

The Effect of Recommendations on Network Structure

Jessica Su
Stanford University
jtysu@stanford.edu

Aneesh Sharma
Twitter
aneesh@twitter.com

Sharad Goel
Stanford University
scgoel@stanford.edu

ABSTRACT

Online social networks regularly offer users personalized, algorithmic suggestions of whom to connect to. Here we examine the aggregate effects of such recommendations on network structure, focusing on whether these recommendations increase the popularity of niche users or, conversely, those who are already popular. We investigate this issue by empirically and theoretically analyzing abrupt changes in Twitter’s network structure around the mid-2010 introduction of its “Who to Follow” feature. We find that users across the popularity spectrum benefitted from the recommendations; however, the most popular users profited substantially more than average. We trace this “rich get richer” phenomenon to three intertwined factors. First, as is typical of network recommenders, the system relies on a “friend-of-friend”-style algorithm, which we show generally results in users being recommended proportional to their degree. Second, we find that the baseline growth rate of users is sublinear in degree. This mismatch between the recommender and the natural network dynamics thus alters the structural evolution of the network. Finally, we find that people are much more likely to respond positively to recommendations for popular users—perhaps because of their greater name recognition—further amplifying the cumulative advantage of well-known individuals.

Keywords

Cumulative advantage, network evolution, social networks, Twitter

1. INTRODUCTION

It is now commonplace for individuals to turn to personalized, algorithmic recommendations to find products and information, including news, music, movies, and books. Significant effort has gone into designing and optimizing recommendation engines, but the aggregate effects of such systems are still poorly understood. In particular, there is debate

over what effect recommenders have on the overall marketplace. Anderson [2] and others have argued that recommenders primarily help individuals discover niche content, raising the fortunes of small, obscure producers at the cost of traditional hit-makers. Fleder and Hosanagar [13] have countered that recommenders, particularly collaborative filters, can lead to “rich get richer” effects [21], with already-popular products accruing most of the benefits.

Here we consider the setting of online social networks, and investigate the theoretical and empirical effects of recommendations on these platforms. Focusing on Twitter, we analyze abrupt changes in the network following the July 2010 introduction of its “Who To Follow” feature [15], one of the largest and most active network recommendation systems. We find that though users across the popularity spectrum benefitted from recommendations, the recommender disproportionately accelerated the growth of already-popular users, corroborating theories of cumulative advantage. We further find the system increased triadic closure and promoted the formation of uni-directional network ties. By treating the recommender’s introduction as a “natural experiment” [10, 25], whose precise timing was largely unrelated to other significant events, we are able to estimate the causal impact of the recommender on network structure, sidestepping concerns that often plague traditional observational analysis.

We attribute the observed structural changes to three subtle and interrelated factors. First, as is typical of network recommenders, Twitter relies on a “friend-of-friend”-style algorithm. We show, both analytically and empirically, that the total number of times an individual is recommended in such systems grows linearly in popularity (i.e., linearly in the number of people who follow the user). Second, we find that users have baseline growth rate that is *sublinear* in popularity. Notably, standard models of network growth such as preferential attachment [6] predict linear growth. This mismatch between the rate at which a user is recommended and their natural growth rate alters the structural evolution of the network, bolstering popular users and increasing the formation of uni-directional ties (since connections to popular users are typically unreciprocated). Finally, we find that individuals are much more likely to respond positively to recommendations for popular users—perhaps because of their greater name recognition—further amplifying the cumulative advantage of well-known users.

2. RELATED WORK

There is an extensive literature on designing and optimizing social network recommendation systems, which we only briefly review here. The friend-of-friend algorithm and similar approaches, such as personalized PageRank, generate recommendations by exploring a user’s local network neighborhood [4, 15]. Such personalized network algorithms are widely used, and have been found to be quite effective in driving recommendations on a variety of social platforms [3, 8, 15]. For a thorough overview of this area, we point the interested reader to some relevant papers and surveys [1, 18, 19, 20], and the references therein.

In the paper most closely related to our work here, Daly et al. [9] conducted a large-scale user study on IBM’s Social-Blue social network site to examine the effects of four different network recommenders, including a friend-of-friend algorithm and content-based recommenders. Users were partitioned into groups of approximately 600 individuals, and each group was exclusively shown recommendations derived from a single algorithm. The authors measured several properties of the resulting subnetworks, including their degree distributions, clustering, and user activity. Notably, and consistent with our results, they find that the friend-of-friend recommender conferred disproportionate gains to popular users, though the paper did not investigate the underlying mechanism for this phenomenon.

Finally, our work touches on research that considers the effects of recommendations on content diversity. In particular, Pariser [24], Sunstein [27], and others have warned that algorithmic recommendations can create “filter bubbles” or “echo chambers,” in which individuals are largely exposed to conforming opinions. There is a variety of work that attempts to measure and address such concerns in network recommendations [26, 29]. The empirical investigations of this issue have found the effects to be real, but often relatively small in magnitude [5, 12]. However, highlighting the complexity of the issue, Hosanagar et al. [17] find that recommenders can in fact *increase* commonalities—rather than segregate individuals—by simultaneously increasing overall consumption and pushing users toward similar products. Pertaining to these findings, the network changes we observe likely alter the diversity of information users consume on the platform, though we leave rigorous analysis of such consequences to future work.

3. DATA

On July 30, 2010, Twitter launched the “Who to Follow” network recommender [15], and rolled out the service to all users over the course of approximately one week. Recommendations appear on the main page of the Twitter website, in a module next to the user’s timeline, and three recommendations are typically shown at a time (see Figure 1). As explained in more detail in Section 5.1, these recommendations are often drawn from one’s “friends-of-friends,” in which case one or more mutual contacts are also named in the recommendation (i.e., in the “Followed by” field). Individuals can immediately follow the suggested user by clicking the “Follow” button. Recommendations are refreshed on each visit to the site. In addition, users can actively request to see a new set of three suggested users, or visit a recommendation page with a longer list of recommendations.




Figure 1: Twitter’s “Who to Follow” network recommendations, introduced in July 2010.

Our analysis is primarily based on two distinct sources of information. First, we use edge creation data to trace the evolution of the Twitter network. Throughout our analysis we use the convention that a directed edge from u to v indicates that u follows v ; equivalently we say that u is a *follower* of v , and that v is a *friend* of u . A user may repeatedly follow and un-follow the same individual. The data, however, only record the time of the last event (either a follow or an un-follow action) in this sequence. It is thus impossible to reconstruct the exact state of the network for each point in history. However, periodic snapshots of this network information are stored, and so we can approximate the creation dates of edges. Specifically, we start with a network snapshot taken on August 7, 2011 (about one year after the recommender was introduced), and approximate the creation time for any edge that was still in existence on that date to be the last time the edge was updated. For the vast majority of edges, this approximate creation date should coincide with the true, first creation date. This approach also discards edges that were created and soon thereafter deleted (i.e., created and deleted before August 7, 2011), which is largely desirable for our analysis. However, for edges that were created, deleted, and then recreated, we would misinfer their original creation times.

The second source of information we use is the timestamped recommendations themselves, as well as clicks on the follow links for these recommendations. Unfortunately, these data are no longer available for the period around the launch of the recommender system in 2010, and so we instead use data from October 6–12, 2015. It is, however, unlikely that any of our qualitative findings are affected by this choice, as discussed in more detail below.

4. THE RICH GET RICHER BUT A RISING TIDE LIFTS ALL BOATS

One of the primary goals of network recommendations is to encourage tie formation, and so we start our investigation by examining whether the recommendation system did in fact spur edge creation. In this context, the effect of the recommender is typically measured by counting clicks on the “Follow” button in the “Who to Follow” module. However, as has been noted previously [25], such an estimation scheme can lead to spurious results. In particular, recommenda-




Figure 2: The number of new edges created on Twitter each day during the four weeks before and after introduction of the “Who to Follow” recommender system. The recommender was rolled out over approximately one week starting on July 30, 2010 (indicated by the dashed, vertical black lines). The solid gray line shows the number of edges added each day, and the solid black line shows this number averaged over successive weeks. The two dashed horizontal lines indicate the average daily number of edges added in the four weeks before the recommender was introduced (7.7M) and the four weeks after (9.5M).

tions may encourage follow actions even in the absence of a click (e.g., by increasing an individual’s awareness of the candidate user, akin to brand advertising), leading one to underestimate the effect of the recommender. Conversely, when individuals click on the “Follow” button, they might have independently followed the recommended user even if they had not seen the recommendation, leading to an overestimate.

To deal with these issues of confounding, we examine aggregate changes in tie formation surrounding the recommender’s introduction. Figure 2 shows the number of new edges created on Twitter during the four weeks before and after the recommender’s rollout period, where the gray line plots the number of new edges added each day, and the black line plots the average over each week. As the figure clearly shows, there is a sudden and dramatic increase in edge growth that coincides with the introduction of the recommender. In the four weeks prior to the recommender, average daily edge growth was 7.7 million; in the four weeks after the rollout period, it was 9.5 million, a jump of 22%. A t-test on the difference in average edge growth rates between the pre- and post-recommender periods confirms the uptick is statistically significant, as does a regression that controls for growth in the user base and day of week effects. Moreover, since other major Twitter products were not released at this time, it appears that the recommender did

Relative boost from recommender




Figure 3: The relative increase in number of followers accrued after the recommender was introduced, as a function of users’ pre-recommender follower count. The change is computed by comparing the number of followers accrued in the four weeks after the recommender’s rollout relative to the four weeks prior to its introduction. Users across the popularity spectrum experienced sizable gains, but the most popular users benefitted substantially more than average.

causally and substantially increase the rate at which edges were created.

We next examine how these gains in edge growth were distributed across users. There are at least three reasonable possibilities. First, one might conjecture that all users benefitted equally, with everyone expanding their follower base by the overall growth rate of approximately 20%. Second, in line with the conventional wisdom popularized by Anderson [2] and others, one might guess that the recommender disproportionately benefitted the least popular users, who have few alternative means for garnering attention. Finally, appealing to theories of cumulative advantage [13, 21], one might hypothesize that the most popular users disproportionately accrued the benefits.

To check, we compute the rate at which users expanded their follower base as a function of their pre-recommender follower count. Specifically, we first subset to users who joined Twitter at least 30 days before the introduction of the recommender system. We then logarithmically bin users based on their follower count 30 days before the recommender was introduced (i.e., the bins contain users with 1–9 followers, 10–99 followers, 100–999 followers, etc.). Next, for each user bin, we compute the number of new followers collectively accrued by those users during the four weeks prior to the recommender’s introduction, and the four weeks after the recommender’s one-week roll out period. Finally, we compute the percent change in follower growth for each set of binned users.

The result of this growth computation is shown in Figure 3. The figure demonstrates that the most popular users—

those with over 10,000 followers—do indeed benefit disproportionately from the recommender, corroborating theories of cumulative advantage. In particular, those with 100,000 followers or more experience gains of 50% in their daily growth rate, more than twice the overall average of 20%. The figure, however, also reveals that users across the popularity spectrum are boosted by the recommender. The least popular users (with 1-9 followers) to the modestly popular (with 1,000 to 10,000 followers) all show sizable gains of roughly 15% in their daily growth rate. It thus appears that while the rich get richer, the rising tide of the recommender also lifts all users.

5. A MECHANISM FOR CUMULATIVE ADVANTAGE

Why is it that the most popular users disproportionately benefit from Twitter’s recommendation system? Fleder and Hosanagar [13] suggest one possibility in the context of traditional product recommendations. Such recommender systems, they argue, rely on historical sales and ratings, and so products with limited exposure are at a disadvantage. As a result, popular products, with known quality, are disproportionately promoted. A similar explanation could hold for social network recommendations. It might be safer and easier to recommend the most popular users, who have established track records and active user bases. Though plausible, we uncover an alternative mechanism that results from a subtle interplay between the recommender, the natural network dynamics, and user preferences, as described below.

5.1 The Effect of Popularity on Friend-of-Friend Recommenders

We start by analyzing the network recommendation algorithm itself. As with most social networks, Twitter relies on a “friend-of-friend”-style recommendation system. The key idea in such systems is that recommendations for any individual u are generated by first constructing a candidate set consisting of the users followed by the users u follows. Since the people one follows are commonly referred to as “friends”, the candidate set consists of u ’s friends-of-friends. Candidates within this set can be ranked by a variety of closely related metrics, such as the number of paths from u to the candidate user, or the likelihood that a random walk started at u lands on the candidate. This basic scheme has proven to be quite effective, as evidenced, for example, by the relatively high rates at which users act on the recommendations. Further, friend-of-friend systems typically outperform content-based recommenders, in part because contextual signals in this domain are relatively weak (e.g., there are no explicit ratings), and in part because social proximity ostensibly has value in and of itself in this setting.

In fact, friend-of-friend recommendations are deterministic, with candidates ranked by a pre-specified metric and the top k (e.g., $k = 3$) candidates shown to the user. In practice, ranking measures are computationally expensive to determine exactly on large networks, so approximations are used [15]. Moreover, the candidate rankings change moment to moment as the network evolves. For our theoretical analysis, we abstract away these implementation details, and instead consider a probabilistic version of the basic friend-of-friend algorithm. Namely, similar in spirit to personalized PageRank [16], we assume recommendations for user u are

generated by performing a two-step random walk starting at u , as stated formally in Definition 1 below.

DEFINITION 1. *Friend-of-friend recommendation algorithm.* Suppose G is a directed graph on n nodes, with each node having at least one out-edge. Consider the simple random walk X_0, X_1, X_2, \dots that moves by selecting an out-neighbor of the current node uniformly at random. For a node x_i , the friend-of-friend algorithm generates recommendations $R_G(x_i)$ by taking two random walk steps starting from x_i . That is,

$$\Pr(R_G(x_i) = x_j) = \Pr(X_2 = x_j \mid X_0 = x_i).$$

In our setting, edges are oriented so that they point toward one’s “friends”; that is, an edge from u to v indicates that u follows v . For simplicity, we allow this stylized friend-of-friend algorithm to select any node in the two-hop neighborhood of x_i , including neighbors of x_i (to which x_i is already connected) or even x_i itself. However, in real-world networks, users typically have many more second-degree contacts than first-degree contacts, so these pathological recommendations constitute a small fraction of the candidate set.

The friend-of-friend algorithm above describes how personalized recommendations are generated for any given user. Our goal is to understand the system from the perspective of the users who are recommended. That is, how likely is any given user to be recommended to someone in the network? Intuitively, popular users (i.e., those with large follower counts) should appear in the candidate sets for a relatively large number of people, and it is thus reasonable to expect that popular users will be recommended more often. However, the precise relationship between popularity and recommendations depends on the structure of the network, and so to rigorously analyze the friend-of-friend algorithm we need to specify the underlying network.

Here we consider the directed configuration model [23], which yields random graphs with a given sequence of node degrees. To generate a random graph under the configuration model, one starts with a valid sequence of in- and out-degrees (i.e., so that the sum of the in- and out-degrees are equal), and then selects a matching on the degree “stubs” (i.e., half-edges) uniformly at random (Definition 2). We note that this procedure is not equivalent to selecting a random graph with the specified degree sequences uniformly at random, since each graph configuration can result from several distinct matchings and not all configurations come from the same number of such matchings. We further note that the resulting graph is allowed to contain self-loops and multi-edges, though such edges usually comprise a small proportion of the total.

DEFINITION 2. *Configuration model.* For any integer $n > 0$, and degree sequences $\vec{d}^{\text{in}} = (d_1^{\text{in}}, \dots, d_n^{\text{in}})$ and $\vec{d}^{\text{out}} = (d_1^{\text{out}}, \dots, d_n^{\text{out}})$ such that

$$\sum_{\ell=1}^n d_{\ell}^{\text{in}} = \sum_{\ell=1}^n d_{\ell}^{\text{out}}$$

the configuration model produces random (multi) graphs \mathcal{G} on n nodes with the specified degree sequences by selecting a matching on the in- and out-degree “stubs” (half-edges) uniformly at random.

With formal descriptions of the friend-of-friend recommendation algorithm (Definition 1) and the network structure (Definition 2) in hand, we are now ready to state and prove our main result on the relationship between popularity and recommendations. Theorem 1 shows that the number of times a node is recommended is approximately proportional to its in-degree (i.e., its number of followers).

THEOREM 1. *For any integer $n \geq 2$ and degree sequences $\vec{d}^{\text{in}} = (d_1^{\text{in}}, \dots, d_n^{\text{in}})$ and $\vec{d}^{\text{out}} = (d_1^{\text{out}}, \dots, d_n^{\text{out}})$ such that $d_\ell^{\text{out}} \geq 1$ and*

$$\sum_{\ell=1}^n d_\ell^{\text{in}} = \sum_{\ell=1}^n d_\ell^{\text{out}}$$

let \mathcal{G} be a random graph generated via the configuration model with the specified degree sequences. Denote the number of edges in the graph by $|E| = \sum_{\ell=1}^n d_\ell^{\text{in}}$. For a realization G of \mathcal{G} and a fixed node x_i , let $R_G(x_i)$ be a random recommendation for x_i generated via the friend-of-friend algorithm. Then,

$$\Pr(R_G(x_i) = x_j) = \begin{cases} \frac{d_j^{\text{in}}(1 - \epsilon_i)}{|E|} & x_j \neq x_i \\ \frac{d_j^{\text{in}}(1 - \epsilon_i)}{|E|} + \epsilon_i & x_j = x_i \end{cases}$$

where

$$\epsilon_i = \frac{d_i^{\text{in}}/d_i^{\text{out}}}{|E| - 1}$$

and the probability captures randomness both in the graph generation and in the recommendation algorithm.

PROOF. We start with the case where $x_j \neq x_i$. Consider the set of events $E_h(u, v)$ that occur when, in the random graph \mathcal{G} , u is connected to v via the half-edge originating from u with index $h \in \{1, \dots, d_u^{\text{out}}\}$. The process of generating a recommendation R_G for x_i gives rise to several intermediate random variables (where we write R_G instead of $R_G(x_i)$ to simplify notation):

1. $X_1 \in \{x_1, \dots, x_n\}$, the first node randomly selected in the 2-hop path from x_i to R_G
2. $H_1 \in \{1, \dots, d_{X_1}^{\text{out}}\}$, the index of the randomly selected half-edge from x_i to X_1
3. $H_2 \in \{1, \dots, d_{X_1}^{\text{out}}\}$, the index of the randomly selected half-edge from X_1 to R_G .

In terms of these random variables and events, we can write the probability of selecting a particular 2-hop path as,

$$\begin{aligned} & \Pr(R_G = x_j, X_1 = x_k, H_1 = h_1, H_2 = h_2) \\ &= \Pr(H_1 = h_1, H_2 = h_2, E_{h_1}(x_i, x_k), E_{h_2}(x_k, x_j)) \\ &= \frac{1}{d_i^{\text{out}} d_k^{\text{out}}} \Pr(E_{h_1}(x_i, x_k), E_{h_2}(x_k, x_j)) \\ &= \frac{1}{d_i^{\text{out}} d_k^{\text{out}}} \Pr(E_{h_1}(x_i, x_k)) \Pr(E_{h_2}(x_k, x_j) | E_{h_1}(x_i, x_k)) \\ &= \frac{1}{d_i^{\text{out}} d_k^{\text{out}} |E|} \Pr(E_{h_2}(x_k, x_j) | E_{h_1}(x_i, x_k)). \end{aligned}$$

The second equality holds since half-edges are selected by the recommendation algorithm independently and uniformly

at random from the number of out-edges of the respective nodes; the fourth equality comes directly from the definition of the configuration model.

We now analyze the conditional probability in the above expression by considering three cases. First, if $x_k \notin \{x_i, x_j\}$ (i.e., all three nodes are distinct), then

$$\Pr(E_{h_2}(x_k, x_j) | E_{h_1}(x_i, x_k)) = \frac{d_j^{\text{in}}}{|E| - 1}.$$

Consequently,

$$\begin{aligned} & \sum_{x_k \notin \{x_i, x_j\}} \Pr(R_G = x_j, X_1 = x_k) \\ &= \sum_{\substack{x_k \notin \{x_i, x_j\} \\ 1 \leq h_1 \leq d_{x_k}^{\text{out}} \\ 1 \leq h_2 \leq d_{x_k}^{\text{out}}}} \Pr(R_G = x_j, X_1 = x_k, H_1 = h_1, H_2 = h_2) \\ &= \sum_{x_k \notin \{x_i, x_j\}} \frac{d_k^{\text{in}}}{|E|} \cdot \frac{d_j^{\text{in}}}{|E| - 1} \\ &= \left(1 - \frac{d_i^{\text{in}} + d_j^{\text{in}}}{|E|}\right) \frac{d_j^{\text{in}}}{|E| - 1}. \end{aligned}$$

Second, if $x_k = x_j$, then

$$\Pr(E_{h_2}(x_k, x_j) | E_{h_1}(x_i, x_k)) = \frac{d_j^{\text{in}} - 1}{|E| - 1}$$

and so

$$\begin{aligned} & \Pr(R_G = x_j, X_1 = x_j) \\ &= \sum_{\substack{1 \leq h_1 \leq d_{x_j}^{\text{out}} \\ 1 \leq h_2 \leq d_{x_j}^{\text{out}}}} \Pr(R_G = x_j, X_1 = x_j, H_1 = h_1, H_2 = h_2) \\ &= \frac{d_j^{\text{in}}}{|E|} \cdot \frac{d_j^{\text{in}} - 1}{|E| - 1}. \end{aligned}$$

Third, if $x_k = x_i$, then

$$\Pr(E_{h_2}(x_k, x_j) | E_{h_1}(x_i, x_k)) = \begin{cases} \frac{d_j^{\text{in}}}{|E| - 1} & h_1 \neq h_2 \\ 0 & h_1 = h_2 \end{cases}$$

and we have

$$\begin{aligned} & \Pr(R_G = x_j, X_1 = x_i) \\ &= \sum_{\substack{1 \leq h_1 \leq d_{x_i}^{\text{out}} \\ 1 \leq h_2 \leq d_{x_i}^{\text{out}}}} \Pr(R_G = x_j, X_1 = x_i, H_1 = h_1, H_2 = h_2) \\ &= \frac{d_i^{\text{out}} - 1}{d_i^{\text{out}}} \cdot \frac{d_i^{\text{in}}}{|E|} \cdot \frac{d_j^{\text{in}}}{|E| - 1} \end{aligned}$$

where the $h_1 = h_2$ term in the sum drops out since the conditional probability is 0. Combining the expressions, we

get

$$\begin{aligned}
\Pr(R_G = x_j) &= \sum_{x_k \notin \{x_i, x_j\}} \Pr(R_G = x_j, X_1 = x_k) \\
&\quad + \Pr(R_G = x_j, X_1 = x_i) + \Pr(R_G = x_j, X_1 = x_j) \\
&= \left(1 - \frac{d_i^{\text{in}} + d_j^{\text{in}}}{|E|}\right) \frac{d_j^{\text{in}}}{|E| - 1} + \frac{d_j^{\text{in}}}{|E|} \cdot \frac{d_j^{\text{in}} - 1}{|E| - 1} \\
&\quad + \frac{d_i^{\text{out}} - 1}{d_i^{\text{out}}} \cdot \frac{d_i^{\text{in}}}{|E|} \cdot \frac{d_j^{\text{in}}}{|E| - 1} \\
&= \frac{d_j^{\text{in}}}{|E| - 1} \left(1 - \frac{d_i^{\text{in}} + d_j^{\text{in}}}{|E|} + \frac{d_j^{\text{in}} - 1}{|E|} + \frac{d_i^{\text{in}}}{|E|} - \frac{d_i^{\text{in}}}{d_i^{\text{out}}|E|}\right) \\
&= \frac{d_j^{\text{in}}}{|E|} \cdot \frac{|E| - 1 - d_i^{\text{in}}/d_i^{\text{out}}}{|E| - 1} \\
&= \frac{d_j^{\text{in}}}{|E|} \left(1 - \frac{d_i^{\text{in}}/d_i^{\text{out}}}{|E| - 1}\right).
\end{aligned}$$

This calculation establishes the result for $x_j \neq x_i$. Finally, for $x_j = x_i$, we note that

$$\begin{aligned}
\Pr(R_G(x_i) = x_i) &= 1 - \sum_{k:k \neq i} \Pr(R_G(x_i) = x_k) \\
&= \frac{d_i^{\text{in}}}{|E|} \left(1 - \frac{d_i^{\text{in}}/d_i^{\text{out}}}{|E| - 1}\right) + 1 - \sum_{k=1}^n \frac{d_k^{\text{in}}}{|E|} \left(1 - \frac{d_i^{\text{in}}/d_i^{\text{out}}}{|E| - 1}\right) \\
&= \frac{d_i^{\text{in}}}{|E|} \left(1 - \frac{d_i^{\text{in}}/d_i^{\text{out}}}{|E| - 1}\right) + \frac{d_i^{\text{in}}/d_i^{\text{out}}}{|E| - 1}.
\end{aligned}$$

□

Theorem 1 shows that users gain exposure from the recommender in proportion to the size of their follower base. Thus, as expected, popular users receive a larger absolute number of recommendations. This finding, however, is at odds with informal and theoretical arguments that suggest recommendation systems (at least collaborative filters) recommend the most popular products disproportionately often (i.e., superlinearly in their market share) [13]. Moreover, the result does not immediately explain why the most popular users saw their growth rate increase by more than 50% compared to the 20% experienced by most users. Indeed, all else equal, the linear relationship between one’s number of followers and number of recommendations suggests all users should see comparable gains.

This theoretical result is validated by empirical evidence on Twitter. Figure 4 shows average daily number of recommendations per follower, for groups of users binned by their follower counts. For confidentiality, only relative rates are shown, with 1 indicating the overall average. If, hypothetically, the overall average were 0.05, then the figure shows that those uses with approximately 10,000 followers would typically be recommended $10,000 \times 0.05 = 500$ times per day. The plot is based on “Who To Follow” recommendations shown October 6–12, 2015 (data on recommendations shown during the 2010 launch of the system were no longer available for analysis).¹ The figure highlights two points.

¹The number of recommendations per follower was likely substantially higher in 2010 than in 2015, in part because of the growing use of mobile devices, on which recommendations are not as readily visible. Nonetheless, it seems rea-

Daily recommendations per follower




Figure 4: The average daily number of recommendations per follower, for groups of users binned by their follower counts. For confidentiality, only relative rates are shown, with 1 indicating the overall average. If the overall average were (hypothetically) 0.05, then users with 10,000 followers would typically be recommended $10,000 \times 0.05 = 500$ times per day. The most popular users receive fewer impressions per follower than less popular users, which indicates that the popularity effect seen in Figure 3 does not stem from popular users simply being recommended disproportionately often. The plot is based on “Who To Follow” recommendations shown October 6–12, 2015.

First, across the spectrum of users with 1 to more than 10 million followers, the daily number of recommendations per follower differs by at most a factor of two, hovering in a relatively narrow band as predicted by Theorem 1. Second, recommendations per follower in fact *decreases* with popularity. This latter result likely stems from modifications in Twitter’s production friend-of-friend recommender that further down-weights recommendations for popular users in an effort to increase diversity. Thus, both in theory and in practice, the friend-of-friend recommender does not appear to confer an obvious advantage to the most popular users.

5.2 A Mismatch Between the Recommender and the Natural Growth Rate

To understand the “rich get richer” effect observed in Figure 3, we now more carefully examine the growth dynamics of Twitter’s network. How does the follower base of a user grow? One possibility is that users attract attention in proportion to their current number of followers, consistent with preferential attachment models of network growth [6]. In these models, edges are probabilistically added from one user to another in proportion to the target user’s degree. Accordingly, over a sufficiently short period of time, the expected number of new followers a user accrues increases ap-

sonable that the qualitative trends would be similar between the two time periods.




Figure 5: The growth rate (i.e., daily percent increase in follower count) before and after introduction of the recommender, for groups of users binned by number of followers. Growth rate is computed for the two weeks before the recommender was introduced, and the two weeks after the rollout period. In contrast to preferential attachment models of network evolution, growth rate decreases with popularity.

proximately linearly in its number of current followers. Such growth that is linear in the number of followers can equivalently be thought of as growth rate (i.e., percent increase in number of followers) that is constant in the number of followers. For example, users might increase their number of followers by a fixed percentage (e.g., 1%) each day, regardless of the number of followers they have.

Figure 5 shows the empirically measured growth rate for users before and after the introduction of the recommender, binned by number of followers. The dotted line indicates average daily growth rate during the week before the recommender’s introduction on July 30, 2010, and the solid line indicates average daily growth for the week after the recommender’s rollout period. As expected, users across the popularity spectrum see a substantial boost in their growth rate after the recommender is introduced. However, importantly, Figure 5 also shows that the most popular users have the *lowest* growth rates, both before and after the recommender. For example, before the recommender was introduced, the most popular users, with more than one million followers, had an average daily growth rate of 0.1% (i.e., their number of followers grew by about 0.1% each day). In comparison, moderately popular users, with follower counts of 100–1,000, had daily growth rate of approximately 0.3%, three times higher.

There are likely a variety of interacting elements that result in the lower growth rate of popular users. Figure 6 suggests one factor: compared to modestly popular users, extremely popular users typically have fewer *distinct* followers-of-followers. For example, users with 1 million followers have on average 10 million followers-of-followers (a




Figure 6: The median number of distinct followers-of-followers, for groups of users binned by follower count, based on the state of the network on July 20, 2010, 10 days prior to the recommender’s introduction. The most popular users have lower relative reach (i.e., number of followers-of-followers divided by followers) than moderately popular users.

multiple of 10), whereas users with 1,000 followers have on average 1 million followers-of-followers (a multiple of 1,000).² This phenomenon, which has been observed previously [22], indicates that a user’s number of followers is not a good proxy for its reach, and in particular, that popular users might grow slower than preferential attachment suggests.³

To generate Figure 6, we first estimated the number of distinct followers-of-followers for each user based on the July 20, 2015 network, 10 days before the recommender was introduced. Given the size of the network, exact computation is prohibitively expensive, and so we used the HyperLogLog algorithm [11] to create a sketch of each user’s set of followers. These neighborhood sketches were then combined to efficiently estimate each user’s number of distinct followers-of-followers [7]. For each bin of users grouped by follower count, the figure shows the median number of followers-of-followers.

The relatively low natural growth rate of popular users helps to explain why popular users experienced such a big boost from the recommender. As shown above, the friend-of-friend recommender suggests users roughly in proportion

²Given the wide range of follower counts, Figure 6 is shown on a log-log scale. However, we caution that concavity on the log-log scale is not equivalent to concavity on the linear scale. The more relevant quantity here is the slope of the line, which is less than 1 starting at the 100 follower-count bin.

³On the Facebook network, there is, somewhat surprisingly, a near-linear relationship between number of followers and number of followers-of-followers [28], perhaps because relationships are symmetric and users are limited to 5,000 connections.

to their number of followers and so, all else equal, the recommender increases each user’s follower count by a fixed percentage (e.g., 0.05%) each day, regardless of their number of followers. However, as shown in Figure 5, the natural growth rate of very popular users (0.1%) is much smaller than for moderately popular ones (0.3%), and so popular users see disproportionate gains. This mismatch between the recommender and the natural growth rate thus results in cumulative advantage for popular users.

Lemma 1 formalizes this intuition, showing that if the natural growth of a user’s follower base is concave in popularity and the recommender is linear in popularity, then growth rate is increasing in popularity (i.e., popular users would have higher growth rates than less popular ones).

LEMMA 1. *Suppose g is a positive concave function on $[0, \infty)$ and is differentiable on $(0, \infty)$. Then for any constant $c > 0$, the function*

$$z(x) = \frac{g(x) + cx}{g(x)}$$

is increasing.

PROOF. First note that for $x > 0$, $z(x) > 1 = z(0)$. Accordingly, we need only show that z is increasing on the open interval $(0, \infty)$. Now, on $(0, \infty)$, the derivative of z is

$$\begin{aligned} z'(x) &= \frac{cg(x) - cxg'(x)}{[g(x)]^2} \\ &= \frac{cx}{[g(x)]^2} \left[\frac{g(x)}{x} - g'(x) \right]. \end{aligned}$$

Consequently, z is increasing if and only if $[g(x)/x] > g'(x)$. By the mean value theorem, there exists $a \in (0, x)$ such that

$$\begin{aligned} \frac{g(x)}{x} &= \frac{g(x) - g(0)}{x} + \frac{g(0)}{x} \\ &= g'(a) + \frac{g(0)}{x} \\ &> g'(a) \\ &\geq g'(x) \end{aligned}$$

where the last inequality follows from concavity. \square

5.3 Popularity and Response Rates

We conclude our investigation of cumulative advantage in network recommendation systems by discussing one final factor, follow-through rate (FTR): the proportion of recommendations that result in a click on the “follow” button in the recommendation module. As discussed above, the follow-through rate is an imperfect proxy for the effect of any one recommendation (e.g., users may be influenced to follow a recommended user even if they do not click on the link). Nevertheless, follow-through rate is useful for understanding qualitative properties of the system.

Figure 7 shows, perhaps surprisingly, that the most popular users have follow-through rates that are more than three times the average. It is not entirely clear why this is the case, but the effect is likely driven by at least two forces. First, in many domains it has been found that people confer higher value to that which they recognize or are familiar with [14], and so individuals may simply believe (accurately or not) that recommendations of popular people are of relatively higher quality. Second, it could be the case that when




Figure 7: *Follow-through rate (i.e., fraction of recommendations that led to an immediate follow), for groups of users binned by follower count. For confidentiality, only relative rates are shown, with 1 indicating the overall average. The plot is based on recommendations shown October 6–12, 2015.*

popular users are recommended, these recommendations are in fact of generally higher quality, perhaps because less data are available on niche users to construct accurate suggestions [13]. Regardless of the root cause, higher response rates for popular users does appear to be a significant source of the observed rich-get-richer effect.

6. CHANGES IN NETWORK STRUCTURE

We have thus far investigated the effect of the recommender on individuals, describing the disproportionate gains experienced by popular users. In aggregate, these consequences for individual users translate to small but observable changes in the global network structure, including the degree distribution, edge reciprocity, and clustering [30].

Figure 8 shows that after the recommender was introduced, a higher proportion of new edges were directed toward popular users, illustrating one noticeable change in how the network evolved. Specifically, for a given number of followers k on the x -axis, the y -axis shows the proportion of new edges directed to users with fewer than k followers. The two lines indicate the distributions for the two weeks before the recommender was introduced (dotted line), and the two weeks after the rollout period (solid line). Both lines are based on follower counts computed 30 days prior to the recommender’s introduction. For example, before the recommender was introduced, 75% of edges were directed toward users with 10,000 or fewer followers, whereas 71% of edges were after the recommender was launched.

Since popular users are less likely to know their followers, they are also less likely to reciprocate network ties. With an increase in connections to popular users, we would thus expect to observe a drop in the proportion of reciprocated edges. Figure 9 indeed shows this pattern, though the ef-

Popularity of followed users




Figure 8: For each k , the proportion of new edges directed to users with k or fewer followers. The two lines indicate the distributions for the two weeks before the recommender was introduced (dotted line) and the two weeks after the rollout period (solid line). After the recommender, a higher proportion of new edges was directed toward popular users.

fect is relatively modest. In the weeks prior to the recommender’s introduction, 38% of new edges were reciprocated, with the number dropping to 36% after the recommender was introduced. For this plot, we consider a new edge from v to w (i.e., v followed w) to have been reciprocated if w also followed v before August 23, 2015 (the date of a recent network snapshot), regardless of who followed whom first.

Finally, Figure 10(a) shows the effect of the recommender on triadic closure. Before the recommender was introduced, 72% of edges closed at least one undirected triangle (i.e., the user followed a second-degree neighbor, ignoring edge direction), and after the recommender, 75% closed triangles. We note that this 3% change is precisely inline with expectations given the high base rate of triangle closing, the fact that recommendations boost edge creation by approximately 20%, and that nearly all recommendations close a triangle. Looking at the number of triangles closed per edge, Figure 10(b) shows a larger effect, with the median number of triangles closed per edge increasing from 5 before the recommender to 7 afterwards. This more pronounced change reflects the fact that friend-of-friend recommenders favor candidates that are reachable from the user via a large number of two-hop paths.

7. DISCUSSION

By theoretically and empirically analyzing Twitter’s “Who to Follow” system, we have uncovered a subtle mechanism for cumulative advantage—and the accompanying structural effects—in network recommenders. Namely, whereas users are recommended in proportion to their current popularity, they naturally grow at a rate that is sublinear in popularity. This sublinear growth likely stems in part from the fact that one’s number of distinct followers-of-followers (a proxy

Fraction of new edges that are reciprocal




Figure 9: The proportion of new edges that are eventually reciprocated, for eight weeks before and after the recommender was introduced. The decline in reciprocation rates after the recommender was introduced is likely due to an increase in individuals following popular users, whom they do not personally know.

for a user’s reach) is also sublinear in popularity. These mismatched dynamics, together with higher follow-through rates for popular users, contribute to the observed rich-get-richer effect.

Cumulative advantage is often thought to hinder the emergence of the best products and ideas, and is hence typically portrayed as an undesirable property of a system. In accordance with this view, one might believe it preferable to design a recommender that mimics the natural network dynamics. Though reasonable, it is not immediately clear that such an approach would lead to better outcomes. For example, some of the effects we observe are driven by individuals actively and disproportionately following recommendations for popular users, which may indicate a latent preference for popularity [26]. More generally, while it is certainly useful to understand the equilibrium behavior of the system in the absence of a recommender, it is not necessarily the case that such non-recommender dynamics are optimal.

We have throughout attempted to estimate causal effects of recommendations by analyzing the abrupt changes in network structure following Twitter’s introduction of a recommender system. Though this approach mitigates many concerns of traditional observational studies, we caution that care must still be taken to interpret our results. In particular, we have estimated only local effects, both in time and population. For example, it is likely that at least some of the edge creation induced by the recommender would have occurred organically, in the absence of a recommender, at a later date. Moreover, as new and different types of users join the service, it is unclear what effect the recommender has on these dynamic and heterogeneous populations. Given the plethora of major and minor changes to the platform

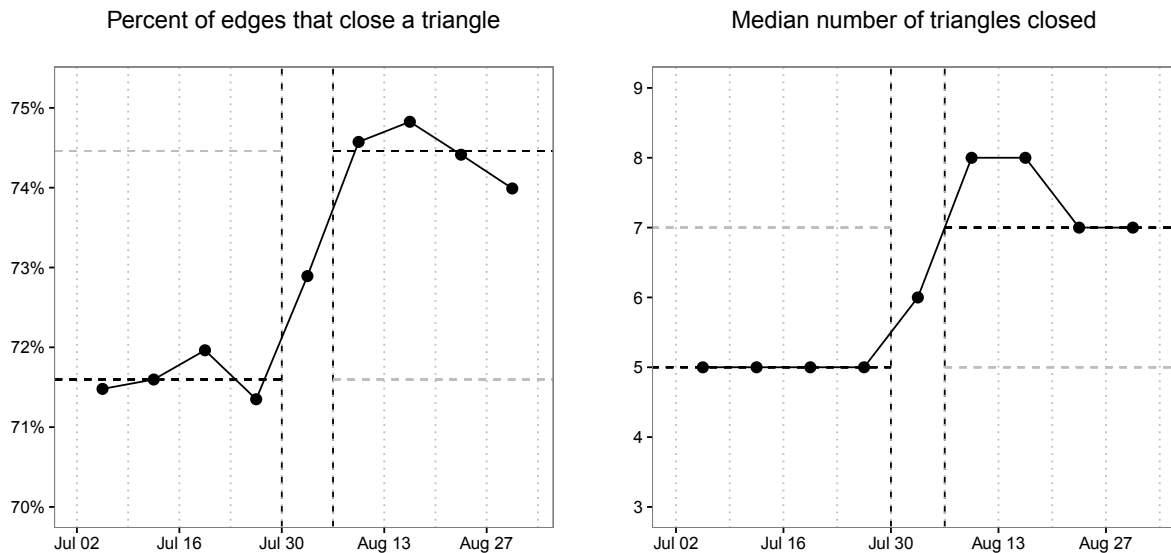


Figure 10: Proportion of new edges that close at least one triangle (left plot), and the median number of triangles closed per edge (right plot), for four weeks before and after the introduction of the recommender.

over the years, it is difficult, and perhaps even impossible, to accurately estimate the long-term effects of the recommendations. Indeed, in a large, complex system such as Twitter, which itself exists in an even larger and more complex ecosystem of online and offline services, it is not obvious how to even define a counterfactual relative to which one can estimate causal effects. For example, if algorithmic recommender systems did not exist, users might turn to non-personalized, third-party lists of suggested users, which could result in even stronger cumulative advantage. Nevertheless, despite these limitations, we believe our theoretical and empirical analysis offers insight into the qualitative, system-wide effects of recommenders on network structure.

Acknowledgements

We thank Ashish Goel and Johan Ugander for helpful comments and suggestions.

References

- [1] M. Al Hasan and M. J. Zaki. A survey of link prediction in social networks. In *Social network data analytics*, pages 243–275. Springer, 2011.
- [2] C. Anderson. *The long tail: how endless choice is creating unlimited demand*. Random House, 2007.
- [3] L. Backstrom and J. Leskovec. Supervised random walks: predicting and recommending links in social networks. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 635–644. ACM, 2011.
- [4] B. Bahmani, A. Chowdhury, and A. Goel. Fast incremental and personalized pagerank. *Proceedings of the VLDB Endowment*, 4(3):173–184, 2010.
- [5] E. Bakshy, S. Messing, and L. Adamic. Exposure to ideologically diverse news and opinion on facebook. *Science*, 2015.
- [6] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- [7] P. Boldi, M. Rosa, and S. Vigna. Hyperanf: Approximating the neighbourhood function of very large graphs on a budget. In *Proceedings of the 20th international conference on World wide web*, pages 625–634. ACM, 2011.
- [8] J. Chen, W. Geyer, C. Dugan, M. Muller, and I. Guy. Make new friends, but keep the old: recommending people on social networking sites. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 201–210. ACM, 2009.
- [9] E. M. Daly, W. Geyer, and D. R. Millen. The network effects of recommending social connections. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 301–304. ACM, 2010.
- [10] T. Dunning. *Natural experiments in the social sciences: a design-based approach*. Cambridge University Press, 2012.
- [11] P. Flajolet, É. Fusy, O. Gandouet, and F. Meunier. Hyperloglog: the analysis of a near-optimal cardinality estimation algorithm. *DMTCS Proceedings*, 2008.
- [12] S. R. Flaxman, S. Goel, and J. M. Rao. Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, 2016.
- [13] D. Fleder and K. Hosanagar. Blockbuster culture’s next rise or fall: The impact of recommender systems on sales diversity. *Management science*, 55(5):697–712, 2009.

- [14] D. G. Goldstein and G. Gigerenzer. Models of ecological rationality: the recognition heuristic. *Psychological review*, 109(1):75, 2002.
- [15] P. Gupta, A. Goel, J. Lin, A. Sharma, D. Wang, and R. Zadeh. Wtf: The who to follow service at twitter. In *Proceedings of the 22nd international conference on World Wide Web*, pages 505–514. International World Wide Web Conferences Steering Committee, 2013.
- [16] T. H. Haveliwala. Topic-sensitive pagerank. In *Proceedings of the 11th international conference on World Wide Web*, pages 517–526. ACM, 2002.
- [17] K. Hosanagar, D. Fleder, D. Lee, and A. Buja. Will the global village fracture into tribes? recommender systems and their effects on consumer fragmentation. *Management Science*, 60(4):805–823, 2013.
- [18] D. Jannach, M. Zanker, A. Felfernig, and G. Friedrich. *Recommender systems: an introduction*. Cambridge University Press, 2010.
- [19] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46(5):604–632, 1999.
- [20] D. Liben-Nowell and J. Kleinberg. The link-prediction problem for social networks. *Journal of the American society for information science and technology*, 58(7):1019–1031, 2007.
- [21] R. K. Merton. The matthew effect in science. *Science*, 159(3810):56–63, 1968.
- [22] S. A. Myers, A. Sharma, P. Gupta, and J. Lin. Information network or social network?: The structure of the twitter follow graph. In *Proceedings of the companion publication of the 23rd international conference on World wide web companion*, pages 493–498. International World Wide Web Conferences Steering Committee, 2014.
- [23] M. E. Newman, S. H. Strogatz, and D. J. Watts. Random graphs with arbitrary degree distributions and their applications. *Physical review E*, 64(2):026118, 2001.
- [24] E. Pariser. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin, 2011.
- [25] A. Sharma, J. M. Hofman, and D. J. Watts. Estimating the causal impact of recommendation systems from observational data. *Economics & Computation*, 2015.
- [26] H. Steck. Item popularity and recommendation accuracy. In *Proceedings of the fifth ACM conference on Recommender systems*, pages 125–132. ACM, 2011.
- [27] C. R. Sunstein. *Republic. com 2.0*. Princeton University Press, 2009.
- [28] J. Ugander, B. Karrer, L. Backstrom, and C. Marlow. The anatomy of the facebook social graph. *arXiv preprint arXiv:1111.4503*, 2011.
- [29] S. Vargas and P. Castells. Rank and relevance in novelty and diversity metrics for recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*, pages 109–116. ACM, 2011.
- [30] M. Zignani, S. Gaito, G. P. Rossi, X. Zhao, H. Zheng, and B. Y. Zhao. Link and triadic closure delay: Temporal metrics for social network dynamics. In *Eighth International AAAI Conference on Weblogs and Social Media*, 2014.