

# Maintaining Ties on Social Media Sites: The Competing Effects of Balance, Exchange, and Betweenness

**Daniel M. Romero**  
Cornell University  
dmr239@cornell.edu

**Brendan Meeder**  
Carnegie Mellon  
bmeeder@cs.cmu.edu

**Vladimir Barash**  
Cornell University  
vdb5@cornell.edu

**Jon Kleinberg**  
Cornell University  
kleinber@cs.cornell.edu

## Abstract

When users interact with one another on social media sites, the volume and frequency of their communication can shift over time, as their interaction is affected by a number of factors. In particular, if two users develop mutual relationships to third parties, this can exert a complex effect on the level of interaction between the two users – it has the potential to strengthen their relationship, through processes related to triadic closure, but it can also weaken their relationship, by drawing their communication away from one another and toward these newly formed connections.

We analyze the interplay of these competing forces and relate the underlying issues to classical theories in sociology – the theory of balance, the theory of exchange, and betweenness. Our setting forms an intriguing testing ground for these two theories, in that it provides a scenario in which their qualitative predictions are largely at odds with one another.

In the course of our analysis, we also provide novel approaches for dealing with a common methodological problem in studying ties on social media sites: the tremendous volatility of these ties over time makes it hard to compare one’s results to simple baselines that assume static or stable ties, and hence we must develop a set of more complex baselines that takes this temporal behavior into account.

## 1 Introduction

In studying the interactions on a social media site, a basic question is to understand what causes relationships among users to be strengthened and what causes them to weaken. This is an issue that is not well understood: there are multiple forces that govern the strengths of social ties and pull in competing directions. It is an important problem to design methods of analysis for these systems that can begin to separate out the effects of these different forces. Existing work in on-line domains has approached this issue by identifying dimensions that characterize the strength of ties (Gilbert and Karahalios 2009), and by incorporating factors such as triadic and focal closure (Kossinets and Watts 2006), similarity among individuals (Anagnostopoulos, Kumar, and Mahdian 2008; Crandall et al. 2008; Aral, Muchnik, and Sundararajan 2009; Kossinets and Watts

2009), and the role of positive and negative relationships (Leskovec, Huttenlocher, and Kleinberg 2010).

Here we develop an analysis framework through which we can use data from social media sites to begin isolating the effects of three distinct social forces on the strengths of relationships: *balance*, *exchange*, and *betweenness*. We begin by describing how these forces operate in a social media context, which will also make clear the sense in which they can produce opposite effects. For this discussion we will focus on undirected links, in which relationships are symmetric.

**Balance and Exchange.** First, we consider the force of balance. Suppose we have a user  $B$  who is friends with users  $A$  and  $C$ . The principle of balance argues that if  $A$  and  $C$  do not have a social tie, this absence introduces latent strain into the  $B-A$  and  $B-C$  relationships, and this strain can be alleviated if an  $A-C$  tie forms (Heider 1958; Rapoport 1953). Hence, balance is a force that causes the formation of an  $A-C$  tie to strengthen the  $B-A$  tie, when  $C$  is also linked to  $B$ .<sup>1</sup>

Counterbalancing this is an equally natural force, which is the principle of *exchange* (Emerson 1962; Willer 1999). Let’s return to the user  $B$  who is friends with users  $A$  and  $C$ . If  $A$  were to become friends with  $C$ , this provides  $A$  with more social interaction options than she had previously. The theory of exchange argues that this makes  $A$  less dependent on  $B$  for social interaction, thereby weakening the  $B-A$  tie.

Figure 1 is a schematic diagram of the forces of balance and exchange as they act on a set of three nodes. Our first set of analyses studies the aggregate effect of these forces on the communication patterns between Twitter users. For this, we say that a *tie* between two Twitter users has formed when they have each sent at least 3 @-messages to the other.<sup>2</sup> We

<sup>1</sup>One sees balance theory applied in two related contexts when we consider scenarios such as this, when  $B$  has positive relations with  $A$  and  $C$ . In one line of argument, the absence of an  $A-C$  link produces stress that needs to be resolved. A related line of argument considers situations in which there is in fact antagonism between  $A$  and  $C$ , which produces even stronger forms of stress (Cartwright and Harary 1956). Both of these situations point to the same conclusions, and both fall under the principle of balance.

<sup>2</sup>@-messages are a basic Twitter mechanism in which one user directs a tweet to another; since they are used between people who know one another as well from users toward celebrities, we require

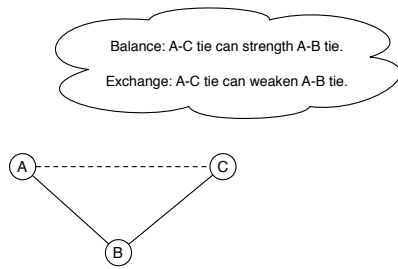


Figure 1: The theories of balance and exchange postulate the effect of A and C forming a relationship on the B-A and B-C relationships.

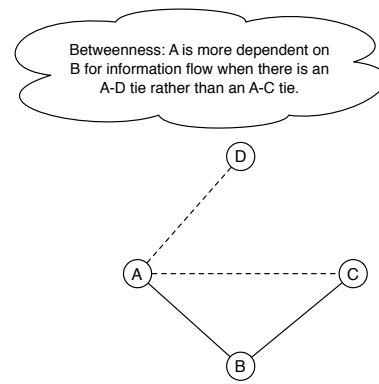


Figure 3: Betweenness postulates that A is more dependent on B for information flow when there is an A-D tie rather than an A-C tie.




Figure 2: Outside influence: The A-B relationship is potentially weakened not only by additional relationships within the online social network, but also by activities that altogether draw users away from the network.

examine ties between users in a large collection of public tweets. We also consider scenarios, such as the one pictured in Figure 1, in which a user  $B$  has ties to users  $A$  and  $C$ , and look at cases in which an  $A-C$  tie does or does not form.

### Decaying Relationships and Outside Opportunities.

We find first of all that the formation of an  $A-C$  link in our Twitter data makes it significantly more likely that the  $A-B$  tie will persist (as measured by the generation of future messages from  $A$  to  $B$ ). At one level, this points to the dominance of balance over exchange in this particular scenario; however, as we investigate the effect of tie formation on tie persistence more closely, a more subtle picture emerges. Going back to users  $A$ ,  $B$ , and  $C$ , suppose that we consider the effect on the  $A-B$  tie of  $A$ 's sending  $k$  messages to arbitrary users other than  $B$ , for some relatively large value of  $k$  — potentially even requiring these messages to go to users not linked to  $B$ . Even in this case, these messages from  $A$  to others lead to an increase in the persistence of the  $A-B$  tie.

This observation underscores the need to be careful in reasoning about how the persistence of ties operates on a social media site. One might suppose, via the principle of

multiple reciprocations before we consider the messaging to constitute evidence of a tie.

exchange, that the  $k$  messages from  $A$  to others divert  $A$ 's attention from  $B$ , to the detriment of the  $A-B$  tie. But we should step back and think about the full set of activities that might draw  $A$  away from  $B$ . Interaction with other users on Twitter is one source of such activities. However, there are many activities completely outside Twitter that might draw  $A$ 's attention away from  $B$  as well. Thus, abstractly, the picture from Figure 1 should be expanded to look more like the larger picture in Figure 2.

In context of Figure 2, the principle of exchange is not irrelevant to the discussion, but we are applying it too narrowly if we view other Twitter users as the only sources of outside opportunities for  $A$  in the  $A-B$  relationship. And the point, then, is that  $k$  messages from  $A$  to many users other than  $B$  still provide strong evidence that  $A$  is actively involved in Twitter, rather than in other activities. This increased involvement makes it easier for  $A$ 's Twitter activity to “spill over” to the  $A-B$  tie.

In Section 4, we consider ways of capturing this spillover effect, and propose a reconceptualization of exchange theory in the particular context of social media to integrate the outside opportunities of a user  $A$  at both the “micro” level (to other users on the site) and the “macro” level (to potentially unobserved activities off the site).

This framework also suggests an important methodological consideration that is underscored by our analyses. Social media sites are domains in which the typical relationship exists in a state of rapid decay, since either user involved in the relationship may begin to rapidly reduce their involvement in the site, or leave it altogether and never return. Such issues are much less of a constraint (even if they are present at lower levels) in analyses of relationships in the physical world — but in on-line settings, they need to be carefully controlled for.

**Balance and Betweenness.** Given these considerations, we explore a further set of questions about social forces and relationships in which we control for  $A$ 's overall level of in-

involvement in the site. Specifically, consider again a user  $B$  who has ties with users  $A$  and  $C$ . Now, let a fixed amount of time pass, and consider two possible scenarios: (i)  $A$  forms a tie with  $C$ , or (ii)  $A$  forms a tie with a user  $D$  who is not connected with  $B$ . In which scenario is the  $A$ - $B$  more persistent? (See Figure 3.) Both (i) and (ii) provide evidence of comparable involvement by  $A$  in the site, and so we must look to the finer structure of the interaction pattern to decide which has a more positive effect on the  $A$ - $B$  tie.

As before, the principle of balance argues that the  $A$ - $B$  tie should be more strengthened in scenario (i). The principle of exchange is a bit tricky to apply here, but we can use the principle of *betweenness* instead to identify a natural argument that says that scenario (ii) should be better for the  $A$ - $B$  link. The principle of *betweenness* is used, for example, by Burt (1992) in his formulation of the theory of structural holes.

The argument for betweenness is as follows. Twitter is an environment in which access to information, and the flow of information, is a crucial force in the shaping of users' activities — consider, for example, the set of social and informational links that are activated whenever a piece of content is extensively retweeted (repeated by users). As a result, when there is no  $A$ - $C$  link, user  $B$  plays an important brokerage role in her relationship with  $A$ :  $B$  provides  $A$  with access to information from  $C$ . If a direct  $A$ - $C$  tie forms, this brokerage role is sharply diminished; on the other hand, the role is not as strongly diminished if  $A$  forms a tie with  $D$ . Thus, considerations of betweenness and brokerage suggest that the  $A$ - $B$  might persist more strongly in scenario (ii), with the formation of an  $A$ - $D$  tie, rather than in scenario (i), with the formation of an  $A$ - $C$  tie.

In Section 3, we carry out a careful analysis of this trade-off, finding significant evidence that the balance argument is operating more strongly than the betweenness argument in the setting of Twitter: the closing of the  $A$ - $B$ - $C$  triangle (as in scenario (i)) has a more positive effect on the  $A$ - $B$  relationship than the formation of ties by  $A$  that leave it open (as in scenario (ii)).

**Persistence of Ties.** Finally, in Section 5, we develop further methodologies for analyzing the persistence of relationships in social media domains such as Twitter, given the rapid rate at which they decay over time. In particular, we identify fundamental asymmetries in the way that relationships ramp up in intensity compared to the way in which they fall off after their peak level of activity, and we show how the closing of triads in the vicinity of a tie can have important effects on its persistence.

## 2 Data Set and Network Extraction

We have collected and processed a large corpus of data from the Twitter social network. From August 2009 until January 2010, we crawled Twitter using their publicly available API. Twitter provides access to only a limited history of tweets through its search mechanism; however, because user identifiers have been assigned contiguously since an early point in time, we simply crawled each user in a comprehensive

range. Due to limitations of the API, if a user has more than 3,200 tweets we can only recover the last 3,200 tweets; all messages of any user with fewer than this many tweets are available. We collected over three-billion messages from more than 60 million users during this crawl.

The primary analysis of this data is to extract all @-messages and build a temporal network of 'attention relationships.' A directed edge exists from user  $A$  to  $B$  if  $A$  sends at least  $k$  @-messages to  $B$ ; the time this edge is created,  $t_D(A, B)$ , is the time at which the  $k$ th @-message is sent. In our analyses we use  $k = 3$ . There are multiple ways of defining a network, and our definition is one way of defining a proxy for the attention that a user  $A$  pays to other users. The resulting network contains 8,509,140 non-isolated nodes and 50,814,366 links.

From this directed, temporal network we extract an undirected, temporal network of ties. An undirected edge between two users  $A$  and  $B$  is formed when  $A$  has sent at least 3 @-messages to  $B$  and  $B$  has sent at least 3 @-messages to  $A$ . The edge  $E = (A, B)$  has time-stamp equal to  $t(A, B) = \max\{t_D(A, B), t_D(B, A)\}$ , the later of the times when the two directed edges were formed. This tie network contains 20,492,393 ties between 3,701,860 users, and although fewer than half of the users remain in the tie network, over 80% of attention relationships contribute to a tie.

We define an *open triad*  $O$  as a graph of three nodes  $A$ ,  $B$ , and  $C$  containing the ties  $(A, B)$  and  $(B, C)$ . The time-stamp of the open triad is  $O_t = \max\{t(A, B), t(B, C)\}$ , the time at which the last of the two ties forms. Open triads  $O = (A, B, C)$  in which the undirected  $(A, C)$  edge eventually forms are said to *close*. We define an open triad that closes  $d$  days after  $O_t$  ( $t(A, C)$  is  $d$  days after  $O_t$ ) to be a *d-closed triad*.

## 3 Balance Vs. Betweenness

We begin by considering the contrast between balance and betweenness discussed in the introduction. We take an open triad  $(A, B, C)$ , and as in Figure 3, we compare the amount of interaction from  $A$  to  $B$  after one of the following two events takes place: (i) the  $A$ - $C$  tie forms, or (ii)  $A$  forms a tie with a user  $D$  who is not connected to  $B$ . Because we have recorded not only the evolution of a triad (whether it closed or not), but also the communication times, we can control for factors such as the delay between triad formation and the creation of the additional tie. Additionally, we will control for  $A$  being 'active'; we make sure that  $A$  was communicating when the triad formed, when the new tie forms, and some time after the new tie formed. In this way, we will not end up studying phenomena that arise primarily because users are immediately leaving the site.

**Representing the competing scenarios.** In particular, we consider the percentage of messages that  $A$  directs to  $B$  in two comparison sets of triads designed to represent scenarios (i) and (ii). First, we choose a value for  $d$  and consider all  $d$ -closed triads; we also want to guarantee that  $A$  had a certain minimum level of activity overall, so we require that

$A$  sent between 200 and 1000 messages in total after the open triad  $(A, B, C)$  was formed, and moreover that  $A$  sent at least one message 1,  $d$ , and  $2d$  days after the open triad was created. This subset of  $d$ -closed triads with these conditions ensuring  $A$  is sufficiently active forms our population for scenario (i).

For scenario (ii), we want an open triad  $(A, B, C)$  where  $A$  sends a message to a node not connected to  $B$ . Thus, for each triad  $O' = (A, B, C)$  that never closes, we look at all of the nodes  $D$  that are not connected to  $B$ , and with which  $A$  forms a tie after  $O'_t$ . We pick such a node  $D$  at random and say that  $O'$  is  $d$ -open, where  $d$  is the number of days after  $O'_t$  that the  $A$ - $D$  tie formed. As before, we also require that  $A$  sent between 200 and 1000 messages after the open triad  $(A, B, C)$  was formed, and that  $A$  sent at least one message 1,  $d$ , and  $2d$  days after the open triad was created. This population of  $d$ -open triads with these conditions on  $A$  forms our population for scenario (ii).

For each population, we measure the percentage of  $A$ 's communication that goes toward  $B$ , as a function of the time since the formation of the open triad. As noted in the introduction, relationships on social media sites have a default tendency to decay, but by observing which scenario provides a slower aggregate decay rate for the  $A$ - $B$  tie, we can begin to learn about the different effects of balance (scenario (i)) and betweenness (scenario (ii)).


**Results.** In Figure 4, we adopt this test with  $d$  chosen to be 1, 3, 5, and 10 days. Each plot shows the average percentage of messages that  $A$  sent to node  $B$  as a function of the number of days after  $O_t$ . The red curve is based on the  $d$ -closed triads, while the green curve is based on the  $d$ -open triads.

We observe first that for all choices of  $d$ , the red curve decreases at a slower rate than the green curve. This indicates that the  $A$ - $B$  tie decays more slowly in the population corresponding to scenario (i). But beyond this, the gap between the two curves is widening: the rate at which they decrease is separating. After day 100 (about three months after the formation of  $A$ 's additional connection), the communication percentage for the open triads decreases at a noticeably faster rate. This suggests that closing the triad benefits communication from  $A$  to  $B$  by slowing the inevitably decreasing amount of online interaction.


In interpreting these results as evidence for the effect of balance, it is important to understand that the formation of the  $A$ - $C$  tie is not causing the extent of  $A$ - $B$  interaction to increase in an absolute sense, but rather for its rate of decay to be slowed. In general, the effect of social forces on relationships in our analysis is ubiquitously modulated by the overall rate of link decay on Twitter.

#### 4 Exchange Theory and Spill-Over Effects


In the previous section, we observed that in the triad  $(A, B, C)$  the communication between  $A$  and  $B$  benefits in the long run from the triad's closing. At a more general level, we will now ask what can be predicted about the  $A$ - $B$  interaction from knowledge of how active  $A$  was with respect to users other than  $B$ .




(a)  $d = 1$



(b)  $d = 3$



(c)  $d = 5$



(d)  $d = 10$

Figure 4: Percentage of message from  $A$  to  $B$  vs. the number of day after creation of open triad. The green curve is based on the  $d$ -open triads and the red curve is based on the  $d$ -closed triads.  $A$  must have sent from 200 to 1000 messages in total after day = 0 and  $A$  must have sent at least one messages on days 1,  $d$ , and  $2d$ .




Figure 5: Number of messages  $A$  sends to everyone but  $B$  vs. number of messages  $A$  sends to  $B$ , 3 days after the creation of the  $A$ - $B$  edge.

Exchange theory posits that as  $A$  has more “outside options” provided by communication partners who are not  $B$ ,  $A$  will spend less time communicating with  $B$ . One hypothesis, then, is that as  $A$  spends more time talking to her friends who are not  $B$ ,  $A$ ’s communication with  $B$  will decrease. Alternatively, we can consider a simple model based on the schematic picture in Figure 2, where  $A$  first decides how much time to spend on Twitter, and then divides that time evenly between all of her friends on Twitter. According to this model, the more time  $A$  spends talking to anyone on Twitter, the more time she will spend talking to  $B$  as well.

We test these two predictions by plotting the number of messages  $A$  sends to everyone but  $B$  vs. the number of messages that  $A$  sends to  $B$  for various points in time after the creation of the  $A$ - $B$  edge. The plots have the same general shape up to several weeks after the creation of the edge. In Figure 5, we present the plot for three days after the creation of the edge. The figure shows a pattern of monotonic increase, which suggests that the second model is a better approximation to the real outcome: the more  $A$  talks to anyone on Twitter, the more she talks to  $B$  as well.

**The Role of Balance in Spill-Over Effects.** This analysis makes precise the sense in which we think of  $A$ ’s activity toward users other than  $B$  as “spilling over” in a positive way toward  $B$ . We now show that the principle of balance can enhance this spill-over effect. To do this, we consider the set-up above, but vary the number of  $A$ ’s messages that go to users with whom  $B$  also has ties.

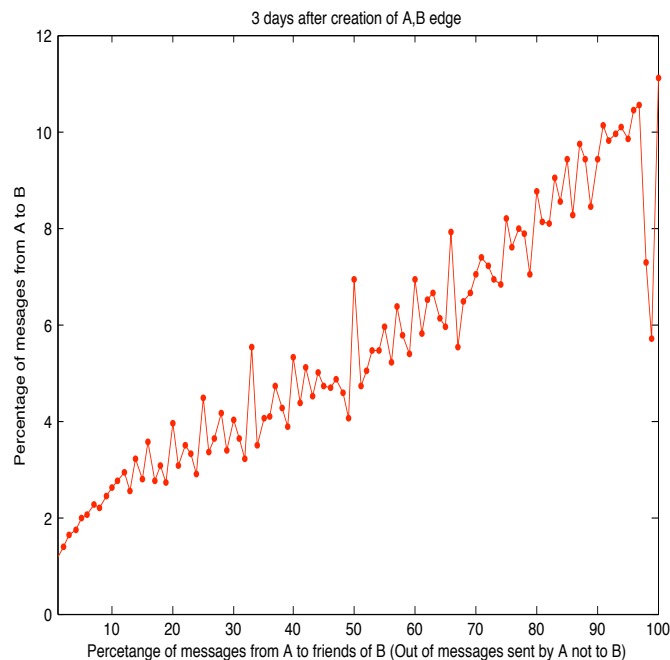


Figure 6: Percentage of messages that  $A$  sends to  $B$  as a function of the percentage of  $A$ ’s non- $B$  messages that go to friends of  $B$ . These messages take place 3 days after the creation of the  $A$ - $B$  edge.

In particular, Figure 6 depicts the following analysis. We consider the messages sent by  $A$  to users other than  $B$ , and ask what fraction of these messages go to users  $C$  with whom  $B$  also has a tie. What Figure 6 shows is that the percentage of messages from  $A$  to  $B$ ’s friends increases: in other words, the spill-over in  $A$ ’s activity toward  $B$  is accentuated when  $A$ ’s activity toward users other than  $B$  takes place with friends of  $B$ .

There is a respect in which Figure 6 can be a bit subtle to interpret, based on the fact that it aggregates many users  $A$  of different activity levels. As a result, we show (in Figure 7) a related analysis in which the set-up is identical except that we require  $A$  to have sent exactly 10 messages to users other than  $B$ . We then ask: how many messages does  $A$  send to  $B$ , as a function of the number (out of 10) of these non- $B$  messages that go to friends of  $B$ ? Again we find that the spill-over in  $A$ ’s activity is enhanced when  $A$ ’s non- $B$  activities include many friends of  $B$ . Indeed, we see in Figure 7 a striking super-linear relationship whereby the spill-over effect ramps up very rapidly once most of  $A$ ’s non- $B$  communication is directed at friends of  $B$ .

**A Situation with Apparent Lack of Spill-Over.** Thus far we have not seen any situations in which  $A$ ’s activity toward users other than  $B$  has had any kind of negative effect on the  $A$ - $B$  tie. Here we identify the possibility of one such situation, leaving the underlying mechanism for it as an open




Figure 7: Number of messages  $A$  sends to friends of  $B$  vs. number of messages  $A$  sends to  $B$ , five days after the creation of the  $A$ - $B$  edge. Node  $A$  sent exactly 10 messages in total to users other than  $B$ .




Figure 8: Zoom-in of figure 4(d). We observe jumps on the green curve at days  $d$  and  $2d$  and on the red curve at day  $2d$  but not on day  $d$ .

question.

The situation is the following. Figure 8 zooms in around the days  $d$  and  $2d$  (in this case 10 and 20) on the curves from Figure 4. We observe that the green curve has jumps on days 10 and 20, while the red curve only has a jump on day 20. The jumps can be explained by the fact that to construct the curves we only take triads in which  $A$  sent messages on days  $d$  and  $2d$  and therefore there is an increased likelihood that a fraction of those messages were sent to  $B$ . However, we do not see such jumps on day  $d$  on the red curve even though node  $A$  was active on that day in the  $d$ -closed triads as well as the  $d$ -open triads. The only difference between the red and the green curves is that on day  $d$ ,  $A$  messaged a neighbor of  $B$  in the red curve, but in the green curve  $A$  messaged a node  $D$ , not connected to  $B$ . The lack of jump on day  $d$  can be observed on all the plots of Figure 4. This suggests that the communication from  $A$  to  $B$  is in some sense suppressed on the day of the triad's closure, and hence points to a possible case in which  $A$ 's actions toward others are reducing the level of activity on the  $A$ - $B$  link. Understanding the extent of this effect and the mechanism behind it is an intriguing open question.

## 5 Basic Properties of Relationship Decay

Since much of our analysis involves the basic fact that interactions on Twitter decay over time, making sustained ties hard to maintain, we now explore the basic properties of relationship decay in more detail.

We begin with a simple question. If we observe an event in which  $A$  communicates with  $B$  on day 0, what is the probability that we will observe another such  $A$ -to- $B$  communication event on day  $d > 0$ ? Figure 9 shows how this probability decreases as a function of  $d$ : note that it starts with a decay rate that is slower than exponential (though also not a very close fit to a power law), and then straightens out into an approximately exponential rate. We note that the rate of decay is faster than for the curves in Figure 4, which were based on  $d$ -open and  $d$ -closed triads. Hence, users  $A$  involved in triads tend to maintain their communication with users  $B$  more than the average. One possible interpretation is that their involvement in triads is indicative of a higher level of activity on Twitter overall, triggering the types of spill-over effects that were the focus of the previous section.

The probability of seeing future communication based on just a single observation is one extreme in this genre of questions. At the other extreme, we can study the dynamics of a strong relationship from one user to another, in which many messages are sent. Specifically, let's consider a pair of users  $(A, B)$  for which  $A$  sends  $B$  at least 100 messages total. We then investigate how the amount of communication from  $A$  to  $B$  changes over time. For each such  $(A, B)$  pair, we partition time into bins of length one week and look for the week during which  $A$  sent the most @-messages to  $B$ . This is the *peak* of the communication. We define the function  $M$  as follows:




Figure 9: Probability that A will send a message to B  $d$  days after having sent her one




Figure 10: Average function  $M$  for pairs  $(A, B)$  in which  $A$  sent  $B$  at least 100 @-messages

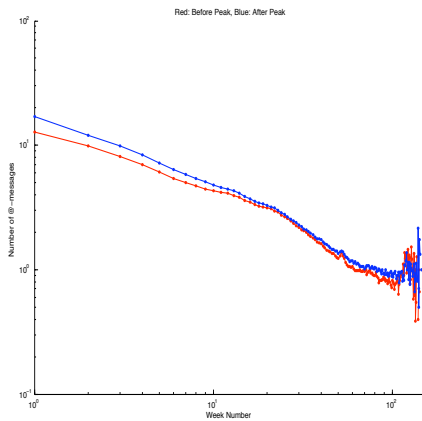


Figure 11: Average of  $\log(M(n))$  as a function of  $\log(n)$  where  $n > 0$  (in blue) and as a function of  $\log(-n)$  for  $n < 0$  (in red)




Figure 12: Average of  $\log(M(n))$  as a function of  $\log(n)$  where  $n > 0$  (in blue for unreciprocated links and green for reciprocated ones) and as a function of  $\log(-n)$  for  $n < 0$  (in red for unreciprocated links and black for reciprocated ones)

$$M(n) = \begin{cases} \text{Num @-mess during peak} & \text{if } n = 0 \\ \text{Num @-mess during } n^{\text{th}} \text{ week before the peak} & \text{if } n < 0 \\ \text{Num @-mess during } n^{\text{th}} \text{ week after the peak} & \text{if } n > 0 \end{cases}$$

Figure 10 shows the average function  $M$  for pairs  $(A, B)$  in which  $A$  sent  $B$  at least 100 @-messages. Figure 11 shows the average of  $\log(M(n))$  as a function of  $\log(n)$  where  $n > 0$  (in blue) and as a function of  $\log(-n)$  for  $n < 0$  (in red). The blue curve is above the red curve which suggests that the communication between  $A$  and  $B$  tends to ramp up to the peak faster than it decays from the peak.

We can refine this analysis a bit further as follows. When a user  $A$  sends a large number of messages to a user  $B$ , there are two possibilities: (i) it could be that  $B$  never send any messages to  $A$  (perhaps because  $B$  is simply a celebrity that  $A$  mentions on a regular basis); or (ii) it could be because  $A$  and  $B$  are actually exchanging messages, suggesting a more overtly social form of interaction. With this in mind, we can consider the plot in Figure 10 broken down separately based on whether the messages from  $A$  to  $B$  are *reciprocated* (with  $B$  messaging  $A$  as well) or *unreciprocated* (with no  $B$ -to- $A$  messages). Figure 12 shows the results for these two categories: we find that the rates of ramp-up and decay do in fact differ between the two, with the curves for unreciprocated links lying slightly above the corresponding curves for reciprocated links. This suggests that the ramp-up and ramp-down for reciprocated links is in fact slightly more abrupt than it is in the unreciprocated case.

## 6 Conclusions and Future Work

There are many forces that affect the strength and longevity of ties on social media sites, and it is a challenge to separate these into their distinct effects. In this paper we have offered a set of data analysis methodologies that lets us begin to isolate the effect of three such forces: balance, in which ties are strengthened when they close triads; exchange, in which ties

are weakened when one end of the tie has other opportunities; and betweenness, in which ties are strengthened when they serve as conduits for information.

Our analyses show the power of balance in the domain we study, Twitter. It also shows that exchange theory should be broadened to conceptually include off-site opportunities for participants in a tie, reflecting the rapid rate at which ties decay. We believe that the framework developed here can be applied to social media settings quite broadly. In particular, it could be used to analyze the differential rates and trajectories by which relationships grow and decay across different domains, and more intriguingly, it could expose contrasting relative extents to which balance, exchange, and betweenness apply across domains. Ultimately, being able to characterize different social applications through the different ways in which these forces operate could provide a useful framework for modeling and reasoning about the behavior of these applications.

## References

- Anagnostopoulos, A.; Kumar, R.; and Mahdian, M. 2008. Influence and correlation in social networks. In *Proc. 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 7–15.
- Aral, S.; Muchnik, L.; and Sundararajan, A. 2009. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proc. Natl. Acad. Sci. USA* 106(51):21544–21549.
- Burt, R. S. 1992. *Structural Holes: The Social Structure of Competition*. Harvard University Press.
- Cartwright, D., and Harary, F. 1956. Structure balance: A generalization of Heider's theory. *Psychological Review* 63(5):277–293.
- Crandall, D.; Cosley, D.; Huttenlocher, D.; Kleinberg, J.; and Suri, S. 2008. Feedback effects between similarity and social influence in online communities. In *Proc. 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 160–168.
- Emerson, R. M. 1962. Power-dependence relations. *American Sociological Review* 27:31–40.
- Gilbert, E., and Karahalios, K. 2009. Predicting tie strength with social media. In *Proc. 27th ACM Conference on Human Factors in Computing Systems*, 211–220.
- Heider, F. 1958. *The Psychology of Interpersonal Relations*. John Wiley & Sons.
- Kossinets, G., and Watts, D. 2006. Empirical analysis of an evolving social network. *Science* 311:88–90.
- Kossinets, G., and Watts, D. 2009. Origins of homophily in an evolving social network. *American Journal of Sociology* 115(2):405–50.
- Leskovec, J.; Huttenlocher, D.; and Kleinberg, J. 2010. Signed networks in social media. In *Proc. 28th ACM Conference on Human Factors in Computing Systems*, 1361–1370.
- Rapoport, A. 1953. Spread of information through a population with socio-structural bias I: Assumption of transitivity. *Bulletin of Mathematical Biophysics* 15(4):523–533.
- Willer, D. 1999. *Network Exchange Theory*. Praeger.