Friendship Paradoxes and the Quora Downvoting Paradox

Shankar Iyer* Massey Cashore

Abstract
The "friendship paradox" refers to the statistical pattern that, for most participants in many social networks, their friends have more friends on average than they do. In recent years, the availability of large volumes of data on online social networks has enabled the study of the friendship paradox in new contexts and on unprecedented scales. Researchers have shown that phenomena similar to the friendship paradox, called "generalized friendship paradoxes," occur in quantities other than friend count: for example, in certain online social networks, the average neighbor of a typical individual is a more active user and content contributor. Furthermore, researchers have also found that online social networks exhibit stronger friendship and generalized friendship paradoxes than are usually measured in the literature: often, for most people in a social network, *most* of their neighbors score more highly on various metrics. This is typically a stronger statement than the usual one about mean of the quantity over neighbors. In this article, we apply these developments in the study of the friendship paradox to Quora, an online knowledge-sharing platform. There is a directed social network of people following one another on Quora, and we first confirm that standard directed-network variants of the friendship paradox occur in this context. We then proceed to investigate a more exotic variant of the friendship paradox that we call "downvoting paradox." This is a variant of the phenomenon that emerges through one of the core interactions on Quora, the "downvote," which people use to give negative feedback on one another's answers. Under certain conditions on the contribution level of the participants, we find that, for most people who got downvoted in a given period of time, most of their downvoters got downvoted more than they did. This is an example of a paradox occurring in a network that represents negative interactions, which is a relatively unexplored context. Furthermore, certain aspects of the product mechanics of Quora make this a particularly interesting setting to explore these types of phenomena. We discuss the implications of the observation of the "downvoting paradox" and suggest opportunities for further investigation.

Key Words: social network analysis, friendship paradox, social media

1. Introduction

"Your friends have more friends than you do": so announces the title of a seminal 1991 paper on social networks by S.L. Feld (Feld, 1991). That paper catalyzed study of the so-called "friendship paradox." Despite its name, the friendship paradox is not really a paradox; instead, it is a term for the statistical pattern that, for the majority of individuals in many social networks, their friends have more friends on average than they do. In the two and a half decades since Feld's paper, versions of this phenomenon have been observed in various contexts, including social networks at schools, academic collaboration networks, and epidemic spreading in real-world contact networks (Feld, 1991; Cohen, Havlin, & Ben-Avraham, 2003; Christakis & Fowler, 2010; Eom & Jo, 2014). Recently, the availability of large volumes of data on online social networks has led to renewed attention on manifestia-tions of the friendship paradox in products such as Facebook and Twitter (Ugander, Karrer, Backstrom, & Marlow, 2011; Hodas, Kooti, & Lerman, 2013, 2014; Bollen, Gonçalves, van de Leemput, & Ruan, 2016). Alongside empirical work, there has also been theoretical work aimed at clarifying when and how certain versions of the friendship paradox emerge

*Quora, Inc., 650 Castro Street #450, Mountain View, CA 94041

(Lattanzi & Singer, 2015; Cao & Ross, 2016) and research into the impact of the friendship paradox on collective decision making (Lerman, Yan, & Wu, 2016; Jackson, 2016).

So-called "generalized friendship paradoxes," which concern statements similar to the friendship paradox but for traits other than friend count, have also been identified in online social networks. For example, in the Twitter network, it has been observed that the people whom a typical user follows are more active than that user on average and see more viral content on average (Hodas et al., 2013). Very recently, it has even been found that, in the mutual follow network on Twitter, most people experience a "happiness paradox": their neighbors are happier than they are on average, at least according to the sentiment encoded in their tweets (Bollen et al., 2016). It has been shown that generalized friendship paradoxes can lead to biases in how individuals in a network perceive global characteristics of the network, as the behavior of popular nodes (who may have more extreme opinions) is overrepresented in the typical node's neighbors (Lerman et al., 2016; Jackson, 2016). The mathematical underpinnings of the generalized friendship paradox have been studied, and stem from the fact that node attributes (e.g., virality of content produced, or citation count) are correlated with the degree of the node (Jo & Eom, 2014; Fotouhi, Momeni, & Rabbat, 2014).

The availability of large quantities of data on online social networks has also enabled researchers to identify that many of these networks exhibit stronger friendship paradoxes than are typically measured in the literature. In particular, Hodas, Kooti, and Lerman have emphasized the distinction between *weak* and *strong* paradoxes (Hodas et al., 2014). Weak paradoxes are those where the mean neighbor of most individuals in the network scores higher on some metric (be it friend count, as in Feld's paper, or some other metric, as in the case of generalized paradoxes). Meanwhile, strong paradoxes are those where the *median* neighbor of most individuals in the network scores higher on a metric. Hodas et al. have shown that, in the cases of Twitter and Digg, weak generalized paradoxes often survive random permutation of the metric in question over the nodes in the network, while the strong paradox disappears. This is because weak paradoxes often inevitably follow from the long-tailed distributions of metrics in social networks, while strong paradoxes can reveal non-trivial aspects of behavioral correlations on the network (Hodas et al., 2014).

In this article, we study strong variants of the friendship paradox in Quora, an online knowledge-sharing platform that is structured in a question-and-answer format. On Quora, people can "follow" other Quora members to indicate interest in the content produced by those people. These follow relationships help determine the questions and answers that get automatically recommended to people in their homepage feeds and digest emails; other important inputs into Quora's recommendation engines include the topics that people follow and the feedback that they have given on previous recommendations. People can give positive and negative feedback on the answers of others in the form of "upvotes" and "downvotes." Many of these interactions (including following, upvoting, and downvoting) produce relationships between people, topics, and questions that can be represented as networks, and in this article, we identify manifestations of the friendship paradox in some of these networks.

In Section 2, we begin by confirming the existence of friendship paradoxes in the follow network (i.e., the network representing follow relationships between people on Quora). The follow network is a canonical example of a directed network, and we confirm that the strong versions of the four "standard" paradoxes on directed networks occur on Quora; for most people who have both followers and followees:

- Most of their followers have more followers than they do.
- Most of their followees (i.e., people whom they follow) have more followers than

they do.

- Most of their followers follow more people than they do.
- Most of their followees follow more people than they do.

We then move on, in Section 3, to study a manifestation of the friendship paradox in a more unorthodox context: a network representing negative interactions between people in a social network. We refer to this as the "downvoting network," and in this network, each directed link represents a unique "downvoter, downvotee pair" within a certain time window. A "downvoter, downvotee pair" is a pair of people on Quora such that the "downvoter" downvoted at least one piece of content by the downvotee during the time window. In the downvoting network, we find the following:

- For most sufficiently frequent writers who have been downvoted by other sufficiently frequent writers, most of the people in their peer group who downvote them get downvoted more often than they do.
- For most people who have cast downvotes, most of the people whom they downvote get downvoted more than they do.

We refer to these observations as the "downvoting paradox." Because this paradox occurs in a network representing negative interactions, it differs markedly from most manifestations of the friendship paradox that have been studied in the past. We discuss the implications of the "downvoting paradox" and, in the Conclusion, provide an outlook for future work.

2. Friendship Paradoxes in the Quora Follow Network

We begin our analysis by studying the network of people following one another on Quora. This is the most "obvious" social network that undergirds the Quora platform and, as such, it is a natural setting to examine friendship paradoxes before looking for these phenomena in more unusual contexts. Furthermore, understanding the friendship paradox in this "standard" context allows us to develop tools that are useful in dissecting the more exotic phenomena to follow in Section 3.

2.1 Definition of the Network and Core Questions

As we alluded to in the Introduction, people typically follow one another on Quora to indicate interest in one another's content. These follow relationships then provide input into Quora's recommendation engines, which are designed to automatically surface relevant content to Quora users. Two contexts in which this occurs are the Quora homepage (where anyone with a Quora account is shown a "feed" of content that is tailored to their interests) and the "digest email." In both of these contexts, people are also shown content based on topics that they follow and their previous activity on the product. Thus, neither homepage feed nor the digest email are purely social, but the follow network *does* constitute one important input into these product features.

In studying friendship paradoxes in the follow network, we restrict our attention to people who visited the product at least once in the four weeks preceding June 1, 2016 and only include follow relationships between these people. This excludes people from the analysis who have Quora accounts, but who are not actively engaged with the product. Amongst the people who meet our criterion for activity, we sample 100,000 random individuals who have at least one follower and one followee; since these people have both incoming and outgoing links, we can actually pose questions about both types of neighbors. For these people, we ask the following questions:

- What is the individual's follower count (i.e., the individual's indegree)?
- What is the average follower count of their followees?
- What is the average follower count of their followers?
- What is the individual's followee count (i.e., the individual's outdegree)?
- What is the average followee count of their followees?
- What is the average followee count of their followers?

Note that the followers and followees of our 100,000 randomly chosen individuals must meet our criterion for activity but need not have both incoming and outgoing links. Note further that the "average" can mean either a mean or median over neighbors, and if we compute both, this gives us a list of ten values for each of our 100,000 randomly sampled individuals: two values of degree for the individual, four means over the person's neighbors, and four medians over those same neighbors. We then take a median, over the 100,000 randomly chosen instances, of each of those ten values to compute a "typical" values for each of the questions above. Since we subsample 100,000 randomly chosen people from the full population, we also compute so-called distribution-free 95% confidence intervals for these "typical" values to ensure that the paradoxes are statistically significant (Hollander, Wolfe, & Chicken, 1999).

2.2 Demonstration of the Paradoxes

Table 1 reports these typical values and suggests the existence of all four paradoxes in the follow network. Moreover, it suggests the existence of all four paradoxes in both their weak and strong versions, because the typical values of both the means and medians over neighbors are greater than the typical values of the degrees of the randomly selected individuals.

Nevertheless, Table 1 is not a completely conclusive demonstration of these paradoxes. This is because, by taking typical values over the ten quantities independently, we have ignored correlations across links in our network. To capture these, we can instead subtract an individual's degree from each of the ten averages over his or her neighbors before taking a median over the 100,000 randomly chosen instances. We do this in Table 2. This table fully accounts for correlations across links, and the values in the table show that the weak and strong versions of the degree-based paradoxes occur in the follow network: for example, for most people in the network, most of the people whom they follow have at least 28 more followers than they do.

It is worth noting that, as Table 2 shows, the weak paradoxes (where we compute the mean over neighbors) are more dramatic than the strong paradoxes (where we compute the median over neighbors): for most people, the people whom they follow have a mean of 76.2 more followers than they do but a median of 28.0 more followers. This difference in magnitude is well-explained by the arguments of Hodas et al. (Hodas et al., 2014). When we take a mean over several neighbors, we give the distribution of degrees over those neighbors more opportunities to sample an extreme value. Sometimes, when we choose a random person in our network and then choose a random followee of that person, we arrive at a followee who has an extraordinarily large follow count. This inflates the mean, resulting in a more dramatic weak paradox. The median over neighbors is less affected by extreme values, and in this sense, the strong paradox is a better reflection of the typical neighborhood of individuals in our network.

Table 1: This table reports statistics for the degrees of 100,000 randomly sampled people in the follow network. The "person" row shows the median values of indegree (follower count) and outdegree (followee count) over these randomly-sampled people. Meanwhile, the "mean follower" and "mean followee" rows show the "typical" (i.e., median) value of the mean degree of the neighbors of the randomly sampled people. Finally, the "median follower" and "median followee" rows show the "typical" (i.e., median) value of the median degree of the neighbors of the 100,000 randomly sampled people. Since we subsample the full population in these estimates, we also report a 95% confidence interval around each of our values ; our estimates are in bold.

	Typical Values of Degree	
	follower count (indegree)	followee count (outdegree)
person	[6.0, **6.0**, 6.0]	[9.0, **9.0**, 9.0]
mean follower	[35.0, **35.5**. 36.0]	[72.7, **73.5**, 74.2]
median follower	[17.0, **17.0**, 17.5]	[42.0, **42.0**, 42.5]
mean followee	[104.7, **106.3**, 108.0]	[63.8, **64.4**, 65.0]
median followee	[51.0, **52.0**, 52.0]	[32.0, **33.0**, 33.0]

Table 2: This table reports statistics for the differences in degree between 100,000 randomly sampled people in the follow network and their neighbors. The "mean follower - person" and "mean followee - person" rows show the typical (i.e., median) values of the difference between the mean degree of the neighbors of P and the degree of P for each of the randomly sampled people P. Meanwhile, the "median follower - person" and "median followee - person" rows show the typical (i.e., median) values of the difference between the median degree of the neighbors of P and the degree of P for each of the randomly sampled people P. Compared to Table 1, averaging over differences better captures correlations in degree across links in the network. Since we subsample the full population in these estimates, we also report a 95% confidence interval around each of our values; our estimates are in bold.

	Typical Values of Differences in Degree	
	follower count (indegree)	followee count (outdegree)
mean follower - person	[16.0, **16.4**, 16.7]	[49.4, **50.0**, 50.7]
median follower - person	[2.0, **2.5**, 2.5]	[20.0, **20.0**, 20.5]
mean followee - person	[75.0, **76.2**, 77.3]	[35.3, **35.8**, 36.2]
median followee - person	[27.5, **28.0**, 28.0]	[9.0, **9.5**, 10.0]

2.3 How Correlations in the Follow Network Account for the Paradoxes

Note that two of the paradoxes that we demonstrate in Table 2 are somewhat easier to reason about than the others. When we choose a random individual in the network and then choose a random follower of that person, it makes sense that the latter person tends to follow more people, since we found this person by randomly selecting one of his or her outgoing links, so people with more outgoing links are intuitively more likely to be selected by this process. Similar logic applies to the observation that a randomly selected followee of a randomly selected individual tends to have more followers, and indeed, to the standard friendship paradox in undirected networks (Feld, 1991). We will ignore the subtleties in this type of argumentation for our present purposes (Lattanzi & Singer, 2015; Cao & Ross, 2016) and instead ask: what accounts for the other two paradoxes? In particular, what accounts for the observation that most followers of most individuals have more followers and the observation that most followees of most indivudals follow more people?

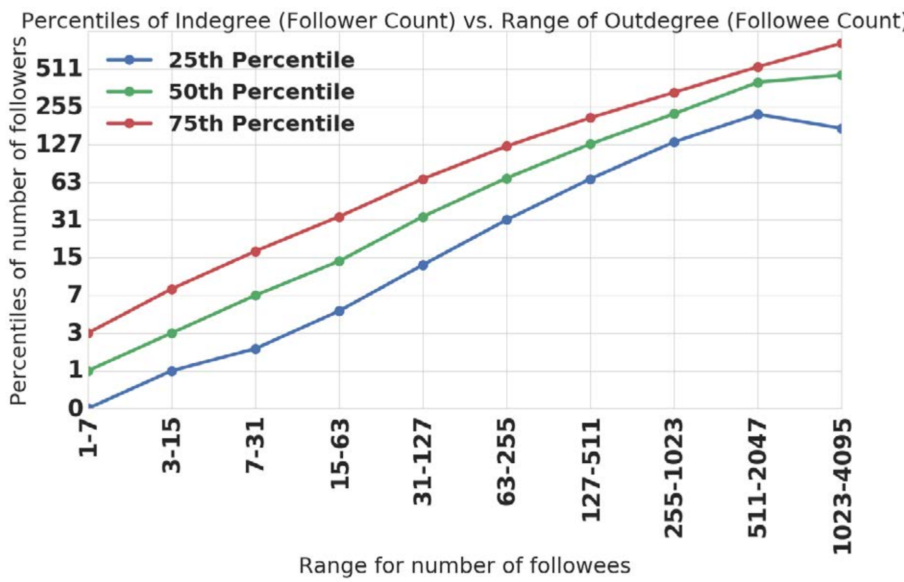

Figure 1: This plot shows percentiles of the overall indegree distribution in the follow network vs. ranges of outdegree for 100,000 randomly sampled people with at least one follower or one followee. We show the 25th, 50th, and 75th percentiles of the distribution. As we consider people who follow more and more people, the distribution of follower counts shifts to higher and higher values. This reveals strong positive correlations between indegree and outdegree in the follow network.

These phenomena are rooted in the statistical correlation between indegree and outdegree in our network. In Figure 1, we visualize these correlations by showing how the distribution of a person's indegree varies as we condition on the person having higher outdegree values. This plot shows that the joint distribution $p(k_{\text{in}}, k_{\text{out}})$ of indegree k_{in} and outdegree k_{out} exhibits strong positive correlations. Suppose *only* these within-node correlations between indegree and outdegree existed and that the network was otherwise wired randomly, thus exhibiting no across-link correlations. Such a scenario can be realized through the approach of the configuration model, where nodes are assigned "stubs" for their incoming and outgoing links and each incoming stub in the network is randomly paired to an outgoing stub (Newman, 2003). In such a scenario, when we sample a random node and then sample a random follower of the node, we would reach a neighbor in proportion to that neighbor's outdegree. It follows that the probability of reaching a neighbor with indegree

k_{in} and outdegree k_{out} would be governed by the distribution:

$$p_{\text{fwr}}(k_{\text{in}}, k_{\text{out}}) = \frac{k_{\text{out}} p(k_{\text{in}}, k_{\text{out}})}{\langle k_{\text{out}} \rangle} \tag{1}$$

Here, $\langle k_{\text{out}} \rangle$ is the mean outdegree in the network. Compared to the distribution $p(k_{\text{in}}, k_{\text{out}})$, the weighting in equation 1 clearly pushes more weight of the distribution out to higher k_{out}, helping to account for the observation that a randomly chosen follower of a randomly chosen individual follows more people. However, the key observation is that, given the positive correlations between indegree and outdegree, this weighting is also likely to push more weight of the distribution out to higher *indegree*, so this can also help explain why the randomly chosen follower typically has more followers.

In practice, the link structure between real people is definitely *not* random. Empirical networks display effects such as *assortativity* (i.e., the tendency of nodes to preferentially connect to others of similar degree). This means that, when we actually sample a random node and then a random follower of that node, we will not find that the joint distribution exactly obeys equation (1). In Figure 2, we compare the overall distribution of indegree in the follow network, the overall distribution of indegree over people who have both types of links, the expected follower distribution of indegree from equation (1), and the actual distribution of indegrees over followers that we find by repeatedly choosing a random follower of a randomly chosen node. Comparing the two-step sampling plot to the inferred plot from equation (1), we see that the shift of the distribution over followers to higher values is diminished by across-link correlations such as assortativity. Nevertheless, this shift is still strong enough so that the median of the two-step sampling distribution is higher than the median of the distribution over people with at least one follower and one followee. This suggests that correlations are likely strong enough for the paradox to survive, a fact that is confirmed by Table 2.

To recap, we have provided evidence that the strong versions of the four degree-based paradoxes occur in the network of people following one another on Quora, and we have shown how these paradoxes are rooted in the correlations between indegree and outdegree. These specific correlations are, perhaps, not surprising: people who have many followers typically have a large following because they are active participants in the product. This means that they are also likely to follow many other participants. However, dissecting this relatively common paradox has allowed us to develop machinery that will be useful in probing friendship paradoxes in a less familiar context in Section 3.

3. Paradox in the Network Induced by Downvoting

Although the network of people following one another provides valuable input into Quora's recommendation systems, in practice, people also see content for reasons that have nothing to do with whom they follow. This is, for example, because they follow topics or directly follow questions. In other words, the network of people following one another is not synonymous with the actual network of interactions on Quora. In this section, we show that friendship-paradox phenomena also exist in "induced networks" of real interactions on the product. We focus on a specific interaction, the answer downvote, for which we identify the special variant of the friendship paradox that we referred to in the introduction as the "downvoting paradox."

3.1 Primer on Upvoting and Downvoting

Before examining the data, we will briefly give some more context on what "downvoting" means. As we mentioned in the Introduction, anyone with a Quora account has the

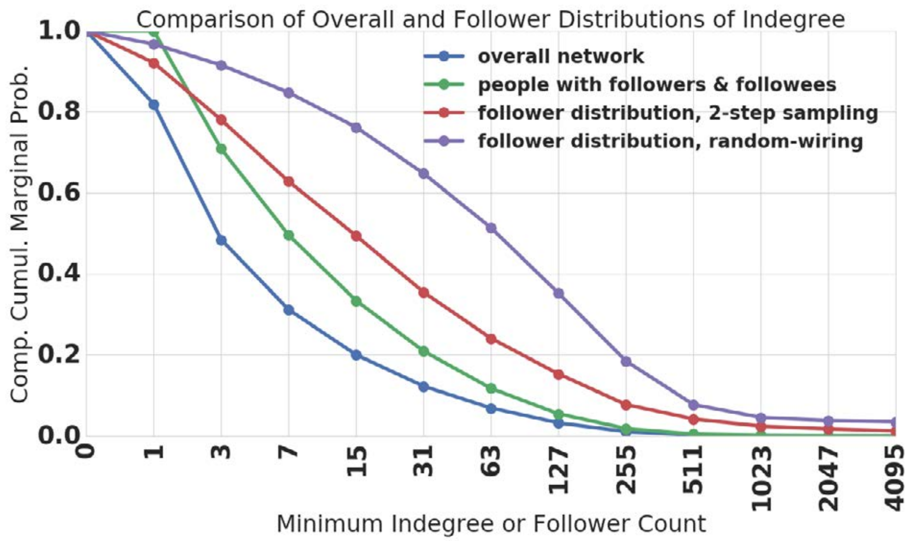

Figure 2: This plot shows four distributions of indegree (i.e., follower count). We plot complementary cumulative marginal distributions, which show probabilities that the indegree is at least the value on the x-axis. In blue, we show the real distribution of indegree over 100,000 randomly selected people in the follow network who had at least one follower or one followee. In green, we show the real distribution of indegree over 100,000 randomly chosen people in the follow network who had at least one follower *and* one followee. In red, we show the distribution of indegree over followers that we find if we repeatedly randomly sample an individual from our 100,000 randomly chosen people with at least one follower and one followee and then randomly sample one of that person's followers (we perform this two-step sampling 10,000 times). Finally, in purple, we show the inferred distribution of indegree over followers that we would expect if we apply the random-wiring assumption in equation (1) to our 100,000 randomly selected people with at least one follower or one followee.

opportunity to provide feedback on any answer by "upvoting" or "downvoting." People typically upvote an answer because they consider it to be factually correct, because they agree with the opinions expressed in the answer, or because they consider it to be otherwise compelling reading. Upvotes are one of a large number of signals that are used to rank an answer, relative to other answers to the same question, on the Quora page for that question. People have the opportunity to cast an upvote anonymously, but in the majority of cases, people do not elect this option. If an answer is upvoted publicly, then it also serves a social distribution role: people who follow the upvoter are more likely to be shown that answer in their homepage feeds and digests. Moreover, when an upvote is cast publicly, the upvoter's name is displayed publicly on the answer to indicate that he or she took that action.

In many ways, "downvoting" is the complementary negative action to upvoting, but there are some important differences. People typically downvote to indicate they believe an answer to be factually incorrect or low quality. This is used as negative signal for ranking answers on question pages. Naturally, downvoting does *not* serve a social distribution function: the downvoter's followers are not more likely to see the answer as a result of the action. In contrast to upvoting, information about who downvoted an answer is never displayed publicly, and even the person who got downvoted (the "downvotee") does not have access to the identities of his or her downvoters.

Both upvoting and downvoting provide compelling settings to study friendship paradoxes on Quora, but there are certain properties of downvoting that make it particularly interesting. First, the relationship between a downvoter and downvotee is a negative one, and friendship paradoxes, as the term "friendship" itself suggests, are rarely studied in these contexts. Moreover, as we have mentioned above, the identities of downvoters are hidden on the product, and this precludes certain types of social explanations for network paradoxes. We will see how this impacts the interpretation of our results later in this section.

3.2 Definition of the Network and Core Questions

We now define our "downvoting network": we begin with all downvotes cast on public, non-deleted answers in the four weeks preceding June 1, 2016. To avoid trespassing upon anonymity, we restrict our attention to downvotes where the downvoter was not anonymous on the question. Even with "public" downvotes (we use quotation marks here because downvotes are not publicly shown even if the person has not elected the anonymous option), we only look at aggregate statistical properties of the data for the purposes of this analysis, never at the names of specific people involved in downvoting interaction). We draw a directed link for every unique downvoter, downvotee pair within the four-week window. A cartoon version of the "downvoting network" is illustrated in Figure 3.

For the network depicted in Figure 3, we ask the following questions:

1. **The "downvoter $\rightarrow$ downvotee" question**: For most downvotees (i.e., people who have been downvoted), do most of their downvoters get downvoted more or less than they do?

2. **The "downvotee $\rightarrow$ downvoter" question**: For most downvoters, do most of their downvotees get downvoted more or less than they do?

Note that the answers that we provide to these questions below are not peculiar to the specific four-week window under consideration; we have checked that they hold for other four-week windows as well.

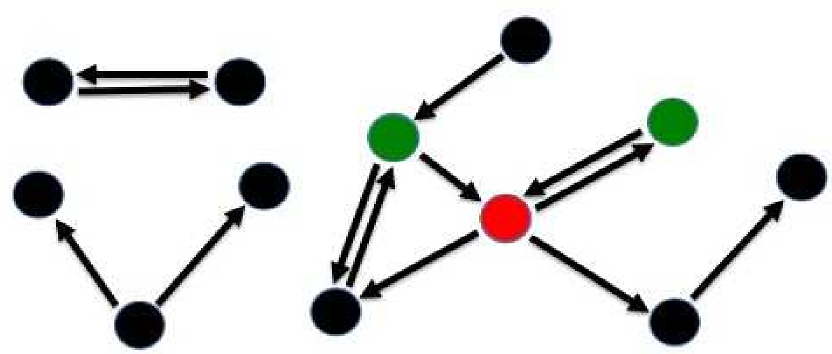

Figure 3: A cartoon illustration of the downvoting network, representing the downvotes within a four-week period on Quora. A directed link exists between two nodes if the person represented by the origin node (the "downvoter") cast at least one downvote on any answer by the person represented by the target node (the downvotee) during the four-week period. In this diagram, the nodes in green represent all the unique downvoters of a particular downvotee, who is represented by the node in red.

Table 3: This table reports the typical differences in the number of downvotes received by people and their average "neighbors" in the "downvotee $\rightarrow$ downvoter" and "downvoter $\rightarrow$ downvotee" questions. The table reveals that the analog of the friendship paradox occurs in the "downvoter $\rightarrow$ downvotee" question (i.e., the right column), but not in the "downvotee $\rightarrow$ downvoter" question. Please see the text of Section 3 for the details of the calculations that lead to this table.

Typical Values of Differences in Downvotes Received		
	downvotee $\rightarrow$ downvoter	downvoter $\rightarrow$ downvotee
mean downvoter - downvotee	-1.0	-39.0
median downvoter - downvotee	-1.0	-29.0

3.3 The Downvoting Paradox Does Not Occur in the Full Downvoting Network

Table 3 provides the answers to the "downvotee $\rightarrow$ downvoter" and "downvoter $\rightarrow$ downvotee" questions over the entire downvoting network. Figure 3 can assist in interpreting this table. Suppose we focus on the "downvotee $\rightarrow$ downvoter" question. In the figure, the node pictured in red represents a downvotee, and by tracing back the links that point at this node, we can find all the downvoters of the downvotee; these downvoters are pictured in green. In the "downvotee $\rightarrow$ downvoter" question, we take an average (mean or median) of the number of downvotes received by the green nodes, subtract the number of downvotes received by the red node, and then repeat this calculation for all downvotees (i.e., all nodes with at least one incoming link - the red node is one example, but there are ten in the figure). We finally take a median over all downvotees to obtain the differences that we report in the "downvotee $\rightarrow$ downvoter" column of Table 3. An analogous calculation leads to the "downvoter $\rightarrow$ downvotee" column, except that we iterate over downvoters and take averages over their downvotees. Note that we do not do any subsampling of downvotees in the "downvotee $\rightarrow$ downvoter" question nor any subsampling of downvoters in the "downvoter $\rightarrow$ downvotee" question. Thus, the statistics reported in Table 3 are population values for the four-week window that the downvoting network represents.

Table 3 shows that the analog of the friendship paradox occurs in the "downvoter $\rightarrow$ downvotee" question (i.e., the right-hand column) but not in the "downvotee $\rightarrow$ downvoter" question. For most downvoters, most of their downvotees get downvoted more than they do; however, for most downvotees, most of their downvoters get downvoted *less* than they do. Is the latter fact surprising? From a product perspective, it may not be: maybe most downvotees get downvoted for understandable reasons (e.g., writing controversial or factually incorrect content), and therefore, we ought to expect them to get downvoted more than their. However, it is instructive to translate the "downvotee $\rightarrow$ downvoter" observation back into the language of the previous section, where we examined friendship paradoxes in the follow network. In that language, this is analogous to finding that, for most people, most of their followers have fewer followers than they do. As we saw, positive correlations between indegree and outdegree actually produce the opposite trend, so where does that break down in the downvoting network?

The nodes in the downvoting network can be characterized by four variables:

- k_{in}: the number of unique downvoters of the person (i.e., that person's indegree in the downvoting network).
- d_{in}: the number of downvotes the person received, which should respect $d_{\text{in}} \geq k_{\text{in}}$.

- k_{out}: the number of unique downvotees of the person (i.e., that person's outdegree in the downvoting network).
- d_{out}: the total number of downvotes the person cast, which should respect $d_{\text{out}} \geq k_{\text{out}}$.

We refer to the joint distribution over these variables as $p(k_{\text{in}}, d_{\text{in}}, k_{\text{out}}, d_{\text{out}})$. If we were to choose a random downvotee in the network and then choose a random downvoter of that person, under a random-wiring assumption, we should expect the joint distribution of these four variables over the downvoter to obey:

$$p_{\text{dvr}}(k_{\text{in}}, d_{\text{in}}, k_{\text{out}}, d_{\text{out}}) = \frac{k_{\text{out}} p(k_{\text{in}}, d_{\text{in}}, k_{\text{out}}, d_{\text{out}})}{< k_{\text{out}} >} \tag{2}$$

If k_{out} is positively correlated with d_{in}, we would expect the marginal distribution of d_{in} over downvoters to be shifted to higher values compared to the overall distribution of d_{in} in the network, potentially resulting in a paradox in the "downvotee $\rightarrow$ downvoter" question.

In Figure 4, we bucket people by their value of k_{out} and plot percentiles of their distribution of d_{in}. The plot shows that, over a large range of downvotee counts, the majority of people received no downvotes at all. We call these people "undownvoted downvoters." This observation allows us to attach a story to the "downvotee $\rightarrow$ downvoter" column of Table 3: the typical values of the differences in this column are -1 because the typical downvotee is someone who got downvoted once during the four-week window by someone who did not get downvoted at all.

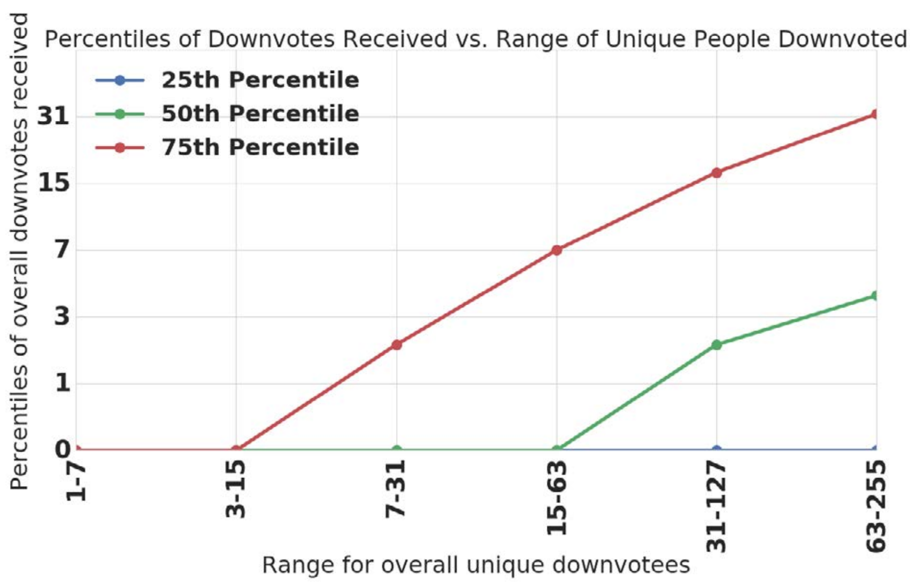

Figure 4: This plot shows percentiles of the number of downvotes an individual received vs. ranges of the number of unique people that individual downvoted. We show the 25th, 50th, and 75th percentiles of the distribution. This plot shows that, over a large range of unique downvotee counts, the median number of downvotes received is zero. In other words, the distribution is strongly affected by the presence of "undownvoted downvoters."

3.4 The Downvoting Paradox Occurs When Focus on Sufficiently Frequent Writers

This observation, in turn, motivates a question about why "undownvoted downvoters" play such an influential role in the averages that are reported in Table 3. One possible explanation is that there is a barrier to being a downvotee that does not exist for being a downvoter:

namely, to be a downvotee, you must have actually written answers. As such, it is interesting to reframe both the "downvotee $\rightarrow$ downvoter" and "downvoter $\rightarrow$ downvotee" questions for solely those people who have performed the prerequisite actions that are necessary to get downvoted:

1. **The revised "downvotee $\rightarrow$ downvoter" question**: For most downvotees who have written at least n recent answers, do most of their downvoters who have written at least n recent answers get downvoted more or less than they do?

2. **The revised "downvoter $\rightarrow$ downvotee" question**: For most downvoters who have written at least n recent answers, do most of their downvotees who have written at least n recent answers get downvoted more or less than they do?

We specifically consider people who wrote at least $n = 3$ non-anonymous answers during the four-week period for which the downvoting network is constructed; however, our findings hold for higher n as well. Note that the variable that we compare between downvoters and downvotees is still *total* downvotes received, not just downvotes received from others meeting the content contribution condition. Furthermore, note that, in the revised questions, to be considered a "downvotee," a person must satisfy the content contribution condition themselves and must have been downvoted by someone who also meets the condition; the same applies to "downvoters."

Before recomputing Table 3 for the revised questions, we first consider how the content contribution condition impacts correlations in our network. We should adjust our definitions of the four variables for each node in light of our modified questions:

- $\tilde{k}_{\text{in}}$: the number of unique downvoters of the person who have written at least n answers.

- d_{in}: the total number of downvotes the person received.

- $\tilde{k}_{\text{out}}$: the number of unique downvotees of the person who have written at least n answers.

- d_{out}: the total number of downvotes the person cast.

Figure 5 probes correlations between $\tilde{k}_{\text{out}}$ and d_{in} and shows that the content contribution restriction has restored positive correlations between these two variables.

We can go further and produce the analog of Figure 2. We do this in Figure 6. The plot shows that the correlations seen in Figure 5 do produce an outward shift of the majority of the distribution of d_{in} over downvoters in the "downvotee $\rightarrow$ downvoter" question, both under the random-wiring assumption and when we account for across-link correlations by randomly sampling a downvotee and then a downvoter. In fact, the random-wiring assumption works very well in this setting, with minimal corrections due to across-link correlations.

In Table 4, we show that the correlations that we explored in Figures 5 and 6 do result in paradoxes in both the revised "downvotee $\rightarrow$ downvoter" and "downvoter $\rightarrow$ downvotee" questions. Thus, both sides of the "downvoting paradox" hold once the content contribution condition is imposed.

3.5 Implications of the Observations

What does this realization of the "downvoting paradox" actually imply? First, it provides an example of friendship-paradox phenomena occurring within networks representing negative interactions, and as we mentioned previously, this is underexplored territory. Second,

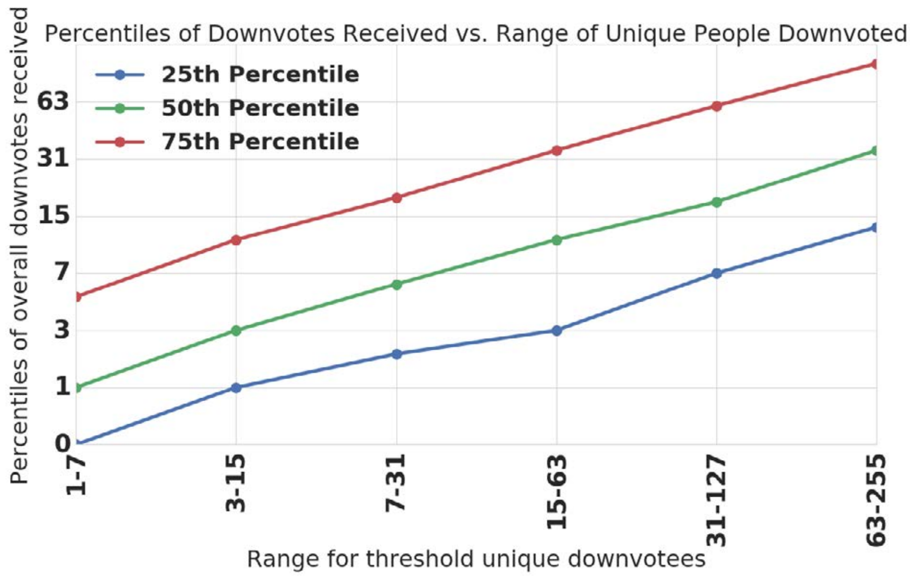

Figure 5: This plot, like Figure 4, shows percentiles of the number of downvotes an individual received vs. ranges of the number of unique people that individual downvoted. The difference with respect to Figure 4 is that we have imposed the content-contribution threshold that we discuss in the test. This means that all people considered for this plot contributed at least $n = 3$ non-anonymous answers during the four-week window represented by the downvoting network. Furthermore, the number of "threshold unique downvotees" for each individual only counts those downvotees who also satisfy the content-contribution criteria. Meanwhile, the number of "overall downvotes received" still includes all downvotes received from any downvoter, not just those who satisfy the content-contribution threshold.

it shows that these type of phenomena can occur in networks representing interactions that are hidden from the actors. This is important because one candidate explanation for the downvoting paradox might be *retaliation*: we could reason that the typical downvoter gets downvoted more frequently in the revised "downvotee $\rightarrow$ downvoter" question because people "get back" at these people by downvoting their answers. However, the product mechanics of Quora essentially rule out this explanation, and we need to consider alternatives.

Table 5 provides evidence that supports one alternative explanation: when we choose a random downvotee and then choose a random downvoter of the downvotee, the downvoter has typically contributed more answers during the four-week window. Thus, a "content contribution paradox" accompanies the "downvoting paradox" and implies that the downvoter, by virtue of contributing more content, had more opportunities to be downvoted. Nevertheless, this is far from a guarantee that the downvoter *will* be downvoted more in the typical case, so other potential explanations (e.g., that the downvoter in the "downvotee $\rightarrow$ downvoter" question typically writes more controversial content) may also be at play. A deeper natural language study would be needed to tease out the extent to which these different factors contribute to the realization of the paradox. Finally, we note that the "downvotee $\rightarrow$ downvoter" side of the downvoting paradox turns the usually demoralizing nature of the friendship paradox ("your friends have more friends than you do") on its head: it may be comforting to content contributors that the people in their peer group who give them negative feedback are no more immune to getting negative feedback themselves.

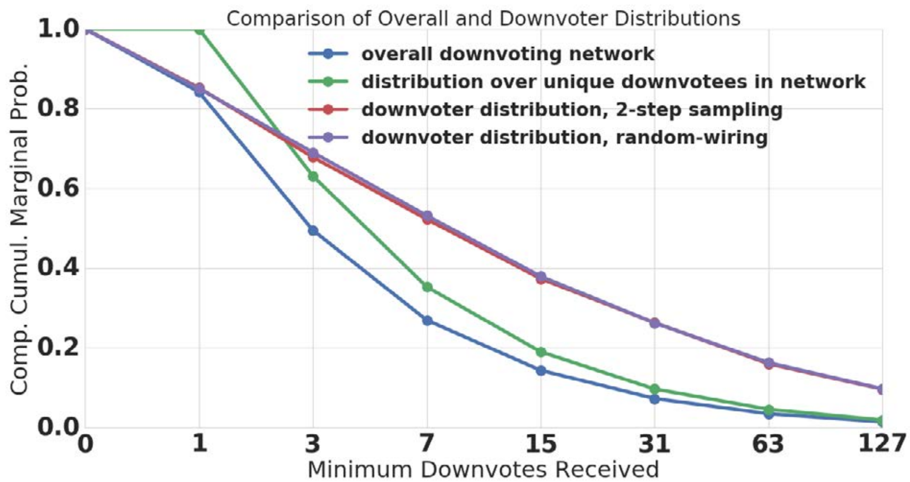

Figure 6: This plot shows four distributions of the number of downvotes received by people in the downvoting network. In blue, we plot the real distribution of downvotes received over all sufficiently active writers in the downvoting network (i.e., those who wrote at least $n = 3$ answers during the time period that the network represents. In red, we plot the real the distribution of downvotes received over sufficiently active writers who received at least one downvote from another sufficiently active writer. In green, we plot the distribution of downvotes received that we find if we repeatedly randomly sample a downvotee then a downvoter (we do this two-step sampling 10,000 times). Finally, in purple, we plot the inferred distribution of downvotes received when sampling a downvotee and then a downvoter under the random-wiring assumption. As in 2, we plot complementary cumulative marginal probabilities. The distributions are computed over people who contributed at least $n = 3$ non-anonymous answers during the four-week window represented by the downvoting network. Furthermore, when we randomly sample a downvotee and then a downvoter, we require that both parties satisfy the threshold. However, the number of downvotes received still includes all downvotes received from any downvoter, not just those who satisfy the content-contribution threshold.

4. Conclusion

In this article, we have examined various manifestations of the friendship paradox on Quora. We first demonstrated that the "standard" directed-network variants of the friendship paradox hold for the network of people following one another. We then took the toolkit that we used to explore friendship paradoxes in the follow network and used it to study the "induced" network of people downvoting one another over a given time period. This revealed the existence, in certain contexts, of a variant of the friendship paradox that we have called the "downvoting paradox." We repeat what the "downvoting paradox" entails here:

- For most sufficiently frequent writers who have been downvoted by other sufficiently frequent writers, most of the people in their peer group who downvote them get downvoted more often than they do.
- For most people who have cast downvotes, most of the people whom they downvote get downvoted more than they do.

Our analysis of the downvoting paradox motivates many follow-up questions. For example, to what extent is the downvoting paradox explained by a correlation between downvoting and increased content contribution, and to what extent is it driven by an increased

Table 4: In this table, we report statistics that we obtain when we repeat the calculations that led to Table 3 but restrict our attention to downvotees and downvoters who contributed at least $n = 3$ non-anonymous answers during the four-week window that our downvoting network represents. Note that the variable that we compare between downvoters and downvotees is still *total* downvotes received, not just downvotes received from active contributors.

Typical Values of Differences in Downvotes Received		
	downvotee $\rightarrow$ downvoter	downvoter $\rightarrow$ downvotee
mean downvoter - downvotee	5.0	-44.5
median downvoter - downvotee	2.0	-23.0

Table 5: In this table, we report statistics that we obtain when we repeat the calculations that led to Table 4, including imposition of the content-contribution threshold. However, the variable that we compare between downvoters and downvotees is now the number of non-anonymous answers contributed during the four-week window represented by the downvoting network.

Typical Values of Differences in Answers Written		
	downvotee $\rightarrow$ downvoter	downvoter $\rightarrow$ downvotee
mean downvoter - downvotee	8.0	-28.7
median downvoter - downvotee	4.0	-17.0

tendency of downvoters to produce controversial content themselves? Moreover, downvoting on Quora represents a very particular type of negative interaction. The identity of downvoters is hidden from downvotees and this can have important consequences for the behavior of these parties: downvoters may feel freer to give negative feedback if they are not publicly identified, and the downvotees cannot retaliate against any specific individual if they believe that they have been downvoted. Does something like the downvoting paradox survive if the underlying product principles are different (e.g., if the identity of downvoters is public), or would such a situation fundamentally alter the dynamics? It may be possible to address these questions by analyzing friendship paradoxes in networks representing other types of negative interactions in online or real-world social networks.

It is also worth noting that, in contrast to many networks in which friendship paradoxes are studied, the downvoting paradox occurs in an "induced network" representing real interactions during some time period. This is *not* an explicit network that persists over time in the product. It is likely that variants of the friendship paradox exist in several other induced networks as well, both on Quora and other social producs. Such paradoxes may even be found in products with no explicit "following" or "friendship" structure. Thus, our study may point the way towards identifying other phenomena of this type in various online social products.

However, our study also indicates some nuances in how these paradoxes need to be measured. In the case of downvoting, there is a class of participant in the downvoting network (the "undownvoted downvoter") who invalidates the paradox for the full network. This is because these people, while participating in one side of the interaction (i.e., downvoting), have typically not performed some other prerequisite action (i.e., answering sufficiently often) that would subject them to the other side of the interaction (i.e., getting downvoted). Understanding these product features and accounting for them (e.g., by considering only interactions between active contributors) can help expose other phenomena of this type.

Acknowledgements: We thank Kristina Lerman and Yaron Singer for helpful discussions on their recent work on the friendship paradox. We also thank the Quora data science team for helpful comments on this blog post. We especially thank William Chen for reviewing important parts of the code, and William, Kao Zi Chong, and Olivia Angiuli for introducing us to the distribution-free method for computing confidence intervals for percentile metrics. Finally, we thank Paula Griffin for helping to coordinate the presentation of this work at the 2016 Joint Statistical Meetings.

References

Bollen, J., Gonçalves, B., van de Leemput, I., & Ruan, G. (2016). The happiness paradox: your friends are happier than you. *arXiv preprint arXiv:1602.02665*.

Cao, Y., & Ross, S. M. (2016). The friendship paradox. *Mathematical Scientist*, *41*(1).

Christakis, N. A., & Fowler, J. H. (2010). Social network sensors for early detection of contagious outbreaks. *PLOS One*, *5*(9), e12948.

Cohen, R., Havlin, S., & Ben-Avraham, D. (2003). Efficient immunization strategies for computer networks and populations. *Physical Review Letters*, *91*(24), 247901.

Eom, Y.-H., & Jo, H.-H. (2014). Generalized friendship paradox in complex networks: The case of scientific collaboration. *Scientific Reports*, *4*, 4603.

Feld, S. L. (1991). Why your friends have more friends than you do. *American Journal of Sociology*, 1464–1477.

Fotouhi, B., Momeni, N., & Rabbat, M. G. (2014). Generalized friendship paradox: An analytical approach. In *International conference on social informatics* (pp. 339–352).

Hodas, N., Kooti, F., & Lerman, K. (2013). Friendship paradox redux: Your friends are more interesting than you. *arXiv preprint arXiv:1304.3480*.

Hodas, N., Kooti, F., & Lerman, K. (2014). Network weirdness: Exploring the origins of network paradoxes. In *Proceedings of the international conference on web and social media* (pp. 8–10).

Hollander, M., Wolfe, D., & Chicken, E. (1999). Wiley series in probability and statistics. *Nonparametric Statistical Methods, Third Edition*, 821–828.

Jackson, M. O. (2016). The friendship paradox and systematic biases in perceptions and social norms. *Available at SSRN*.

Jo, H.-H., & Eom, Y.-H. (2014). Generalized friendship paradox in networks with tunable degree-attribute correlation. *Physical Review E*, *90*(2), 022809.

Lattanzi, S., & Singer, Y. (2015). The power of random neighbors in social networks. In *Proceedings of the eighth acm international conference on web search and data mining* (pp. 77–86).

Lerman, K., Yan, X., & Wu, X.-Z. (2016). The" majority illusion" in social networks. *PLOS One*, *11*(2), e0147617.

Newman, M. E. (2003). The structure and function of complex networks. *SIAM review*, *45*(2), 167–256.

Ugander, J., Karrer, B., Backstrom, L., & Marlow, C. (2011). The anatomy of the facebook social graph. *arXiv preprint arXiv:1111.4503*.