Incentivized Sharing in Social Networks

Joseph J. Pfeiffer III*
Department of Computer Science
Purdue University, West Lafayette, IN, USA
jpfeiffer@purdue.edu

Elena Zheleva
LivingSocial
Washington, D.C., USA
elena.zheleva@livingsocial.com

ABSTRACT

With the proliferation of online social networks in recent years, there has been an increasing interest in studying social phenomena at a larger scale than it was ever possible before. Many companies have started to create social media strategies to keep up with the buzz around social networks. These strategies consider how to increase the companies' digital presence and how to positively impact the adoption of their products and services. Meanwhile, the question of how monetary incentives for social sharing affect the behavior of social network users remains largely unexplored. Here, we present a case study in which a particular type of incentive changes the structural properties of a social network and shifts the power-law curve of sharing. We distinguish between altruistic and incentivized shares, and we look at the impact of different incentive amounts on the sharing behavior of users. We also propose future directions for this type of research.

Keywords

social networks, experiments, incentives, sharing

1. INTRODUCTION

As a myriad of websites centered around social media and social networks have appeared online over the last few years, a revolutionary change in the way people share information with one another has occurred. Companies such as LivingSocial[1], Amazon[2], Facebook[3] and Google[4] have websites where millions of users can communicate with one another in a variety of ways, including emails, maintaining friendships and "blogging" to one another. Content shared between users through these website channels can range from media such as photos and music to information on news, products and services. Many companies encourage such sharing behavior between individuals; widespread viral sharing of content can increase traffic to their websites, potentially resulting in increased revenue through ads or product sales.

In addition to offering functionality for users to share content, certain companies offer monetary incentives to their users to invite friends to join their website. Thus, they utilize users' social networks to increase the reach of new user acquisition efforts. For example, Bloomspot[5], a daily deal company, gives \$10 to a user whose friend becomes a member and purchases from the website. Similarly, Fab[6], an online marketplace for designer items, has tiered incentives for inviting friends to join and/or purchase from their company. It gives \$30 credit when 10 friends have joined, \$30 more when 25 have joined, free shipping when 50 have joined, and \$25 when a friend makes their first \$25 purchase.

Besides priming users with incentives to gather new user acquisitions, several sites tie monetary incentives into sharing products with new and existing customers, specifically, with customers who *purchase* deals from the site. For example, users who share (or recommend) a Bloomspot deal with their friends get credit equal to the price of that deal towards future purchases if at least two of their friends buy the recommended deal. LivingSocial has a more conservative, but immediately applicable, post-purchase incentive. When a user persuades three (or more) friends to purchase a particular deal, the user is refunded their own purchase of the item. The latter deal sharing mechanism of LivingSocial is the focus of our paper, and we discuss it in more detail in the next section.

Such incentives can have measurable impacts on the behavior of users. For instance, a LivingSocial user with only two recommendation recipients in mind may choose to expand the number of shares, with the intent of convincing one additional person to purchase the deal and thus getting a free deal. Such behavior can be of benefit to both the user and LivingSocial – should the user manage to convince additional individuals to purchase an item they get the value of the deal. Additionally, if several more users purchase the deal LivingSocial can recoup the costs of the free deal, as well as potentially add new users and grow their potential pool of buyers for future deals.

These sharing incentives are given to all users, which makes them mass marketing incentives rather than direct marketing ones. This is in contrast with viral marketing strategies

*Work performed while interning at LivingSocial.
[1]www.livingsocial.com
[2]www.amazon.com
[3]www.facebook.com
[4]www.google.com
[5]http://www.bloomspot.com
[6]http://www.fab.com

VLDB workshop on Online Social Systems (WOSS) 2012

in which the goal is to identify a subset of individuals whom to target, so that they adopt a service or a product and in turn influence others to do the same, producing an effect which spreads in the social network as much as possible [2, 5, 8]. While some argue that there are users who inherently can influence others [1] and should be the target of direct marketing campaigns, others argue that everyone can be an influencer, thus mass marketing efforts are worthwhile. For example, in an experiment with music adoption, Salganik et al. show that people who are given the opportunity to adopt and rate a certain product early are the ones that determine the trends [10], presenting an argument against the existence of intrinsic personal characteristics of influencers.

The deal recommendations between users reveal a social network structure with the communication patterns between users. One important structural property of social networks is the degree of a user, the number of users someone is connected to, and there have been numerous studies which show that degree distributions often follow a power law distribution [3, 4, 7, 9, 11].

In this paper, we present a case study of the LivingSocial deal sharing network, showing how LivingSocial's unique sharing incentive impacts the distribution of shares by *shifting* the power-law distribution and increasing the volume of social shares. We distinguish between altruistic and incentivized shares, and look at the impact of different incentive amounts on the sharing behavior of users.

2. INCENTIVIZED DEAL SHARING AT LIVINGSOCIAL

LivingSocial is an e-commerce company whose main business revolves around connecting local businesses with customers online by running deals on its website. The majority of the deals are location-specific – LivingSocial negotiates with local businesses in different cities around the world, determining a deal price for customers that is considerably lower than the usual item price. The motivation for merchants is exposure; using LivingSocial to advertise their products brings in clients that might have never heard about or considered the merchant without such a deal, which (from the merchant's perspective) should hopefully result in expansion of the merchant's customer base. Customers are motivated to try a new business without having to pay full price and to treat the deal as a trial run.

The primary method LivingSocial uses for social influencing between users is the *Me+3* program[7]. To participate in this program, a user is first required to purchase a particular deal. Then, the user is given the option to share the deal with as many friends as desired, by posting on the user's Facebook Timeline, by emailing friends through the LivingSocial interface, or by posting a customized link through other means. Lastly, should three (or more) of the friends of the user purchase the shared deal, the user will get the deal for free rather than pay the purchase price. Thus, LivingSocial can sell more deals and potentially acquire more users, while the current user is incentivized to share in order to get their own deal for free.

The data that we use for our study is coming from customer-to-customer deal recommendations sent by email through the LivingSocial interface. Each recommendation (or share) data point consists of a unique triple $(sender_i, recipient_j, deal_k)$, and whether the recipient bought the deal or not. We define *outdegree* of a $(sender_i, deal_k)$ pair as the number of recipients with which $sender_i$ shared $deal_k$ This means that individual users can share different deals at different degrees, and each separate $(sender_i, deal_k)$ pair is included as a different outgoing degree sample. In contrast, *indegree* of a $(recipient_j, deal_k)$ pair is defined as the number of senders who shared $deal_k$ with $recipient_j$.

3. CHANGES IN SHARING BEHAVIOR

As users of LivingSocial send deals between one another (as well as to outside users), a *directed social network* is formed with users acting as *senders* and *receivers* of recommended deals. As the users are incentivized to share the deals they purchase with other users in the network, we can observe interesting and unusual structural changes in the network that distinguish the LivingSocial social network from other social networks. Additionally, we explore the sharing behaviors at different price points, as well as the predictive power of the shared recommendations in terms of recipient preferences.

3.1 Shifting the power-law

To start, we examine the degree distribution of *non-incentivized* individuals who share LivingSocial deals. LivingSocial users can share any deal posted on the website with their friends, through email, Twitter, and Facebook. If the user hasn't purchased the deal prior to sharing, then the share cannot be a part of the *Me+3* program, and thus it has has no monetary incentive. This implies that their sharing is inherently *altruistic*; the users are likely sharing because the receiving party might be interested in the deal. We call these *pre-purchase shares*.

Figure 1.a shows the outdegree distribution for a sample of 100,000 pre-purchase shares. The distribution was fitted to a power law using maximum likelihood [4], scaled to fit the occurrences. A power-law distribution for a random variable X is one where:

$$P(X = x) \propto x^{-\alpha}.$$

It is seen that the altruistic sharing distribution follows the power law quite well – such behavior mirrors the degree distributions found in many other social networks. The implication is that we have a large number of senders who share with a handful of recipients and far fewer senders that share with a large number of recipients. A notable exception of the LivingSocial share network compared to other networks is the *sharpness* or *steepness* of the distribution. The majority of networks typically have an power law exponent α where $2 \leq \alpha \leq 3$ [4]; in contrast, the exponent for LivingSocial share data is nearly 3.5, which is considerably larger. This implies that users are emailing deals to subsets of their friends rather than everyone they know. This is partly due to the time-sensitive nature of LivingSocial deals. All deals through LivingSocial are for a fixed amount of time, most for less than a week, which gives users considerably less opportunity to share links with friends. This contrasts with email and Facebook friendships, where the outgoing connections of users accrue over a long time period.

We also examine the patterns of users who share a deal after they purchase a deal, i.e., *post-purchase* shares. These users have been incentivized through the *Me+3* program

[7]http://help.livingsocial.com/articles/how-does-the-me-3-promotion-work

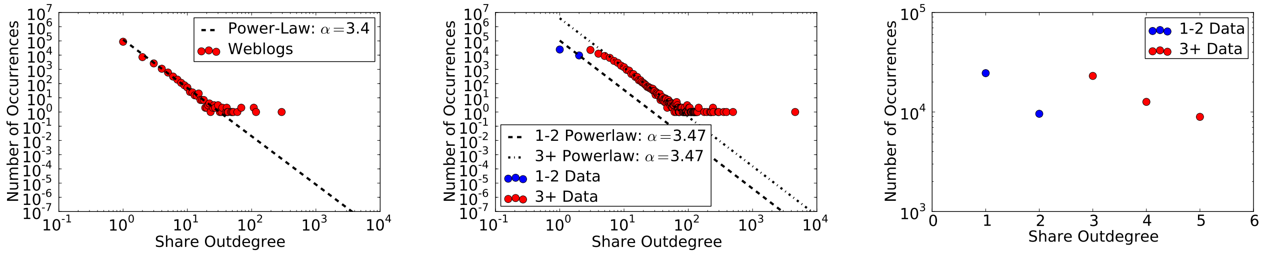

(a) (b) (c)

Figure 1: (a) Non-incentivized share distribution. (b) The degree distribution of shares through email. (c) Closer view on the distribution of degree 1 through 5.

– that is, they can receive their deal free from LivingSocial if they can get three of their friends to purchase the deal as well. The incentivized sharing distribution is shown in Figure 1.b, where from a sample of $100,000$ shares we can see a noticeable 'dip' behavior in the distribution at degrees 1, 2, and 3 (Figure 1.c highlights this behavior). Most interestingly, we see that the power law beginning around degree 3 has nearly the same exponent as found in the non-incentivized case, both having an exponent around 3.5. Thus, we can assume the distribution of incentivized users as having been *shifted* – or *moved over*, resulting in approximately the same distribution, the only difference being that the distribution has added approximately 2 shares to every point. The average outdegrees in the two share scenarios differ widely. While the non-incentivized shares have an average outdegree of 1.3, the incentivized shares have an average degree of 4.1, with the distribution tail contributing more heavily in the incentivized case. This means that the incentive does drive a significant amount of additional share volume.

An interesting question is what is the effect of the incentivized shares on whether or not the recipient makes a purchase. Next, we analyze the behavior of recipients and their share *indegree*.

3.2 Social pressure and adoption

To assess the effect of social pressure on adoption, we next examine the probability of a user purchasing a deal given that a certain number of people recommended the deal to them. This can be found in Figure 2.a, while the share indegree distribution can be found in Figure 2.b. This is performed on a sample of 100,000 $(recipient_j, deal_k)$ pairs, and indegree is the number of emails that were sent to $recipient_j$ for $deal_k$.

Notice the beginning part of Figure 2.a, where we can see that the higher the number of recommendations for a *particular* deal, the higher the probability of a user to purchase that deal, up to degree 4 which has the highest probability of 0.4. It appears that after this instance we have little improvement as the number of recommendations increases, but this can be attributed to the low number of sample points for these higher degrees (e.g., only two sample points at 9). As Figure 2.b shows, the number of shares with more than 4 recommendations drops significantly. However, the points between 1 and 4 indicate that users send deals according to whether their friend is likely to buy the product, and a high volume of incoming recommendations for a product is a good predictor for whether the user is likely to purchase the product.

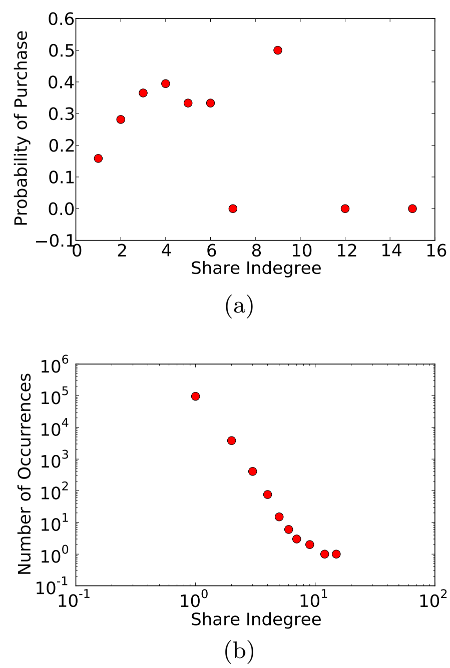

Figure 2: (a) Probability of purchase given an incoming recommendation degree for a *single* product (b) Number of instances with incoming degree for a single product

A complementary finding, based on Figure 2.b, is that unlike the incentivized share outdegree distribution, the indegree distribution does not have a dip at 2. This means that while the *outgoing* share distribution has been shifted for 3 and above, the effect on the incoming share distribution remains unclear. This can be explained by the fact that the $Me + 3$ incentive does not apply to the recipient, therefore it does not shift the behavior of recipients in the same manner as the behavior of senders.

As the recipient purchase probability is correlated with the number of recommendations that user receives, a key question is which *type* of recommendation contributed to the purchase of the product. More specifically, did the sender recommend the deal to the recipient because they believed

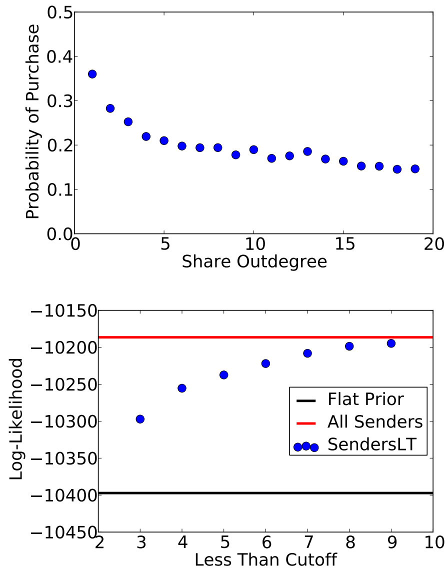

Figure 3: (a) Probability of purchase when sender shares deals with varying numbers of recipients. (b) Predictive ability using various amounts of recipients per recommendation.

the recipient had an interest in the deal, or was the sender simply trying to get their own free deal and sending the deal at random in hopes of a fortunate purchase? We explore these questions next.

3.3 Shares and recipient purchase preferences

People have intrinsic motivations to share information online with others, and sharing deals is no exception. While certain users may only share deals with other users due to the monetary incentive, many users share LivingSocial deals without such an incentive, as discussed in Section 3.1. This type of behavior has been observed also in peer-to-peer network sharing where some users serve files to others because of the group's common welfare, while others respond only to monetary incentives [6].

The distribution of shares per sender showed that a large proportion of shared deals are shared *only* with one or two other individuals (Figure 1.b). As these deal shares can never reach the incentive, it is safe to assume that the sender has little interest in the incentive and believes the recipient truly has some interest in the recommended deal. We categorize this type of sharing as *altruistic*. In contrast, any sender who shares with three or more recipients has the potential to be sharing solely for the incentive, or being an *incentivized* sender.

In order to learn whether low volume senders are more accurate with their assessments on which friends might purchase a deal, we first examine how often a recipient of a deal will purchase the deal, based on the number of recipients the sender shared with. Here, we only consider *recipients who are registered* with LivingSocial with the email address that the sender used.

In Figure 3.a, we can see that senders who only send to a handful of users are considerably more accurate when choosing which friends may be interested in a deal. Here, the *probability of purchase* is the proportion of users who purchased the recommended deal from a sender. This figure shows that if only a single person is sent the email, the probability of making that purchase is 36 percent, as compared to 5 or more friends, which tails down from 20 percent. However, as the probability of any user making any particular purchase at random is much lower than 20 percent, *there is still useful information provided by higher degree senders*. This implies that higher volume senders, while less accurate than when sending to a low volume number of recipients, are still being somewhat discriminative when determining which users might be interested in particular deals.

In addition to whether or not a deal recommendation results in a purchase given the volume of targeted recipients, we analyze whether senders know their friends well and whether shares reflect the recipient preferences. Every deal at LivingSocial is hand classified in a particular *category*: for example, a dentist visit may be placed in a category named Healthcare while a diner deal might be classified under Eateries. We pick a sample of $5,000$ users with whom at least one deal was shared in 2011 and who have bought at least one deal in 2012. Using the recommendations sent across 2011, we build simple preference models which learn the multinomial distribution over the shared deal categories for *each recipient*. A flat Dirichlet prior is then applied over the categories for each user to avoid category probabilities of 0 if a user has never been recommended a deal of that category. Then we check the log-likelihood of future purchases made in 2012, given the models.

The performance of the model which considers all incoming deal shares is shown as a red line in Figure 3.b. We can see that it performs better than a flat multinomial distribution which is shown as a black line.

An interesting question is whether the deals shared with a high number of recipients are too noisy to extract any meaningful recipient preference information, and whether they can be excluded from the model. In other words, do the high volume deal shares contain information about the recipients' categorical preferences that is not already present and available through the low volume deal shares? If the high volume deal shares are mostly noise, then at best they add little value to models which build categorical preferences of recipients, while at worst they could potentially *decrease* the models' performance by washing out the relevant information found from low volume deal shares.

To test whether these high volume deal shares interfere with our categorical model's ability to predict future deal purchases, we build models which include varying levels of share volumes (Figure 3.b). The x-axis indicates the model's "cutoff" point, that is, the maximum number of recipients per share which will be included in the recipient's categorical model. For example, the point corresponding to the '3' on the x-axis indicates a model where only deals shared with 1, 2, or 3 recipients will be incorporated in the recipients' models, omitting all deal shares sent to 4 or more recipients from any recipient's model. We then gradually increase the volume of deal shares to be incorporated in the recipients' models, up to 9. We first observe that each of these points outperforms a 'flat' distribution (the black line), meaning they never do worse than random. Furthermore, we can see that as the additional deal shares are included, the predictive power of the recipients' model increases, indicating better performance when the higher volume deal shares are

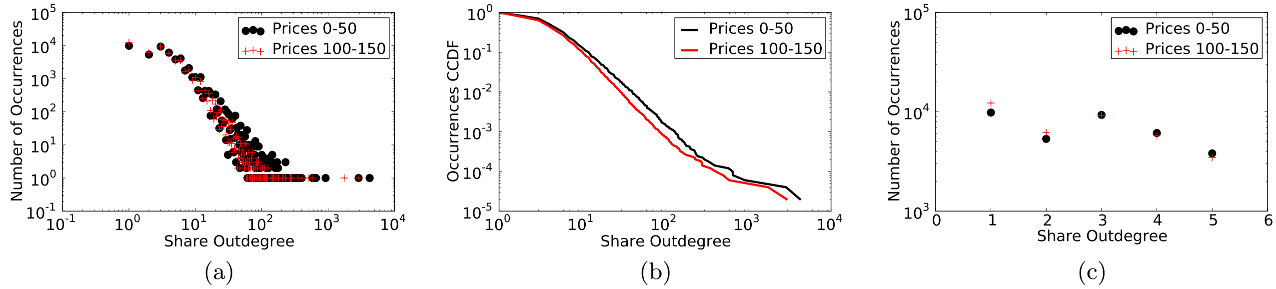

Figure 4: (a) The distribution of the number of recipients per share for varying prices. (b) The complementary cumulative degree function for the number of recipients per share for varying prices (c) Close view of the volume of senders who shared with 1 through 5 recipients.

incorporated into the recipients' categorical models. Finally, we compare against the model which incorporates all deal shares sent, regardless of volume (the red line). As this line outperforms all the other various points, it shows that while a limited amount of information is available from these high volume shares, the noisy data does not interfere with the lower volume shares. Furthermore, it does add some additional performance, indicating slight amount of information can be gleaned from such high volume shares.

This shows that the flat prior alone (black) performs worse than any of the recipient informations, indicating that senders are using some amount of discrimination when determining which users might be interested in a deal. When comparing the likelihoods between models as the number of recipients allowed per share is increased, we find that the amount of information found in the high volume senders stagnates, indicating less discrimination by the sender as to whether the receiver has interest in the present deal. However, the information is not completely random, as the predictions continue to slowly increase in accuracy. High degree senders are likely highly motivated by the incentive, and they also target individuals they believe have a chance at purchasing the deal.

Next, we discuss the impact the deal price has on the sharing behavior, in particular whether the price can affect on how desirable the incentive is.

3.4 Impact of the incentive amount

So far, we discussed how the incentive and volume of sharing affect the behavior of both recipients and senders. An interesting question to explore next is the effect of the incentive amount. One hypothesis is that the higher the incentive, the more additional shares it drives. However, this turned out not to be the case.

To understand the effect price has on shares, we examine the shift in the share outdegree distribution for two distinct price ranges, $(\$0, \$50]$ and $[\$100, \$150]$. We sample 100,000 $(sender_i, deal_k)$ pairs from the corpus of data (50,000 per price range). Figure 4.a shows the share outdegree distribution, i.e., the number of shares with a certain number of recipients. Part (a) shows the distribution of send volume and part (b) shows the complementary cumulative distribution function (CCDF) of the distribution of send volume - at every point along the x-axis we see the number of points greater than this one in the distribution. Contrary to our initial intuition, when the price is higher, the sender typically shares with *fewer* people. As $100 - 150$ remains below $0 - 50$ in 4.b, this shows that fewer senders are flooding the inboxes of their friends when the items are in a higher price range, but are more comfortable asking around to see if a random friend may be interested in the cheaper deals.

This is further magnified in Figure 4.c. Here, we have focused on just the senders who send to 5 or fewer recipients. We note a pivot point at 3 shares - for low priced items the senders more frequently send to 4 or 5 recipients. For high priced items the senders are more focused on the individuals they believe are interested in the items, sending more frequently to just one or two individuals. The volume of additional incentivized sharing appears to be lower at the higher price point.

4. CONCLUSIONS

We presented an initial study of one type of monetary incentive strategy for sharing in social networks. Due to the unique nature of the incentive, namely *"if three of your friends buy a deal, you get the deal for free"*, users have interest in sharing with at least 3 friends and we are able to see a clear shift in the behavior due to the incentive. The share outdegree distribution follows a power law for users who share with 3 or more friends, while users who only share with 1 or 2 friends do so much less frequently than expected. In addition, we showed the number of shares greatly increases in the incentivized case, in contrast to data compiled where users were not incentivized.

While our study suggests that monetary incentives can potentially change the behavior of social sharing, many questions remain unanswered. For example, we expect that one fruitful area of future research is studying models which optimize the incentive price points given the potential benefit from sharing. Another interesting direction is in attribution modeling to determine which adoption behavior can be attributed to the monetary incentive and which one to altruistic motivation. This seems particularly hard, given that the goal of both types of incentives is to make a good recommendation and drive adoption. Comparisons between different types of monetary incentive strategies and positioning them according to the scenarios in which they work best is another interesting research direction.

5. REFERENCES

[1] N. Agarwal, H. Liu, L. Tang, and P. S. Yu. Identifying the influential bloggers in a community. In *WSDM*,

2008.
[2] A. Anagnostopoulos, G. Brova, and E. Terzi. Peer and authority pressure in information-propagation models. In *PKDD*, 2011.
[3] A. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, October 1999.
[4] A. Clauset, C. Shalizi, and M. Newman. Power-law distributions in empirical data. *SIAM Review*, 51(4):661–703, 2009.
[5] P. Domingos and M. Richardson. Mining the network value of customers. In *KDD*, 2001.
[6] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge. Incentives for sharing in peer-to-peer networks. In L. Fiege, G. M§hl, and U. Wilhelm, editors, *Electronic Commerce*, volume 2232 of *LNCS*, pages 75–87. Springer, 2001.
[7] S. Lattanzi and D. Sivakumar. Affiliation networks. June 2009.
[8] J. Leskovec, L. A. Adamic, and B. A. Huberman. The dynamics of viral marketing. *ACM Trans. Web*, 1(1), May 2007.
[9] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. pages 177–187, 2005.
[10] M. Salganik, P. Dodds, and D. Watts. Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science*, 311(5762):854–856, 2006.
[11] E. Zheleva, H. Sharara, and L. Getoor. Co-evolution of social and affiliation networks. 2009.