

Changes in Engagement Before and After Posting to Facebook

Nir Grinberg^{1,2}, P. Alex Dow², Lada A. Adamic² and Mor Naaman¹
nir@cs.cornell.edu, adow@fb.com, ladamic@fb.com, mor@jacobs.cornell.edu
¹Jacobs Institute, Cornell Tech ²Facebook Inc.

ABSTRACT

The asynchronous nature of communications on social network sites creates a unique opportunity for studying how posting content interacts with individuals' engagement. This study focuses on the behavioral changes occurring hours before and after contribution to better understand the changing needs and preferences of contributors. Using observational data analysis of individuals' activity on Facebook, we test hypotheses regarding the motivations for site visits, changes in the distribution of attention to content, and shifts in decisions to interact with others. We find that after posting content people are intrinsically motivated to visit the site more often, are more attentive to content from friends (but not others), and choose to interact more with friends (in large part due to reciprocity). In addition, contributors are more active on the site hours before posting and remain more active for less than a day afterwards.

Our study identifies a unique pattern of engagement that accompanies contribution and can inform the design of social network sites to better support contributors.

Keywords

Computer-mediated communication; Social Media; Information Sharing; Social Participation; Engagement; User Behavior.

Categories and Subject Descriptors

H.1.2 [Models and Principles]: User/Machine Systems—*Human factors, Human information processing*; H.5.m [Information Interfaces and Presentation (e.g. HCI)]: Miscellaneous

1. INTRODUCTION

In this work, we seek to shed light on behavior and engagement practices accompanying posting on Facebook. Affordances of information sharing on Social Network Sites (SNS) [3] determine the experience for contributors, their community and the dynamics of the network as a whole. As a result, much research has focused on people's motivations to post on social media [10, 19, 28, 29, 30, 32]. However, to date, little research has examined posters' behavior and engagement *directly* after (or before) the act of posting.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI'16, May 07-12, 2016, San Jose, CA, USA

© 2016 ACM. ISBN 978-1-4503-3362-7/16/05...\$15.00

DOI: <http://dx.doi.org/10.1145/2858036.2858501>

We draw on existing theories from communication and social psychology to formulate hypotheses about contributors' behavior on SNS. We address three different questions in this work. First, we test whether contributors (those who post their own content at a given point in time) are intrinsically feedback-seeking and visit the site more often after contribution even when no knowledge of feedback exists. Second, we examine whether content consumption practices of contributors change both in quantity and selectivity. Lastly, we investigate changes in interaction rates with others' content, and quantify the effect of reciprocity in interactions with friends.

Better understanding of the mechanisms behind contribution is important for both theoretical and practical reasons. The underlying processes that accompany contribution to SNS are not yet well understood [5, 7, 22]. Studying the relation between contribution and user engagement in large-scale observational datasets can provide a new perspective for understanding individuals' behavior in context, and complement previous research that relied on self-reported measures (e.g. [6, 17, 27]). Examining user engagement around posting can identify changing needs and preferences of contributors, as well as indicate expectations for feedback from others. Practically, better understanding of contributors' behavior can help encourage posting, better support users at times of contribution, and may even be used to improve personalized recommendations.

We devise a within-subject, observational data analysis of de-identified log data of Facebook activity from a sample of 2.4 million people over a period of 9 days. In our design, we observe individuals' actions on Facebook around times of contribution (without any intervention) and another comparable activity, like liking or commenting on another's post. Specifically, we consider when an individual posts a piece of content, e.g. writes a post or posts a photo, and compare her activity around that time to a different time when she gives feedback on someone else's content. We use measures of activity such as site visits, number of stories read and number of stories interacted with in the 48 hours surrounding contribution, in order to learn about the relation between posting and contributors' behavior.

Our contributions are therefore:

- First large-scale evidence for within-subject differences in engagement around times of contribution, e.g. when posting content to Facebook rather than commenting on others' posts.
- Empirical evidence for an increase in site visits, reading more stories from friends and interacting more with friends in the 24 hours after posting.
- Potential design implications for better supporting contributors on social network sites.

To further motivate this study, we describe the theoretical frame-

work used to draw hypotheses about changes in contributors' behavior.

2. BACKGROUND

We build on theories from various fields to examine behavioral changes of contributors in SNS. These theories help us reason about the ways in which posting content can affect how individuals use Facebook, consume content, and interact with others on it. But first, we need to describe the motivating factors for contribution on SNS.

Previous research identified key motivating factors for participation in online communities, and gratifications contributors draw from it. For example, Dholakia et al. [10] identified five motivating factors for contribution online: purposive value (exchange of information), self-discovery (acquiring knowledge), entertainment, enhancing social status and maintaining relationships. Other studies [29, 32] examined the motivations for active participation on Wikipedia, finding similar motivations and gratifications. Preece and Shneiderman [31] describe contributors' recognition and ability to build reputation as a major motivating factor for social contribution. Several studies examined contribution to SNS, and Facebook in particular. Both Joinson et al. [19], and Papacharissi and Mendelson [30] provided evidence that Facebook contribution helps support expressive information sharing and maintaining relationships.

While previous research mostly relied on self-reported measures for studying *why* people contribute online, we focus in this work on the ways in which contribution may affect user behavior, using a large-scale dataset of contributors' actions that are free of any intervention.

2.1 Feedback Expectations and Site Activity

Feedback is a key component of any social exchange: it is important both for motivating contributions in the first place [5, 8, 21] and for evaluating social relationships over time [16, 25, 36]. Most, if not all, of the motivating factors for contribution identified by Dholakia et al. [10] depend on feedback from the online community, which suggests that contributors will expect some feedback. For example, *purposive value* is the value people derive from achieving a pre-determined purpose with the help of the community such as planning a trip or selling items. Similarly, if people post on Facebook to maintain relationships as suggested by previous research [19, 30] then it is reasonable that contributors expect responses. In anticipation of new interactions, contributors may visit Facebook more frequently after posting. We refer to site visits that are not initiated by a notification (e.g. email sent by Facebook) as *self-motivated site visits* and hypothesize that:

H1 *Following a post, self-motivated site visits will increase.*

2.2 Shifts in Content Consumption Patterns

Contrasting theoretical explanations can be argued for changes in consumption of content from others after posting. On the one hand, contributors already spent time crafting their message, which may directly compete with the limited amount of time they have to spend online after posting. On the other hand, contribution may take place at times when people are more free in the first place, and posting may be associated with a further increase in their consumption of content. The later argument is consistent with an account of participation taking place in a more active state [31] or aroused state in psychological terms, which was shown to be associated with increased levels of activity [14, 23, 34, 38].

At the same time, alertness or arousal may also mean more selective distribution of attention. Easterbrook hypothesized, based on

studies of cue utilization, that arousal would lead to narrowing of attention [12], a finding that was later verified in an eye movement experiment [26]. If the act of posting makes one more selective, it is feasible that contributors would focus more on content from friends, as opposed to pages or other broadcast sources that are less specific to them.

The fact that habitual time-passing behavior is a major motivation for social media use [19, 30] leads us to believe that contribution would not come at the expense of content consumption, but rather enhance and make it more selective. Therefore, our hypotheses for content consumption are:

H2.a *Following a post, contributors will consume more content.*

H2.b *Following a post, contributors will consume more content from friends.*

2.3 Interaction Rates and Reciprocity

Contributing content is likely to have an effect on subsequent interactions with others, but different factors may positively or negatively affect the overall rate of interactions over time. On the one hand, higher interaction rates after posting may occur due to greater time availability, more active state or reciprocity. On the other hand, fatigue or a fixed-quota for interactions may result in a lower interaction rate after posting. We describe each of these arguments next and consider how these factors may affect interaction rates jointly.

Two of the arguments presented before, regarding contribution happening at more flexible times and more active state, can also explain an increase in the rate of interactions. For example, if people post when they have more free time then they may continue to interact more with content after posting. If contributors are more active and selective, as suggested before, they may choose to interact more in general, and with friends in particular.

In addition, *reciprocity* as the social norm of returning a favor, can also lead to higher interaction rate with friends after posting. In the realm of computer-mediated communication, even simple one-way communications such as a like or short "composed communication" (as defined by Burke et al. [4]) bare value. Therefore, receiving feedback from friends on a post, perhaps similarly to receiving a gift, creates indebtedness and calls for reciprocation. Reciprocity in social exchanges can take one of two forms: direct or indirect (also known as generalized reciprocity) [24, 33]. Direct reciprocity in our settings implies that contributors would interact more with the friends who responded to their post, while indirect reciprocity suggests more interactions with friends in general. In both cases, reciprocity results in more interactions with friends after posting.

In contrast to the above theories, fatigue or a fixed-quota policy may explain a decrease in interaction rates after contribution. If contributors consume more content, as postulated in the previous section, they may experience fatigue over time and engage in fewer interactions. Similarly, if people have a fixed amount of interaction they can engage in, and more content is consumed, the rate of interaction would decrease. We believe that the additional amount of content consumed would be relatively small and thus neither fatigue nor interaction limits would be dominant in our case.

Therefore, our hypotheses are:

H3.a *Following a post, contributors are more likely to give feedback to friends.*

H3.b *Contributors are more likely to give feedback to those who responded to their content than to other friends.*

3. METHODS

To test the hypotheses listed above, we devised a quantitative,

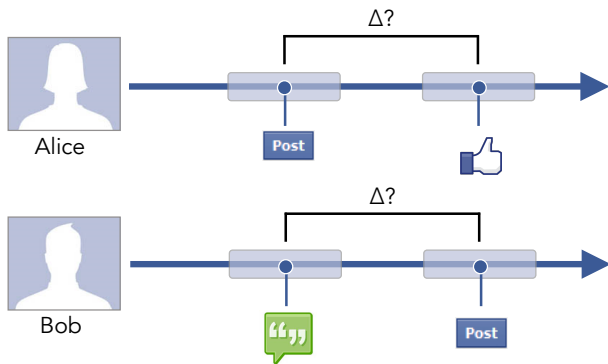


Figure 1: Research design: observational analysis comparing individuals’ activity in the 48 hours centered around either a contribution action C (e.g. posting a status update) or feedback action F (e.g. a like or a comment). We chose pairs of anchoring actions C and F that took place a week apart, with equal number of pairs having contribution followed by a feedback action (as in Alice’s case) and vice versa (as in Bob’s case).

within-subject, observational data analysis of Facebook activity logs. We wanted to isolate the effect of contribution as much as possible while controlling for other variables. To that end, we devised a comparative analysis of activity before and after posting on Facebook with a baseline of activity from the same individual at another time. We used feedback actions such as liking or commenting on someone else’s content as our baseline because those are similar times where people are on Facebook and actively engage with others. As we will show in the results section, there are no material differences in the context in which feedback and contribution actions take place. But first, we describe the dataset and the measures used in our analysis.

3.1 Dataset

Our dataset consists of the activity a sample of Facebook users engaged in, without any intervention, around two types of actions: contributing content (C), and providing feedback to others (F). The data were de-identified and content of posts was not analyzed. **Contribution** is defined as the act of posting content to Facebook, for example, an individual posting a status update, sharing of a link or uploading a photo. **Feedback** is defined as reacting to someone else’s content on Facebook: a like, a comment or re-share of others’ content. Identifying such pairs of actions from the same individual allows us to compare behavior around contribution with a baseline of activity around feedback.

Figure 1 illustrates the setup of our dataset. Each individual had one contribution action C and one feedback action F that happened on the same day-of-week, one week apart from each other, in any order, using Facebook’s web interface on a desktop device. In Figure 1, Alice posted a status update first and liked a friend’s photo a week later, while Bob commented on a friend’s post first and posted his own photo a week later. Both such sequences were included in our study.

We wanted to control, as much as possible, for external factors driving changes in individuals’ engagement other than contribution. In cases where individuals had multiple pairs of actions we randomly selected one pair in order to equally represent people in our dataset. We further balanced the dataset such that there is an equal number of pairs with contribution happening first (like Al-

ice) and feedback first (like Bob). We required both actions to have been performed on a mid-week workday (some time during the 24 hour span of Wednesday Pacific Standard time) to reduce bias from day-to-day variation. Our comparison of activity around actions included Facebook use through any device (mobile or not), but we required posting and feedback actions to have happened on Facebook’s web interface using a desktop device. We focused on contributions happening on the web interface in order to reduce bias stemming from differences in device capabilities, screen resolutions, and versioning, all of which vary more on mobile.

Given the selected actions C and F for each individual, we compared their behavior 24 hours before and after each action. We chose a window of 48 hours around actions in order to respect the natural and regular periodicity of human behavior. The matching of actions did not exclude the other type of action from occurring around that same time. For example, it is possible that a given individual posted content some time before or after the feedback action F selected for the analysis, and vice versa. Stricter filtering, requiring no contribution by the user around the time of the F action selected for analysis, would have resulted in a much smaller dataset, which would have been less representative of the general population of contributors. Our non-strict selection criteria are noisier, but provide a less biased lower bound on the actual effect size of contribution versus feedback.

Our selection criteria of two actions per contributor yield a sample of individuals who are slightly more active than a reference population (RP) who used Facebook’s web interface to post that week. The median person in our dataset is 37 years old (RP median=35), has 400 friends (RP median=344), has been using Facebook for 4.2 years (RP median=4.0), and has logged into Facebook 26.8 days out of the last 28 (RP mean = 24.3). Our sample is 55.7% female (RP: 51.9%).

In summary, our dataset includes C and F actions for 2.4 million individuals who posted content to Facebook or gave feedback to others using the web interface on two specific dates, February 11th and February 18th of 2015. The dataset is balanced in terms of the order in which contribution and feedback actions appear in it. Each individual included in the analysis has exactly one contribution action and one feedback action, where actions took place on the same day-of-week, interface and device. This set of individuals and actions is a sample of all users with actions that aligned with the selection criteria for those dates. Except for the analysis of self-motivated site visits that uses a subset of contributors, the rest of analysis uses the complete dataset.

3.2 Measures

We now turn to define the key measures used in our analysis.

Self-Motivated Site Visits

The measure of self-motivated site visits refers to the number of site visits that are not initiated by a notification, before any knowledge of feedback is available to contributors. We count site visits in terms of sessions, where each session is a sequence of actions of a logged-in user where actions are less than 30 minutes apart; if the individual was not active for 30 minutes, we count a subsequent action as a new session and a “site visit”¹.

When measuring site visits and sessions we want to ignore those visits that are due to offline notifications – users getting e-mail, SMS or mobile push notifications about Facebook activity that invites them to come back to the site. Therefore, we examined a

¹We chose relatively long (30 minutes) sessions in order to enhance resilience for short-term attention shifts. We experimented with shorter spans and found similar results.

subset of contributors for whom Facebook did not generate any of-line notifications in the two days preceding an action and the day following it. This subset of contributors did not receive notifications because they disabled offline notifications explicitly in their profile preferences or there was no activity that led to a notification being generated for them.

Stories Read

We measure content consumption by examining the number of News Feed stories read by contributors in the 24 hours preceding or following an action. Facebook’s News Feed is the landing page for people browsing to facebook.com or opening the mobile app, where content from friends and followed accounts is algorithmically ranked. A story is considered read if it was visible in the central portion of the user’s screen for a least two seconds. Note that we explicitly exclude stories that originated from the contributor herself as this may appear in her News Feed. In addition, our measure of stories read is not directly impacted by notifications because stories read as a result of clicking on a notification (on any platform) are logged separately and thus not counted towards our measure of stories read².

Interaction Rate

We define interaction rate as the proportion of likes or comments given per News Feed story read by the contributor in 24 hours before or after activity. Interaction rate is the portion of stories read from others (as defined above) that contributors liked or commented on directly from the News Feed. In other words, our measure of interaction rate excludes likes and comments that occur in other parts of Facebook such as Timeline or groups. Here again, any interactions with the contributor’s own content (reads, likes, comments) were excluded.

3.3 Statistical Analysis

For most of the analyses described below, we use Difference in Differences (DID) analysis in order to estimate the effect size of contribution while accounting for exogenous variation external to contribution. DID is a common statistical analysis technique used in observational data analysis to mimic a random assignment experimental design. DID estimates the effect of “receiving treatment” (in our case choosing to post) by controlling for a trend evident in the control group (feedback action in our case). In particular, DID analysis for our measure of stories read would be calculated as follows:

$$DID_{reads} = (R_C^{after} - R_C^{before}) - (R_F^{after} - R_F^{before}) \quad (1)$$

Where R is our measure of stories read in this case, and indices of after/before designate period relative to contribution C and feedback F actions for which the measure was computed. The underlying assumption in DID is that the treatment and control groups are comparable in every respect other than the assignment to treatment or control. Recall that we compare activities from the same individuals, day-of-week, interface, device, and comparable context as we will show in the next section. Therefore, we believe DID approach is particularly adequate for our settings since it highlights differences in engagement after contribution and contrasts them with the trend in engagement around comparable feedback action from the same person.

²Notifications may affect the number of stories read *indirectly* by encouraging people to visit their Facebook profile more often even if they do not directly follow the link on the notification. However, these changes in engagement are moderated by the individual and therefore an integral part of the behaviors we wish to study.

Two elements in the way we apply the DID help reduce selection bias and bias due to ordering effects. First, DID is often suspected for a selection bias in the assignment of individuals into treatment and control groups. In our analysis, however, both control and treatment groups include the *same* people, which eliminates individual differences between groups by design. Second, we reduce bias due to ordering effects by choosing a long gap in between the actions we examine (C and F) and balance the occurrence of actions in any particular order (contribution or feedback first). While we cannot rule out that one action may effect another action over a long period of time, our preliminary analysis suggest a diminishing difference in activity after 24 hours from posting or giving feedback. We use a much longer gap, of one week in between actions C and F , to further eliminate such interactions. In addition, the balanced order at which contribution and feedback actions appear in our dataset reduces the bias that observed effects are due to a one-time external event that affects only one of the conditions, or other time-based trends like increase in use over time.

All of our statistical tests were done using the standard technique of bootstrapping, with 10,000 replicas. We estimated means and 95% confidence intervals around them using the bootstrapped samples. Bootstrapping is more stable, asymptotically more accurate than estimates of confidence intervals based on a single empirical sample, and do not require normality assumptions [11]. We also favor bootstrapping over traditional paired t-tests since the latter tends to yield highly-significant p-values in all cases due to the sheer size of the sample (hundreds of thousands people in our smallest sample).

4. RESULTS

In this section we present the results of our comparative analysis of individuals’ behavior around contribution and feedback actions. Before we address the hypotheses described in the Background section, we first establish the validity of comparing activity around feedback and posting actions to each other.

4.1 Preliminary Analysis

We performed a series of descriptive and comparative analyses to better understand the data, and verify that there are no material differences between the contexts in which people performed the different actions (contribution and feedback).

A central question to our analysis is how active people are before and after different actions. Figure 2 addresses exactly this question by presenting on its top panel the percentage of people in our sample who were active on Facebook as a function of time, for 24 hours before and after each of the two actions that they took. The figure shows activity around contribution action (solid red line) and feedback action (dashed black line). Data points in the figure correspond to the percentage of the 2.4M people in our dataset that had an active session during each 20 minute time bin on the x-axis. For example, at the exact time of an action (time 0) all of the individuals in our dataset were active on Facebook since they either posted content or gave feedback. As a result, the plot spikes for both conditions at exactly 100%. The bottom panel shows the differences between the percent of active sessions around contribution and feedback (in other words, the difference between the solid red and dashed black lines on top). Figure 2 clearly shows that except for the 20 minutes immediately following an action, contributors are more active for several hours both before and after posting content compared to their activity around feedback at the same time frame. The only exception is the 20 minutes shortly after feedback where people are more likely to continue to engage with News Feed content rather than leave Facebook, as 7% of contributors do im-




Figure 2: Percent of active users (top panel) in the 48 hours around contribution and feedback actions, with the differences ($A_C - A_F$) visible on the bottom panel. The 95% confidence intervals were too small to be visible.

mediately after posting content.

The discontinuity observed at around zero in Figure 2 informed our decision to exclude the hour immediately following or preceding an action from our analysis. The fluctuation visible in the differences panel about an hour before the action and about an hour afterwards indicate short-term differences, probably stemming from the different sequence of user interactions at which feedback and contribution occur in. Therefore, for the rest of our analysis we use a window of 48 hours around an action, but exclude the 2 hours centered around an action.

Three interesting findings emerge from Figure 2 regarding the higher activity levels around contribution, and its return to baseline levels at the ± 24 hour period. First, we see that higher activity levels start as early as six hours before contribution and last more than 12 hours afterwards. The fact that contributors are more active even six hours before contribution is interesting and cannot be simply explained by the additional time necessary to conceive and articulate a post. The higher levels of activity after contribution are likely to be driven, at least in part, by notifications that contributors get due to feedback on their content. Below, when we address hypothesis H1, we show that notifications are not the only factor that explains higher level of user engagement after contribution. Second, Figure 2 shows uptick in activity in the 24 mark before and after each action. The increased activity indicates regular patterns in user activity and justifies the choice of 24 hours for analysis. Lastly, the diminishing differences at the ± 24 hours relative to actions demonstrate that the effect is largely dissolved in a day.

We further examined the data to make sure posting sessions are not fundamentally different than feedback sessions. We looked at the length of sessions and the position within a session where actions were recorded. As before, our definition for a session is a sequence of actions that are less than 30 minutes apart from each other. Figure 3 shows the average duration of contribution and feed-




Figure 3: Average duration of sessions around feedback/contribution actions, with the percent of time spent before the action.

back sessions. While a contribution session lasts more than 80 minutes on average, feedback sessions are significantly shorter, lasting only close to 68 minutes (95% confidence intervals were 20 seconds long, too short to be visible on the relevant scale). The long duration of sessions is likely to be a result of the long sessionization window used, but the relative position of actions within sessions are more robust. The figure shows that contributions are positioned similarly within a session, with 43.8% of the session time passing by before feedback occurs and 44.7% for contribution. A one percent increase in the relative position of contribution within session is equivalent to ~ 50 seconds, which is relatively small and could potentially be explained by the extra time required to compose a post.

We also verified that contribution and feedback actions occur at comparable time of day. For example, we wanted to make sure our dataset is not biased such that contribution takes place in the morning and feedback at night. By computing the difference in time of day for each pair of user actions, we find no statistically significant difference. The average difference in time of day is bound by a 95% confidence interval of $(-3.2, 3.0)$ minutes. No difference (difference of zero) is well within the 95% confidence interval. Therefore, we conclude that contribution and feedback actions occur at roughly the same time of day.

In summary, the preliminary analysis provided evidence for the adequacy of our comparative analysis of individuals' engagement in the 24 hours before and after contribution and feedback actions. This initial analysis informed our decision to exclude the hour right before and after an action for the rest of the analysis and established that contribution and feedback actions are positioned comparably within sessions and within the day.

4.2 Site Visits

We test hypothesis H1 about an increase in self-motivated site visits by conducting DID analysis on our measure of site visits. Recall that for this analysis, we wish to neutralize the effect of notification. To this end, we focus on a sub-sample of 150,000 people for whom Facebook did not send any offline notifications in the 48 hours preceding an action and 24 hours after. These people either chose not to receive offline notifications or there was no activity that led Facebook to generate a notification for them.

Figure 4 shows the average number of site visits for the same set of people before and after contribution and feedback actions. For instance, we see that in the 24 hours before feedback individuals had an average of 3.9 self-motivated site visits, while closer to 4 site visits after taking a feedback action (excluding the feedback/posting session itself). The dashed line designates the projected number of site visits, if the general trend apparent in the feedback condition occurred at times of contribution.

The evidence from Figure 4 is that the number of self-motivated site visits after contribution exceeds the projection by 0.11 site vis-

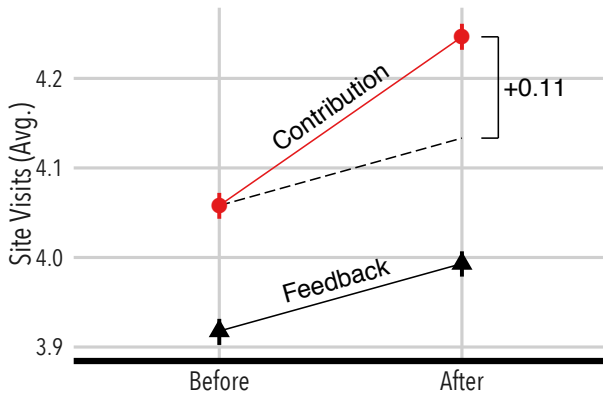


Figure 4: Difference in Differences analysis shows a significant increase in self-motivated site visits in 24 hours before/after activity (95% CIs). The dashed line designates the DID prediction for levels after contribution based on the trend evident in the Feedback condition. Square brackets highlight the significant increase in site visits of +0.11 on average.

its on average, the difference is statistically significant, and is in line with hypothesis H1. The figure shows that in the 24 hours after both feedback and contribution actions, people are visiting Facebook more often even without getting any offline notifications. On top of the projected increase in site visits from giving feedback, individuals visit Facebook 0.11 (+2.6%) more often on average when posting. These findings show that there is a small increase in site visits not stemming from notifications or from merely taking an action on the site.

4.3 Content Consumption

We now examine how contribution affects an individual’s attention to content. Hypotheses H2.a and H2.b postulate that contributors will consume more content overall and particularly more content from friends, respectively. We test these hypotheses using a DID analysis on the measure of stories read, counting stories viewed for at least two seconds in the central portion of the user screen. The measure of stories read will increase if people reading more pieces of content or decrease if they are skipping content.

Figure 5 presents the DID analysis of, separately, stories read from friends and stories read from other sources like Facebook Pages. While fewer stories are read on average after giving feedback (evident in the decreasing trend in black), the trend for contribution is positive for content from friends and neutral for pages. Similar patterns emerge when we do not distinguish between friend and page content – people read (on average) slightly fewer stories after engaging in feedback actions and about three more stories (+2.1%) after contribution. The number of stories read from pages also increases on average by 1.2 (+1.8%) compared to the DID projection, and are statistically significant as explained above. These findings support an increase in overall content consumption and consumption of friend content.

Interestingly, only the number of stories read from friends increases compared to pre-contribution levels, suggesting a shift in attention towards friends but not others. These results are consistent with the previous observation that contributors remain active for longer periods of time after contribution, but also indicate that the additional time is spent more selectively on friends’ content.

We performed further analysis to verify that the above changes in

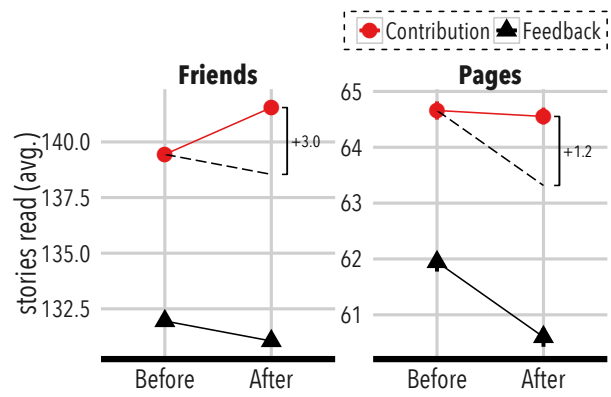


Figure 5: Difference in Differences in number of Newsfeed stories read in 24 hours before/after activity (95% CIs). Dashed lines designate the DID prediction for levels after contribution based on the trend evident in the Feedback condition. Square brackets highlight the significant increase in the number of stories read.

individual consumption habits do not simply stem from differences in the content available to people in the News Feed around feedback actions and contribution actions. As a crude measure of content availability, we test whether the distribution of content available from weak and strong ties changes before and after the *C* and *F* actions in our dataset. We use a measure of tie strength that is based on the frequency of past communication between two individuals and we simply associate the tie strength of the friend authoring the post with the content viewed by the contributor. We conducted DID analysis on the tie strength associated with content and found no significant difference. Therefore, we conclude that the content available for consumption around contribution is not significantly different than the content around feedback actions.

In conclusion, we find that contributors consume more content before posting, but increase consumption even further after posting, particularly of friends’ content. We rule out an explanation that those changes in consumption habits simply arise from News Feed ranking or other differences in the availability of content at different times.

4.4 Interaction Rates and Reciprocity

Previous sections established that contributors are more engaged around contribution and consume more content, even though the content itself remains the same. We now examine whether posting content affects individuals’ decisions to interact with others as postulated by hypotheses H3.a and H3.b.

Figure 6 shows our DID analysis of interaction rates with content from friends and pages. The bottom left panel, for example, shows that before posting, users comment on 0.74% of the stories they read from friends and that this rate significantly increases to 0.77% after posting. DID that were statistically significant are highlighted in the figure by square brackets as can be seen in the left two panels.

Several interesting findings should be noted about Figure 6. Across the board, the interaction rate before and after posting is significantly higher than the rate when simply giving feedback to others. After giving feedback, there is no significant change in the interaction rate or even a slight decrease compared to pre-feedback levels. In contrast, the interaction rate with friends after posting

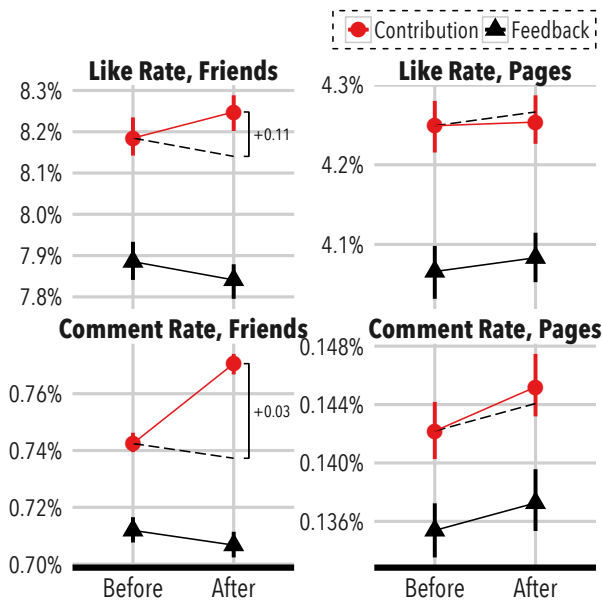


Figure 6: Difference in Differences in Liking and Commenting rates with friend/page content in 24 hours before/after activity (95% CIs). Dashed lines designate the DID prediction for levels after contribution based on the trend evident in the Feedback condition. Square brackets highlight the effect size when significant.

increases significantly for both likes (an absolute gain of +0.11%, which is a 1% gain relative to the “before” level) and comments (+0.03%, 4% gain). The changes in interaction rates are statistically significant, substantial, and even more interesting given that there are no significant changes in interaction rate for pages (right side of Figure 6). These findings provide supporting evidence for hypothesis H3.a about increase in interaction rate with friends, and provide counter evidence to the idea of a “fixed-quota” or decision fatigue over time.

Next, we provide a deeper examination of the interaction rate with friends to understand the role of reciprocity in these interactions. For example, consider an individual, Anna, who posted a status update on Facebook and later saw stories from two of her friends, Brian and Colin. Hypothesis H3.b suggests that if Colin gave feedback to Anna’s original post, she would be now more likely to comment on Colin’s post than on Brian’s. Of course, it is possible that Anna and Colin are simply more likely to interact with each other in general, for example, because they are closer friends. To control for this difference in relationship, we use Propensity Score Matching (PSM) with a score based on tie strength (as described in the previous section). For every person posting and friend who commented/liked their post (designated as indebted), we match an equally close friend who did not comment/like the person’s post (control). We verified that the average tie strength in the indebted and control groups is not significantly different. We can then compare the interaction rates of contributors with content viewed from the two groups, where the only difference between group is whether the friend previously responded to the contributor’s post or not.

Figure 7 shows liking and commenting rates for indebted and control groups of friends. On the left part of the figure, we see the

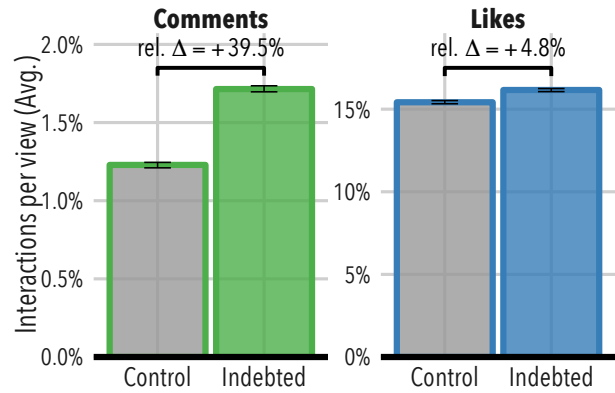


Figure 7: Commenting and Liking rates (95% CIs) on friends’ content who responded to the contributor’s post (indebted) or not (control), controlling for tie-strength.

rate in which contributors commented on content they saw, split to friends who responded to the contributor’s post (indebted condition) and equally close friends who *did not* respond to contribution (control group). As the figure shows, commenting on content from friends who responded to contribution (indebted condition) occurs at an average rate of 1.71 comments per 100 stories read (1.71%, solid green bar, second from the left in Figure 7). In contrast, contributors only commented on content from friends who did not respond to their contribution (control condition) at a rate of 1.22 comments per 100 stories read (1.22%, first bar from the left).

The relative change in commenting rate in the indebted condition is a large 39.5% increase over the control group and a more modest 4.8% increase for likes. These findings are highly significant, align well with the theory regarding direct reciprocity, and supportive of hypothesis H3.b. As a side note, notice that rates of interaction in Figure 7 are much higher than on the left side of Figure 6. This observation is reasonable since the interactions in Figure 7 are with friends who responded to contribution, which are more likely to be friends one frequently interacts with, and thus results in higher rates.

We note that the interaction rates increase also for friends who did not respond to contribution, but at smaller rates than those in Figure 6. This result is consistent with indirect (generalized) reciprocity as we described in the background section. However, more complex analysis is needed to substantiate indirect reciprocity in this case because it requires careful control for the activity of others in addition to the actions taken by the contributor herself.

In summary, we see an overall higher interaction rate around times of contribution, with further increase after posting, especially for friends and not others. We find that whether equally close friends respond to contribution or not affects the likelihood of future interactions with their content, resulting in substantially more likes and comments for friends who responded to contribution.

5. DISCUSSION

In this study we examined individuals’ behavior when posting to Facebook and found significant changes in engagement both before and after contribution. We discuss here why we think these shifts in individuals’ engagement occur and what SNS can do to better support contributors.

5.1 Contribution and Changes in Engagement

Higher Engagement Before Posting

A salient theme across all of our findings is that contribution is associated with more active engagement even before contribution takes place. These findings can be explained by external factors that influence both posting and engagement, or by higher engagement leading to contribution. External factors may include contribution taking place when people have more free time to spend on Facebook, being in a more active and alert state, or simply when people attend an event together with their friends. All of the above can increase engagement and be associated with posting. Alternatively, higher engagement can also lead, through various means, to contribution. For example, being exposed to interesting content from others can inspire or simply remind people to post.

The fact that posting is positioned similarly within session to feedback actions suggests that people often spend considerable amount of time on others' content before posting their own content. These findings are consistent with the notion of influence from Social Learning Theory [1], which posits that people learn by observing others and gradually behave more similarly to them, even without any external incentives. Whether increased engagement leads to posting remains an open question, with implications for newcomers [5] as well as contributors in general.

Higher Engagement After Posting

Our findings show that contributors are not only more engaged before posting, they also increase their engagement after posting at a higher rate than they do after feedback activities.

The result showing an increase in site visits after posting (without notifications) supports the idea of self-motivated changes in individuals' engagement. We believe that some of the additional site visits are motivated by anticipation of feedback and that similar changes occur for people who do get notifications. On platforms at the scale of Facebook, an effect of 2.6% increase in site visits translates into hundreds of thousands of additional site visits each day that are presumably motivated by anticipation of feedback.

Most consistent with all of our findings, both before and after contribution, is that posting is associated with a more active and alert state. These results interact with ideas from attention theories looking at how we allocate attention [12]. Key recent theories of attention deal with selection processes (what do people pay attention to) and vigilance (how do we sustain attention over time) [9].

Other alternative explanations for the increased engagement after posting are consistent with some of our findings, but not all of them. Some of this higher level of activity can simply be tied to contributors interacting with the responses on their post. However, it was not established until now that other activities on Facebook, unrelated to contribution, also rise. In addition, if people post when they have more time to spend on Facebook it is feasible that they will continue to engage even after posting. However, looser time constraints around contribution do not immediately explain the changes in selectivity of consumption and interactions with content. Similarly, attending an event with friends and posting about it on Facebook could explain some of the increases in interaction rates with friends, but not the persistently high levels of engagement with non-friend content. Reciprocity, as we will discuss in greater detail next, does not explain the high engagement levels before posting or with page content afterwards.

Contribution and Reciprocity

Once contribution is made and responses come in, it is reasonable that contributors will reciprocate, but the magnitude and speed at which it occurs is somewhat surprising. In the 24 hours after contribution, commenting rate on content from friends who responded to contribution increased close to 40% more than the control, compared to a more modest (but still substantial) 4.8% increase for likes over the control. While reciprocity is a well-documented and replicated phenomenon, this is the first time the immediacy of the effect is shown in social media settings and at large scale.

An important question is whether the reciprocity effect is deliberate. In other words, do contributors seek out opportunities to comment or like the content from those who gave feedback on their content? Or are they unconsciously inclined to reciprocate because they have positive feelings towards those who gave them feedback? In offline settings, a well-established result shows that we are more likely to like people who evaluate us positively [2, 35], or in other words, "we like those who like us." [15]³. These previous findings may suggest that individuals develop more positive feelings towards those who give them positive feedback, and as a result may be more inclined to like or comment on their content. Our working assumption is that both deliberate and more implicit mechanisms are in effect here, perhaps demonstrating a dual process mechanism that is known to apply in social settings [13].

Under the assumption that at least some of the feedback is due to a deliberate attempt at reciprocity, these findings are in line with the claims of the hyperpersonal model in Social Information Processing [37]. In particular, the model captures how interactions in CMC get amplified over face-to-face communication, which can then turn into greater indebtedness and reciprocity. This theory aligns well with the more substantial increase in comments versus likes; the different time investment and meaning for comments over likes has been well documented, and the fact that contributors choose to comment more than like content may indicate a greater sense of indebtedness on their part. These findings are in line with the changes in tie strength highlighted in [4] and the preference for "composed communications".

5.2 Limitations

While we attempted to carefully design our analysis and control for key factors, the study still has several limitations.

First, as a purely observational study our findings are only suggestive of the causal relations between posting and user engagement. We believe that posting does lead to an increase in overall activity and changes the composition of actions contributors take on the site. Similarly, we think that seeing more engaging content can encourage, inspire, or simply remind people to post their own content. Nevertheless, by merely observing the actions people take on Facebook we cannot definitely discern these causal explanations from other alternative explanations that were mentioned before.

Second, by focusing on aggregate measures of activity over a period of 24 hours we average out some of the behaviors that only occur at shorter time spans and lose the ordinal aspect of activity. For example, our measures are likely to smooth effects that happen on the next session immediately following a post, especially since we are excluding the one hour before and after posting.

Lastly, including in our analysis contributors who were active on Facebook at two different times a week apart introduces some selection bias. While we did work with a sample of millions of people, our methodology is not suitable for drawing inferences about

³As [15] shows, we even like those who positively evaluate *others* – "everybody likes a liker".

less active contributors and may not generalize for contributions on other SNS.

5.3 Future Work

Future studies could examine the role of feedback in modulating contributors' behavior and the time-range for these effects. While most feedback received on SNS is positive, even in sites with a weaker sense of identity and friendship than Facebook (see Cheng et al. [7] for details), the question remains as to how feedback affects behavior. Even more challenging is the fact that the effect of feedback is likely to depend heavily on contributors' expectations, which are subjective and not directly observable. A closer investigation can examine the temporal aspect of the behavioral changes we identified and try to link the short-term changes in engagement with long-term effects on relationships.

Other extensions of our work can investigate how engagement changes as a result of individual differences as well as differences in form and substance of the posted content. Different populations (e.g. women and men, young and old) engage differently with SNS [18, 20] and analyzing the effect for different sub-populations can reveal additional differences. Posted content may very well differ in form, style, content, effort and intent embedded in it, which all call for further exploration of their effect on contributors' behavior.

5.4 Design Implications

Our findings suggest a potential for designing adaptive systems that encourage social participation, help contributors focus on the content that is important to them, and recommend content based on the context of actions. First, the observation that contributors are more active six hours before posting opens possibilities for researchers to design nudges for contribution at times of high engagement and evaluate whether these are perceived as beneficial. Second, the importance of feedback from friends may call for refinement of user experiences around feedback interactions and rethinking how to surface these to contributors. Lastly, we demonstrated that individuals' engagement with content depends on the context in which it occurs (e.g. posting on Facebook), a finding that recommendation systems can use to differentially value explicit feedback from people. Further research is needed to better serve the naturally-changing needs, expectations and preferences of contributors.

5.5 Conclusions

In this work, we examined the short-term engagement of individuals when posting to Facebook and contrast it with their activity at another time when they give feedback to others. Our within-subject comparative analysis resulted in two major findings. First, we found the people are more active around posting actions than feedback actions for about six hours before posting and more than 12 hours afterwards. The deeper engagement happens both before and after the time of posting and across all the measures we examined: self-motivated site visits, stories read and interaction rates with content. Second, contributors not only start more engaged, but also further increase their engagement after posting at a higher rate than any other feedback action. Self-motivated site visits increase after posting as well as the consumption and interaction with friends' content, but not others.

We highlighted a few areas in which interface design can better support contributors, encourage social participation and possibly improve content ranking in recommendation systems. Taken together, our findings identify an important pattern of engagement that is consistent with key behavioral and social theories. It is pos-

sible that underlying all of these is a distinct cognitive state that is associated with contribution, greater desire for social connection and more willingness to engage with friends. However, we believe that additional evidence needs to accumulate before a more holistic theory could emerge, explaining individuals' engagement in the complex social context in which it is embedded.

6. REFERENCES

- [1] Albert Bandura. 1977. Social learning theory. *New York: General Learning Press* (1977).
- [2] Carl W. Beckman and Paul F. Secord. 1959. The effect of perceived liking on interpersonal attraction. *Human Relations* (1959).
- [3] danah m. boyd and Nicole B. Ellison. 2007. Social Network Sites: Definition, History, and Scholarship. *Journal of Computer-Mediated Communication* 13, 1 (2007), 210–230.
- [4] Moira Burke and Robert E. Kraut. 2014. Growing Closer on Facebook: Changes in Tie Strength Through Social Network Site Use. In *proceedings of the 32nd international conference on human factors in computing systems (CHI '14)*.
- [5] Moira Burke, Cameron Marlow, and Thomas Lento. 2009. Feed me: motivating newcomer contribution in social network sites. In *proceedings of the 27th international conference on human factors in computing systems (CHI '09)*.
- [6] Jaclyn Cameron and Nick Geidner. 2014. Something old, something new, something borrowed from something blue: Experiments on dual viewing TV and Twitter. *Journal of Broadcasting & Electronic Media* 58, 3 (2014), 400–419.
- [7] Justin Cheng, Cristian Danescu-Niculescu-Mizil, and Jure Leskovec. 2014. How Community Feedback Shapes User Behavior. In *Proceedings of the 8th international conference on weblogs and social media (ICWSM '14)*.
- [8] Coye Cheshire and Judd Antin. 2008. The social psychological effects of feedback on the production of Internet information pools. *Journal of Computer-Mediated Communication* 13, 3 (2008), 705–727.
- [9] Marvin M. Chun, Julie D. Golomb, and Nicholas B. Turk-Browne. 2011. A taxonomy of external and internal attention. *Annual review of psychology* 62 (2011), 73–101.
- [10] Utpal M. Dholakia, Richard P. Bagozzi, and Lisa K. Pearo. 2004. A social influence model of consumer participation in network-and small-group-based virtual communities. *International journal of research in marketing* 21, 3 (2004), 241–263.
- [11] Thomas J. DiCiccio and Bradley Efron. 1996. Bootstrap confidence intervals. *Statistical science* (1996), 189–212.
- [12] James A. Easterbrook. 1959. The effect of emotion on cue utilization and the organization of behavior. *Psychological review* 66, 3 (1959), 183–201.
- [13] Jonathan St. B.T. Evans. 2008. Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology* 59 (2008), 255–278.
- [14] Hans J. Eysenck. *The biological basis of personality*. Charles C. Thomas, Publisher, Chapter Activation, Arousal, and Emotion, 227–263.
- [15] Valerie S. Folkes and David O. Sears. 1977. Does everybody like a liker? *Journal of Experimental Social Psychology* 13, 6 (1977), 505–519.
- [16] Mark S. Granovetter. 1973. The strength of weak ties. *American journal of sociology* (1973), 1360–1380.

- [17] J. Brian Houston, Joshua Hawthorne, Matthew L. Spialek, Molly Greenwood, and Mitchell S. McKinney. 2013. Tweeting during presidential debates: Effect on candidate evaluations and debate attitudes. *Argumentation and Advocacy* 49, 4 (2013), 301.
- [18] David J. Hughes, Moss Rowe, Mark Batey, and Andrew Lee. 2012. A tale of two sites: Twitter vs. Facebook and the personality predictors of social media usage. *Computers in Human Behavior* 28, 2 (2012), 561–569.
- [19] Adam N. Joinson. 2008. Looking at, looking up or keeping up with people?: motives and use of facebook. In *proceedings of the 26th international conference on human factors in computing systems (CHI '08)*.
- [20] Funda Kivran-Swaine, Samuel Brody, and Mor Naaman. 2013. Effects of gender and tie strength on Twitter interactions. *First Monday* 18, 9 (2013).
- [21] Cliff Lampe and Erik Johnston. 2005. Follow the (slash) dot: effects of feedback on new members in an online community. In *Proceedings of the 13th international conference on supporting group work (GROUP '05)*.
- [22] Cliff Lampe, Rick Wash, Alcides Velasquez, and Elif Ozkaya. 2010. Motivations to participate in online communities. In *proceedings of the 28th international conference on human factors in computing systems (CHI '10)*.
- [23] Randy J. Larsen and David M. Buss. 2002. Personality psychology: domains of knowledge about human behavior. (2002).
- [24] Claude Lévi-Strauss. 1969. *The elementary structures of kinship*. Beacon Press.
- [25] Nan Lin. 1999. Building a network theory of social capital. *Connections* 22, 1 (1999), 28–51.
- [26] Elizabeth F. Loftus, Geoffrey R. Loftus, and Jane Messo. 1987. Some facts about "weapon focus". *Law and Human Behavior* 11, 1 (1987), 55.
- [27] Misa T. Maruyama, Scott P. Robertson, Sara K. Douglas, Bryan C. Semaan, and Heather A. Faucett. 2014. Hybrid media consumption: How tweeting during a televised political debate influences the vote decision. In *Proceedings of the 17th conference on computer supported cooperative work (CSCW '14)*.
- [28] Mor Naaman, Jeffrey Boase, and Chih-Hui Lai. 2010. Is it really about me?: message content in social awareness streams. In *Proceedings of the 13th conference on computer supported cooperative work (CSCW '10)*.
- [29] Oded Nov. 2006. What motivates wikipedians? *Communication of the ACM* 50, 11 (2006), 60–64.
- [30] Zizi Papacharissi and Andrew Mendelson. 2011. Toward a new (er) sociability: uses, gratifications and social capital on Facebook. In *Media perspectives for the 21st century*, Stylianos Papathanassopoulos (Ed.). Routledge, New York, 212–230.
- [31] Jennifer Preece and Ben Shneiderman. 2009. The reader-to-leader framework: Motivating technology-mediated social participation. *AIS Transactions on Human-Computer Interaction* 1, 1 (2009), 13–32.
- [32] Sheizaf Rafaeli, Tsahi Hayat, and Yaron Ariel. 2009. Knowledge building and motivations in Wikipedia: participation as "Ba". *Cyberculture and new media* (2009), 51–68.
- [33] Marshall D. Sahlins and Michael Banton. 1965. On the Sociology of Primitive Exchange in The Relevance of Models for Social Anthropology. (1965).
- [34] Tali Sharot and Elizabeth A. Phelps. 2004. How arousal modulates memory: disentangling the effects of attention and retention. *Cognitive, Affective, & Behavioral Neuroscience* 4, 3 (2004), 294–306.
- [35] J. Sidney Shrauger and Stephen C. Jones. 1968. Social validation and interpersonal evaluations. *Journal of Experimental Social Psychology* 4, 3 (1968), 315–323.
- [36] Dalmas A. Taylor and Irwin Altman. 1987. Communication in interpersonal relationships: social penetration processes. (1987).
- [37] Joseph B. Walther. 1996. Computer-Mediated Communication: Impersonal, Interpersonal, and Hyperpersonal Interaction. *Communication Research* 23, 1 (1996), 3–43.
- [38] Robert M. Yerkes and John D. Dodson. 1908. The relation of strength of stimulus to rapidity of habit-formation. *Journal of comparative neurology and psychology* 18, 5 (1908), 459–482.