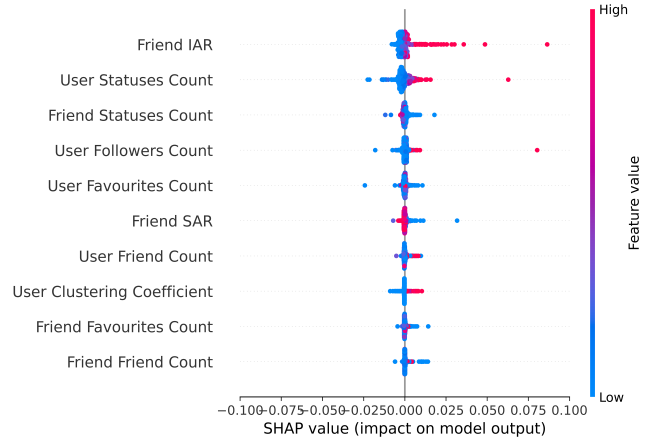


Figure 11: SHAP value distributions for IAR predictions in the ELECTION dataset. Each distribution illustrates the direction and magnitude of the impact of individual features on the model’s output, with features ordered by importance from top to bottom.

### Sensitivity Analysis with Varying Thresholds for Target User Identification

To address potential sampling biases introduced by the network construction method, we experiment with different URL adoption thresholds to identify target users, specifically testing thresholds of 5 and 20 URLs, in addition to the previously used threshold of 10 URLs. Similarly, we construct three distinct friendship networks based on the types of reciprocal interactions between these target users: *Interaction*, *Retweet*, and *Mention* networks. The resulting node and edge counts for the friendship networks in both datasets, based on these new thresholds, are presented in Table 6.<sup>1</sup>

Consistent with previous findings, the results indicate that both susceptibility metrics (IAR and SAR) exhibit homophily within friendship networks, constructed using thresholds of 5 and 20 shared URLs (Table 7). Additionally, IAR demonstrates a stronger degree of homophily compared to SAR. Furthermore, as shown in Table 8, the *Generalized Friendship Paradox* holds for influence-driven adoption (IAR) but not for spontaneous adoption (SAR) under both thresholds, corroborating our earlier conclusions. These results suggest that our main findings are robust across different sample sizes.

<sup>1</sup>The results for the 10 URL threshold ( $G_{10}(V, E)$ ) is available in the main paper.

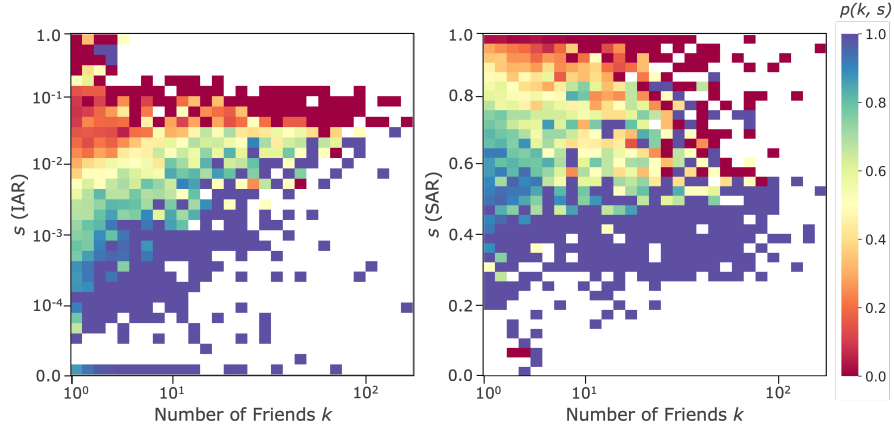


Figure 9: Paradox holding probability  $p(k, s)$  at varying degrees  $k$  using the *Interaction* network in the ELECTION dataset.

Table 6: Number of nodes ( $|V|$ ) and edges ( $|E|$ ) in the user networks  $G_5(V, E)$ ,  $G_{10}(V, E)$ , and  $G_{20}(V, E)$ , constructed with URL thresholds of 5 and 20, respectively.

$G_5(V, E)$	ELECTION		COVID	
	$ V $	$ E $	$ V $	$ E $
Interaction	98,592	242,856	134,409	300,432
Retweet	56,701	106,893	119,456	198,654
Mention	94,702	148,114	115,584	173,022

$G_{10}(V, E)$	ELECTION		COVID	
	$ V $	$ E $	$ V $	$ E $
Interaction	49,794	117,384	73,960	163,974
Retweet	25,077	49,679	53,161	90,800
Mention	45,289	72,554	56,060	81,969

$G_{20}(V, E)$	ELECTION		COVID	
	$ V $	$ E $	$ V $	$ E $
Interaction	45,845	111,402	67,702	154,068
Retweet	23,824	48,588	50,192	87,513
Mention	41,536	67,632	50,695	75,227

Table 7: The susceptibility metrics (IAR and SAR) exhibit homophily within the social network, as evidenced by the weighted average correlation (Wgt. Avg. Corr.) between users and their friends. All correlation coefficients are significant ( $p < 0.001$ ).

$G_5(V, E)$	ELECTION		COVID	
	IAR	SAR	IAR	SAR
Wgt. Avg. Corr.				
Interaction	0.512	0.344	0.545	0.352
Retweet	0.520	0.371	0.569	0.356
Mention	0.458	0.302	0.346	0.288

$G_{10}(V, E)$	ELECTION		COVID	
	IAR	SAR	IAR	SAR
Wgt. Avg. Corr.				
Interaction	0.541	0.356	0.564	0.369
Retweet	0.541	0.391	0.577	0.364
Mention	0.483	0.315	0.372	0.299

$G_{20}(V, E)$	ELECTION		COVID	
	IAR	SAR	IAR	SAR
Wgt. Avg. Corr.				
Interaction	0.579	0.354	0.577	0.365
Retweet	0.586	0.389	0.589	0.367
Mention	0.506	0.324	0.399	0.294

Table 8: Empirical results for the *Generalized Friendship Paradox* in three friendship networks derived from the ELECTION dataset are presented. For each susceptibility metric  $s$  (IAR and SAR), we calculate the Spearman correlation coefficient  $\rho_{ks}$  between degree  $k$  and  $s$  ( $p < 0.001$  for all  $\rho_{ks}$ ), average paradox holding probability  $P$  and the average susceptibility  $\langle s \rangle$  of nodes and their friends  $\langle s \rangle_{nn}$ .

$G_5(V, E)$	Metric $s$	$\rho_{ks}$	Individual-level GFP	Network-level GFP	
			$P$	$\langle s \rangle$	$\langle s \rangle_{nn}$
Interaction	IAR	0.188	0.527	0.005 <	0.006
	SAR	-0.174	0.385	0.889 >	0.876
Retweet	IAR	0.302	0.493	0.012 <	0.013
	SAR	-0.296	0.419	0.826 >	0.812
Mention	IAR	0.115	0.532	0.005 <	0.006
	SAR	-0.103	0.316	0.902 >	0.889
$G_{10}(V, E)$	Metric $s$	$\rho_{ks}$	Individual-level GFP	Network-level GFP	
			$P$	$\langle s \rangle$	$\langle s \rangle_{nn}$
Interaction	IAR	0.179	0.541	0.007 <	0.008
	SAR	-0.166	0.409	0.880 >	0.868
Retweet	IAR	0.298	0.501	0.013 <	0.014
	SAR	-0.294	0.425	0.821 >	0.809
Mention	IAR	0.111	0.544	0.005 <	0.006
	SAR	-0.099	0.383	0.896 >	0.880
$G_{20}(V, E)$	Metric $s$	$\rho_{ks}$	Individual-level GFP	Network-level GFP	
			$P$	$\langle s \rangle$	$\langle s \rangle_{nn}$
Interaction	IAR	0.170	0.548	0.007 <	0.009
	SAR	-0.159	0.419	0.874 >	0.863
Retweet	IAR	0.295	0.504	0.013 <	0.014
	SAR	-0.294	0.428	0.817 >	0.805
Mention	IAR	0.099	0.552	0.006 <	0.007
	SAR	-0.089	0.393	0.890 >	0.874