

Structural Predictors of Tie Formation in Twitter: Transitivity and Mutuality

Scott A. Golder
Department of Sociology
Cornell University
Ithaca, NY, 14850
sag262@cornell.edu

Sarita Yardi
School of Interactive Computing
Georgia Institute of Technology
Atlanta, GA, 30308
syardi3@gatech.edu

Abstract — New ties are often formed between people who already have friends in common. Though the social sciences have addressed the effects of existing structural patterns on the formation of new ties, less attention has been given to ties in directed networks. Drawing from the microblogging service Twitter, we conducted a web-based experiment in which subjects were asked to rate their interest in forming ties to other people, blind to existing network connections between them. We show that two structural characteristics, transitivity and mutuality, are significant predictors of the desire to form new ties. Our findings shed light on tie formation, especially in online networks.

Keywords: social networks, Twitter, triadic closure.

I. INTRODUCTION

Social life, whether online or offline, presents people with a problem – how to meet other people, make new friends, and build professional contacts.

People form new social ties for many reasons and under many conditions. Exposure to people of similar age, socioeconomic status, or education level in neighborhoods, schools, and workplaces means that people are likely to encounter others who share interests or opinions. We evaluate others positively when they have characteristics in common with ourselves, thus driving the strengthening of these connections. This is the principle of *homophily*, that “similarity breeds connection” [18] and that people’s social networks are therefore relatively homogenous with respect to personal characteristics, beliefs, and values. Homophily can be seen as a psychological as well as structural phenomenon; we have a psychological predisposition to forming ties to those like us, but at the same time we are also structurally more likely to encounter those people in our everyday lives, both online and offline.

Homophily arises through the process of *triadic closure*, or the formation of new ties to people who are friends of existing friends. That is, when A and B are friends, and B and C are friends, triadic closure occurs when A and C become friends.

In many circumstances, it is enough to say “A and B are friends.” This is because the relationship can be considered *symmetrical*; if A is friends with B, B must be friends with A. However, many times it is useful to think of relationships as *asymmetrical*. This may be because the actors have different roles (employee-employer) or because the power dynamics in the relationship call for it (e.g. B would like to be good friends with A, but A is content having B as an acquaintance).

Symmetrical relationships can be modeled with undirected graphs, and asymmetrical relationships with directed graphs.

We are interested in *directed triadic closure*, or triadic closure when individuals’ relationships are asymmetric. Triadic closure has been examined extensively in undirected networks, but less in directed networks. In Section III we look to develop an understanding of how homophily and triadic closure work in directed networks.

Following this, we describe our web-based experiment in which we measured how the structural properties of subjects’ networks could predict their interest in other people. We found that, despite having no knowledge about the ties between themselves and those they were asked to rate, subjects’ ratings of others were significantly associated with the structural characteristics of their egocentric networks.

This study was conducted using the microblogging service Twitter. Begun in 2006, Twitter supports posting very short messages (“tweets”). A person’s tweets are visible to his or her “followers,” or those others who have chosen to pay attention to that person.

Subjects in the experiment were shown sequences of other Twitter users who they did not already follow, and were asked to rate their interest in following each of these others. We used a hierarchical regression model to estimate the effects of structural characteristics on subjects’ appraisals, and also elicited descriptive feedback from the subjects.

II. TWITTER

Twitter is a *conversational microblog*, a service for posting short, quasi-public messages up to 140 characters in length. People create lists of others and are shown a reverse-chronological list of all of the posts of those people. The substantive nature of the social tie on Twitter is attention-based. In addition to paying attention to one another by “following,” Twitter users can *address* tweets to other users and can *mention* others obliquely in their tweets [13]. Another common practice is “retweeting,” or rebroadcasting someone else’s message (with attribution) so as to direct attention toward that person’s tweets [1].

Twitter differs from other online social networking services in that ties are asymmetric. Consider friendship ties in LinkedIn, Facebook, or MySpace; in these services, when two people share a friendship tie, the tie is symmetrical; A being friends with B implies B is friends with A. This is not so in Twitter; A

can "follow" B, but B need not follow A.¹ People who are popular, such as basketball player Shaquille O'Neal (Twitter name: @THE_REAL_SHAQ) or actor Ashton Kutcher (@aplusk), can be "followed" by millions of others, but can pay attention to as many or few as they like.

III. SOCIAL TIES

Modeling ties as directed networks introduces complexity but offers significant analytical benefits. When a tie is symmetric, there are only two states: the tie is present or absent. When ties are asymmetric, there are four states. Three are shown in figure 1: A is connected to B, B is connected to A, or A and B are mutually connected. The fourth state, not shown, is the absence of a tie between A and B.

If A has a directed tie to B and B does not have such a tie to A, we might say B has a power or status advantage over A, since B is more important to A than A is to B [6]. In the context of Twitter, A follows and pays attention to B, but B does not follow – and therefore pays no attention to – A.

In many cases, the directed link is an indicator of the direction in which attention flows. Prominent cases include the hyperlink graph of the web, in which directed edges identify authoritative pages [17], and citation networks in which scholarly articles reference those that came before [15].

Another way of thinking about the directed tie is that it is a conduit of information. In a sense, all social networks are information or communication networks [10, 2, 20]. This is why "too much homophily" can result in networks in which there is a lack of access to diverse viewpoints, resulting in an "echo chamber" (e.g. [24]) or provincialism.

Though we use arrows to indicate the flow of attention, such as $A \rightarrow B$ when A follows B, such an arrow can be read backwards, because information flows from B, to his follower A. We can think of any social tie as representing an exchange relationship: when A follows B, A is in effect exchanging his attention, a valuable resource, for B's information, which is presumably of value to A. For a social tie to be maintained, it must be *rewarding* to the parties involved [12, 6].

We now move from dyads to triads, or from structures of two actors to structures of three actors. Directed ties greatly increase the number of configurations available to groups of three actors. As detailed in figure 2, instead of one kind of path connecting A to a neighbor X to the target B ($A-X-B$), there are nine such paths.

A. Structural Balance & Triadic Closure

Starting with the observation that friendship choices are interdependent, early work in sociometry and network analysis examined the prevalence of various configurations of triads.

Heider's theory of cognitive balance (see [3,12,11] for further discussion) describes how stability and consistency arise in network configurations. Consider three individuals – let us call them A, X and B – in an undirected network. If A is friends



Figure 1. Directed Ties.

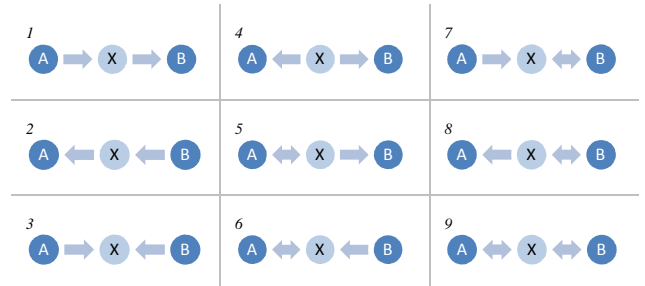


Figure 2. Directed two-step paths.

with X and X is friends with B, then what do we make of the relation between A and B? If A is friends with² B then the triad is "balanced." If A is not friends with B, it is "unbalanced." As the strength of the $A-X$ and $X-B$ ties increase, the likelihood of the $A-B$ tie existing grows; a configuration lacking the $A-B$ is empirically least likely and has been called a "forbidden" triad [10]. If the $A-B$ tie does form, it is an example of triadic closure, or the formation of an $A-X-B$ triad in which all three "legs" are present.

The concept of balance is connected to the idea of psychological strain being placed on unbalanced relationships. In everyday social life, if X's two friends A and B were openly hostile to one another, then A or B might suffer emotionally when the other shows up at X's party, or X might suffer from having to choose which friend to exclude. In order to bring about balance in the network, X might have to sever his tie to either A or B, or A and B would have to change their opinions about one another.

In the context of Twitter, the strain of unbalanced triads might be due to the presence of diverse fractions of conversations rather than unified, coherent ones. It may also relate to having to choose carefully what to tweet when distinct groups of followers make mutually inconsistent normative demands (e.g. both conservative and liberal followers).

Holland and Leinhardt [11] considered closure in directed rather than undirected networks using direction of tie to model choice, the $A \rightarrow B$ tie indicating A's choice to be friends with B. They examined the presence of transitivity, or the degree to which, when $A \rightarrow X$ and $X \rightarrow B$, the $A \rightarrow B$ tie was present. They considered transitivity to be an essential element of structure; since social life does not map perfectly onto the expectations of mathematical models, they relaxed the transitive/intransitive distinction by proposing an index of transitivity. Empirical social networks do not exhibit perfect transitivity (when all triads are transitive), but they also exhibit far less intransitivity than expected in random networks.

More recently, Romero and Kleinberg [23] made the observation that determining whether or not transitivity plays a

¹ To clarify: "friends" is the Twitter term to describe the people who the individual in question is following, or those who the individual's outlinks point towards. "Followers" are those who are following the individual in question, whose outlinks point into him.

² We can replace "is friends with" with "has positive feelings towards," "evaluates favorably," or some other positively-valenced relation.

role in directed triadic closure in Twitter requires knowledge of the temporal ordering of the ties. For example, in the example in the previous paragraph, if the $X \rightarrow B$ tie forms *after* the $A \rightarrow B$ tie does, then transitivity could not play a role, since the path from A to X to B did not exist at the time the $A \rightarrow B$ tie formed. Romero and Kleinberg showed that the presence of directed triadic closure is higher than chance for egocentric networks, further supporting the idea that transitivity is an important precursor to directed triadic closure.

B. Attention-Information Networks

One of our goals is to describe a way to think about the real-world meaning behind the numerous structural configurations given in Figure 2. The conceptual categories we describe below each draw on the observation that social networks are inherently directed, regardless of how they are modeled, and that all social networks are *attention-information* networks in which attention is exchanged for information. Therefore, the relationships between people can be described as their positions with respect to flows of information and attention.

1) Shared Interests

Social ties can be established around shared interests, organizations, and activities [7]. Sharing many interests with another person is one kind of similarity, and forms the basis for collaborative filtering algorithms, such as book and movie recommendations. In network terms, a shared interest is best represented by the $A \rightarrow X \leftarrow B$ triad, where A and B each pay attention to X. This is called *outlink equivalence* – A and B are equivalent in who their outlinks point towards, to the extent that the sets of out-neighbors overlaps [25]. As A and B become more outlink-equivalent, the likelihood of them being interested in one another is expected to increase.

2) Shared Audiences

An analogous argument can be made for people who share inlinks in common. The kind of tie $A \leftarrow X \rightarrow B$ describes a scenario in which some person X is interested in both A and B. A and B are therefore *inlink-equivalent* to the extent that the set of people who link to A overlaps with that of B [25].

We may say that the people called X form the *audience* paying attention to A and/or B. If a large number of people find both A and B worthy of attention, A and B may have some shared traits driving this, such as tweeting about the same or similar topics. It is entirely possible that for a given X, A and B are interesting for different reasons; perhaps X and A live in the same city, but X and B share an interest in comic books or moose hunting. Nevertheless, as the number of X's grows, the likelihood of A and B being interested in one another is expected to grow as well.

3) Transitivity & Filtration

Twitter users perform a valuable service as curators; they read the tweets of those who they follow, and pass the good ones along to their own followers by "retweeting" them [1]. If A follows X, then A sees the tweets by B which X has thought high enough quality (along any of a number of possible dimensions) to rebroadcast it.

Considering the path $A \rightarrow X \rightarrow B$, A may then *not* want to follow B, because he already receives the best of B's tweets

without doing the filtering work that X does. However, this is predicated on A's belief that X is doing a good job of filtering and rebroadcasting B's tweets.

In contrast, suppose many such people X are following B. A may then be susceptible to the influence of his neighbors and likewise follow B. The idea is that X's provide *social proof* [4] that it is a good idea to follow B. That is, people judge actions as good, prestigious, or valuable to the extent that they see others doing them. Indeed, A finding it valuable to follow B would indicate attention is transitive and being a curator is of comparatively little value.

4) Mutuality & Reciprocity

Reciprocity means behaving toward someone in the manner in which they behave toward you. Paying back a favor or returning a smile with a smile are examples of reciprocal behavior. Reciprocity is a source of social cohesion [9]; when two individuals attend to one another, the bond is reinforced in each direction and both people will find the tie rewarding [6]. Reciprocal exchange relationships, in which individuals give something of value (e.g. attention) to one another in turns – also leads to stronger affective ties [19], and giving support is strongly associated with receiving it [22].

Though the $B \rightarrow A$ tie could be evidence of some basis for an $A \rightarrow B$ tie, there are two reasons why this may not be the case. First, the cost of feigning attention is low. Since follower count is one visible metric by which people measure status, spam bots are known to add spurious "follower" relations in the hope that others will follow back. This is an attempt to gain legitimacy. Second, status differentials persist; the essence of popularity is that attention is a scarce resource, and more people pay attention to a popular person than he could possibly reciprocate to.

IV. METHOD

We ran a web-based experiment from July-August 2009 with randomly-selected active Twitter users. This section describes the experimental design, including the suggestion generation process, subject recruitment process, and front-end interface.

In brief, our study invited each subject to rate the profiles of 14 randomly-selected people in their 2-degree networks on a 1-5 Likert scale, in terms of how much the subject would be interested in following that person. For the purposes of clarity, we will refer to these 14 other people as "alters," though they are actually only potential alters.

The alters were selected at random from the 200 people with the most 2-step paths to the subject, provided they were not already followed by the subject. Selecting randomly from these 200 did not necessarily produce many highly-connected alters; the Twitter network is sparse, and by the 200th-ranked person, the number of connections was often quite small. Nevertheless, these people are by design more connected than people chosen at random from the 2-degree network. We conducted a pilot run of the experiment in which alters were selected at random. The results were inconclusive due to sparsity of network data and are not discussed further.

We deliberately did not allow subjects to rate anyone with more than 5,000 friends or followers. Our intent was to filter out celebrities, news outlets, and others who might be found

Citation: Scott A. Golder and Sarita Yardi. "Structural Predictors of Tie Formation in Twitter: Transitivity and Mutuality." *Proceedings of the Second IEEE International Conference on Social Computing*. August 20-22, 2010. Minneapolis, MN.

through a broadcast medium. In any case, following @THE_REAL_SHAQ is unlikely to be evidence of a substantive social tie. Note that in this respect our population of interest is different from that of Romero and Kleinberg [23], who focused specifically on high-indegree individuals with tens of thousands of followers, or "micro-celebrities."³

We also excluded as alters those whose accounts were set to "private" because our experimental interface requires that individuals' tweets be visible to the subject so that the subject can decide whether the target person is interesting. Our estimates suggest that, at the time of the experiment, 8-14% of Twitter accounts were private.⁴

A. Subject Recruitment

Our sample frame consisted of all Twitter users appearing at least once in the public timeline during the last week of May 2009. The public timeline is an RSS feed of tweets from randomly-selected Twitter users, updated every minute.⁵ Though users may appear multiple times in the public timeline, we sampled by name, not by tweet, so highly active users had no additional selection advantage.

By sampling from the public timeline, we are more likely to sample on active users than if we had sampled from all accounts; like most web sites, many Twitter accounts are used rarely or even abandoned, and such accounts are less likely to exhibit social behavior of any kind. It is not clear whether active users are more likely to want to follow suggested others (desire for new contacts growing with use) or less likely (they are, in some sense, saturated already).

Using an account specifically created for this project, we sent recruitment messages to 2,085 users:

@user Hi! Would you be willing to help us with a short Twitter experiment? It's easy and takes less than 5 minutes. Thanks!

Over 250 people responded to our initial request. Some agreed to participate immediately, others asked for clarification first. We generated back-end data (next subsection) and passwords for 157 subjects and sent each a tweet containing a link to the experiment and their password.

Of these 69 fully completed the experiment, 32 participated in the pilot and, as detailed in the Results section, 37 participated in the actual run of the experiment. Though overall the attrition rate was high, it is not out of the ordinary for unsolicited web experiments [5], but it is worth emphasizing that our subjects are likely fairly active Twitterers.

B. Back-end Data Collection

Before running subjects in the experiment, we required their profile data and a comprehensive map of their 2-degree networks: all their friends and followers, all those people's friends and followers, and all the links among them.

Despite excluding high degree individuals, crawling 2-degree networks was computationally intensive – the average number

of friends and followers in our subject pool was 441 and 500, respectively, and some 2-degree networks contained over 1 million edges – requiring us to run the back-end software as a separate intermediate step, rather than doing it in real time when subjects visited our site. Our back-end software read subjects' 2-degree networks into memory and computed the number of each kind of 2-degree path for each person in that network, as seen in Fig. 2. It then generated profile pages for 14 alters who were 2 steps from the subject.

C. Experimental Interface

Subjects were given an introduction to the experiment on the login page and a brief set of directions:

In the next few pages, we will show a series of different Twitter users. We want to know if you think you would want to follow the user and why. We will first ask you to take a look at the user's most recent tweets and tell us what characteristics you find interesting or not. Then we would like you to rate whether or not you want to follow the user.

After logging in, they were shown each of the 14 suggested alters in succession and asked to rate them and, optionally, to give a free-response explanation for their rating. Each page showed the alter's username, profile picture, location and bio (if present), and friend, follower, and tweet counts. This was followed by a list of their 20 most recent tweets (see Fig. 3).

At the top of the page, subjects were given a Likert rating scale to respond to the question: "Would you want to follow this person?" The wording choice "would you" was based on wording used in common large survey polls (e.g. [14]). The response options provided were labeled, "Definitely!", "Probably", "I'm not sure", "Probably Not" and "Definitely not!", and a button to report missing or broken items. Subjects were given a text box for the optional free response. Subjects were shown only the profiles of the suggested people; they were *not* shown any information about the network ties that link them to these people.

D. Validity Check

Subjects' stated preferences in laboratory experiments do not

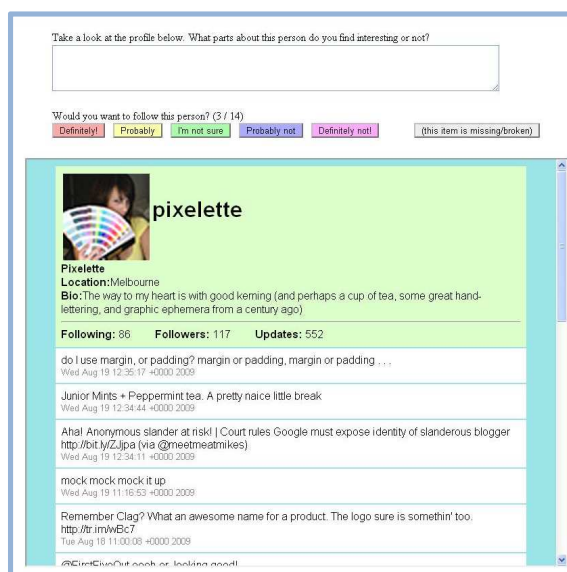


Figure 3. The experimental interface.

³ Incidentally, @TheRealShaq is a mega-celebrity, one of the small number of Twitter users who has multiple millions of followers.

⁴ This estimate is derived from random sampling of valid User IDs.

⁵ Every person (subjects and alters) in our study had a public account. Data collection involved white-listed accounts and the public data API.

always match their subsequent real-world behavior. To test the validity of responses, at the conclusion of the experiment, we offered subjects the opportunity to actually follow with their own Twitter accounts the 14 alters that they had just rated. Given a list of alters' photos and names, clicking on the alter's "follow" button made an API call to add the alter to the subject's followees. This last step allows a comparison of subjects' self-reported and actual behavior.

E. Model Specification

A problem in the analysis of networks is that ties are not independent; individuals' friendship choices are affected by the choices their friends make, and therefore estimates in the resulting models are subject to bias. We control for this by using a hierarchical regression model in which the subjects are clusters, and the 14 people they rated are cases nested within those clusters. We estimate a random intercept for each subject, reflecting the variation in their basic interest in adding new friends. Since the subjects (and the people they are shown) are chosen at random and the graph is sparse, we did not expect alters to appear in the data for more than one subject, and indeed none did.

The response data are on a five-option scale, ranging from very positive to very negative, but are analyzed as continuous, since the underlying construct, desire to follow that person, is continuous. This requires an assumption that the responses are evenly spaced [16]. If we were not willing to make this assumption, the appropriate model would be a hierarchical ordered logistic regression model. We ran all models shown in Table 2 both as linear and ordered-logistic models, and the results are qualitatively similar.

V. RESULTS

37 subjects completed the web-based experiment. Subjects were considered to have completed the experiment if a rating was given for at minimum 9 of the 14 alters. Due to a software bug, responses for the 14th alter were never recorded, so subjects have a maximum of 13 responses. Since the alters are chosen randomly, we do not expect any systematic errors are introduced by leaving the final one out.

For both subjects and alters, the number of friends, followers, and statuses ("tweets") is log-normally distributed. That is, each distribution has a long right-hand tail such that when the log is taken, the resulting transformed distribution is approximately normal. Therefore, in all the data analysis and statistical models, the logs of these counts will be used.

Table 1 presents some descriptive statistics about the subjects and alters. For some characteristics, 2-tailed t-tests were performed to examine whether any appreciable differences exist between the subjects and their alters. Alters were significantly more well-connected than subjects, with more friends as well as followers. This is to be expected; Feld [8] points out that highly-connected people are by definition more likely to appear in lists of alters since they have more chances to be selected; to put it another way, friends are likely to be friendlier than average [21]. Likewise, subjects were more prolific than alters, which was also to be expected since, as described earlier, our recruitment method biased toward active

Twitterers. However, longevity did not vary significantly between subjects and alters.

We examine the results incrementally, by fitting four regression models: subject characteristics alone, alter characteristics alone, subject and alter characteristics together, and all characteristics plus network structure (Table 2).⁶ In general, coefficients are stable across models. The alters-only model (Model 2) performs least-well; besides that, goodness of fit (pseudo- R^2) increases as additional predictors are included, and the model that considers network structure (Model 4) performs considerably better than the equivalent model without those predictors (Model 3).

The results of the validity check suggest that subjects' responses were consistent with their ratings. 15 of the 37 subjects chose to follow one or more alters; for the other 22, we cannot tell if they actively chose not to follow anyone, or if they stopped participating at this stage because it was optional. The 15 subjects who chose to follow one or more alters followed a total of 32 alters. Of the 32, 18 had been given a rating of 5, and 12 a rating of 4. Since the alters who were followed were generally given high scores, we are confident that higher ratings are a good measure of more desirability.

A. Subject-Specific Predictors

We begin with subject-level predictors, which measure aspects of the subjects' networks, irrespective of the alters in question.

Subjects' interest in adding new friends does not appear to satiate. There is not a significant relationship between the number of friends subjects already have and their ratings of alters. A commonsense belief is that adding new friends imposes a cost of time and attention; this is because more friends generate more tweets which requires more time to read or filter them all, and because new friends' tweets might draw one's attention away from existing friends. Subjects who already feel they have enough friends might therefore be less inclined to add new friends and thus impose more stringent standards. However, we did not find such a result; subjects with many and few friends were not significantly different in their ratings of alters.

Another commonsense belief is that Twitter users construct their egocentric network early in their time on Twitter and, having done so, do not seek to add more friends. A finding consistent with this belief would require a negative, significant effect to account age. However, we did not find such a result; new and old subjects were not significantly different in their ratings of alters.

However, more active subjects do appear to be less interested in adding friends. The number of tweets a subject has written is significantly negatively associated with the desire to add new friends; a one standard-deviation increase in the log count of tweets written is associated with a decreased alter rating of about 0.3. Note that this is net of the effect of account age, and that account age and number of tweets have a low correlation,

⁶ We used standardized coefficients for each continuous variable, and the outcome variable is subjects' rating of alter from 1-5. Coefficients should be interpreted in the manner of the following example: according to Model 4, a one standard deviation increase in the log-count of an alter's followers is associated with an increase in the estimated rating of that alter by 0.258.

TABLE 1. SUBJECTS' AND ALTERS' CHARACTERISTICS.

	SUBJECTS		ALTERS		p^3
	M	SD	M	SD	
Friends ¹	5.228	0.968	5.958	1.512	0.004
Followers ¹	5.424	0.982	6.170	1.440	0.002
Tweets ¹	7.146	1.241	6.421	2.000	0.031
Default Photo ²			0.005		
Includes a Bio ²			0.849		
Includes a Location ²			0.896		
Account Age (days)	326.149	231.231	320.367	247.415	0.891
N	37		443		

Notes:

- (1) natural log
- (2) Coded 1/0. The mean is interpretable as a fraction of subjects.
- (3) For each 2-tailed t-test, $df=478$. $N_{\text{subjects}}=37$, $N_{\text{alters}}=443$.

so it is activity and not longevity that generates this effect. In short, more active Twitterers are less interested in adding new friends, as demonstrated by lower ratings of alters. Finally, network density⁷ and number of followers both had positive, significant effects on alter ratings.

Since we would expect individuals' egocentric networks to be more sparse the larger they grow, if subjects who are popular (many followers) and subjects who are part of closer-knit communities (more dense) are both associated with higher ratings of others, then two separate processes might be underlying these effects, and further investigation is needed.

B. Alter-Specific Predictors

1) Friend and Follower Counts

Alters were rated more desirable to the extent they had more followers and fewer friends. That is, people who are popular already appear more likely to draw new followers. This was consistent across models.

It is worthwhile to take a moment and reflect on this. Alters' number of followers is a signal of status, of being desirable to many others. The key to being popular is having the attention of many other people, more than one can pay attention to oneself; this may be why having many followers is positive, but having many friends (following many others) is negative.

It cannot be the case that high-indegree alters are simply more active; number of tweets is arguably a better proxy for activity and was not significant. We observed that more popular alters are rated more highly simply because they are more popular.

Popularity has a cumulative effect; social psychologists call this *social proof* [4] and network analysts call it *preferential attachment* [23]. However, these models imply social learning – that people observe what others do and choose from those actions; more algorithmically, they select a node to link to randomly, weighted by the in-degree. Instead, what we observe is a status effect; it is not the case that high-indegree people were more likely to be seen, but that *the very fact* of their high in-degree accounts for some of their desirability.

TABLE 2. ESTIMATED EGO AND ALTER EFFECTS ON EGO'S RATING OF ALTER.

PREDICTORS ¹	MODEL 1	MODEL 2	MODEL 3	MODEL 4
(Intercept)	3.078 *** (0.102)	3.385 *** (0.278)	3.399 *** (0.259)	3.364 *** (0.266)
Subject's Friends ²	0.131 (0.213)		0.148 (0.215)	-0.039 (0.228)
S.'s Followers ²	0.799 *** (0.247)		0.730 ** (0.250)	0.803 ** (0.273)
S.'s Tweets ²	-0.308 * (0.147)		-0.293 * (0.148)	-0.245 † (0.147)
S.'s Network Density	0.357 ** (0.125)		0.360 ** (0.125)	0.354 * (0.149)
S.'s Account Age	-0.024 (0.131)		-0.026 (0.133)	0.024 (0.139)
Alter's Friends ²		-0.130 (0.086)	-0.166 * (0.084)	-0.222 * (0.098)
A.'s Followers ²		0.273 ** (0.094)	0.235 ** (0.093)	0.258 * (0.103)
A.'s Tweets ²		-0.045 (0.081)	-0.013 (0.080)	0.002 (0.079)
A.'s Account Age		0.073 (0.077)	0.019 (0.076)	0.03 (0.077)
A.'s order in expt.		-0.033 * (0.015)	-0.036 * (0.015)	-0.016 (0.016)
A. has default photo		-0.067 (0.881)	0.002 (0.875)	-0.011 (0.847)
A. includes bio		0.358 † (0.194)	0.336 † (0.193)	0.24 (0.192)
A. includes location		-0.423 † (0.226)	-0.400 † (0.225)	-0.439 * (0.22)
<i>Subject-Alter Paths</i>				
Reciprocity (A → S) ³				0.348 (0.221)
(1) S → X → A ⁴				-0.053 (0.067)
(2) S ← X ← A				0.142 * (0.067)
(3) S → X ← A				0.068 (0.087)
(4) S ← X → A				-0.132 (0.101)
(5) S ↔ X → A				0.222 * (0.107)
(6) S ↔ X ← A				-0.155 † (0.091)
(7) S → X ↔ A				0.215 ** (0.072)
(8) S ← X ↔ A				-0.175 * (0.088)
(9) S ↔ X ↔ A				0.249 * (0.106)
<i>Random components⁵</i>				
Random effect	0.515 (0.088)	0.793 (0.113)	0.516 (0.089)	0.466 (0.086)
Residual	1.195 (0.042)	1.170 (0.041)	1.171 (0.041)	1.130 (0.040)
Log Likelihood	-729.117	-732.748	-720.386	-685.754
LR vs Model 1	-	-	17.41 *	86.56 ***
LR vs Model 2	-	-	25.54 ***	94.69 ***
LR vs Model 3	-	-	-	69.15 ***
Pseudo-R ²	0.205	0.072	0.233	0.301
Clusters	37	37	37	36 ⁶
N	443	443	443	432

† $p < 0.10$ * $p < 0.05$ ** $p < 0.01$ *** $p < 0.001$

Notes:

- (1) Standard errors shown in parentheses. Standardized coefficients for all predictors except A's order, photo, bio, location.
- (2) Natural log.
- (3) Binary-valued predictor (1/0)
- (4) All tie types are measured as the number of paths of that type from Subject (S) through some other person X to Alter (A), divided by the number of people in Subject's network.
- (5) The subject-level random effects and alter-level residuals are normally distributed.
- (6) One case was dropped due to missing data.

⁷ "Density" here is the clustering coefficient of the 1.5-degree network; it measures the proportion of possible ties that exist: $e/(n*(n-1))$

The mass media-driven claim that having more followers than friends is socially desirable in Twitter⁸ turns out to be true.

2) Self-Presentation

Self-presentation-related predictors, such as having the default photo, including a bio or a location, had mixed effects for alters. Using the default photo (rather than uploading a personal profile photo) had a non-significant effect that was nearly zero; however, as noted in Table 1, only 0.5% of alters had the default photograph, so perhaps there is not enough variation in the data to pick up any effect. Though about 85% of alters included a bio and 90% of alters included a location, having a bio did not have a consistently significant effect, and location had a consistently *negative* effect, associated with a drop in rating by about 0.4.

One reason why location might have a negative effect is that it highlights a salient difference between subject and alter. We turn to free-response data to examine why. The following subjects were motivated by shared geographic region, giving 4's and 5's in these cases:

Definitely, from houston where I'm from. tweets worth reading.

In same city as me, Dallas. That's why I will follow him.

He's funny, lives in my area... Definitely adding him.

The fact is, more people live far away from any given person than live near that person. We suggest that alters' location makes geographic distances salient to the subject and when the locations match the effect is positive, but in the larger number of mismatched cases, the differences between subject and alter are highlighted, and the expected score decreases:

I would follow this if I were interested in New York. I could see following a twitter dedicated to my neighborhood in Chicago.

C. Relational Predictors

We now turn to the predictors of greatest importance, the existing network connections between subject and alter.

1) Reciprocity

Let us start by considering reciprocity, or the directed tie from alter to subject. That is, is the alter following the subject already? We did not observe a statistically significant effect.

The meaning of this is both counterintuitive and subtle. In other studies of reciprocity, it has been unclear whether the existence of the network tie ($A \rightarrow S$) is motivating S 's desire to link to A , or whether S is motivated to link to A due to the same reasons A linked to S , namely, interpersonal similarity. Note that in our experiment, S was not aware of the tie to A , so any desire to link to A cannot be based on the structural existence of the incoming tie. Therefore, we can conclude that in studies in which the $A \rightarrow S$ tie has a significant effect on the formation of the $S \rightarrow A$ tie, the visibility of the $A \rightarrow S$ tie is having a positive effect on S 's choice to link to A , separate from any choice based on similarity or attraction.

In short, because we do not observe an effect for the $A \rightarrow S$ tie when it is not visible to the subject, we support the hypothesis that $S \rightarrow A$ ties form *due* to reciprocity for its own sake: the

establishment of a tie to the alter, specifically because the alter linked to the subject first.

2) Mutuality and Transitivity

We now turn to the nine two-step path types. Each numbered item in the "Subject-Alter Paths" section in Table 2 corresponds to a kind of path from the subject to the alter.⁹ Each is normalized by dividing by the number of nodes in the subjects' network. The reason for this is simple: 10 paths from S to A would be a lot more meaningful if S has 100 neighbors than if S has 10,000.

Following Holland and Leinhardt [11], we call the $S \leftrightarrow X$ tie *mutual*. In ties exhibiting *mutuality*, the two actors mutually attend to one another.¹⁰ We observe that all of the two-step paths with at least one mutual tie (5-9 in Figure 2) have a statistically significant effect on the rating, and those that do not have a mutual tie (1-4) do not, with the exception of (2).

A dyad exhibiting mutuality is one in which each actor pays attention to and receives information from the other. As described earlier, such behavior leads to more social cohesion [9] and stronger affective ties [19]. Crucially mutuality affords Twitterers the ability to engage in conversation, which may also be a proxy for relational strength. Since each gives attention to the other, mutuality may also be a proxy for equal status. More investigation is needed. In any case, our results show that mutuality of at least one of the two steps in the path is an important criterion in determining whether a path will be influential to the subject.

Notice that, of the paths exhibiting mutuality, (5,7,9) have positive coefficients and (6,8) have negative coefficients. Paths (6,8) do not exhibit transitivity – there is no directed flow of attention from S to A – while paths (5,7,9) do. That is, transitivity is associated with an increased desire to form a tie; this supports findings by Holland and Leinhardt [11] and Romero and Kleinberg [23]. The finding that lack of transitivity has a significant, negative effect (rather than no effect) has not, to our knowledge, been observed before.

Again, status offers an explanation. A transitive path like $S \leftrightarrow X \rightarrow A$ indicates a consistent status hierarchy; A is higher-status than X , who is equal status to S . Since A is higher status than S , S might like to pay attention to A . In contrast, paths like (6) and (8) indicate A is lower status than S ; A is either lower status than someone (X) who is a status equal of S , or A is the status equal of someone (X) lower status than S (8).

An illustrative example is a request for an introduction. Suppose, given a relationship like (5) or (7), S would like to know A , and would like X to introduce them. Examples include a graduate student asking her advisor for an introduction to a senior scholar, or one entrepreneur asking another for an introduction to a potential investor. Though the "S" would be happy to know the "A", flipping the arrows around we see that status differences indicate the "A" might not be interested in meeting the "S".

⁸ <http://www.techcrunch.com/2009/08/26/twitters-golden-ratio-that-no-one-likes-to-talk-about/>

⁹ Instead of the $A-X-B$ notation used in Figure 2, we refer to paths like $S-X-A$ to denote (S)ubject and (A)lter.

¹⁰ Note that our use of "mutuality" is unrelated to its use in [21], which uses it to describe incidence of having friends in common.

VI. DISCUSSION

Earlier, in trying to develop an intuitive sense behind directed two-step paths, we identified concepts like shared audiences, shared interests, transitivity and filtration.

Paths (3) and (4), which describe shared interests and shared audiences, respectively, did not have statistically significant effects. Therefore, simply sharing a neighbor X in these ways do not appear to be sufficient bases for new tie formation. However, the conditions of (3) are satisfied by (6,7) and the conditions of (4) by (5,8). Paths (5,7) also satisfy mutuality and transitivity, and we can think of (5) as a special case of a shared audience tie, and (7) as a special case of a shared interest tie. In future work, we hope to use text analysis to examine how text content affects closure in the context of the (5) and (7) path types.

Regarding transitivity, we considered two alternative outcomes: first, that the middle individual in the $A \rightarrow X \rightarrow B$ triad would filter B 's tweets, which would redound to the benefit of A ; in effect X is a curator for B . Second, A would learn about B via X and, taking X 's link to B as a vote of confidence, begin to follow B himself; this is the link-copying mechanism behind preferential attachment [23]. We can rule out filtering as a mechanism, since filtering would imply negative, not positive coefficients for (5,7,9).

However, we can rule out link copying as well. Our experimental design made subjects *blind* to existing connections, thus ruling out observational learning as an underlying mechanism. We can also rule out purely structural effects (i.e. increased exposure); though our selection process selected for alters two steps away, after that selection, all alters were equally likely to be selected.

Though positive coefficients for (5,7,9) would be expected if a copying mechanism were in effect, we have found positive effects for (5,7,9) even in an experimental design in which copying is deliberately impossible. Though preferential attachment generates similar results, we caution that a different underlying mechanism may be at work.

VII. CONCLUSION

We designed our experiment so that subjects would be blind to the connections between them and the alters they were rating. We wanted to isolate the psychological and structural effects of homophily and our results show that while structure is a good proxy for measuring the desire to form ties, since subjects could not see that structure, it could not be the mechanism driving the desire. Though all the selections made by our subjects were driven purely by choice, it is important to recognize that structure provides a useful means of estimating choice-driven results.

We have thus been able to observe that reciprocity does not appear to take place when the ties individuals would be reciprocating are not visible, and that transitivity and mutuality are important conditions which, together, are associated with an increased desire to form ties; we suggest that a consistent status hierarchy and some level of tie strength drive this effect.

This experiment is suggestive rather than definitive, and future work is called for – examination of these patterns in other

domains, and in naturalistic in addition to experimental data. For those doing applied work and system design, our work can inform the improvement of recommendations offered by "friend suggestion" algorithms.

REFERENCES

- [1] danah boyd, Scott A. Golder, and Gilad Lotan. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. *Proc. HICSS-43.*, 2010.
- [2] Ronald S. Burt. *Structural Holes: The Social Structure of Competition*. Harvard University Press, 1992.
- [3] Dorwin Cartwright and Frank Harary. Structural balance: A generalization of heider's theory. *The Psychological Review*, 63(5):277–293, 1956.
- [4] Robert Cialdini. *Influence: Science and Practice*. Allyn and Bacon, Needham, MA, 2001.
- [5] Mick Couper. *Designing Effective Web Surveys*. Cambridge University Press, 2008.
- [6] Richard M. Emerson. Power-dependence relations. *American Sociological Review*, 27:31–41, 1962.
- [7] Scott Feld. The focused organization of social ties. *American Journal of Sociology*, 86(5):1015–1035, 1981.
- [8] Scott Feld. Why your friends have more friends than you do. *American Journal of Sociology*, 96(6):1464–1477, 1991.
- [9] Alvin Gouldner. The norm of reciprocity: A preliminary statement. *American Sociological Review*, 25(2):161–178, 1960.
- [10] Mark Granovetter. The strength of weak ties. *American Journal of Sociology*, 78:1360–1380, 1973.
- [11] Paul W. Holland and Samuel Leinhardt. A method for detecting structure in sociometric data. *American Journal of Sociology*, 76(3):492–513, 1970.
- [12] George C. Homans. *Social Behavior: Its Elementary Forms*. Harcourt Brace Jovanovich, New York, 1974.
- [13] Courtenay Honeycutt and Susan C. Herring. Beyond microblogging: Conversation and collaboration via twitter. *Proc. HICSS-42*, 2009.
- [14] John Horrigan. Home broadband adoption Technical report, Pew Internet and American Life Project, June 19, 2009 2009.
- [15] Norman P. Hummon and Patrick Doreian. Connectivity in a citation network: The development of DNA theory. *Social Networks*, 11:39–63, 1989.
- [16] James F. Jekel, David L. Katz, and Joann G. Elmore. *Epidemiology, Biostatistics and Preventive Medicine*. Elsevier, 2nd edition, 2001.
- [17] Jon Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46(5):604–632, 1999.
- [18] Miller McPherson, Lynn Smith-Lovin, and James M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444, 2001.
- [19] Linda D. Molm. Theoretical comparisons of forms of exchange. *Sociological Theory*, 21(1):1–17, 2003.
- [20] Peter R. Monge and Noshir S. Contractor. *Theories of Communication Networks*. Oxford University Press, 2003.
- [21] Mark E. J. Newman. Ego-centered networks and the ripple effect. *Social Networks*, 25(1):83–95, 2003.
- [22] Gabriele Plickert, Rochelle R. Cote, and Barry Wellman. It's not who you know, it's how you know them: Who exchanges what with whom? *Social Networks*, 29:405–429, 2007.
- [23] Daniel M. Romero and Jon Kleinberg. The directed closure process in information networks with an analysis of link formation on twitter. In *International Conference on Weblogs and Social Media (ICWSM)*, 2010.
- [24] Cass Sunstein. *Republic.com*. Princeton University Press, 2001.
- [25] Stanley Wasserman and Katharine Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, 1994.