Intrinsic versus Image-Related Motivations in Social Media: Why do People Contribute Content to Twitter?

Olivier Toubia
Professor of Business
Columbia Business School
522 Uris Hall, 3022 Broadway, New York, NY 10027-6902
ot2107@columbia.edu

Andrew T. Stephen
Assistant Professor of Business Administration & Katz Fellow in Marketing
University of Pittsburgh, Joseph M. Katz Graduate School of Business
318 Mervis Hall, Pittsburgh, PA 15260
astephen@katz.pitt.edu

Intrinsic versus Image-Related Motivations in Social Media: Why do People
Contribute Content to Twitter?

We empirically study the motivations of users to contribute content to social media in the context
of the popular microblogging site Twitter. We focus on non-commercial users who do not
benefit financially from their contributions. Previous literature suggests two main possible
sources of motivation to post content for these users: intrinsic motivation and image-related
motivation. We leverage the fact that these two types of motivation give rise to different
predictions as to whether users should increase their contributions when their number of
followers (audience size) increases. In order to address the issue that the number of followers is
endogenous, we conducted a field experiment in which we exogenously increased this number
for a set of users (treatment group) over a period of time and observed the impact on their
posting activities compared to a control group. We find evidence for both types of motivation.
We estimate each treated user's utility function using a dynamic discrete choice model.

Keywords: Social Media, Field Experiments, Dynamic Discrete Choice Models.

1. Introduction

In recent years, social media has emerged as major channel for broadcasting information. For instance, by late 2011 there were over 173 million public blogs,[1] and 250 million messages ("tweets") were sent each day through the popular microblogging platform Twitter.[2] While it is clear that millions of people around the world are contributing content to social media platforms on a regular basis, it is less clear what motivates them to do so. Although some contributors are able to derive advertising revenue from their content (using for example platforms such as Google's AdSense; cf. Sun and Zhu 2011), social media platforms predominantly rely on the benevolent contributions of individuals as "content providers." While publishers' incentives in traditional media are well understood and are typically a function of the number of "eyeballs" reached by their content, motivations to benevolently contribute content in social media are not well understood. Understanding consumers' motivations to post content to social media is not only important theoretically, it is also relevant managerially. Indeed, one may argue that a firm cannot devise an efficient social media strategy and optimize the ways it engages with consumers through social media if it does not understand what motivates consumers to be active on such platforms in the first place. Moreover, for the platforms themselves, understanding what motivates their users to contribute is important since the viability of these platforms as businesses depends not only on how many users they have but also on how active their users are as content contributors.

---

[1] Nielsen Blogpulse, October 7, 2011. http://www.blogpulse.com.
[2] Mediabistro, October 18, 2011. http://www.mediabistro.com/alltwitter/costolo-future-of-twitter_b14936.

This paper considers what motivates individuals to contribute content to social media in the absence of explicit economic incentives. In this context, the incentives literature suggests two relevant types of motivations beyond financial or other extrinsic economic incentives: intrinsic motivation and image-related motivation. *Intrinsic motivation* assumes that users receive direct utility from posting content, and is defined as "the doing of an activity for its inherent satisfactions rather than for some separable consequence" (Ryan and Deci 2000). *Image-related motivation*, on the other hand, assumes users are motivated by the perceptions of others (see Fehr and Falk 2002 for a review of the psychological foundations of incentives).[3] Image-related motivation is also related to status seeking or prestige motivation (e.g., Glazer and Konrad 1996; Harbraugh 1998a, 1998b; Fershtman and Gandal 2007; Lampel and Bhalla 2007).

Intrinsic and image-related motivations have been studied quite extensively in the domain of prosocial behavior (see for example Glazer and Konrad 1996; Harbaugh 1998a, 1998b; Bénabou and Tirole 2006; Ariely, Bracha and Meier 2009). In a domain closer to social media, Lerner and Tirole (2002, 2005) contrast the intrinsic pleasure open-source developers derive from working on "cool" projects with the (image-related) desire for peer recognition (see for example Lakhani and Wolf 2005 for empirical evidence for intrinsic motivation in open-source development and Fershtman and Gandal 2007 for empirical evidence for image-related motivation). Evidence for intrinsic and image-related motivations has also been found in the context of electronic knowledge repositories (see for example Kankanhalli, Tan and Wei 2005; Wasko and Faraj 2005; Lampel and Bhalla 2007; Nov 2007).

In the context of social media, to the best of our knowledge, research to date on individuals' motivations has been survey-based. The existing evidence suggests that intrinsic and image-

---

[3] Fehr and Falk also discuss reciprocity as a psychological source of motivation, which depends on whether an agent perceives the action of another agent as hostile versus kind. This is less relevant in our context.

related motivations are indeed critical drivers of content posting in social media. Bughin (2007) surveys users of an online video-sharing site and finds that their primary motivations to upload videos are image-related ("I seek fame") and intrinsic ("It is fun"). Hennig-Thurau et al. (2004) survey the motivations of web-based opinion-platform users to post content to these platforms. Besides some motivations specific to their particular context, the motivations found by these authors tend to be either intrinsic or image-related (e.g., "It is fun to communicate this way with other people in the community," "My contributions show others that I am a clever consumer").

Therefore, based on the extant literature it appears that intrinsic and image-related motivations are both plausible and realistic reasons for people to contribute content in social media. In this paper we empirically compare these two types of motivation, focusing specifically on the context of the popular microblogging platform Twitter.[4]

Two recent related papers are Kumar (2009) and Shriver, Nair, and Hofstetter (2011). Kumar (2009) studies consumers' purchase of ring-back tones for their mobile phones (a ring-back tone is not consumed by the user purchasing it but rather by those who call that user), and estimates the utility consumers derive from having a high status (i.e., more recently updated tones), from consuming the tones purchased by their peers, and from expressing themselves through the tones they purchase. Shriver, Nair, and Hofstetter (2011) study the causal relations between content generation and the number of social ties in an online windsurfing community, using a set of instrumental variables. However, neither paper studies specifically the motivations to post content in social media. Kumar (2009) uses a context slightly different from social media, and Shriver et al. (2011) do not study motivations directly. Our approach focuses on one of the

---

[4] Whereas a blog is a website or part of a website that displays entries or elements of content (text, graphics, video, etc.) usually posted by an individual, a microblog is a type of blog that allows users to exchange smaller elements of content (e.g., short sentences, individual images, links).

most popular social media platforms, and directly studies individuals' motivations to contribute content. Moreover, a unique aspect of the present paper is the use of a field experiment as a source of identification.

The rest of the paper is organized as follows. In Section 2, we provide an overview of the Twitter platform, discuss how the concepts or intrinsic and image-related motivation are operationalized in this context, and describe our empirical strategy. We describe our data in Section 3, provide some model-free analysis in Section 4, and analyze the data using a dynamic discrete choice model in Section 5. We conclude in Section 6.

2. Background

2.1. Twitter

Twitter allows users to share "tweets," which are short text messages of up to 140 characters. Users have the ability to "follow" other users, i.e., subscribe to their tweets. This creates a directed social network (unlike other social networks such as Facebook or Linkedin which are undirected networks, user A following user B on Twitter does not automatically imply that B follows A). A user's home page (as seen by that user) contains a "timeline" that captures all the tweets posted by the users this user follows (ordered in a reverse chronological order), a text box labeled "what's happening" that allows the user to post a tweet, and a reminder of the number of users following the user and the number of users followed by the user. Twitter users may be split into non-commercial and commercial users. Commercial users may be classified into celebrities, media organizations, non-media organizations, and brands (Wu et al. 2010). In this paper we focus on non-commercial users for whom there exists no apparent financial incentives to contribute content.

Besides typing a message in the "What's happening" window of their home page, users may also post tweets as "replies" or "retweets." A "reply" to a previous tweet is a text message of up 140 characters that will be seen by users who follow both the user who posted the initial tweet and the user replying to that tweet. A "retweet" forwards a previous tweet to a user's followers. Other features of Twitter include the ability for a user to "unfollow" another user (i.e., stop following a user that he or she had been following), to "block" another user (i.e., prevent that other user from following him or her), the ability to create "lists" of users, and the ability for users to "protect" their accounts. Accounts that are not protected are called "public" (this is the default setting) and may be followed and accessed by any user (i.e., the account's number of users followed, number of followers, cumulative number of tweets, text of recent tweets are public information). If a user elects to protect his or her account, the text of his or her tweets may only be accessed by that user's followers. However, the number of users followed, number of followers, and cumulative number of tweets of protected accounts are still public information. In addition, requests by other users to follow a user with a protected account need to be approved by that user. According to Cha et al. (2010), approximately 8% of Twitter accounts are protected.

One feature of Twitter that is critical to our analysis is that posting content is a way to attract new followers. This claim is supported by our data (i.e., the state transition probabilities reported in Section 5), and is consistent with Shriver et al. (2011) who find a positive causal effect of content generation on the number of social ties a user has in an online windsurfing community. Note that unlike other directed social networks, reciprocity (i.e., A follows B and B follows A) on Twitter is only moderate. Kwak et al. (2010) report that of all user pairs on Twitter with at least one link between them, only 22.1% have a reciprocal relationship (i.e., each user in the pair

follows the other user).[5] In other words, a user's number of followers is not simply a by-product of his or her following activities, and posting content in the form of tweets is one way for users to attract new followers.

Twitter usage has been steadily growing. The number of unique US visitors to twitter.com in September 2011 was estimated at 35 million, up from 28 million in September 2010.[6] The average number of tweets per day grew from an average of 5,000 in 2007 to 300,000 in 2008 to 2.5 million in 2009, and in October 2011 this number reached 250 million tweets per day.[7] Even if each tweet takes only a few seconds to write, with 250 million tweets written per day, the equivalent of multiple decades of one person's life are spent each day posting content on Twitter (250 million tweets times 5 seconds per tweet = 39.6 years).

Given the scale and relevance of Twitter in society, it is not surprising that academic research on Twitter has started to emerge, mostly from Computer Science and Information Systems. Extant research has focused primarily on studying the structure and the nature of the Twitter social network, and on issues related to influence and information diffusion on this network (see for example Cha et al. 2010; Kwak et al. 2010; Weng, Lim and Jiang 2010; Bakshy et al. 2011; Romero et al. 2011; Wu et al. 2011). However, to the best of our knowledge academic research on Twitter in marketing and other social sciences has been extremely limited so far. Exceptions include Ghose, Goldfarb, and Han (2011) who compare user search costs in online versus mobile platforms using data from a microblogging site comparable to Twitter; and

---

[5] In comparison, this proportion has been found to be around 70% in the Flickr social network (Kumar, Novak and Tomkins 2006; Cha, Mislove and Gummadi 2009) and 84% in the Yahoo! 360 social network (Kumar, Novak and Tomkins 2006).

[6] Source: http://siteanalytics.compete.com/twitter.com/, accessed November 15, 2011.

[7] Mediabistro, October 18, 2011. http://www.mediabistro.com/alltwitter/costolo-future-of-twitter_b14936.

Stephen, Dover, Goldenberg, and Muchnik (2011) who study how user activity on Twitter affects the extent to which URLs posted by users in tweets spread through the Twitter network.

2.2. Intrinsic versus image-related motivations on Twitter

2.2.1. Intrinsic Motivation

Intrinsic motivation to post content in Twitter should translate into a utility from posting that is monotonically non-decreasing in the user's number of followers (i.e., the size of the user's audience). Such motivation is consistent with Twitter's initial positioning as "a real-time information network powered by people all around the world that lets you share and discover what's happening now" (Twitter home page, accessed 02/2010). Twitter further states that "Twitter asks 'what's happening' and makes the answer spread across the globe to millions." The public nature of Twitter and the claims that the information spreads "to millions across the globe" are consistent with the intrinsic value from posting being monotonically non-decreasing in a user's number of followers. Put simply, a user who is intrinsically motivated to broadcast content derives more utility from doing so as the size of his or her audience increases. This is similar to the case where content publishers receive explicit financial incentives, which are typically monotonically non-decreasing in the size of the publisher's audience.

2.2.2. Image-related motivation

There is some evidence suggesting that image-related motivation on Twitter is related to a user's number of followers. While any user is able to contribute as much content and follow as many users as he or she wants, followers need to be "earned" and a user's number of followers is an informative signal. The number of followers has been used as a measure of influence by academics (Cha et al. 2010; Kwak et al. 2010) and is often associated with popularity by the

general public (e.g., www.twitaholic.com, wefollow.com, championist.com).[8] There have been

several reports in the press of Twitter users attaching a lot of importance to their number of

followers. According to Poletti (2009), Twitter has become an avenue for self-promotion, and

one's number of followers is becoming "the new barometer of how we gauge our self worth".[9]

Leonhardt (2011) claims that the number of followers on Twitter is "just how people keep score

on the site and compare themselves to friends and colleagues."[10] Teitell (2011) reports on the

social pressures to achieve high numbers of followers on Twitter and high scores on sites such as

www.klout.com and www.peerindex.net that rate all Twitter users based on their number of

followers (as well as other metrics, using proprietary scoring rules).[11] The importance for many

Twitter users of having a large number of followers is further revealed by the plethora of

websites that offer advice on how to increase that number (a partial list may be obtained by

searching for "increase Twitter followers" on Google). Therefore it seems appropriate to

measure the stature or prestige of a Twitter user by his or her number of followers.

It is reasonable to model utility from stature as a non-decreasing concave function of the

number of followers (Baumeister and Leary 1995). As mentioned above, posting content on

Twitter is a way to attract new followers. Therefore, posting content on Twitter is a way for

image-conscious users to increase their stature on Twitter, as measured by their number of

followers.

---

[8] See also Beck, Howard. 2009. "New Way to Gauge Popularity." *The New York Times*, October 21.
[9] Poletti, Therese. 2009. "What if your friends won't follow you on Twitter?" *MarketWatch*, November 12.
[10] Leonhardt, David. 2011. "A Better Way to Measure Twitter Influence." *The New York Times*, March 27.
[11] Teitell, Beth. 2011. "Ascent of the Social-Media Climbers." *The Boston Globe*, February 18.

2.3. Empirical strategy

Twitter offers a unique opportunity to study and contrast intrinsic and image-related motivations in social media for at least two reasons. First, by focusing on non-commercial users we are able to study contributions to social media in a context in which financial or other extrinsic incentives are minimal, if present at all. This allows us to focus on contrasting intrinsic and image-related motivations. Second and most importantly, Twitter provides a context for empirically comparing intrinsic versus image-related motivations, because they give rise to different predictions as to how users should react to an increase in their number of followers. If users contribute content to Twitter because of the intrinsic value they derive from broadcasting information to their followers (with the utility derived from posting being monotonically non-decreasing in a user's number of followers), we should expect users to increase their posting activities as they receive additional followers. On the other hand, image-related motivation gives rise to the opposite prediction: posting activities should be reduced as the number of followers increases. As mentioned above, posting content on Twitter is a way to attract new followers. If the image-related utility derived by a user is concave in that user's number of followers, i.e., the marginal utility from an additional follower is decreasing with the number of followers, then the motivation to post content in order to attract new followers should decrease with that user's current number of followers. Therefore, image-related motivation for contributing content on Twitter suggests that users should post less as they have more followers.[12] Interestingly, image-related motivation also leads to the prediction that users should post less content as the structure

---

[12] One may think that users who are motivated by image would feel compelled to actually post more as they amass more followers, in order to maintain their number of followers. However, our data suggest that the expected change in the number of followers when no posting occurs is not negative (see state transition probabilities reported in Section 5.1.4), so this scenario seems unlikely. More generally, our model in Section 5 will enable us to take such scenarios into account, by explicitly capturing and quantifying the impact of posting on a user's future number of followers.

of the network becomes stable (i.e., a non-evolving static structure of connections is achieved) and as posting activities become less likely to lead to additional followers, which raises questions on the longer-term sustainability of the Twitter platform. This issue will be addressed using counterfactual analyses in Section 5.3.4.

We note that there are conditions related to the way posting affects one's future number of followers under which intrinsic and image-related motivations could in fact have the opposite effects to those just described. These conditions have low face validity and are discussed at the beginning of Section 5. Notwithstanding, our model in Section 5 enables us to quantify intrinsic versus image-related utility even under such conditions. In particular, the identification of our model does not rely on the assumption that intrinsic (image-related) motivation always gives rise to an increase (decrease) in posting activity following an increase in the number of followers.

In order to empirically compare intrinsic versus image-related motivations to post content on Twitter, we collected panel data on a random set of 2,493 non-commercial Twitter users, which include a daily count of their number of followers and posting activities. In order to address the issue that the number of followers is endogenous, we conducted a field experiment in which we randomly selected 100 of these users (the treatment group) and exogenously increased each of their numbers of followers by 100 over time. We observed the impact of this intervention on the posting activities of the treated users compared to users in a control group.

To preview our results, we find evidence of both types of motivation being present. The predominant source differs across users and depends on how many followers they initially had prior to our experimental treatment. Treated users with a moderately small initial number of followers became more active compared to a control group after an increase in their number of followers, which is consistent with intrinsic motivation. In contrast, treated users with a

moderately large initial number of followers became less active compared to a control group after an increase in their number of followers, which is instead consistent with image-related motivation. These results are based on a model-free analysis of the data, presented in Section 4. In Section 5 we estimate a dynamic discrete choice model that allows us to quantify intrinsic versus image-related utility for each treated user while allowing for heterogeneity across users' preferences. This model also enables us to perform some counterfactual analysis of how posting activity would be affected if the network's structure became stable.

3. Data

Our data were collected directly from Twitter using Twitter's application programming interface (API; see dev.twitter.com). We selected a random set of 2,493 non-commercial Twitter users from an initial database of approximately 3 million user accounts. The users in our dataset are a mix of public and protected accounts. We ensured that our users were non-commercial by checking account names, and checking against lists and classifications on sites such as wefollow.com and twitterholic.com. We collected data daily on the following variables for each user in our sample: (i) the number of followers, (ii) the number of users followed, and (iii) the cumulative number of tweets posted by that user since that user profile was created (including retweets and replies). Neither the text of the tweets posted by these users nor the structure of the social network to which these users belong are available to us.

3.1. Initial calibration dataset

We first collected data daily for these 2,493 users for 52 days, between May 8, 2009 and June 28, 2009. This initial dataset allowed us to identify active users among the set. We classified a user as "active" if he or she increased his or her cumulative number of tweets or number of users

followed at least once during this screening observation window. Out of all users, 1,355 were classified as active.

3.2. Field experiment

We collected daily data again from the same set of 2,493 users for 160 days, between September 14, 2009 and February 20, 2010 (our main observation window). In order to introduce exogenous variations in the number of followers, we selected 100 users randomly from the set of 1,355 active users, and exogenously and gradually increased their number of followers by 100 over a 50-day period (days 57 to 106). We refer to these 100 active users as the "treated" users and the other 1,255 active users as the "control" users.

   In order to exogenously increase the number of followers of the treated users we created 100 synthetic Twitter users (50 males, 50 females). With the help of two undergraduate research assistants who were avid Twitter users, we attempted to make these users as realistic as possible. The names of these synthetic users were generated using the name generator available at www.fakenamegenerator.com. Before creating any link to the treated users, profile pictures were uploaded to the synthetic users' profiles and each synthetic user followed an average of five other synthetic users as well as some celebrities and media organizations (as is typical for many Twitter users). The synthetic users also posted tweets on a regular basis. In order to increase the credibility of the exogenous links to the treated users from the synthetic users, we started by creating one link (i.e., adding one synthetic follower) per day to each treated user.[13] After doing so each day for four days, we increased the daily number of exogenous links per treated user to two per day, and so on until the rate increased to five per day for four days, after which it was

---

[13] If the treated user had a protected account, the synthetic user sent a request to follow that user.

decreased to four per day for four days, and so on. By day 106 each synthetic user had created

one link to each treated user. Figure 1 shows the number of exogenous links created to each

treated user on each of the 160 days in our main observation window.[14] Note that our

experimental procedure respects Twitter's Terms of Service (available at twitter.com/tos).

[Insert Figure 1 Here]

4. Model-free analysis

4.1. Descriptive Statistics

We first report some key descriptive statistics. Figure 2 shows histograms and log-log plots of

the distribution of the number followers on the first day of the main observation window, for all

2,493 users and for all 1,355 users who were active during the screening period (i.e., the set of

users from which our treated users were drawn). The distribution of the number of followers is

close to a truncated power-law (the log-log plots are close to linear), which is typical of social

networks. Networks with power-law degree distributions (the degree of a node is its number of

connections), often referred to as "scale-free" networks, have been found to describe connections

as diverse as the collaboration between film actors, the citation patterns of scientific papers, links

between pages on the World Wide Web, and links between social commerce stores (e.g., de

Solla Price 1965; Barabási and Albert 1999; Stephen and Toubia 2009). Figure 3 shows the

distribution across all 1,355 active users of the average daily posting rate during the main

observation window. The average daily posting rate is measured as the total number of posts

during the window divided by the number of days. We see that the distribution is heavily

skewed, with many users posting very little and few users posting heavily. Figure 4 shows the

---

[14] The gaps in Figure 1 are due to our RAs needing to take breaks from this labor-intensive activity.

evolution of the median number of followers over time for treated versus control users. Note that these groups' median numbers of followers were very comparable before the start of the intervention (days 1-57). Also note that the actual increase in number of followers for treated users is not necessarily larger than or equal to 100 for all treated users. This is because some treated users had protected accounts and did not accept all synthetic users' follow requests, and because users have the ability to block any of their followers. Nevertheless, by the end of the main observation window, the median number of followers for treated users was greater than the median number for non-treated users by a margin of 85.00.

[Insert Figures 2-4 Here]

4.2. Impact of increase in number of followers on posting activity

We now consider the posting behaviors of treated versus control users. We compute the average daily posting rate for treated and control users before the intervention (days 1 to 56) and after the intervention (days 107 to 160). As seen in Figure 3, the distribution of average daily posting rates is heavily skewed. Therefore, instead of comparing summary statistics across groups of users, we compare within-user the average daily posting rate after the intervention to the average daily posting rate before the intervention. We find that the proportion of users for whom the average daily posting rate after the intervention was higher than the rate before the intervention is somewhat greater among treated users than it is among the control users. Specifically, 40.82% of treated users had a greater posting rate after the intervention than before, compared to 34.19% of control users. However, the difference between these two proportions is not statistically significant ($Z = 1.32, p = 0.19$).

A richer pattern of results emerges if we compare changes in average daily posting rates conditional on a user's initial number of followers (on day 1 of the main observation window). Indeed, if intrinsic and image-related utility from posting vary differently as a function of one's number of followers, we should expect the relative prominence of these two sources of motivation to change as well. In other words, while the behavior of a user may be driven primarily by one source of motivation when that user has few followers, the other source of motivation may be more prevalent as his or her number of followers increases. In addition, there is likely heterogeneity in users' motivations to post content on Twitter, and to the extent that this heterogeneity is reflected in the number of followers, different groups of users with different numbers of followers may have different motivations.

Figure 5 (respectively Figure 6) plots the probability that a user increased (decreased) his or her posting rate after the intervention versus before, as a function of the log of that user's initial number of followers measured on day 1 of the main observation window, for treated and non-treated users. These figures were obtained by smoothing the raw data using a Gaussian kernel function (bandwidth = 1). Treated users with lower initial numbers of followers tended to increase their posting rates relative to control users. However, treated users with higher initial numbers of followers tended to decrease their posting rates relative to control users.

[Insert Figures 5-6 Here]

In order to statistically compare the impact of the intervention on posting behavior as a function of the initial number of followers, we split our treated users into quintiles based on their number of followers on day 1 of the main observation window. The five quintiles are described in Table 1. Table 2 reports the proportion of users with increased average daily posting rates and with decreased average daily posting rates (after versus before the intervention) in each quintile.

Treated users in the 2nd quintile were significantly more likely to increase their posting rates compared to users in the control group (z=2.42, p<0.02), and marginally significantly less likely to decrease their posting rates compared to users in the control group (z=-1.85, p=0.06). Users in the 4th quintile, however, show the opposite result: treated users in that group were significantly less likely to increase their posting rates compared to users in the control group (z = -2.18, p < 0.03), and significantly more likely to decrease their posting rates compared to users in the control group (z=1.94, p=0.05). The differences in the other quintiles are not statistically significant. Therefore, our results suggests that exogenously increasing the number of followers made some users post more (users in the 2nd quintile), made some users post less (users in the 4th quintile), and had little effect on the others.

[Insert Tables 1-2 Here]

The fact that the intervention had no effect on the 1st and 5th quintiles is not very surprising. The 1st quintile is composed of users with very few followers who are marginally active and may hardly visit the Twitter platform. The 5th quintile is composed of users with over one thousand followers on average, for whom the addition of 100 followers over 50 days may have gone largely undetected. The results in the 2nd quintile are consistent with intrinsic motivation being predominant. As discussed above, intrinsically motivated users contribute content for the value they derive from broadcasting information to their followers, and this value should be monotonically non-decreasing in the number of followers (i.e., the size of their audience). If utility derived from posting is monotonically non-decreasing in the number of followers, posting activity should increase following an increase in the number of followers. The results in the 4th quintile, however, are consistent with image-related motivations being prevalent. Image-related motivations lead users to post content in order to attract new followers and increase their stature

on the platform (recall that posting content on Twitter is a way to attract new followers). As argued above, if a user's utility comes from having a large number of followers and if the marginal utility from additional followers is decreasing, then we should expect posting activity to decrease following an increase in the number of followers. The results in the 3rd quintile are consistent with the effects of the two sources of motivation canceling each other on average for users with an initial number of followers within the corresponding range.

Note that one simple alternative explanation for the results in the 4th quintile could be that some users feel comfortable posting on Twitter when their followers are limited to immediate relations and when some level of intimacy is preserved, but become less comfortable as their posts become more public. However users in the 4th quintile had 125.550 followers on average on day 1 (and 148.151 on day 57), i.e., these users were not posting only to a small intimate group of followers. Therefore it seems unlikely that the decrease in their posting activity was due to a switch from a private, intimate context to a public one.

5. Dynamic Discrete Choice Model

The model-free analysis provides support for both intrinsic and image-related motivations being present, with the predominant motivation being different across users depending on their initial numbers of followers. We now formulate and estimate a dynamic discrete choice model of posting behavior on Twitter. This complements and extends the model-free analysis in at least three important ways.

First, it enables us to quantify intrinsic versus image-related utility for each treated user. Our model allows for heterogeneity across users in the shape of their utility function and in the relative weight on intrinsic versus image-related utility.

Second, while we have argued that intrinsic motivation should lead to an increase in posting activities as the number of followers is increased and that image-related motivation should lead to a decrease in posting activities, there exist conditions under which these theoretical predictions would be reversed. These conditions, which do not have high face validity, are related to the way posting affects one's future number of followers. For example, if attracting new followers became much harder as users have more followers, it might be possible that although a user motivated by image receives less marginal value from each additional follower as his or her number of followers is increased, he or she has to post much more heavily in order to attract new followers, leading to a net increase in posting activities.[15] Conversely, it is possible that users driven by intrinsic motivation would post less after having more followers, if for example their early posting activities were targeted toward building an audience to which they would broadcast later and if the only way to build such an audience was to post heavily early on. By capturing the effect of posting on the number of followers, through the state transition probabilities, our model allows quantifying intrinsic versus image-related utility irrespective of whether these conditions are satisfied. In other words, the identification of our model does not rely on the assumption that intrinsic (image-related) motivation always gives rise to an increase (decrease) in posting activity following an increase in the number of followers.

Third, our model enables us to perform counterfactual analysis of how posting activity would be affected under various conditions. In light of the recent public debate on the sustainability and the future growth of social media platforms such as Twitter,[16] we use a counterfactual experiment to study how posting activities would be affected if the network's

---

[15] This would go against the popular notion of preferential attachment, which is often believed to govern the evolution of "scale-free" social networks like this one (Barabási and Albert 1999).

[16] See for example Hagan, Joe. 2011. "Twitter Science." *New York Magazine*, October 02.

structure became stable and if posting content did not lead to additional followers. Such counterfactual analysis would not be possible based solely on the model-free analysis presented in the previous section.

We first describe our model, then the estimation procedure, and finally the results.

5.1 Model

We index users by $i=1,\ldots I$. We index time (day) by $t=1,\ldots\infty$. In each time period, each user chooses one of four possible actions: (i) follow at least one new user and post content, (ii) follow at least one new user and post no content, (iii) follow no new users and post content, and (iv) follow no new users and post no content. We model each user's decision in each time period as a multinomial choice over these four possible actions. The utility derived by a user from each action in each time period is a function of the number of followers this user has at that time. Because the user's action will have an impact on his or her future number of followers, it will also have an impact on the utility offered by each possible action in the future. In order to capture these dynamic effects, we build a dynamic discrete choice model in which users take into account how their current actions will impact their future decisions. Like with every dynamic discrete choice model, developing and estimating our model involves (i) defining the states, (ii) defining the actions, (iii) defining the utility function, (iv) modeling the state transition probabilities, and (v) specifying the likelihood function. We describe each of these steps next.

5.1.1. States

We define user i's state at time t by his or her number of followers on that day, s_{it}.

5.1.2. Actions

We denote the action taken by user i at time t by $a_{it} = \{n_{it}, p_{it}\}$, where n_{it} is a binary variable equal to 1 if user i followed at least one new user at time t and p_{it} is a binary variable equal to 1 if user i posted content at time t.[17] As mentioned above, each user faces a choice between four possible actions in each period: $\{1,1\}, \{1,0\}, \{0,1\}, \{0,0\}$. We next describe the costs and benefits associated with each action.

5.1.3. Utility function

We model the utility derived by user i in period t as:

$$(1) \qquad u(s_{it}, a_{it} \mid \theta_i) = \theta_{i1}(1 + s_{it})^{\theta_{i2}} + p_{it} \cdot \theta_{i3}(1 + s_{it})^{\theta_{i4}} + \theta_{i5}n_{it} + \theta_{i6}p_{it}$$

where $\theta_i = \{\theta_{i1}, \theta_{i2}, \theta_{i3}, \theta_{i4}, \theta_{i5}, \theta_{i6}\}$ and the following constraints are imposed:

$\theta_{i1}, \theta_{i3} \geq 0$; $\{\theta_{i2}, \theta_{i4}\} \in [0,1]^2$; $\theta_{i5}, \theta_{i6} \leq 0$ (when estimating the model, we further constrain θ_i to a compact set Θ).

The specification of our utility function is driven by the discussion in Section 2.2. The first term, $\theta_{i1}(1 + s_{it})^{\theta_{i2}}$, captures image-related utility. As discussed above, the stature or prestige of a Twitter user may be measured as a monotonically non-decreasing and concave function of that user's number of followers. This term does not depend on the action chosen in period t, but it does depend on the current state which is the result of past actions. The next term, $p_{it} \cdot \theta_{i3}(1 + s_{it})^{\theta_{i4}}$, captures intrinsic utility from posting content. This term is positive only if the user posts content in period t, and is equal to 0 otherwise. When positive, this term is also

---

[17] An alternative formulation would consider the *number* of posts and the *number* of new users followed. However this would make the action space unbounded, and would pose significant computational challenges. Therefore, we treat these actions as binary.

monotonically non-decreasing and concave in the number of followers. Finally, the last two terms capture the cost of following a new user and of posting, respectively.

5.1.4. State transition probabilities

We denote the state transition probabilities by $f(s'|s,a)$, the probability of reaching state s' in the next period given a state s and an action a in the current period. Because we have access to a substantial amount of data (daily states and actions of 2,493 users over 160 days) and because our action space only includes four possible actions, we are able to use the observed transition frequencies as state transition probabilities instead of estimating them parametrically (see for example Bajari, Benkard and Levin 2007). These empirical state transition probabilities are based on all observations from all 2,493 users in our main observation window. Each observation consists of a triplet $\{s_{it}, a_{it}, s_{i(t+1)}\}$, i.e., a starting state, an action and a resulting next state. The empirical transition probability $f(s'|s,a)$ is simply the proportion of times states s' was observed among all the observations where action a was taken in state s. Using observed transition frequencies implies that we need to limit our analysis to the largest set of states in which all four actions were observed and in which these actions always led to states in that same set of states. This results in a state space containing 539 states with numbers of followers ranging from 0 to 1618, which covers 93.36% of the initial observations (only 2.08% of the initial observations involve states with more than 1618 followers).

Across all states and observations, we find that the actions $\{follow, post\}$, $\{follow, don't\ post\}$, $\{don't\ follow, post\}$ and $\{don't\ follow, don't\ post\}$ give rise to expected changes in the number of followers $E(s'-s|a)$ of 1.209, 0.487, 0.197, and 0.022 respectively, and that the probability of an increase in the number of followers, $\text{Prob}(s'-s>0|a)$, is 0.511, 0.335, 0.257,

and 0.075 respectively. This confirms that both posting content and following new users are ways to attract new followers.

5.1.5. Likelihood function

Following Rust (1987), we assume an unobservable shock ε_{it} to utility for consumer i at time t which follows a double-exponential distribution. The value function for user i in state $\{s, \varepsilon\}$ is the solution to the following Bellman equation:

(2) $V(s, \varepsilon \mid \theta_i) = \max_{a \in A}[u(s, a \mid \theta_i) + \varepsilon(a) + \beta . \sum_{s'} f(s' \mid s, a). V(s' \mid \theta_i)]$

where β is a discount factor, A is the action space, and $V(s \mid \theta_i) = \underset{\varepsilon}{E}(V(s, \varepsilon \mid \theta_i))$. Note that we make the standard assumption that all users have correct beliefs regarding the state transition probabilities. We define action-specific value functions as:

(3) $V_a(s \mid \theta_i) = u(s, a \mid \theta_i) + \beta . \sum_{s'} f(s' \mid s, a). V(s' \mid \theta_i)$

Rust (1987) showed that:

(4) $V(s \mid \theta_i) = \log(\sum_{a \in A} \exp(V_a(s \mid \theta_i)))$

and that the probability that user i chooses action a_{it} at time t in state s_{it} is given by:

(5) $P(a_{it} \mid s_{it}, \theta_i) = \dfrac{\exp(V_{a_{it}}(s_{it} \mid \theta_i))}{\sum_{a' \in A} \exp(V_{a'}(s_{it} \mid \theta_i))}$

Equation (5) defines our likelihood function. Maximizing this likelihood function poses great computational challenges, due to the fact that the likelihood involves the solution to Bellman's

equation. We use a method recently proposed by Norets (2009a) to estimate the set of parameters $\{\theta_i\}$. We describe our estimation procedure next.

5.2. Estimation

Although we use the full set of users to estimate the state transition probabilities, we estimate the parameters of the model $\{\theta_i\}$ for treated users only, mainly for tractability and identification purposes. In particular, we estimate the model on the set of treated users over the main observation window ($t=1,\ldots160$) and remove from the analysis users with fewer than 10 usable observations. The estimation method proposed by Norets (2009a) is closely related to the method proposed by Imai, Jain and Ching (2009), and uses a Bayesian approach to simulate the posterior distribution of the parameters conditional on the data. Given our particular dataset and model, we specifically adapt the approach proposed in Corollary 2 of Norets (2009a). This method offers the combined benefits of being computationally tractable, of being based on the full solution of the dynamic program, and of allowing us to capture heterogeneity in the estimated parameters.

Our estimation algorithm relies on a hierarchical Bayes model that consists of a likelihood, a first-stage prior and a second-stage prior. As mentioned above, our likelihood function is given by Equation (5). We then impose a first-stage prior on the parameters, according to which θ_i comes from a truncated normal distribution with mean θ_0 and covariance matrix D_θ, where the truncation ensures that θ_i remains in a compact set Θ: $\theta_i \sim TN(\theta_0, D_\theta)$. The arrays θ_0 and D_θ are themselves parameters of the model, on which we impose a second-stage prior. Following standard practice (Allenby and Rossi 1998), we use a non-informative second-stage prior, in order to let these parameters be driven by the data. We fix β=0.995.

We draw values of all the parameters using a Gibbs sampler. The estimation algorithm proposed by Norets (2009a) relies on an approximation of the value function. Let m index the iterations of the Gibbs sampler. The approximation of the value function at iteration m leverages the value function approximations from previous iterations. Let $\{\theta_i^{*(1)},...,\theta_i^{*(m-1)}\}_{i=1,...,I}$ be the parameter draws from the first m-1 iterations. For each user i, the approximation at iteration m uses only the past $N(m)$ iterations, where $N(m)=[m^{\gamma_1}]$ (we set γ_1=0.6). Among these iterations, the approximation relies on the ones at which the draws of the parameters were the closest to the current draw of θ_i. In particular, the $\widetilde{N}(m)$ closest neighbors are considered, where $\widetilde{N}(m)=[m^{\gamma_2}]$ (we set γ_2=0.4). At iteration m, we consider the following approximation of V_a for user i:

$$(6) \quad \widetilde{V}_a^{(m)}(s\,|\,\theta_i) = u(s,a\,|\,\theta_i) + \beta \cdot \frac{1}{\widetilde{N}(m)} \sum_{l=1}^{\widetilde{N}(m)} \sum_{s'} \widetilde{V}^{(k_l)}(s'\,|\,\theta_i^{*(k_l)}) \cdot f(s'\,|\,s,a)$$

where $\{k_l\}_{l \in \{1,...\widetilde{N}(m)\}}$ index the closest neighbors to θ_i among the $N(m)$ past draws, and $\widetilde{V}^{(k_l)}$ refers to value function approximations computed in previous iterations. The likelihood is then approximated by:

$$(7) \quad \widetilde{P}^{(m)}(a_{it}\,|\,s_{it},\theta_i) = \frac{\exp(\widetilde{V}_{a_{it}}^{(m)}(s_{it}\,|\,\theta_i))}{\sum_{a' \in A} \exp(\widetilde{V}_{a'}^{(m)}(s_{it}\,|\,\theta_i))}$$

At the end of each iteration, a set of new value function approximations are generated to be used in future iterations. These new approximations are obtained by applying the Bellman operator:

$$(8) \quad \widetilde{V}^{(m)}(s\,|\,\theta_i^{*(m)}) = \log(\sum_{a \in A} \exp(\widetilde{V}_a^{(m)}(s\,|\,\theta_i^{*(m)})))$$

Norets (2009a) proves theoretically that these approximations converge uniformly and completely to the exact values (see Theorem 1, p. 1677). We run the Gibbs sampler for 2,000 iterations, using the first 500 as burn-in. Convergence was assessed visually using plots of the parameters. Further detail of the estimation procedure is presented in the Appendix.

5.3. Results

5.3.1. Model fit

A common measure of model fit in Bayesian statistics is the marginal density of the data according to the model, defined as: $LMD = \int_{\Theta^I} P(data | \{\theta_i\}) P(\{\theta_i\}) d\theta$, where

$$P(data | \{\theta_i\}) = \prod_{i=1}^{I} \prod_{t=1}^{T} P(a_{it} | s_{it}, \theta_i)$$ and $P(\{\theta_i\})$ is the prior distribution on the parameters. The log marginal density is approximated by the harmonic mean across the Gibbs sampler iterations of the likelihood of the data (see for example Sorensen and Gianola 2002). To make the results more intuitive, we report the marginal density to the power of the inverse of the total number of observations, i.e., we report the geometric mean of the marginal density per observation. We obtain a value of 0.4726. For comparison, a null model that assumes that all four actions are equally likely would achieve a per-observation marginal density of 0.2500. We also assess the fit of the model using posterior predictive checks (Gelman, Meng, and Stern 1996). Posterior predictive checks assess how well the model fits key aggregate statistics of the data, by comparing the posterior distribution of these statistics based on the model with the observed values. We consider the proportion of observations in our data in which posting was observed (i.e., $a_{it} = \{0,1\}$ or $\{1,1\}$). The observed proportion is 0.2470. The predicted value of this quantity is computed at each iteration of the Gibbs sampler, as

$$\sum_{i=1}^{I}\sum_{t=1}^{T}P(a_{it}=\{1,1\}\,|\,s_{it},\theta_i)+P(a_{it}=\{0,1\}\,|\,s_{it},\theta_i).$$ The posterior distribution (across iterations of the

Gibbs sampler) of the predicted value of this quantity is shown in Figure 7. The mean of the

posterior distribution is 0.2604, and the 95% credible interval ([0.2405;0.2797]) contains the

observed value of 0.2470.

We also consider the distribution (across users) of the proportion of observations in which

posting was observed. We compute a point estimate of the posting frequency for each user (this

point estimate for user i is the average across posterior draws of

$$\sum_{t=1}^{T}P(a_{it}=\{1,1\}\,|\,s_{it},\theta_i)+P(a_{it}=\{0,1\}\,|\,s_{it},\theta_i)).$$ Figure 8 provides a scatter plot of the posting

frequency as predicted by the model versus observed, across users. We see that the model is able

to recover these frequencies very well at the individual level.

[Insert Figures 7-8 Here]

5.3.2. Parameter estimates

Table 3 reports the point estimates and 95% credible intervals of the average parameters across

users, and Figure 9 plots the distribution across users of the point estimates of the parameters.

We see that although the parameter θ_3 tends to be larger than θ_1 (this is the case for 88.30% of

the users), θ_4 tends to be smaller than θ_2 (this is the case for 75.53% of the users). This suggests

that intrinsic utility from posting tends to be larger than image-related utility when the number of

followers is small, with the reverse being true when the number of followers is larger. This is

consistent with the findings of the model-free analysis.

[Insert Table 3 and Figure 9 Here]

5.3.3. Intrinsic versus image-related utility

We now estimate the proportion of intrinsic versus image-related utility derived by each user, according to the model. As discussed above, the image-related utility derived by user i at period is $u_{image_{i,t}} = \theta_{i1}(1 + s_{it})^{\theta_{i2}}$ and the intrinsic utility is $u_{intrinsic_{i,t}} = p_{it} \cdot \theta_{i3}(1 + s_{it})^{\theta_{i4}}$. We define the total image-related and intrinsic utility for user i as: $u_{image_i} = \sum_{t=1}^{T} u_{image_{i,t}}$ and $u_{intrinsic_i} = \sum_{t=1}^{T} u_{intrinsic_{i,t}}$. We estimate the proportion of intrinsic utility for user i as: $proportion_i = \dfrac{u_{intrinsic_i}}{u_{intrinsic_i} + u_{image_i}}$.

Figure 10 plots the distribution across users of the estimate of this proportion. The average across users is 0.2533 and the median is 0.1313. Therefore, according to the model most treated users derived more image-related utility from Twitter than they did intrinsic utility during our main observation window.

[Insert Figure 10 Here]

5.3.4. Counterfactual analysis: change in posting activity if the network's structure became stable

Finally, we consider the question of how posting frequency would be affected if the network's structure were stable, i.e., if a user's actions did not influence his or her future states. If actions did not influence future states, actions would be chosen according to the utility derived in the current period only. For each observation in our data, we estimate the probability of each four actions based only on the immediate utility provided by that action:

$P^{stable}(a_{it} \mid s_{it}, \theta_i) = \dfrac{\exp(u(s_{it}, a_{it} \mid \theta_i))}{\sum_{a' \in A} \exp(u(s_{it}, a' \mid \theta_i))}$. For each user we estimate the posting frequency under these probabilities and compare these frequencies with those obtained under the initial model (used to construct the scatter plots in Figure 8). Figure 11 plots the distribution across users of

the change in posting frequency resulting from the assumption that actions do not influence future states. The average predicted change across users is -0.0502, the median is -0.0347, and the predicted change is negative for 94.68% of the users. In other words, this counterfactual analysis suggests that if users' posting activities did not influence their future number of followers, posting frequency would decrease by an average of approximately 5%.

[Insert Figure 11 Here]

6. Discussion and Conclusion

While publishers' incentives in traditional media are well understood, the incentives of individuals to contribute content on social media platforms such as Twitter are not yet well understood. Recently some researchers have considered the psychological consequences of various behaviors related to social media use. It has been found for example that contributing content to social media can affect self-esteem, well-being, and self-control (Gonzales and Hancock 2011; Valkenburg, Peter, and Schouten 2006; Wilcox and Stephen 2011). However, to the best of our knowledge, individuals' motivations for contributing content in social media have been explored so far only using survey-based approaches (e.g., Bughin 2007, Hennig-Thurau et al. 2004).

Previous literature suggests that two primary types of motivation to contribute content to social media in the absence of financial or other extrinsic incentives are intrinsic and image-related. On Twitter, intrinsic motivation to post content leads to the prediction that users should post more as their number of followers (i.e., their audience) increases. On the other hand, image-related motivation leads to the prediction that users should derive less marginal utility from additional followers as their number of followers (a measure of stature) increases, and therefore

should have less motivation to post content. These opposing predictions have allowed us to empirically identify and compare these two sources of motivation. In order to address the issue that the number of followers is endogenous, we conducted a field experiment in which we exogenously increased the number of followers of a set of users (treatment group). Our results are consistent with the co-existence of intrinsic and image-related motivations, with the relative effects of these two sources of motivation varying across users with different initial numbers of followers. We estimate a dynamic discrete choice model that quantifies intrinsic versus image-related utility for each treated user. Counterfactual analyses based on this model suggest that posting frequency is likely to slow down as the network's structure becomes stable (i.e., as it becomes harder to attract new followers by posting content).

Our findings are relevant not only theoretically but also managerially. Understanding what motivates consumers to be active on social media platforms like Twitter is a prerequisite for marketers interested in devising efficient social media strategies and optimizing the ways they engage with consumers on these platforms. Moreover, our results, and in particular our counterfactual analysis, inform how social media platforms based on information sharing may be expected to mature. We find that non-commercial users should be expected to contribute less content as the network becomes stable. Non-commercial users would, however, still derive value from consuming content contributed by others, in particular commercial users. In this sense, Twitter might become more and more like another media channel than a social network.

Our research also offers several areas for future research. First, our data were limited to the number of tweets and number of users following and followed by each user. Future research may enrich our findings by using data that would include the text of the tweets as well as the structure of the social network. Second, beyond the context of the present paper, future research may also

explore further the use of field experiments as a way to address endogeneity issues in social networks. Endogeneity has been a perennial issue in the study of social interactions (see for example Hartmann et al. 2008; Manski 1993, and Moffitt 2001). Given that social media platforms such as Twitter exist in the public domain and are easily accessible, future research on social media and networks may adopt identification strategies similar to ours. Finally, future research may study motivations in other related contexts. In Twitter, a user's stature is based on how many followers they have, the ability of a user to directly affect this metric is limited, and contributing content is a way to increase stature. These characteristics are found in a number of other social media platforms and we would expect our results to apply to these platforms. For example, in YouTube users can have "subscribers" (akin to followers) and, like Twitter, reciprocity is not automatic and users need to contribute content in order to attract new subscribers. Conversely, our results may be less applicable to other platforms such as Facebook where the social network is undirected and therefore all ties are reciprocal (a Facebook user can directly increase his or her number of friends simply by "friending" more people). Note however that other aspects of Facebook are more directly relevant to our findings. For instance, recent changes to Facebook have introduced a Twitter-like "subscription" feature that allows users to follow others without being followed back.

Appendix: Details of the estimation procedure

At each iteration m of the Gibbs sampler:

-A proposed value of θ_i, $\theta_i^{*(m)}$, is drawn for each user. Let $\theta_i^{(m-1)}$ denote the value retained at iteration m-1. For each i, the new proposed value is drawn such that: $\theta_i^{*(m)} = \theta_i^{(m-1)} + d_i$ where $d_i \sim N(0, \gamma D_\theta)$ and γ is adjusted such that the acceptance rate is around 30%. Using rejection sampling (see for example Allenby, Arora and Ginter 1995), we constrain $\theta_i^{*(m)}$ to a compact set Θ=[0,10]*[0,0.6]*[0,10]* [0,0.6]*[-10,0]^2 (Note: we use 0.6 instead of 1 as an upper bound for $\theta_{i,2}$ and $\theta_{i,4}$ for numerical stability issues – this constraint does not appear to be binding). The new value is retained (i.e., $\theta_i^{(m)} = \theta_i^{*(m)}$) with probability:

$$\min\{\frac{P(\theta_i^{*(m)} \mid \theta_0, D_\theta).\widetilde{L}^{(m)}(\{a_{it}\}_{t=1,...T} \mid \theta_i^{*(m)})}{P(\theta_i^{(m-1)} \mid \theta_0, D_\theta).\widetilde{L}^{(m)}(\{a_{it}\}_{t=1,...T} \mid \theta_i^{(m-1)})}, 1\}, \text{ where } \widetilde{L}^{(m)}(\{a_{it}\}_{t=1,...T} \mid \theta_i) = \prod_{t=1}^{T} \widetilde{P}^{(m)}(a_{it} \mid s_{it}, \theta_i),$$

where $\widetilde{P}^{(m)}(a_{it} \mid s_{it}, \theta_i)$ is defined in Equation (7). Otherwise we set $\theta_i^m = \theta_i^{(m-1)}$.

-A new approximation of the value function is generated for $\theta_i^{*(m)}$, to be used in future iterations. This new approximation is obtained by applying the Bellman operator in Equation (8). As recommended by Norets (2009b), we apply the Bellman operator more than once. In our implementation we apply the Bellman operator until the new values are close enough to the previous ones (maximum absolute deviation less than 1), with a minimum of 10 iterations.

-The first-stage prior parameters θ_0 and D_θ are updated. Our second stage prior on θ_0 is uniform on the compact set Θ, and our second-stage prior on D_θ follows an inverse-Wishart distribution: $D_\theta^{-1} \sim W(v_0, V_0)$ with $v_0 = k+3$ (where k is the number of elements in θ_i) and $V_0 = 0.001I$ (where I is the identity matrix). These second-stage priors have the attractive property of being conjugate with the likelihood for the first-stage prior parameters, which allows us to draw the parameters directly from their respective conditional posterior distributions: $\theta_0 \sim TN(\dfrac{\sum_{i=1}^{I} \theta_i}{I}, \dfrac{D_\theta}{I})$ and

$$D_\theta^{-1} \sim W(v_0 + I, V_0 + \dfrac{\sum_{i=1}^{I}(\theta_i - \theta_0)(\theta_i - \theta_0)^T}{I})$$

References

Allenby, Greg M., and Peter E. Rossi. 1998. "Marketing Models of Consumer Heterogeneity." *Journal of Econometrics*: 57-78.

Allenby, Greg M., Neeraj Arora, and James L. Ginter. 1995. "Incorporating Prior Knowledge into the Analysis of Conjoint Studies." *Journal of Marketing Research*: 152-162.

Ariely, Dan, Anat Bracha, and Stephan Meier. 2009. "Doing Good or Doing Well? Image Motivation and Monetary Incentives in Behaving Prosocially." *The American Economic Review*: 544-555.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. "Estimating Dynamic Models of Imperfect Competition." *Econometrica*: 1331-1370.

Bakshy, Eytan, Jake M. Hofman, Winter A. Mason, and Duncan J. Watts. 2011. "Everyone's an influencer quantifying influence on twitter." *Proceedings of the fourth ACM international conference on Web search and data mining - WSDM '11*: 65-74.

Barabási, Albert-László and Réka Albert. 1999. "Emergence of Scaling in Random Networks." *Science*: 509-512.

Baumeister, Roy F., and Mark R. Leary. 1995. "The Need to Belong: Desire for Interpersonal Attachment as a Fundamental Human Motivation." *Psychological Bulletin*: 497-529.

Bénabou, Roland, and Jean Tirole. 2006. "Incentives and Prosocial Behavior." *American Economic Review*: 1652-1678.

Bughin, Jacques R. 2007. "How Companies Can Make the Most of User-Generated Content." *McKinsey Quarterly*: 1-4.

Cha, Meeyoung, Hamed Haddadi, Fabricio Benevenuto, and Krishna P. Gummadi. 2010. "Measuring User Influence in Twitter: The Million Follower Fallacy." *Proceedings of the fourth international AAAI conference on weblogs and social media*: 10-17.

Cha, Meeyoung, Alan Mislove, and Krishna P. Gummadi. 2009. "A Measurement-driven Analysis of Information Propagation in the Flickr Social Network." *Proceedings of the 18th international conference on World Wide Web*: 721-730.

de Solla Price, Derek J. 1965. "Networks of Scientific Papers." *Science*: 510-515.

Fehr, Ernst, and Armin Falk. 2002. "Psychological Foundations of Incentives." *European Economic Review*: 687-724.

Fershtman, Chaim, and Neil Gandal. 2007. "Open Source Software: Motivation and Restrictive Licensing." *International Economics and Economic Policy*: 209-225.

Gelman, Andrew, Xiao-Li Meng, and Hal Stern. 1996. "Posterior Predictive Assessment of Model Fitness Via Realized Discrepancies." *Statistica Sinica*: 733-807.

Ghose Anindya, Avi Goldfarb Sang-Pil Han. 2011. "How is the Mobile Internet Different? Search Costs and Local Activities." New York University working paper.

Glazer, Amihai, and Kai Konrad. 1996. "A Signaling Explanation for Charity." *American Economic Review*: 1019-1028.

Gonzales, Amy L. and Jeffrey T. Hancock. 2011. "Mirror, Mirror on my Facebook Wall: Effects of Facebook Exposure on Self-Esteem." *Cyberpsychology, Behavior & Social Networking*, (January/February), 79-83.

Harbaugh, William T. 1998a. "The Prestige Motive for Making Charitable Transfers." *American Economic Review*: 277-282.

Harbaugh, William T. 1998b. "What Do Donations Buy? A Model of Philanthropy Based on Prestige and Warm Glow." *Journal of Public Economy*: 269-284.

Hartmann, Wesley, Puneet Manchanda, Matthew Bothner, Peter Dodds, David Godes, Kartik Hosanagar, and Catherine Tucker. 2008. "Modeling Social Interactions: Identification, Empirical Methods and Policy Implications." *Marketing Letters*: 287:304.

Hennig-Thurau, Thorsten, Kevin P. Gwinner, Gianfranco Walsh, and Dwayne D. Gremler. 2004. "Electronic Word-of-Mouth via Consumer-Opinion Platforms: What Motivates Consumers to Articulate Themselves on the Internet?" *Journal of Interactive Marketing*: 38:52.

Imai, Susumu, Neelam Jain, and Andrew Ching. 2009. "Bayesian Estimation of Dynamic Discrete Choice Models." *Econometrica*: 1865-1899.

Kankanhalli, Atreyi, Bernard C.Y. Tan, and Kwok-Kee Wei. 2005. "Contributing Knowledge to Electronic Knowledge Repositories: An Empirical Investigation." *MIS Quarterly*: 113-143.

Kumar, Ravi, Jasmine Novak, and Andrew Tomkins. 2006. "Structure and Evolution of Online Social Networks." *Proceedings of the 12th ACM SIGKDD International conference on knowledge discovery and data mining*: 611-617.

Kumar, Vineet. 2009. "Why do Consumers Contribute to Connected Goods? A Dynamic Game of Competition and Cooperation in Social Networks." Tepper School of Business working paper.

Kwak, Haewoon, Changhyun Lee, Hosung Park, and Sue Moon. 2010. "What is Twitter, a social network or a news media?" *Proceedings of the 19th international conference on World wide web - WWW '10*: 591-600.

Lakhani, Karim R. and Robert Wolf. 2005. "Why Hackers Do What They Do: Understanding Motivation and Effort in Free/Open Source Software Projects." In *Perspectives on Free and Open Source Software*, edited by J. Feller, B. Fitzgerald, S. Hissam and K. Lakhani. Cambridge, Mass: MIT Press.

Lampel, Joseph, and Ajay Bhalla. 2007. "The Role of Status Seeking in Online Communities: Giving the Gift of Experience." *Journal of Computer-Mediated Communication*: 434-455.

Lerner, Josh, and Jean Tirole. 2002. "Some Simple Economics of Open Source." *The Journal of Industrial Economics*, 197-234.

Lerner, Josh, and Jean Tirole. 2005. "The Economics of Technology Sharing: Open Source and Beyond." *Journal of Economic Perspectives*: 99-120.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies*: 531-42.

Moffitt, Robert A. 2001. "Policy Interventions, Low-Level Equilibria, and Social Interactions." In *Social Dynamics*, edited by S. Durlauf and P. Young, 45-82 Washington, D.C.: Brookings Institution Press and MIT.

Norets, Andriy. 2009a. "Inference in Dynamic Discrete Choice Models with Serially Correlated Unobserved State Variables." *Econometrica*: 1665-1682.

Norets, Andriy. 2009b. "Implementation of Bayesian Inference in Dynamic Discrete Choice Models." Princeton University working paper.

Nov, Oded. 2007. "What Motivates Wikipedians?" *Communications of the ACM*: 60-64.

Romero, Daniel M., Wojciech Galuba, Sitaram Asur, and Bernardo A. Huberman. 2011. "Influence and passivity in social media." *Proceedings of the 20th international conference companion on World wide web - WWW '11*: 113-114.

Rust, John. 1987. "Optimal Replacement of GMC Bus Engines: an Empirical Model of Harold Zucker." *Econometrica*: 999-1033.

Ryan, Richard M., and Edward L. Deci. 2000. "Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions." *Contemporary Educational Psychology*: 54-67.

Shriver, Scott K., Harikesh S. Nair, and Reto Hofstetter. 2011. "Social Ties and User-Generated Content: Evidence from an online Social Network." Stanford GSB working paper.

Sorensen, Daniel, and Daniel Gianola. 2002. *Likelihood, Bayesian, and MCMC Methods in Quantitative Genetics*. New York: Springer.

Stephen, Andrew T. and Olivier Toubia. 2009. "Explaining the Power-Law Degree Distribution in a Social Commerce Network." *Social Networks*: 262-270.

Stephen, Andrew T., Yaniv Dover, Jacob Goldenberg, and Lev Muchnik. 2011. "Social Sharing by Social Pumps: The Effect of Transmitter Activity on Information Diffusion Over Online Social Networks." Working paper, University of Pittsburgh.

Sun, Monic, and Feng Zhu. 2011. "Ad Revenue and Content Commercialization: Evidence from Blogs." Stanford GSB working paper.

Valkenburg, Patti M., Jochen Peter, and Alexander P. Schouten. 2006. "Friend Networking Sites and their Relationship to Adolescents' Well-Being and Social Self-Esteem." *Cyberpsychology & Behavior*, 9 (5), 584-90.

Wasko, Molly McLure, and Samer Faraj. 2005. "Why Should I Share? Examining Social Capital and Knowledge Contribution in Electronic Networks of Practice." *MIS Quarterly*: 35-57.

Weng, Jianshu, Ee-Peng Lim, and Jing Jiang. 2010. "TwitterRank: Finding Topic-sensitive Influential Twitterers." *Proceedings of the third ACM international conference on Web search and data mining - WSDM '10*: 261-270.

Wilcox, Keith and Andrew T. Stephen. 2011. "Are Close Friends the Enemy? Online Social Networks, Narcissism, and Self-Control." Working paper, Babson College.

Wu, Shaomei, Jake M. Hofman, Winter A. Mason, and Duncan J. Watts. 2011. "Who Says What to Whom on Twitter." *Proceedings of the 20th international conference on World wide web - WWW '11*: 705-714.

Tables and Figures

Table 1—Five quintiles based on number of followers on day 1.

Quintile	Range of number of followers	Median number of followers	Average number of followers
1	0-12	7	6.499
2	13-26	19	18.941
3	27-61	39.5	40.988
4	62-245	109	125.550
5	246-18940	704	1378.949

Notes: The quintiles (i.e., range of the number of followers) are determined based on the treated users to ensure an equal spread of these users across the five groups. The median and average numbers of followers reported in the table are for the entire set of active users.

Table 2—Proportion of users with increased average daily posting rate and with decreased average daily posting rate (after versus before intervention).

	Increased average daily posting rate		Decreased average daily posting rate	
Quintile	Treated	Control	Treated	Control
1	0.286	0.205	0.238	0.185
2	0.632	0.350	0.158	0.370
3	0.556	0.370	0.278	0.426
4	0.200	0.451	0.700	0.474
5	0.400	0.375	0.500	0.574

Notes: Treated users in the 2^{nd} quintile were significantly more likely to increase their posting rates compared to users in the control group ($z=2.42$, $p<0.02$), and marginally significantly less likely to decrease their posting rates compared to users in the control group ($z=-1.85$, $p=0.06$). Treated users in the 4^{th} quintile were significantly less likely to increase their posting rates compared to users in the control group ($z = -2.18$, $p < 0.03$), and significantly more likely to decrease their posting rates compared to users in the control group ($z=1.94$, $p=0.05$). The differences in the other quintiles are not statistically significant.

Table 3— Point estimates and 95% credible intervals of the average parameters across users.

Parameter	Point estimate	Credible interval
θ_1	0.309	[0.170;0.497]
θ_2	0.329	[0.241;0.388]
θ_3	0.690	[0.586;0.791]
θ_4	0.265	[0.226;0.309]
θ_5	-5.008	[-5.301;-4.648]
θ_6	-5.075	[-5.380;-4.775]

Notes: The utility derived by user i at period t is*: $u_{sit}, ait\theta i = \theta i1(1+sit)\theta i2 + pit\theta i3(1+sit)\theta i4 + \theta i5nit + \theta i6pit$,* where s_{it} is user i's number of followers in period t, p_{it} is a binary variable equal to 1 if user i posts content in period t, and n_{it} is a binary variable equal to 1 if user i follows at least one new user in period t.
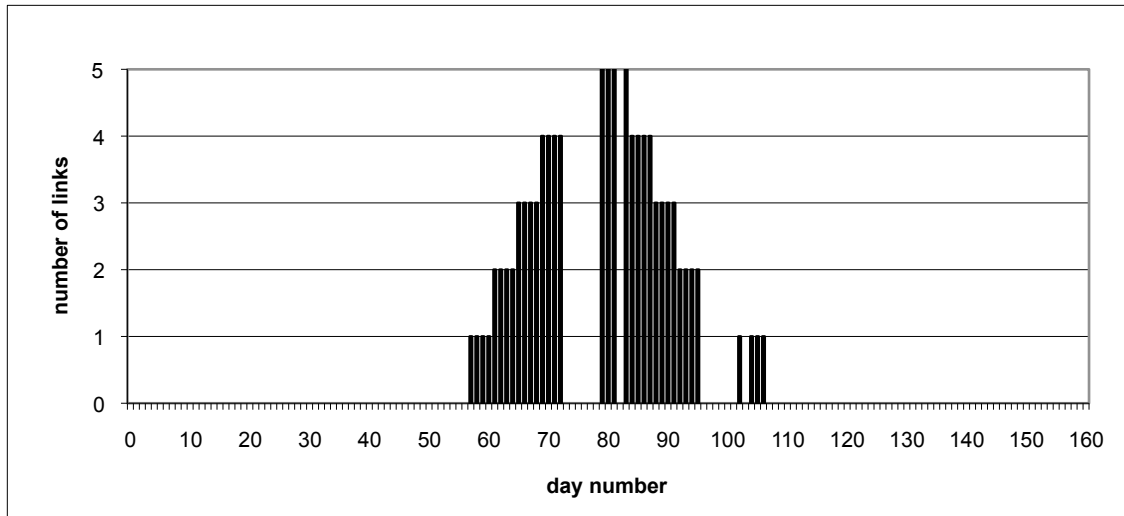
Figure 1. Daily number of exogenous links created per treated user over the main observation window.
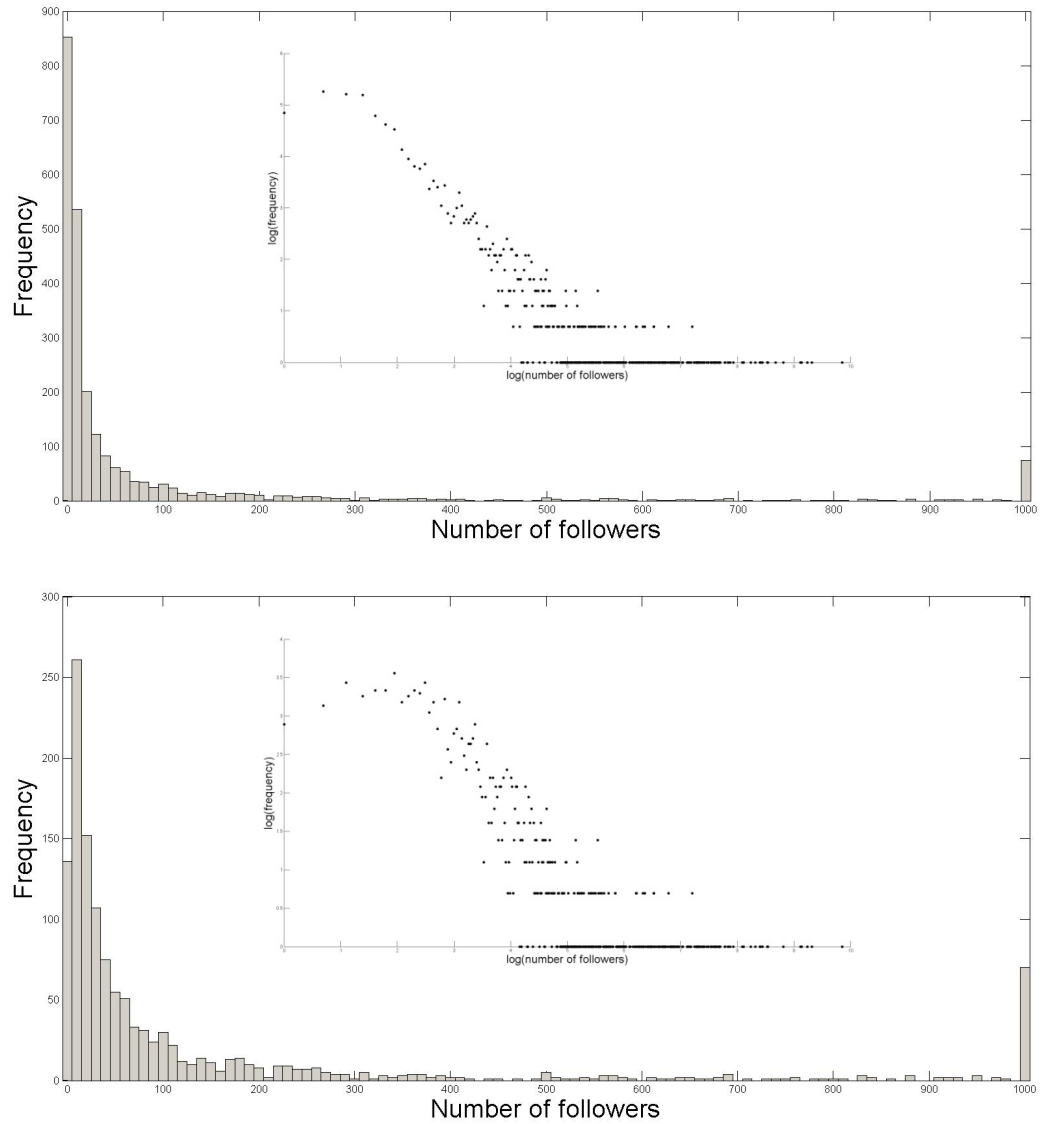
Figure 2. Histograms and log-log plots of the distribution of the number of followers on day 1 for all users (top panel) and all active users (bottom panel).
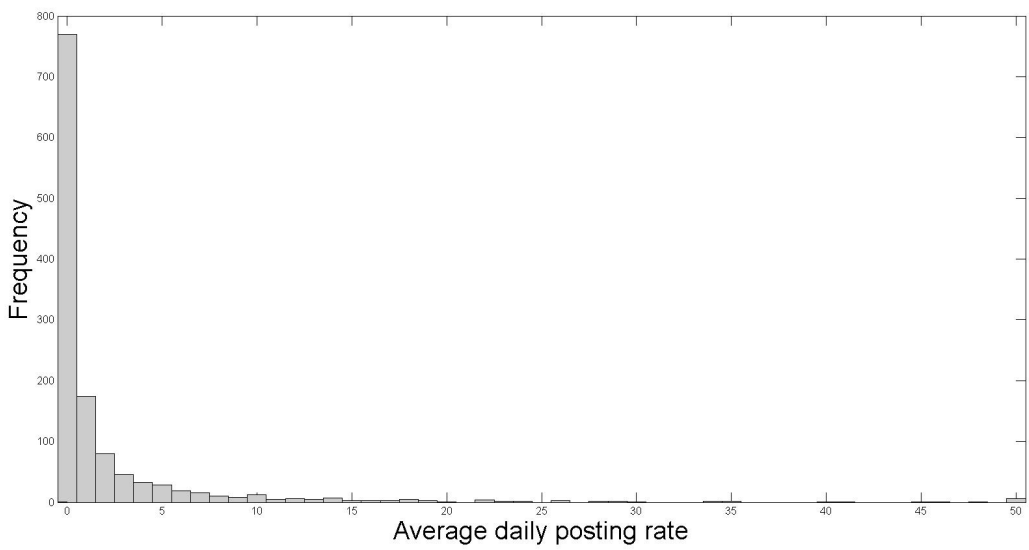
Figure 3. Distribution of average daily posting rate among all active users.
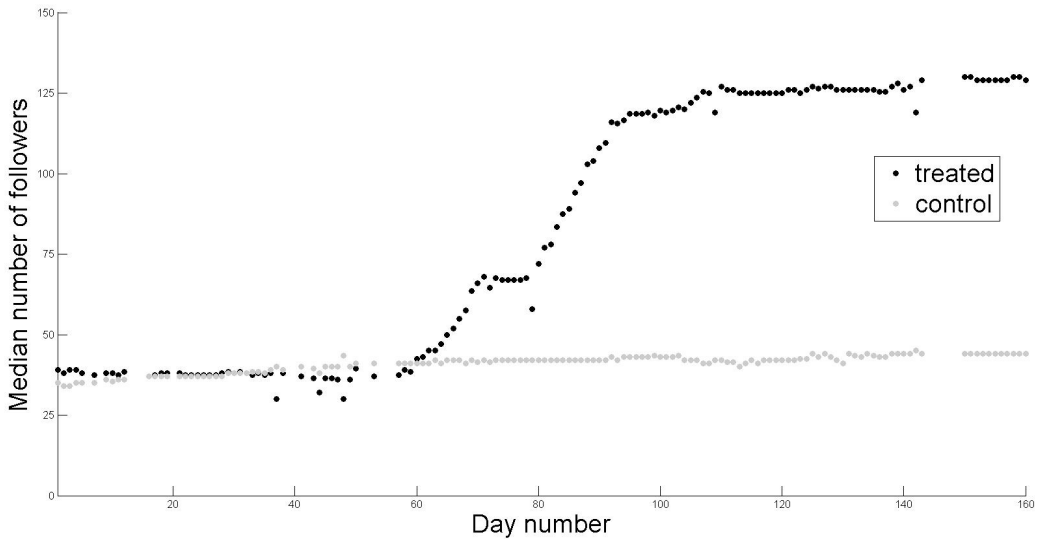
Figure 4. Median number of followers as a function of time for treated and control users.
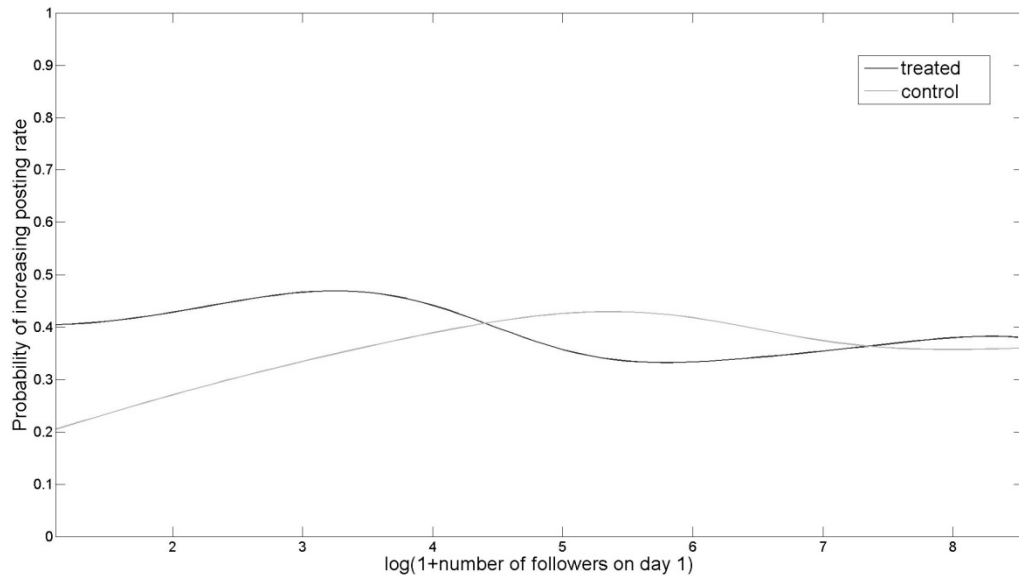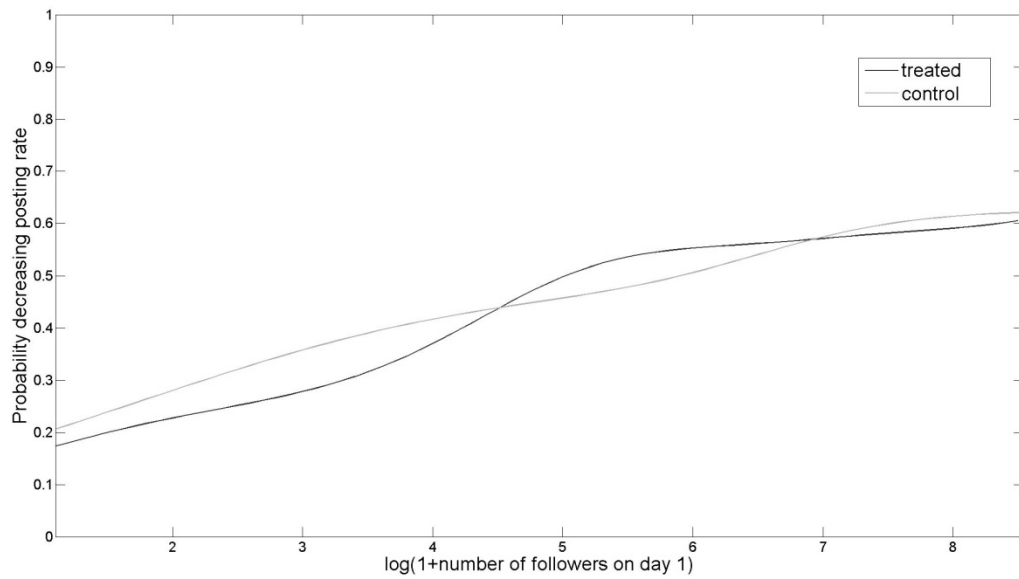
Figure 5. Probability of increasing posting rate after versus before the intervention as a function of the log of the number of followers on day 1.

*Notes:* This Figure is obtained by smoothing the raw data using a Gaussian kernel function (with a bandwidth of 1).

Figure 6. Probability of decreasing posting rate after versus before the intervention as a function of the log of the number of followers on day 1.

Notes: This Figure is obtained by smoothing the raw data using a Gaussian kernel function (with a bandwidth of 1).
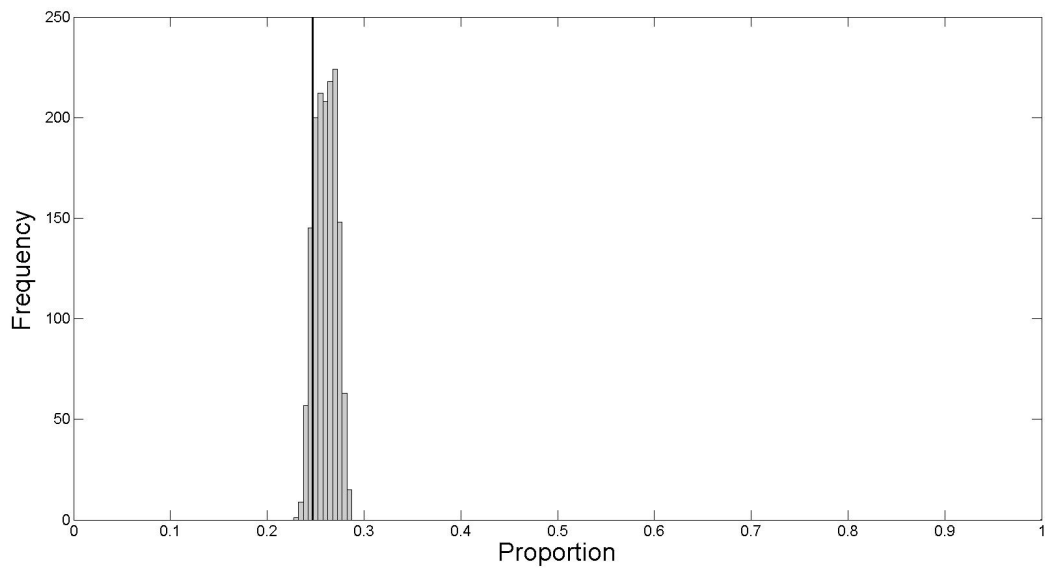
Figure 7. Posterior check of proportion of observations that include posting.

Notes: The histogram plots the posterior distribution of the proportion of observations in which users post content, as predicted by the model. The solid line represents the observed proportion.
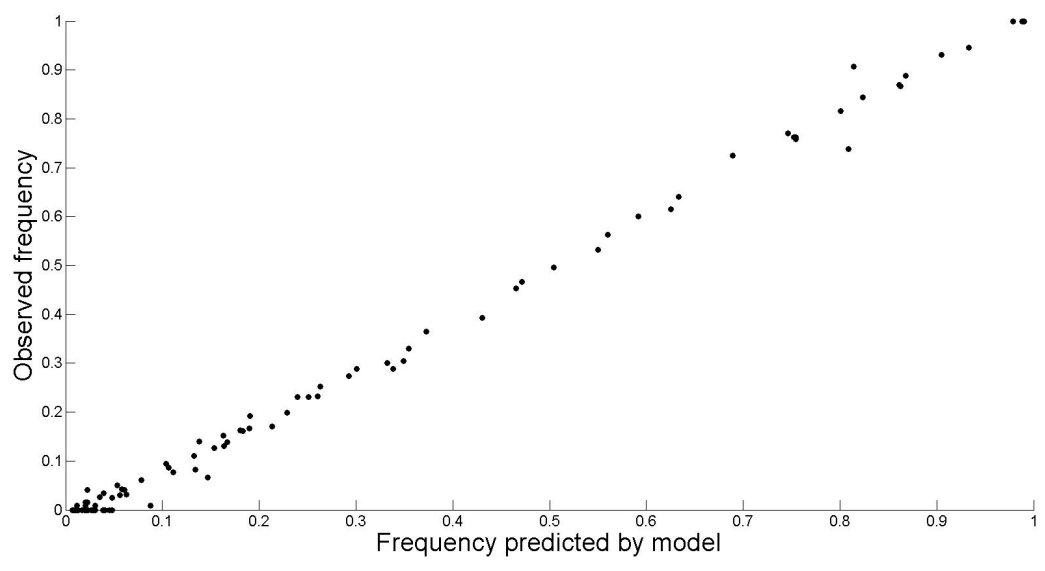

Figure 8. Posting frequency as predicted by the model versus observed

Notes: The x-axis corresponds to the posting frequency (proportion of observations that involve posting) as predicted by the model, the y-axis corresponds to the observed frequency. Each dot corresponds to one user.
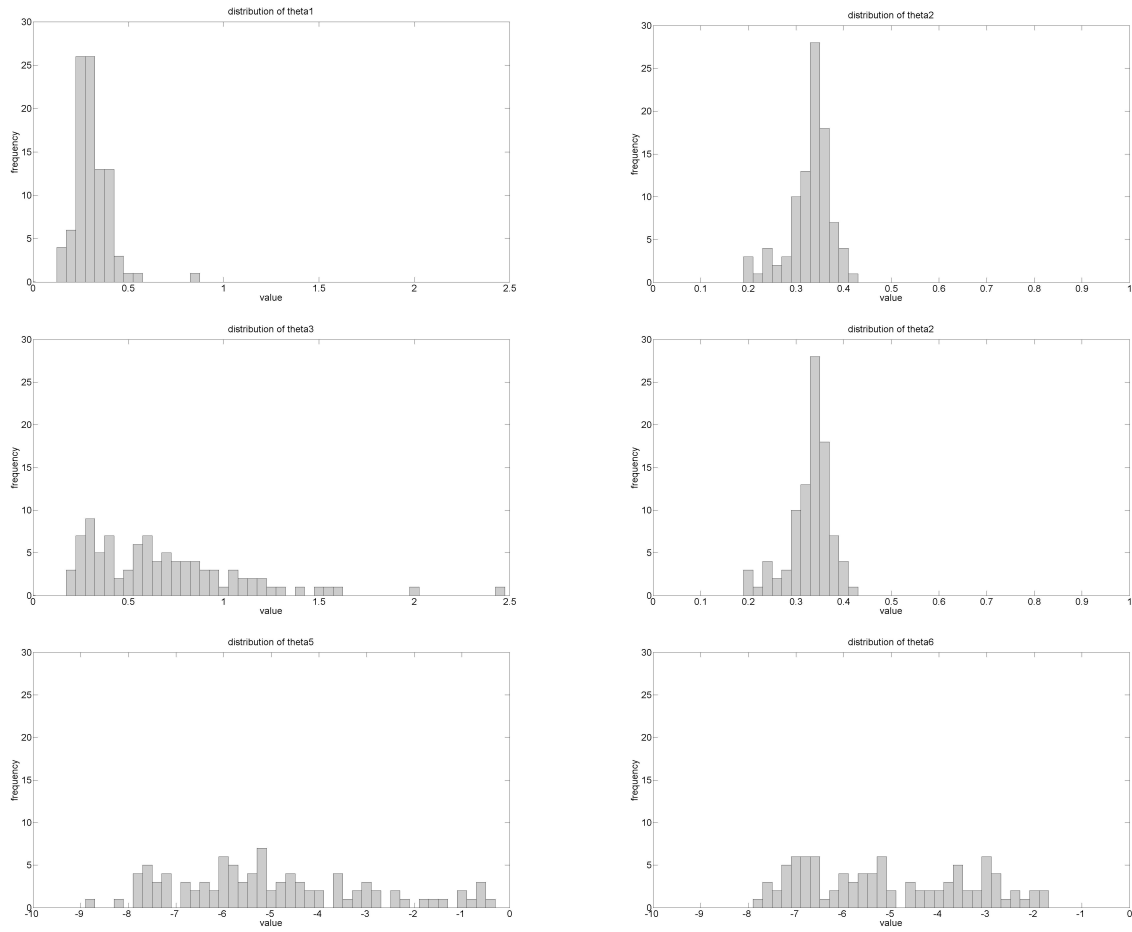
Figure 9. Distribution of the parameters across users.

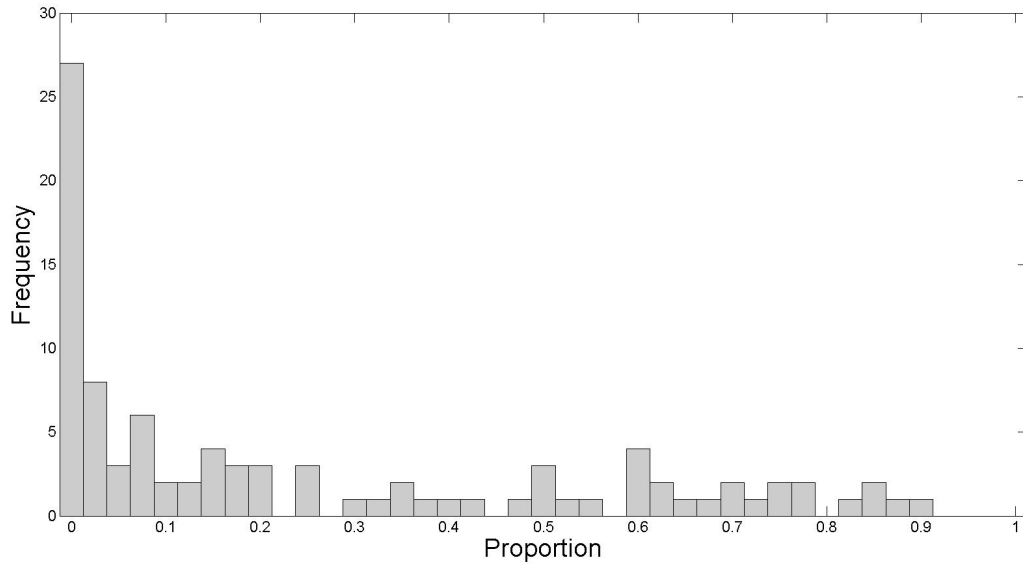*Notes:* Distribution of θ_1 to θ_6 across users (from left to right and from top to bottom).

Figure 10. Distribution across users of the proportion of intrinsic utility (intrinsic / (intrinsic + image-related)).
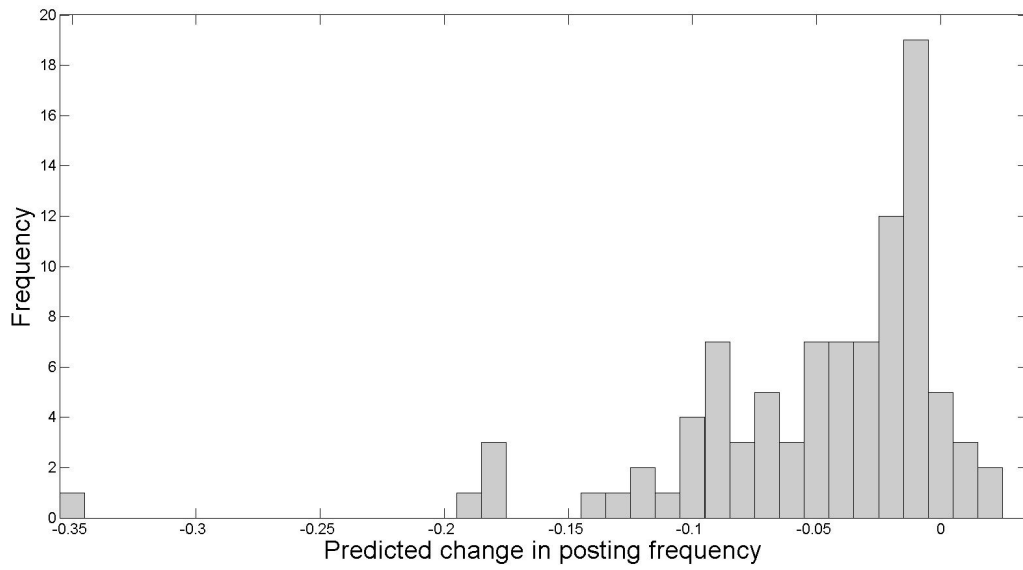

Figure 11. Distribution across users of the predicted change in posting frequency if network's structure became stable.