The Friendship Paradox: An Analysis on Signed Social Networks with Positive and Negative Links

Catherine Yang
catherine.y.yang@vanderbilt.edu
Vanderbilt University

Yuying Zhao
yuying.zhao@vanderbilt.edu
Vanderbilt University

Tyler Derr
tyler.derr@vanderbilt.edu
Vanderbilt University

ABSTRACT

Given the ubiquity and significance of social network systems, comprehending the network topology is essential for a deeper understanding of these networks. One notable phenomenon in social networks, Friendship Paradox (FP), has been extensively studied and has led to the Generalized Friendship Paradox (GFP), which states that an individual's neighbors, on average, have more of some measurable characteristic or quantity than the individual (e.g., friends/degree in the original FP). However, most of the existing works on FP and GFP naturally focus on positive relationships while in the real world, negative relations are also ubiquitous. To bridge this crucial gap, we investigate (G)FP in signed networks which contain both positive and negative relationships (e.g., friends and foes). Specifically, we propose a first-order signed neighbor metric based on the traditional (G)FP that not only considers undirected homogeneous link relations (e.g., comparing an individual's foes to the foes of their foes), but also directed heterogeneous link relation (e.g., comparing an individual's friends to the friends of their foes). Furthermore, we develop a second-order metric to further study the relationship between an individuals positive and negative neighborhood sets (e.g., comparing the average number of friends from an individuals set of foes to that of their friends). Finally we perform an empirical analysis of these proposed metrics in signed networks across a representative set of real-world datasets.

CCS CONCEPTS

• **Information systems** → **Data mining**.

KEYWORDS

friendship paradox, signed social networks, negative links

ACM Reference Format:
Catherine Yang, Yuying Zhao, and Tyler Derr. 2018. The Friendship Paradox: An Analysis on Signed Social Networks with Positive and Negative Links. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 6 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Social network systems have become ubiquitous in our daily lives, providing us with platforms to connect with others and exchange information. Given their significance, this has led to a plethora of studies to better understand our society at both the macro-level (e.g., social stability [12, 13]) and at the micro-level (e.g., individual or subgroup interactions [22]) [2, 28]. At the core of much of this research is the network topology [23], which defines the relationships between the individuals involved in the system. This has led to inherently multi-disciplinary research at the intersection of sociology, graph theory, data science, etc. [8]. One such direction is the well studied Friendship Paradox (FP) [11], which states, on average, an individual's friends have more friends than that individual.

This phenomenon of the friendship paradox has been shown to exist due to a combination of factors in a network, specifically the sampling bias where nodes are counted proportional to their node degree [11] (i.e., the overrepresentation of popular individuals appearing in the averaging among the neighborhood set of many others [17]). While the phenomenon is only truly linked with the degrees of a node, it has led to the proposed Generalized Friendship Paradox (GFP) [9], that generalizes and extends beyond node degree/popularity. More precisely, GFP states that an individual's neighbors, on average, tend to have more than the individual, which has been discovered to exist across many other individual characteristics and rooted in the fact that they have positive correlation with degree/popularity [9, 17]. For example, GFP has been shown to exist empirically for individual characteristics including viral content on social media [17], citations in collaboration networks [9], etc., while also theoretically in the direction of popularity stepping beyond degree to more complex measures, such as eigenvector centrality [16]. However, nearly all existing work has focused on networks where the links are defined according to positive relationships. Hence, in this work we seek to study the friendship paradox in signed networks [3], which models complex social systems containing both positive relations (e.g., friends or followers) and negative relations (e.g., foes or unfollowers) together in one network [4, 22, 25].

In today's world while perhaps most of our social interactions are positive, users in both physical and virtual social systems ultimately also develop negative relations [22]. Hence, to fully understand and uncover the complexities of our society, signed networks are becoming more prevalent [5, 15, 25]. However, to the best of our knowledge, work has yet to focus on the friendship paradox in signed networks with an empirical emphasis on negative links and the generalized friendship paradox from the perspective of the interactions between positive and negative links/degrees. Dedicated efforts are especially desired on the study of negative links, since while they follow some commonalities with positive links, such as power-law degree distributions [7, 24], they also have significant differences across other properties, such as homophily [26] or reciprocity, which has been shown to be related to the level of visibility of negative links within the system [7]. Furthermore, the behavior

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/XXXXXXX.XXXXXXX

that constitutes a negative link is also highly variable depending on the network, which emphasizes the need for a comprehensive study across a representative set of signed networks.

In this work, we first provide a review of related work from the perspective of the friendship paradox and signed network analysis. Thereafter, we propose a first-order signed neighbor paradox based on the traditional friendship paradox, but instead of independently studying negative links and treating them fully independent from positive links, we further calculate the interactions among positive and negative links for both signed and unsigned networks. Expanding on the first-order definition, we then propose a deeper second-order definition, inspired by recent second-order unsigned network statistics studying multistep friendship paradox [19] and second-order homophily [10]. Empirical analysis of the proposed definitions are studied across a representative set of online signed networks covering a bitcoin exchange network [7], product review website [14], technology news website [14], and voting/election network [22], along with a physical signed network collected from a village in Honduras [18]. Our contributions can be summarized as follows:

- We for the first time study the (generalized) friendship paradox in signed networks containing both positive and negative links.
- We propose both first-order and second-order signed neighbor metrics to study the complex phenomenon related to positive and negative node degrees in signed networks.
- We provide a comprehensive analysis of the developed metrics across a representative set of signed networks, which stimulates further research in this direction to deeper understand the preliminary findings of this work.

2 PRELIMINARIES

In this section, we first discuss related work on the Friendship Paradox and signed network analytics, then we provide basic notations used throughout the paper.

2.1 Related Work

2.1.1 Related Studies on the Friendship Paradox. The Friendship Paradox (FP) was first discovered by Scott L. Feld in 1991 as a form of sampling bias[11]. It implies that an individual is more likely to be friends with someone who is popular, and less likely to be friends with someone who has fewer friends. In network terms, a node is more likely to be a neighbor of a node with many neighbors (i.e., high degree), compared to being linked to a node with only a few edges (i.e., low degree). However, researchers later studied/observed this phenomenon not only according to the theoretically grounded node degree, but in a more generalized form where individuals' neighbors, on average, have more of some other measurable individual characteristic/quantity than the individual. This generalized form has been designated as the Generalized Friendship Paradox (GFP)[9]. For instance, GFP occurs in collaboration networks where a researcher's collaborators, on average, tend to be more productive[9]. GFP has also been discovered to be associated with user activity engagements on online social media [17]. Nevertheless, these discourses on FP/GFP primarily concentrate on positive relationships [11, 17, 27, 29], while the seemingly unavoidable and prevalent negative relationships have yet to be explored.

2.1.2 Signed Network Analytics. Although negative relations in social networks have been relatively underexplored, signed networks seek to model these complex systems inherently having both positive and negative links [3, 25]. For example, positive links may represent relations associated with trust/friendships, while negative links might encode distrust or even animosity. Generally, the major directions in signed network analytics are: 1) network theories and analysis [3]; 2) prediction tasks on networks [6, 21]; 3) network models [4, 24]; and 4) network measurements [1, 7]. In this work, we seek to expand the frontier on signed network measurements and theories, by investigating the (G)FP in signed networks by developing first-order and second-order network measurements.

2.2 Signed Network Notations

Formally, a signed network $\mathcal{G} = (\mathcal{V}, \mathcal{E}^+, \mathcal{E}^-)$ is composed of a set of n nodes (e.g., users in a social network) $\mathcal{V} = \{v_1, \ldots, v_n\}$, along with the sets of positive links $\mathcal{E}^+$ and negative links $\mathcal{E}^-$ that exist between the nodes in the network. Furthermore, we define the set of immediate positive neighbors of a node v_i as $\mathcal{N}_+(v_i)$ and similarly as $\mathcal{N}_-(v_i)$ for the set of negative neighbors.

3 SIGNED NEIGHBOR PARADOX METRICS

In this section, we first propose a first-order measurement towards not only understanding the FP applied to negative links (instead of positive links), but furthermore to measure the complex relationships between positive and negative links, which also aligns with the GFP assuming positive/negative degrees are correlated. Thereafter, we introduce a second-order signed neighbor metric that directly studies the potentially paradoxical relationship among an individuals set of friends compared to their set of foes, from the perspective of the average number of friends/foes associated with the individuals in those two neighborhood sets.

3.1 First-order Measurements

In order to study the first-order signed neighbor paradox, we propose calculating each node's four possible relations: incoming positive edges, outgoing positive edges, incoming negative edges, and outgoing edges. Then, we take a set of the node's neighbors and calculate the average number of incoming positive edges, outgoing positive edges, incoming negative edges, and outgoing edges those neighbors have. The set of neighbors is chosen based on its relationship to the node. We analyze the number of each type of neighbor for a node and its set of each neighbor type. More formally, for each node v_i, we calculate:

$$\psi_{i,1}^{\alpha\beta} = \frac{1}{\mathcal{N}_\alpha(v_i)} \sum_{v_j \in \mathcal{N}_\alpha(v_i)} \mathbb{1}\big(|\mathcal{N}_\beta(v_j)| > |\mathcal{N}_\beta(v_i)|\big)$$

where $\mathbb{1}$ is an indicator function such that it equals 1 when the neighbor v_j has more neighbors of β type than it, and 0 otherwise. This expression is summed and then divided by the number of neighbors we compared it against so that $\psi_{i,1}^{\alpha\beta}$ represents the ratio of a node v_i's α type neighbors that have more relations of β type than it. For example, in a simplified undirected case, when α = + and β = -, this measures to what extent a node's friends have more foes than it. Then, we aggregate the results of each node, and calculate the ratio of nodes whose friends have more foes than them. A visualization of the undirected first-order is shown in Figure 1.

3.2 Second-order Measurements

Expanding upon the above first-order measurement, we seek to develop a deeper second-order signed theory, which are even less explored in the literature with only recently a few works focused in unsigned networks [10, 19]. Specifically, we propose to investigate a more advanced potential signed neighbor degree paradox considering second-order information, while also developing a computationally efficient global aggregated measure to avoid the inherent nested all-pairs comparison. Formally, for a node v_i the second-order signed neighbor paradox can be calculated as follows:

$$\psi_{i,2}^{\delta} = \frac{1}{|\mathcal{N}_{+}(v_i)|\,|\mathcal{N}_{-}(v_i)|} \sum_{v_j \in \mathcal{N}_{+}(v_i)} \sum_{v_k \in \mathcal{N}_{-}(v_i)} \mathbb{1}\left(|\mathcal{N}_{\delta}(v_j)| > |\mathcal{N}_{\delta}(v_k)|\right)$$

where $\delta \in \{+,-\}$ for undirected signed networks. However, this results in a $O(n^3)$ time complexity for a naïve global aggregation. Therefore, to avoid the computational costs, we instead propose:

$$\psi_{2}^{\delta} = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}\Bigg(\Big(\frac{1}{|\mathcal{N}_{+}(v_i)|} \sum_{v_j \in \mathcal{N}_{+}(v_i)} |\mathcal{N}_{\delta}(v_j)|\Big) > \Big(\frac{1}{|\mathcal{N}_{-}(v_i)|} \sum_{v_k \in \mathcal{N}_{-}(v_i)} |\mathcal{N}_{\delta}(v_k)|\Big)\Bigg).$$

Here we are interested in comparing the amount of friends and foes between an individuals set of friends and foes. Friends refer to the set of nodes a node has a positive outgoing edge to, and enemies refer to the set of nodes a node has a negative outgoing edge to.

For example, in the positive case where $\delta = +$, for each node v_i, we take its set of friends $\mathcal{N}_{+}(v_i)$ and foes $\mathcal{N}_{-}(v_i)$, and measure the average amount of friends in each set. We sum this over all nodes, and the resulting value represents the ratio of nodes in the network in which its set of friends have more friends in the network than its enemies have friends. The negative case is similarly defined.

3.3 Temporal First-order Measurements

The temporal first-order measurements aim to capture how the measurements for the first-order paradox trend change over time. We are interested in seeing if there is a particular value that it may converge to that represents the average ratio of friends more popular than us compared to friends with less friends. Therefore, this paper will focus on measuring the number of positive outgoing edges each node has and the average number of positive outgoing edges a node's outgoing positive edged neighbors have in relation to how long the node has existed and participated in the network (and similarly for outgoing negative).

In order to measure the temporal behavior of the data, we first convert the temporal graph into a sequence of static graphs so that we can measure the first-order values on each graph in the sequence. More specifically, in every snapshot, we add all new nodes and edges to the graph. The graph is incrementally built up, and its state at each snapshot represents the network at that point in time. Each snapshot constitutes all transactions occurring over up to that point in time since the start of data collection. However, since nodes enter the graph at different points, the amount of time each node has existed also differs. Since we want to measure the first-order paradox in relation to how long a node has existed, we calculate the first-order value, and place it in the proper time bin according to the difference in the current time and the time of the node's creation. Ultimately, first-order values are averaged per bin.

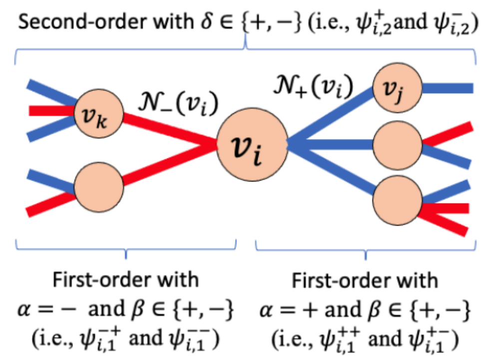

Figure 1: An example visualizing how the undirected first-order and second-order measurements are calculated.

4 EMPIRICAL ANALYSIS

To comprehensively investigate the signed neighbor paradox via our proposed first-order and second-order measurements, we evaluate a representative set of signed networks covering a variety of negative relationships including both online and physical signed networks with their detailed descriptions in Section 4.1 and Table 1. We explore both directed and undirected settings in Section 4.2 and Section 4.3, respectively. Furthermore, in Section 4.4, we explore the dynamic setting where we examine the node-level evolution of directed first-order measurement. In the end, we provide the results and conduct analysis for second-order measurement.

4.1 Signed Network Datasets

4.1.1 Bitcoin Alpha. The Bitcoin Alpha dataset [7] is a directed weighted temporal signed network taken from the Bitcoin Alpha platform, which is a trust network associated with users who anonymously transact with each other for products/services for Bitcoin. Each user can rate their transaction partners with either a positive score (ranging from 1 to 10) or a negative score (ranging from -10 to -1), where the platform offers guidance to standardize the positive/negative scoring. The collective ratings a user receives determine their reputation within the platform, and affect their future ability to transact within the community.

4.1.2 Wiki Elections. Wiki elections [22] is a directed temporal signed network that models Wikipedia's community votes that get cast during elections of new administrators. In this network, positive edges represent users/admins vote in support of the admin candidate while negative edges represent the voters' opposition.

4.1.3 Honduras Village. The Honduras Village dataset [18] is a directed signed network, which represents friendly (positive), antagonistic (negative), or stranger (no-link) real-world physical relationships between individuals living in western Honduras.

4.1.4 Slashdot. The Slashdot dataset [20] is a directed signed network modeling a technology news site, where specifically the Slashdot Zoo feature allows users to tag each other as their friends or foes (i.e., positive or negative neighbors). Here we condense this directed signed network into an undirected network, where for a bi-directional relationship if at least one is negative we treat this as a negative undirected edge.

Table 1: Basic statistics of the signed networks datasets.

Network Type	Dataset Name	# Nodes	# Pos. Edges	# Neg. Edges
Directed	Bitcoin Alpha	3,784	22,650	1,536
	Wiki Elections	7,116	78,440	22,253
	Honduras Village	149	1,252	187
Undirected	Slashdot	82,141	380,933	119,548
	Epinions	131,580	589,888	121,322

4.1.5 Epinions. The Epinions dataset [14] is a directed signed network created from Epinions.com, a product review platform that allows users to assign trust (i.e., positive) or distrust (i.e, negative) links to other users based on their provided reviews. We note that negative links are totally invisible to others, but are provided by Epinions staff for research purposes. Here we condense this dataset to an undirected signed network following the same procedure as on the Slashdot dataset.

4.2 Directed First-order Analysis Results

One observation is that the 4 values for $\alpha \in \{+in, +out\}$, $\beta \in \{+in, +out\}$ are relatively higher than the other sets of four in Table 2. This indicates that the set of users that view a node positively and the set of users the node views positively are likely to have on average even more incoming and outgoing positive reviews. In social networks, this can represent a node's involvement with a popular person who is generally well liked and positive towards others. In online trading networks such as Bitcoin Alpha where a positive review represent trust, the results indicate that users that a user reviews positively is likely to be a trustworthy user who gives and receives more positive reviews than the user. These users are potentially highly active and well known traders that act as hubs that increases the overall average number of positive edges observed in a node's positive neighbor set.

From Table 2, $\alpha = -in$ and $\beta = -out$ is the highest value in its respective row and column for all three datasets, meaning that the users that view a node negatively, are on average, likely to leave more negative reviews than the node does. This is explainable by the fact that we are taking a set of users who have given the node a negative review or vote, and may be more prone to giving negative reviews/votes to others beyond that node.

Another point of interest is that the $\alpha = -out$ and $\beta = -out$ value is consistently amongst the lowest values measured, and the only value that remains below 0.5 in all three datasets. In all the social networks studied, more than half of users have more outgoing negative edges than the set of users it has a negative edge to. This contrasts with the traditional friendship paradox, and our findings in $\alpha = +out$ and $\beta = +out$, where a node's positive neighbors have more positive neighbors than it for most nodes. The low value observed for $\alpha = -out$ and $\beta = -out$ indicates that most users in these networks give a greater amount of negative edges than the nodes it has a negative edge to do.

4.3 Undirected First-order Analysis Results

Table 3 describes the undirected first-order results on the Slashdot and Epinions datasets. It can be observed that the diagonal values where $\alpha = \beta$ are higher than the $\alpha \neq \beta$ values. In other words, for a greater portion of nodes, it has less friends than its friends, and less enemies than its enemies. However, all values are relatively high, showing that most nodes have both less friends and less enemies than its friends and enemies do.

Table 2: First-Order Results on Directed Signed Networks

α \ β	+ in	+ out	− in	− out
Bitcoin Alpha				
+ in	0.95	0.95	0.74	0.59
+ out	0.96	0.95	0.75	0.59
− in	0.48	0.45	0.39	0.88
− out	0.85	0.84	0.91	0.38
Wiki Elections				
+ in	0.84	0.97	0.76	0.85
+ out	0.78	0.64	0.71	0.58
− in	0.72	0.86	0.74	0.95
− out	0.57	0.45	0.79	0.47
Honduras Village				
+ in	0.54	0.67	0.51	0.61
+ out	0.79	0.65	0.59	0.64
− in	0.49	0.57	0.39	0.85
− out	0.54	0.49	0.72	0.31

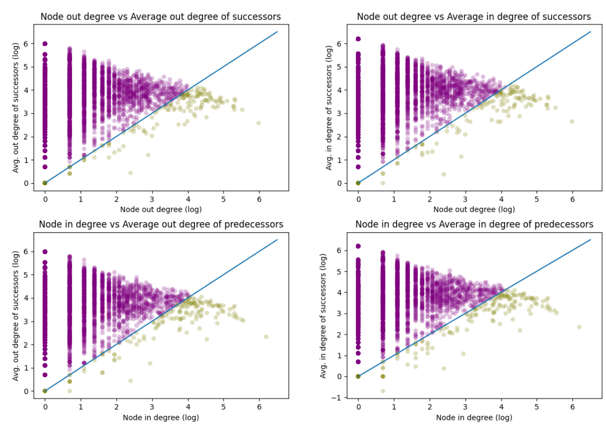

Figure 2: A detailed node-level (log-log) visualization on the Bitcoin Alpha dataset of the (G)FP according to a directed unsigned perspective where positive and negative links are merged together towards a single (in/out) degree. Note that the purple nodes are users who adhere to the theory's expectation, while olive nodes below the diagonal line are those that have a higher (in/out) degree than the average of their (in/out) neighbors.

Table 3: Slashdot (Epinions) Undirected First-Order Results

α \ β	+	−
+	0.97 (0.88)	0.83 (0.71)
−	0.87 (0.81)	0.96 (0.86)

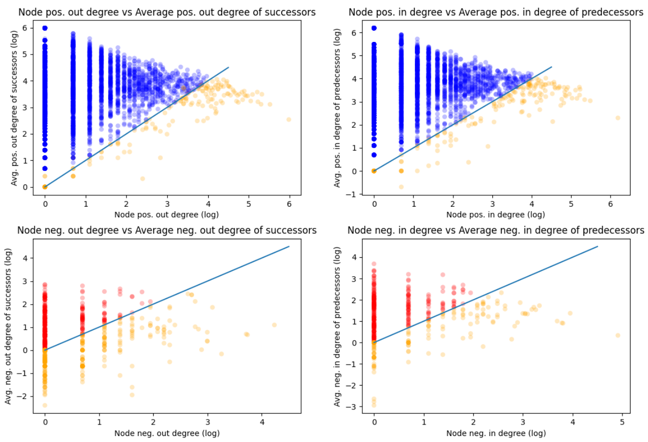

Figure 3: A detailed node-level (log-log) visualization on the Bitcoin Alpha dataset of the (G)FP according to a directed signed perspective where positive and negative links are analyzed separately and specifically focused on Bitcoin Alpha's diagonal of Table 2. Note the blue(red) nodes are users who adhere to the theory's expectation according to their positive(negative) perspective, while the orange nodes do not.

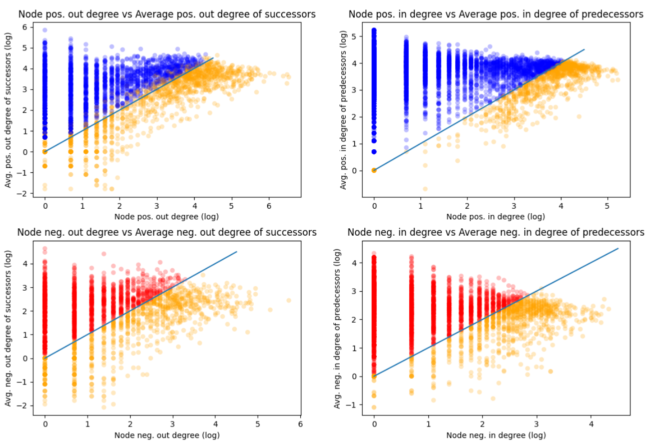

Figure 4: Visualizing the same setting as Figure 3, but on the Wiki Elections dataset.

4.4 Examining the Node-level Evolution of Directed First-order Measurements

In order to measure the temporal behavior of the data, we first decide the regularity at which we create static graphs to measure the first-order values. In this paper, we elected Bitcoin Alpha as a representative dataset and set the snapshots to be every month, or 30 days, giving us 72 first-order measurements over the period the data was collected. We note however that although the data contains 72 months, an insufficient number of nodes exist and have positive relations for over than 64 months, and an insufficient number of nodes have negative relations for over 47 months, so we omit those insignificant data points in Figures 6 and 7, respectively.

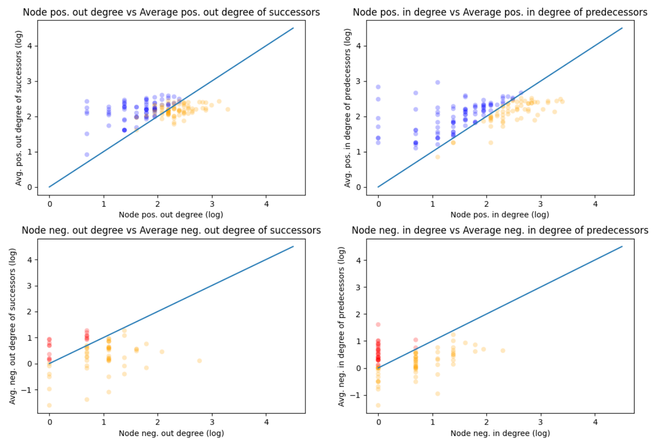

Figure 5: Visualizing the same setting as Figure 3, but on the Honduras Village dataset.

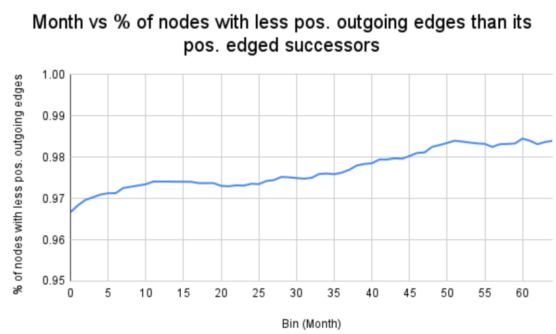

Figure 6: Bitcoin Alpha Positive Directed First Order Paradox over 64 Months describing the relationship between the age of nodes and the average positive friendship paradox value observed at that age.

The first observation, from Figure 6, is that as time goes by, the proportion of nodes with less positive outgoing edges than their successors is steadily increasing. It indicates that more users are experiencing the positive friendship paradox over time.

Similar to the non-temporal first-order results, the range of values are high at above 0.95. 96% of nodes observe that its positive edged neighbors have more positive outgoing edges than it does when it first joins the network. As nodes age, more and more of them observe that on average, its friends have more friends. As the Bitcoin Alpha dataset describes online ratings and transactions between users, it is logical that over time, a user would transact with more popular, active, and trusted users that interact with and review other legitimate users which may increase the overall friend average as well as the % of nodes that fall in this category.

Contrary to the positive results, the % of nodes with less negative outgoing edges than the nodes it has a negative outgoing edge seems to oscillate around a much lower range around 0.3-0.4, indicating that a majority of users leave a greater amount of negative reviews than their enemies do. In the context of bitcoin, someone that a user leaves a negative review for may be an illegitimate trader, and may leave more positive reviews in an effort to appear trustworthy.

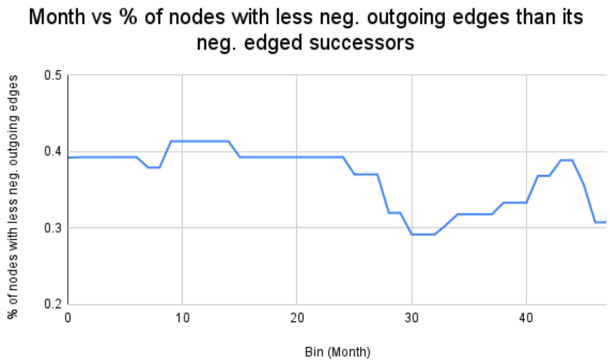

Figure 7: Bitcoin Alpha Negative Directed First-Order Paradox over 47 Months describing the relationship between the age of nodes and the average negative friendship paradox value observed at that age.

Table 4: Second-Order Results

Dataset	$\delta = +$	$\delta = -$
Bitcoin Alpha	0.8852	0.8948
Slashdot	0.8012	0.6461
Epinions	0.8383	0.7557
Wiki Elections	0.7376	0.6039
Honduras Village	0.7718	0.7383

4.5 Second-order Analysis Results

The second-order results are shown in Table 4. We draw three observations as follows: (1) *Positive Trend*: the second-order positive results are similarly high for all datasets with an average of 0.8068, indicating that for most nodes, their friends have more friends than their enemies do; (2) *Negative Trend*: the negative results have a lower average at 0.7278 with a slightly higher variation. This indicates a similar observation with the positive results: for most nodes, their friends also have more enemies than their enemies do; (3) *Positive vs Negative*: when comparing the positive and negative scores, we observe that the negative result is consistently equal to or lower than the positive. It indicates that the negative trend is slightly weaker than the positive trend.

5 CONCLUSION

In this study, we bridge the gap between the Friendship Paradox (FP)/Generalized Friendship Paradox (GFP) and signed networks by introducing corresponding metrics that take into account not only positive but also negative relationships. We propose both first-order and second-order signed neighbor paradox metrics based on the traditional FP metric to investigate at multiple levels along with providing aggregated statistics, node-level visualizations, and examine the node-level evolution of the first-order measurement. We plan to continue this study on more signed network datasets, with emphasis on diverse negative link meanings (e.g., unfollowing), along with a deeper investigation on the temporal signed networks. Additionally, since the current study focuses on the sign of links, we may consider edge weights as a method of accounting for the strength/polarization between nodes when measuring FP and GFP.

REFERENCES

[1] Samin Aref and Mark C Wilson. 2018. Measuring partial balance in signed networks. *Journal of Complex Networks* 6, 4 (2018), 566–595.

[2] Peter J Carrington, John Scott, and Stanley Wasserman. 2005. *Models and methods in social network analysis*. Vol. 28. Cambridge university press.

[3] Dorwin Cartwright and Frank Harary. 1956. Structural balance: a generalization of Heider's theory. *Psychological review* 63, 5 (1956), 277.

[4] Tyler Derr, Charu Aggarwal, and Jiliang Tang. 2018. Signed network modeling based on structural balance theory. In *Proceedings of the 27th ACM international conference on information and knowledge management*. 557–566.

[5] Tyler Derr, Cassidy Johnson, Yi Chang, and Jiliang Tang. 2019. Balance in signed bipartite networks. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 1221–1230.

[6] Tyler Derr, Yao Ma, and Jiliang Tang. 2018. Signed graph convolutional networks. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 929–934.

[7] Tyler Derr, Chenxing Wang, Suhang Wang, and Jiliang Tang. 2017. Signed node relevance measurements. *arXiv preprint arXiv:1710.07236* (2017).

[8] David Easley and Jon Kleinberg. 2010. *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge university press.

[9] Young-Ho Eom and Hang-Hyun Jo. 2014. Generalized friendship paradox in complex networks: The case of scientific collaboration. *Scientific reports* 4, 1.

[10] Anna Evtushenko and Jon Kleinberg. 2021. The paradox of second-order homophily in networks. *Scientific Reports* 11, 1 (2021), 1–10.

[11] Scott L Feld. 1991. Why your friends have more friends than you do. *American journal of sociology* 96, 6 (1991), 1464–1477.

[12] Danielle German and Carl A Latkin. 2012. Social stability and health: exploring multidimensional social disadvantage. *Journal of Urban Health* 89 (2012), 19–35.

[13] Itzhak Gilboa and Akihiko Matsui. 1991. Social stability and equilibrium. *Econometrica: Journal of the Econometric Society* (1991), 859–867.

[14] Ramanthan Guha, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. 2004. Propagation of trust and distrust. In *Proceedings of the 13th international conference on World Wide Web*. 403–412.

[15] Yixuan He, Xitong Zhang, Junjie Huang, Mihai Cucuringu, and Gesine Reinert. 2022. PyTorch Geometric Signed Directed: A Survey and Software on Graph Neural Networks for Signed and Directed Graphs. *arXiv preprint arXiv:2202.10793*.

[16] Desmond J Higham. 2019. Centrality-friendship paradoxes: when our friends are more important than us. *Journal of Complex Networks* (2019).

[17] Nathan Hodas, Farshad Kooti, and Kristina Lerman. 2013. Friendship paradox redux: Your friends are more interesting than you. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 7. 225–233.

[18] Alexander Isakov, James H Fowler, Edoardo M Airoldi, and Nicholas A Christakis. 2019. The structure of negative social ties in rural village networks. *Sociological science* 6 (2019), 197–218.

[19] Josh Brown Kramer, Jonathan Cutler, and AJ Radcliffe. 2016. The multistep friendship paradox. *The American Mathematical Monthly* 123, 9 (2016), 900–908.

[20] Jérôme Kunegis, Andreas Lommatzsch, and Christian Bauckhage. 2009. The slashdot zoo: mining a social network with negative edges. In *Proceedings of the 18th international conference on World wide web*. 741–750.

[21] Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. 2010. Predicting positive and negative links in online social networks. In *Proceedings of the 19th international conference on World wide web*. 641–650.

[22] Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. 2010. Signed networks in social media. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1361–1370.

[23] Alan Mislove, Massimiliano Marcon, Krishna P Gummadi, Peter Druschel, and Bobby Bhattacharjee. 2007. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM IMC*. 29–42.

[24] Pradumn Kumar Pandey, Bibhas Adhikari, Mainak Mazumdar, and Niloy Ganguly. 2020. Modeling signed networks as 2-layer growing networks. *IEEE Transactions on Knowledge and Data Engineering* 34, 7 (2020), 3377–3390.

[25] Jiliang Tang, Yi Chang, Charu Aggarwal, and Huan Liu. 2016. A survey of signed network mining in social media. *Comput. Surveys* 49, 3 (2016), 1–37.

[26] Jiliang Tang, Xia Hu, and Huan Liu. 2014. Is distrust the negation of trust? The value of distrust in social media. In *Proceedings of the 25th ACM conference on Hypertext and social media*. 148–157.

[27] Johan Ugander, Brian Karrer, Lars Backstrom, and Cameron Marlow. 2011. The anatomy of the facebook social graph. *arXiv preprint arXiv:1111.4503* (2011).

[28] Stanley Wasserman and Katherine Faust. 1994. Social network analysis: Methods and applications. (1994).

[29] Ezra W Zuckerman and John T Jost. 2001. What makes you think you're so popular? Self-evaluation maintenance and the subjective side of the" friendship paradox". *Social Psychology Quarterly* (2001), 207–223.