

The effects of rewarding user engagement

- The case of Facebook apps

Jörg Claussen

Ifo Institute – Leibniz Institute
for Economic Research at the
University of Munich
Poschingerstr. 5
81679 Munich, Germany
claussen@ifo.de

Tobias Kretschmer

University of Munich and Ifo
Institute
Schackstr. 4/III,
80539 Munich, Germany
t.kretschmer@lmu.de

Philip Mayrhofer

CDTM, LMU & TU Munich
Arcisstr. 21
80333 Munich, Germany
pmayrhofer@lmu.de

Abstract:

We study the market for apps on Facebook, the dominant social networking platform, and make use of a rule change by Facebook by which highly engaging apps were rewarded with further opportunities to engage users. The rule change led to new applications with significantly higher user ratings being developed. Moreover, user ratings became more important drivers of app success. Other drivers of app success are also affected by the rule change; sheer network size became a less important driver for app success, update frequency benefitted apps more in staying successful, and active users of Facebook apps declined less rapidly with age. Our results show that social media channels do not necessarily have to be managed through hard exclusion of participants, but can also be steered through “softer” changes in reward and incentive systems.

Keywords: app markets, social media, platform management, Facebook.

JEL Classification: L1, L50, O33

1. Introduction

The way firms interact with their consumers has changed significantly with the advent of social media channels. In particular, communication channels based on links among consumers, e.g. word-of-mouth or (active or automated) recommendations are supported through use of social media. In the empirical context of Facebook in-site applications (or apps), we study how firms and users transform their behavior in response to a change in the effectiveness of social media channels. More specifically, Facebook started rewarding apps that engage users more by enabling them to send out more notifications. Based on this change, we ask the following questions: i) Will developers start publishing apps with higher user ratings? ii) Will app ratings become a more important driver for app success? iii) Are other drivers of app success affected by the change in the effectiveness of social media channels?

This study is closely linked to the literature on word-of-mouth. Most of the current literature on word-of-mouth deals with the causal identification of peer effects¹—among others by exploiting the structure of the underlying network (e.g. Bramoullé et al. 2009) or by using (field) experimental designs (e.g. Aral and Walker 2011). We investigate how changing the effectiveness of social media channels changes the behavior of firms and users. Our work is related to Aral and Walker (2011), who compare the effectiveness of active-personalized invitations and passive-broadcast notifications. In contrast to comparing the relative effectiveness of different communication channels however, we study the effects of changes in the effectiveness of communication channels in general. Given the significant effect of social media channels (Aral and Walker 2011), providers of these channels have to find the right balance between spam-like information overload and allowing enough information being sent out to attract and retain users. A further tradeoff exists between giving new channel participants an initial chance and rewarding established participants who already successfully engage users.

¹ The main challenges in identifying peer effects come from Manski's (1993) reflection problem and unobserved heterogeneity of consumers.

The increased importance of social media is reflected by a growing body of literature on this topic. For example, Susarla et al. (2012a, 2012b) examine interdependencies between network structure and the diffusion of user-generated content as well as word-of-mouth cascades on YouTube. Tucker (2011) uses a similar identification strategy to ours and analyzes how a change in Facebook's privacy rules affected the effectiveness of personalized advertising. Other papers using Facebook as their empirical setting deal with the impact of social influence on advertising success (Bakshy et al. 2012), the impact of Facebook communities on transactions away from the original site (Goh et al. 2012), or with the role of social networks in hiring decisions (Acquisti and Fong 2012).

We study the early stages of the app market for Facebook, the largest social network on the Internet and a platform provider for add-on programs. The market for Facebook apps was highly dynamic in the early stages, but suffered from a flood of low-quality apps, which was detrimental to user experience. Facebook undertook a number of changes in reaction to this. Hence, we ask if a change in the rules in February 2008 by which Facebook apps could attract and engage users changed developers' incentives to design engaging apps and ultimately more successful apps.

Managing a platform has the goal of maximizing its monetization opportunities. As revenues for platform owners are often generated through advertising or transaction-based charges, managing active user engagement (which in turn increases advertising effectiveness) is often at the core of platform management. This may be done through non-price instruments imposing rules and constraints, creating incentives and shaping demand and supply behavior (Boudreau and Hagiu 2009). Hagiu (2011) identifies a quality/quantity tradeoff, since both higher quality and quantity are attractive to consumers, but higher quality presents an entry barrier to complementary goods providers, thus reducing quantity. Relatedly, Casadesus-Masanell and Halaburda (2011) argue that platform owners may limit the number of apps on a platform to realize larger benefits from app-level direct network effects. We contribute to

this literature by showing that platforms can be managed through “soft” quality incentives and hard exclusion of low-quality participants may not be necessary.² We also argue and demonstrate empirically that a change in “soft” incentives changes other drivers of app success.

We observe the usage of Facebook apps between September 2007 and June 2008.³ On Facebook, the amount of information an app can send out to users critically influences an app’s success. In February 2008, Facebook implemented a rule change regarding the amount of notifications apps could send out: before February 2008, all apps could send out the same amount of messages per user, while thereafter the amount of notifications permitted was determined by how frequently these messages were clicked on, a useful proxy for an app’s ability to attract and retain users. This increased incentives for producing highly engaging apps and punished apps that sent out information deemed useless or even annoying by users. To isolate the effects of this change, we focus on a twenty-week time window around the change. We use this change (assumed to be endogenous to the platform operator but exogenous for app developers) to analyze how potential drivers of app success changed in response. This natural experiment-like change in the effectiveness of social media channels therefore allows for a similar identification as for field experiments (Aral and Walker 2011, Goldfarb and Tucker 2011, Animesh et al. 2011, Tucker and Zhang 2011, Tucker 2011).

We use a rich, longitudinal data set on 7,784 Facebook apps. This setting is useful for our purposes for several reasons. First, we have data on apps soon after the launch of the platform, which lets us examine the dynamics of a nascent and dynamic market. Second, our complete listing of apps on the platform avoids selection and survivor biases. Third, Facebook is one of the largest and most successful platforms for apps, making it relevant for the entire industry.

² Evaluating quality of platform participants and excluding low-quality apps is costly and prone to error. Circumstantial evidence for this is given by regular debates revolving around apps rejected on Apple’s market for iOS apps. Conversely, not imposing any quality restrictions may lead to a flooding of low-quality apps as observed in the 1983 Atari shock (Coughlan 2004).

³ The app platform was opened in May 2007.

We observe a strong increase in users' ratings of apps once the rule change is implemented, which lends strong support for the intuition that the rule increased developers' incentives to develop more engaging applications. In a second step, we analyze how app success, measured as the number of an app's active users, is influenced by the rule change. Estimating fixed-effect OLS models we find that app rating matters more for the number of active users after the change, in line with Facebook's stated goals of the change. Conversely, we find that the rule change led to the number of installations becoming less important for app success. The frequency with which apps are updated (a proxy for the degree to which an app is managed and maintained) gains in importance as a driver for app success. Further, while the number of active users declines as apps become older the decline is less severe after the rule change, which implies that the intervention was successful in keeping adopters engaged for longer.

2. Facebook

2.1. Apps on Facebook

Facebook is the largest social networking website (other examples are Google+ or LinkedIn). Consumers use social networking services to interact with friends, family members, and increasingly business partners. Core components include personal mini-homepages with which a user creates a digital representation of him-/herself (Boyd and Ellison 2007) as well as different means to communicate (personal messages, boards, chats) and to exchange different media.⁴ Facebook is the largest and fastest-growing social network with over 900 million active users, of which 80% are outside the United States (as of March 2012).⁵ However, direct social networking functions are just the core product of Facebook; other social media functions have been added subsequently. In May 2007, Facebook launched an app platform which allows third parties to develop software that integrates into the social network and adds an additional dimension to Facebook usage not covered by Facebook's core

⁴ For example, Facebook is the largest online photo sharing utility.

⁵ Source: <http://newsroom.fb.com/content/default.aspx?NewsArealid=22>, accessed July 23, 2012.

components. In May 2008, one year after the platform launched, more than 30,000 apps had been developed, with more than 900 million installations in total.

2.2. Entry of app developers

Facebook does not control the market for third-party apps. To generate indirect network effects however, platform providers like Facebook want to encourage a wide variety of apps and experimentation in parallel (Church and Gandal 2000, Boudreau et al. 2008). Hence, Facebook provides developers with a set of tools that decrease development costs and lower entry barriers. This triggers high entry rates both from de-novo entrants as well as from developers with multiple apps. It also affects both users' experience and developers' incentives. On the one hand, a large variety of apps presents opportunities for consumers to discover and adopt apps (Oestreicher-Singer and Sundararajan 2006, Hervas-Drane 2010). On the other hand, high rates of entry could result in particularly fierce competition, which in turn would diminish profits and incentives around the platform (Boudreau 2008).

Facebook encouraged wide entry by developers. The company offered strategic subsidies to third-party developers (Shapiro and Varian 1998) by providing open and well-documented app programming interfaces, multiple development languages, free test facilities, as well as support for developers through developer forums and conferences. Facebook also imposes minimal requirements for apps to be listed in the official directory and it does not "police" or discourage developers imitating or producing "copy-cats" of existing apps.

2.3. Adoption and usage of apps

Having a large variety of apps has consequences for consumers' product search and adoption. On Facebook, app discovery, adoption and usage depends strongly on a user's social context, encouraging developers to build apps designed to intensify social interactions (Boudreau and Hagiu 2009). Users are passively influenced through the visibility of usage patterns such as reviews, ratings or matching

mechanisms (Oestreicher-Singer and Sundararajan 2006, Hervas-Drane 2010). Active forms of social influence take the form of recommendations which are conveyed directly via predominantly digital or online word-of-mouth processes (Katz and Lazarsfeld 1955). IS and Marketing scholars have examined the conditions under which consumers are likely to rely on others' opinions in their purchase decisions, the motives for people to spread the word about a product, and the variation in strength of influence people have on their peers in word-of-mouth communications (Dellarocas 2003, Phelps et al. 2005, Bampo et al. 2008, Agarwal et al. 2009). It is widely accepted that in such contexts bandwagon processes—positive feedback loops where adoption decisions by some increase incentives or pressures to adopt for others—are common (Katz and Shapiro 1985, 1986, Abrahamson and Rosenkopf 1993).⁶

In Facebook's case, the vast majority of installations take place through direct recommendation (word-of-mouth) and automated notifications by Facebook (Ermecke et al. 2009)⁷. That is, in contrast to other platforms where users actively enter "app stores" and search for particular apps or specific functionalities, Facebook apps are predominantly installed in response to receiving a direct recommendation or a notification. This difference matters as the spread and ultimately the success of apps strongly depends on the effectiveness of social media channels. Alternative marketing channels play a much less important role in our setting.

3. Facebook's rule change

Facebook users adopt apps through two main channels. First, users of an app can directly invite friends who are not currently users of the app (invites). Second, Facebook users get regular updates on friends'

⁶ Another feature relates to the costs that users incur in installing and using apps. Due to the dominant business model of indirect monetization, the vast majority of apps are free to use. Also, due to technical and design features, users can install and use multiple apps in parallel, thus "multi-home" (Rochet and Tirole 2003).

⁷ Less than 0.5% of app installations were initiated from Facebook's app directory (<http://developers.facebook.com/blog/post/523>, accessed July 23, 2012).

activities from the built-in “News Feed”. These updates are messages sent by apps to the news feed and signal a friend’s activity in this particular app (notifications).

Both channels have been modified by Facebook. In the very first phase of the app platform (from launch in May to August 2007) invites and notifications could be sent almost without restrictions. Many app developers used this to “spam” their users’ friends. In September 2007 Facebook imposed a set of restrictions (the number of invites and notifications sent out on behalf of a user was limited). In the following months these rules remained unchanged.⁸

However, after a period of steady growth, on February 6th 2008 Facebook announced a rule change such that notifications would be allocated based on user feedback. Apps whose users react more heavily to notifications that are sent out (a measure for relevance of the notifications) would be able to send out more notifications. One week later,⁹ feedback allocation was launched for notifications,¹⁰ Facebook motivated this change by the expectation that the new system “provide[s] users with more compelling notifications and fewer notifications that they are likely to ignore or consider spam” (Figure 1). Further, they “hope this change incentivizes developers to improve the quality of their notifications”.

----- INSERT FIGURE 1 HERE -----

We want to understand how this rule change affected Facebook’s market for apps. Did the rule change lead to the expected increased quality levels of apps? And was increased quality from then on a more effective driver of app usage? Finally, how were other drivers of app usage affected by the rule change?

⁸ To the best of our knowledge based on the official announcements of Facebook to its developers.

⁹ I.e. in week 7 of 2008.

¹⁰ This same change was also implemented later this week for requests and invites. (<http://developers.facebook.com/blog/post/83/>, accessed July 23, 2012). As notifications also have an inherent invite-role as they not only retain but also attract new users, we do not distinguish in this paper between notifications and invitations but just treat notifications as a general term for communication flowing through the social media channel that can attract new users and retain existing ones.

What could have motivated Facebook to initiate these changes? If app quality had always been Facebook's main goal it would be surprising to see no "hard" quality vetting process at the start of the app market and if that was not feasible: Facebook initially did not even have a rule rewarding (notification and/or app) quality directly, but rather sheer size by fixing the number of notifications per user per day. While we cannot infer Facebook's aims (apart from the ones stated in the announcement) with certainty, the strategic subsidies to developers and the ease of getting apps listed suggests that inviting entry may have been on top of Facebook's agenda in the early stages of the apps market. This was reinforced through highly publicized success stories which created a gold rush among developers. For example, the music app iLike grew to several million users within days.¹¹ Within weeks, several thousand app developers had signed up for access credentials to the platform and had started to launch a wide variety of apps. For users, the myriad apps launched early on helped "educate" them about installing add-on apps to make the service fresh and exciting. Users learned quickly. Through invites and a flood of notifications in their news feed, the vast majority of users had installed at least one app within weeks. Many users installed dozens of apps at the same time (multi-homing is comparatively costless here), sometimes even several with largely identical functionality (e.g. within the first month there were several enhanced "walls" that allowed posting and exchanging multi-media items).

After the initial enthusiasm, user sentiment towards apps changed. With a larger installed base of apps and the increasing professionalization of developers in terms of exploiting the opportunities to use the "viral channels", the volume of notifications and invites grew exponentially. Users became annoyed by constant updates about their friend's activities and apps. For both Facebook as the platform operator and the developers this threatened to lead to adverse effects because instead of adopting and using apps, users would start ignoring notifications and requests.

¹¹ See <http://mashable.com/2007/06/11/ilike-facebook-app-success/> (accessed July 23, 2012).

Facebook's rule change came precisely during the time when notifications became annoying in the eyes of users.¹² While the change did not increase entry barriers as such (development kits were still free, developers blogs were still supported etc.), it became harder for low-quality apps to succeed because it was difficult to attract a significant number of (active) users through spam-like notifications. Quality, as announced by Facebook itself, was rewarded more directly. Our econometric analysis therefore looks at the actual effects to assess how this rule change affected the dynamics of apps and their usage.

4. The economics of Facebook apps

4.1. A model of app monetization

When Facebook launched its platform for third-party apps in May 2007, in addition to the challenge of integrating their apps in Facebook's services, developers had a profit motive. Given there is no fee for users to subscribe to Facebook or to install any of the apps, developers looked for other ways of monetization. Facebook left it open to developers to monetize their app pages through advertising or other "in-app" transactions. Facebook also did not impose restrictions on the form of advertising. The most common form of advertising is next to the website's content and core functionality. Facebook also did not take a share of "in-app" transaction revenues at first, leaving developers to capitalize on those.¹³

Facebook's objectives are largely aligned with the objectives of their third-party developers and rely on capitalizing their active user base. Revenues come from selling advertising space to brands, advertisers or Facebook apps that target specific users. Next to each app's canvas page (the space allocated to an app), Facebook can place its own advertising. Consequently, the more users engage with apps, the more page impressions or time Facebook can sell to advertisers. Revenues that can be realized are directly

¹² See for example <http://www.techdirt.com/articles/20071130/110033.shtml> (accessed July 23, 2012).

¹³ Due to the (open) installation process and the lack of a payment system, Facebook could not take a revenue cut from developers without further development. In contrast, Apple takes a 30% cut from all sales in its iTunes store.

determined by the number of active users of the platform and apps.¹⁴ Thus, growing the platform through apps and keeping existing users active (thereby generating transactions or looking at and clicking on ads) is one of Facebook's most important objectives.

Developers' monetization opportunities predominantly depend on the amount of time users spend on an app site. Hence, developers will want to develop apps that engage users and keep them on the site for longer. Figure 2 presents a causal chain that governs the profitability of an app. i) Developers exert effort to increase app quality as perceived by users (we will discuss the meaning of quality for Facebook apps in the following subsection), ii) apps with higher perceived quality are more successful in engaging users, and iii) more engaged users translate into higher advertising revenue and hence profits.

----- INSERT FIGURE 2 HERE -----

In our empirical setting, we do not observe the first and last parts of this chain. However, given the quality of an app as perceived by users is likely to depend at least to some extent on a developer's deliberate efforts and that monetization opportunities for any ad-based business model depend strongly on the level of engagement with the base product, observing perceived quality and app success provides us with a meaningful snapshot of the economics of app monetization.

We reiterate our research questions and illustrate how they relate to the process outlined above. Our first question asks for the effect of rewarding user engagement on the app rating by users. This captures the deliberate efforts by developers to design apps that appeal to and engage users. Our second question captures the effect this user rating has on app success as measured by the number of active users, a meaningful proxy for overall user engagement. Our third question then asks if other

¹⁴ Gnyawali et al. (2010) show a performance-increasing effect of opening a platform to third-party codevelopment.

determinants of user engagement have been affected by the rule change, especially to determine if app rating has increased in importance relative to other potential drivers of user engagement.

In the following subsection, we discuss the two measures we observe as proxies for app quality respectively app success; app rating and the number of active users.

4.2. App rating as a proxy for quality

The stated aim of Facebook's rule change was to reward high-quality apps, incentivizing developer efforts to design such apps. Although we do not observe developers' efforts, we do observe the apps' user ratings. Given the monetization chain introduced above, an app will likely be considered high-quality by users if it keeps their interest and engagement. Hence, perceived app quality will be closely linked to the ability to engage users and ultimately result in monetization opportunities for developers.

Note that by using app rating as a proxy for quality, we focus on quality from the user's perspective. This implies that this may be due to improved "engineering quality", i.e. a vertical dimension that refers to better programming and technical quality, as well as a horizontal element referring to better matching of app features to a group of users. Important for our purposes is that developers will take actions aimed at increasing user perception of quality (i.e. attractiveness) and consequently user engagement. Another assumption is that the relevance of notifications is correlated with the extent of user engagement in general. While we do not observe the success rate of notifications directly, our estimations in Section 7 investigate this potential link empirically: Did the rule change (which, in its most narrow interpretation, affects developers' incentives to design engaging notifications) change the number of active users (which measures user engagement and consequently commercial potential)?

4.3. Number of active users as a proxy for app success

We study how Facebook's rule change affected app success. Finding the most relevant measure for a free-to-install app is not trivial. Conventional measures like profits or revenues are either not meaningful

(as far as direct sales revenues are concerned) or not obtainable (as far as advertising revenues are concerned). The success of free software or goods more generally is therefore often measured by the number of users or installations. However, this may also be misleading for at least two reasons: First, the number of installations (typically measured by the number of installation packs downloaded from the app's website) may double-count adopters if they repeatedly install software on different hardware or if they install upgrades to the basic version they already have. Hence, the number of installations may overstate the actual number of active copies of a piece of software – especially if it is frequently updated, requiring repeated installations. Second, even if the number of technically active copies is accurately measured, not all potential users will use it regularly. This is particularly relevant if a good comprises access and use elements, and most revenues arise from use, not access (Grajek and Kretschmer 2009). For Facebook apps this would seem especially true as access is free (installing an app is costless) and only active users will notice the advertisements from which the app derives its revenues. Hence, rather than the total count of installations, we use the number of daily active users as a metric for app success, which ultimately will translate to advertising revenues (unobservable to us).

4.4. Drivers of app success

We also study how Facebook's rule change may have affected the drivers of app success (measured by the number of active users) in the market for apps.

App rating

Facebook's press release explicitly states the intention that "this change incentivizes developers to improve the quality of their notifications" and "to reward compelling notifications" (Figure 1). That is, notifications and users' app experience are linked both through the new invitation process and the retention effect of in-app notifications. Highly rated apps are more likely to be installed if a user receives an invitation, and they will be used more intensively if the notifications generated are meaningful. The

notification process prior to the change introduced noise in this process by letting all apps issue the same number of notifications per user. Post-change, more successful apps could send out more notifications, leading to higher user engagement.

Update Activity

The frequency of updates to an app's "About" page is a choice variable by app developers. As Facebook apps usually have fairly limited functionality that cannot be changed post-introduction, an app can gradually become less attractive for an individual who has already installed the app. However, user interest can be retained if an app's "About" page in the directory is updated regularly. Apps that are actively managed and updated regularly are therefore expected to better retain and engage their customers. Facebook's rule change aimed at improving user experience by making notifications more relevant, and notifications are more likely to be relevant if an app is managed proactively, thus leading to higher user engagement. Note that updates do not trigger a notification to users but can only be observed by the user when using the app again or when visiting the app's "About" page in the directory.

Installed base effects

The installed base of users of an app can proxy for the number of active users through network effects as well as the composition of users at different points in time (Cabral 2006, Grajek and Kretschmer 2009). If there are app-wide network or peer effects, more installations should result in disproportionately more active users (as engaging with the app becomes more attractive with more users). However, for local network effects, i.e. if an app only becomes more attractive the more friends of a user have installed it, we do not expect a disproportionate effect of an app's (global) installed base on the number of active users. A countervailing mechanism of installed base is that the composition of users changes over time. With heterogeneous adopters and high-intensity users adopting early, the number of active users is expected to increase less than proportionally with a growing installed base

(Cabral 2006). Absent micro-(user-app-level) data, we cannot disentangle these effects empirically, but we can compare their relative strength by observing the net effect (Grajek and Kretschmer 2009).

How did Facebook's rule change affect the role of installed base on the number of active users? As we observe the net effect of user composition and network effects, the influence on both individual forces will make up the change in the net effect. User composition is unlikely to be affected strongly by the rule change as it affects the supply side (app developers), but not the demand side (app users). In contrast, the rule change reduced the advantage of apps with a large installed base since notifications were not proportional to the number of installations but depended on the relevance of notifications.

App age

App age, i.e. the time since which the app has been launched, also drives the number of active users. Older apps are expected to be used less intensively as usage follows a fad, i.e. users are only interested in the app for a short time. One goal of Facebook's rule change was to reward more compelling apps, so that we expect apps to engage and retain their users for longer after the rule change.

Portfolio effects

Portfolio effects matter in cultural industries in which artists or developers create several products. For example, Hendricks and Sorensen (2009) find that spillovers between artists' albums exist. Similarly, on Facebook most developers have a portfolio of (often functionally similar) apps. This may lead to users splitting their time across different apps, or to users developing a taste for a specific developer's apps. The net effect of belonging to a large portfolio is therefore an empirical question. It is interesting to speculate how Facebook's rule change may have changed the role of belonging to such a portfolio. Post-intervention, "new" apps had to attract users by "earning" credit through successful past notifications. As younger apps had less of a history to fall back on, alternative promotion channels mattered more

both in attracting new users and in keeping existing users' interest, so that access to an alternative channel in cross-promotion through developers' app portfolios may become more relevant.

5. Data

We use a unique dataset from Facebook's public directory of apps which included all apps available on the Facebook platform.¹⁵ All app-specific "About"-pages in this directory have been crawled and parsed daily to extract the variables described below.

Even though our data covers the period of September 1, 2007 to June 30, 2008, we focus on a period of twenty weeks around the rule change for most of our analysis. Our observation period falls in the early phase¹⁶ of the app platform and was characterized by strong growth in terms of users of Facebook's service as well as the number of apps and their users. The number of apps on the platform grew immensely from around 2,000 in early September 2007 to over 18,000 in June 2008.

We obtained records for 18,552 apps, of which 7,784 were active and contained all variables in the 20-week time window around the rule change. The records include data on an app's entry to the platform, its usage by Facebook members, its developer and finally an assignment to certain categories. Further, we computed a number of measures by observing changes made to the directory page as well as from clustering apps by developer name.

Our data is particularly well-suited for our analysis for several reasons. First, we have precise measures for an app's active users. Facebook continuously reports how many users have interacted with the app within the previous 24 hours. It also specifies the number of all users who have installed the app at any given time. Hence, we observe both an app's installed base and the number of active users. Second, the number of active users indicates the potential for economic success, as explained in Section 4.1. Third,

¹⁵ Facebook's app directory has been suspended in July 2011 (as announced on <http://developers.facebook.com/blog/post/523/>, accessed July 23, 2012).

¹⁶ Facebook's platform for apps was launched on May 24, 2007.

our data mitigates selection problems originating from minimum size thresholds or launched, but failed products. Developer entry to the platform is frequent due to low entry barriers. Finally, the dataset includes apps that were successful and apps that never reached a meaningful user base. Since data is recorded from the first day an app appears in the directory, information is available independent of the app's subsequent success. This is rare for studies on Internet-based industries where observing entry is often hard due to poor documentation of the early history of a category or firm. Published accounts on the entities often only appear upon them reaching a "threshold scale" (Eisenmann 2006, p. 1193).

Variable definitions, summary statistics, and pairwise correlations are given in Table 1 to 3.

----- INSERT TABLES 1-3 HERE -----

App quality ($AppRating_i$)

Users can rate apps on a five-point Likert scale. The variable $AppRating_i$ captures the average user rating of app i . We can only construct a time-invariant measure of app rating as Facebook's app directory did not report user ratings before February 2008. Facebook users were informed about an app's rating when installing it through the app directory, but not when installing it through other channels (notifications, invites, ads). As Facebook reported that less than 0.5% of app installations originated from Facebook's app directory and even shut down the app directory because of being so unsuccessful¹⁷, this measure is basically unobserved by users. This makes it unlikely that the introduction of app ratings is an alternative driver of the observed changes.

¹⁷ See <http://developers.facebook.com/blog/post/523>, accessed July 23, 2012.

App Success ($ActiveUsers_{it}$)

Our measure of an app's success is app i 's active user base in week t measured as the average number of daily active users.¹⁸ Hence, we observe the number of users of an app's installed base of users that use the app on a given day and form the weekly average.¹⁹ Given the skewness in $ActiveUsers_{it}$, we use the variable's logarithm for our regressions.

Rule change ($RuleChange_t$)

As discussed in Section 3, Facebook changed the rules in how far apps can send out notifications in February 2008. We construct a dummy variable $RuleChange_t$ which takes a value of zero before the change (until the sixth week of 2008) and a value of one thereafter.

Update activity ($UpdateFreq_{it}$)

We observe several events on the app's "About" page indicating an update of the underlying app: we check if the name or one of the descriptive elements of an app has changed, if the screenshot was updated, and if the app category has changed. For each point in time, we calculate an app's update intensity $UpdateFreq_{it}$ as the cumulative number of updates divided by app age ($WeeksActive_{it}$).²⁰

Installed base ($NumUsers_{it}$)

At the app level, we observe the number of Facebook users that have installed an app on their profile page. As the number of installations of an app is highly skewed, we use the natural logarithm of $NumUsers_{it}$.

¹⁸ The exact way how Facebook measures the number of active users is described in this entry on Facebook's developer blog: <http://developers.facebook.com/blog/post/30/> (accessed July 23, 2012).

¹⁹ All other time-dependent variables are also observed on a daily basis and then aggregated up to the weekly level. We aggregate data from daily to weekly levels to average out weekday-dependent differences in the number of active users.

²⁰ Calculating update activity by discounting older updates with weekly discount rates of 10% to 30% leads to qualitatively unchanged results (available from the authors).

App age ($WeeksActive_{it}$)

We measure the age of an app as the weeks since the app has appeared in the app directory for the first time. We use the natural logarithm of $WeeksActive_{it}$ to allow for a non-linear decay in usage over time.

Firm portfolio ($NumSisterApps_{it}$)

As developers can release multiple apps, we measure a firm's portfolio in the form of sister apps. $NumSisterApps_{it}$ is the number of an app i 's sister apps at time t .

6. How did the rule change affect app ratings?

6.1. Descriptive evidence

Going back to the causal chain introduced in Section 4, we first study whether Facebook's rule change led to increased app ratings. Given that the rule change is considered exogenous for app developers, we can use this quasi-experimental setting to draw some inference from changes in average levels of app rating once the rule change is enacted. Figure 3 plots average quality ratings of Facebook apps against their launch date.²¹ We see that apps launched after the rule change (in week 7 of 2008) immediately achieved a significantly higher average rating of around 0.4 points.²² This is strong suggestive evidence that the rule change resulted in higher app ratings. One might wonder why the changes in rating levels took place immediately after the rule change and did not develop slowly over time. In this stage of the app market, apps were fairly simple and easy to program. This might explain why app quality changed immediately following the rule change although it has only been announced several days earlier.

----- INSERT FIGURE 3 HERE -----

²¹ As app ratings are only observed in cross-section, we cannot observe how the rating of a particular application changes over time.

²² Measured on a five-point Likert scale.

6.2. Econometric evidence

Although the rule change came as an exogenous shock to the app developers, part of the observed effect could also be explained by changes in the composition of apps. If the rule change rendered one type of apps with inherently lower ratings unattractive and made app developers move to other types of apps which achieve higher ratings with the same amount of effort, our results could also be driven by these composition effects. To rule out this alternative explanation, we run a simple OLS regression in which we explain the rating of an app by the rule change, but control for the app's category²³ as well as for potential portfolio effects created by sister applications:

$$AppRating_i = \beta_1 RuleChange_t + Category_{i,1} + Category_{i,2} + \beta_2 \ln(NumSisterApps_{it}) + u_{it}$$

Results for this regression are reported in Table 4. In specification (4-1) we simply include a before/after dummy for the rule change, while we create a set of ten dummies in specification (4-2) to recreate the econometric equivalent to Figure 3. Coefficients can be directly interpreted as changes in app ratings against the base category of the first week in our sample, e.g. specification (4-1) suggests that the policy change led to apps with 0.346 higher ratings. Both specifications clearly support our initial findings and we thus exclude the alternative explanation of rating levels being changed by changes in app types.

----- INSERT TABLE 4 HERE -----

6.3. Robustness checks

In a next step we run two sets of additional robustness checks. First, we want to rule out the possibility that our results just pick up an ongoing upward trend in the level of app ratings. If app developers just become better over time in programming apps with high app ratings, our results could pick this up. Similarly, if old apps systematically command lower ratings due to self-selection for early reviews (Li and

²³ Each app is assigned to up to two out of 22 categories. Example categories are “just for fun”, “photo”, “chat”, “dating”, or “gaming”.

Hitt 2008), this would create an upward time trend as younger apps are rated higher. We run an additional regression in which we include a linear week variable in addition to the dummy capturing the rule change. Results are given in the Appendix (Table A.1, specification A1-1) and show no significant continuous time trend, making the alternative explanation of picking up a general upward trend even less likely. As a second robustness check, we enlarge the observed time period around the rule change, first moving to a time window of +/- 10 weeks and then to the whole observed period from 09/2007 to 06/2008. Results are again given in the Appendix (Table A.1, specifications A1-2 and A1-3) and show that our results persist over these longer time horizons, albeit weaker.

6.4. Distribution of app ratings

We have not yet identified the mechanism leading to our clean result of increased app ratings following the rule change. At least two possible explanations could generate this outcome. First, if the rule change worked as Facebook intended and really “incentivizes developers to improve the quality of their notifications” (Figure 1), we would expect a general upward trend across different levels of app ratings. Second, the result could also be generated by a scenario in which the rule change leads to an increase in entry barriers, which in turn leads to the exclusion of those apps with the lowest quality levels while keeping the quality of all other apps unchanged. We explore this question by analyzing the distribution of app ratings before and after the rule change. The first two bars in Figure 4 plot the percentage of apps in the different rating bins.²⁴

----- INSERT FIGURE 4 HERE -----

We can see that the share of apps with a rounded rating of 1 to 4 declined and only apps with a rounded rating of 5 increased their share. The fact that not only apps with the lowest ratings lose shares (and the

²⁴ Average ratings are on a continuous scale from 1 to 5. We round average ratings to allocate each app to one of the five bins.

largest relative loss in is for apps with a rounded rating of 2, not 1) lends support to increased quality incentives as the underlying mechanism. We confirm this intuition by running a counterfactual, given in the third bar in Figure 4. Here, we simulate increased entry barriers by using the population of apps released before the rule change and excluding those apps with the lowest quality until the average quality rating of the population is equal to the average rating level after the rule change. In other words, how would a distribution look like with the same average rating but generated through a process of successively forcing the weakest apps out of the market? Given the stark differences between the resulting hypothetical distribution of app ratings (right bar) and the actual distribution after the rule change (middle bar), we feel the second mechanism is unlikely to drive our result.

In summary, we have gathered strong evidence that the rule change indeed incentivized developers to program apps generating higher app ratings. The results are robust to alternative explanations, including changes in app types, general time trends, and exclusion of low-quality apps.

7. How are the drivers of active users affected by the rule change?

7.1. Empirical specification

Turning to the next step in the causal chain, we now examine how app ratings and other drivers of app success were affected by the rule change. To do this, we use a set of fixed-effect regressions and test how the rule change moderated drivers of an app's active user base. Here, app-specific fixed effects control for all time-invariant characteristics that might drive the number of active users. As the app rating measure is time-invariant, it is absorbed by the app-specific fixed-effects and cannot be identified directly.²⁵ We thus run the following regression:

²⁵ The time-invariant category assignment is also absorbed by the fixed effects.

$$\begin{aligned}
ActiveUsers_{it} = & a_i + \beta_1 RuleChange_t + \beta_2 RuleChange_t * AppRating_{it} \\
& + \beta_3 UpdateFreq_{it} + \beta_4 RuleChange_t * UpdateFreq_{it} \\
& + \beta_5 \ln(NumUsers_{it}) + \beta_6 RuleChange_t * \ln(NumUsers_{it}) \\
& + \beta_7 \ln(WeeksActive_{it}) + \beta_8 RuleChange_t * \ln(WeeksActive_{it}) \\
& + \beta_9 \ln(NumSisterApps_{it}) + \beta_{10} RuleChange_t * \ln(NumSisterApps_{it}) + u_{it}
\end{aligned}$$

All time-invariant, app-specific heterogeneity is absorbed by the app fixed-effects a_i . We then include all potential drivers of active users as main effects as well as in interaction terms with the rule change. As noted, the main effect of app ratings is absorbed by the fixed effects. However, the interaction term with the rule change is still identified and lets us answer if app ratings became a more important driver for the number of active users after the rule change. The other interaction terms let us identify the moderating effect of the rule change on the relative importance of the drivers of the active user base.

Another concern relates to other possibly unobserved shocks affecting the platform. If these shocks drive the effectiveness of the identified drivers of active users, the dummy for the rule change could also capture these shocks and not only the actual rule change. To mitigate this possibility as far as possible, we restrict our analysis to a short time period around the rule change.²⁶

7.2. Results

We now present results for the moderating effects of the rule change on the drivers of active users. We do not estimate the periods before and after the rule change separately but include a timeframe from ten weeks before until ten weeks after the rule change and identify the moderating effects from interactions with the dummy variable $RuleChange_t$. The first column of Table 5 reports the main effects, while interaction effects are included in the second specification.

²⁶ For the main results, we use a time window of ten weeks before the rule change to ten weeks thereafter. In the robustness checks, we restrict the sample to five weeks before and after the rule change and also run an unrestricted regression.

----- INSERT TABLE 5 HERE -----

The coefficients of the main effects show that the rule change led to an average reduction of active users of 11.7%. In addition, we see that increased update activity is beneficial, that each additional installation of an app leads on average to 0.709 additional active users, and that the number of active users declines over time (app age). The coefficients of the main effects maintain their sign and significance when the interaction terms are added in specification (5-2). The main effect of app rating is absorbed by the fixed effect, but the interaction term with the rule change shows an increased importance of app rating after the rule change. The benefits from higher update frequency increase with the rule change, while the number of an app's users becomes an even stronger negative driver of usage intensity. The coefficient for the interaction of rule change with app age is also positive and significant, which suggests that apps stay attractive for longer. We do not find a significant effect for sister apps.

7.3. Robustness checks

We run three checks to assess the robustness of our results. The results are in the Appendix in Table A.2.

First, we include a continuous time trend to again exclude the explanation that we just pick up a general time trend with the rule change dummy. Specification (A2-1) shows no significant effect of adding a continuous time trend and the dummy for the rule change stays nearly identical.

Second, we run a regression including only applications launched before the rule change to make sure that our results are not driven by changed characteristics of newly launched apps. We see in specification (A2-2) that the coefficients for the interaction terms are very similar in size between the full sample (specification (4-2)) and the sample of apps launched before the rule change.

Finally, we restricted our analysis to a short timeframe around the rule change to avoid confounding the effects from the rule change with other contemporaneous trends. We restrict the observation window

to five weeks around the rule change in specification (A2-3) and run the regression without restrictions in specification (A2-4). The results with these different set of restrictions still hold. When restricting our observation period to +/- 5 weeks around the rule change, the effect of app rating becomes approximately 30% weaker, suggesting that the full effect of the rule change has not yet set in. This suggests that users do not simply observe app rating and use highly rated apps more, but are “drawn in” gradually by more engaging apps, making the short-term effect weaker than the long-term one.

8. Discussion and conclusion

8.1. Interpretation of the results

Our empirical results show that the rule change initiated by Facebook, the platform owner, had a profound impact on the determinants of success in Facebook’s app market. All drivers of the active user base were affected in their impact on active user base by Facebook’s rule change.

Facebook’s move was first and foremost designed to “incentivize developers to improve the quality of their notifications”. Given the nature of notifications—a term capturing both invitations for new users (designed to encourage new installations) and activity reports to other existing users (designed to keep engagement high)—notifications are an important part of a user’s app experience. The first robust result we find thus confirms Facebook’s stated aim of improving perceived app quality. This is visible in Figure 3, where there is a marked jump in quality ratings for apps released after the change, but also reflected in our subsequent regressions which show that firms deliberately changed their apps to achieve higher app ratings. Our most demanding results control for time-invariant app characteristics in Table 5. Here, while the coefficient on app rating is not identified (it is absorbed by the app fixed-effect), we find that post-change app rating has a stronger positive impact on app success. Thus, rule change rewarded the apps which successfully kept users engaged through compelling notifications.

The frequency of updates increases the active user base. This supports the intuition that actively managed apps (i.e. frequently updated ones) enjoy higher usage intensity after the rule change as user engagement is rewarded more. This is especially relevant for older apps that were introduced under a regime which did not incentivize high quality greatly. These apps can “catch up” by actively maintaining user interest through updates after the change.

We then consider the net impact of network effects and user composition. As expected, network effects matter less for usage intensity after the change. There are two related explanations for this. First, post-rule change it is more difficult for widely installed, but not particularly engaging apps to leverage their sheer size to keep users’ interest. That is, if users receive plenty of notifications from a particular app, they may eventually start using the app more since the notifications suggest intense activity. Second, notifications are another mechanism of engaging users. In this case, users are not simply triggered to engage through friends using the app, but also through notifications. Thus, the two substitute for each other to some extent, rendering sheer network size relatively less important.

Another direct implication of the increased quality incentives for apps after the change is the intuition that post-change, apps “age well”. That is, the (negative) fad effect is less pronounced after Facebook’s intervention. Note that this is relevant even though we control for other factors like update frequency and app rating (and other time-invariant characteristics), so this result indeed suggests that apps decay more slowly as a consequence of the rule change. “Compelling notifications” are one driver of this, a change in the composition of active apps is another.

Finally, we do not find that a developer’s app portfolio has a consistent effect on active usage across our specifications in Table 5 and Table A.2.

Our results confirm widespread changes in the success factors and consequently the market structure of the market for Facebook apps. We do not know the precise goals Facebook had in mind when initiating

this change, but our results are in line with an incentive to reward high-quality apps and edge low-quality ones out of the market. This may also trigger a concentration towards one dominant app in every genre as higher quality is rewarded with further opportunities to capitalize on it. So while keeping entry barriers low helps keeping up the “long tail” (Anderson 2006), bandwagon behavior (Dellarocas et al. 2010, Duan et al. 2009) may also lead to a “superstar effect” (Rosen 1981).

One can only surmise that this should eventually lead to increased monetization opportunities for Facebook, which is supported anecdotally by the fact that Facebook later implemented a change in the way monetary transactions are channeled through the Facebook website (and that lets Facebook keep a share of revenues).

8.2. Managerial implications

Our results have a number of implications for practice. First, we find that quality can be incentivized through “soft”, i.e. non-excluding rules. This is an alternative to the costly (and error-prone) “quality threshold” rule under which the platform owner exerts direct control over which software appears on the platform through a vetting and quality assessment process. While such an approach may increase the average quality of active apps, it may also be counterproductive in a nascent market in which consumer preferences are not (yet) settled and there may be innovative apps that would fail the established quality criteria. Second, we find that the drivers of app success are contingent on the environment set by the platform owner. This includes the promotion channels available to apps and the opportunity to implement minor technical (or layout) changes after an app’s launch.

8.3. Limitations and further research

Our study has some limitations and is exploratory in several respects. First, our data is aggregated at the app level. Thus, we can observe changes in the aggregate behavior of users of an app, but not how an individual’s behavior changes over time. Especially as post-change app usage decays less rapidly it would

be interesting to gather individual-level data to see what drives this result. Second, we do not observe profits or revenues, neither by app developers or the platform owner, Facebook. Hence, we cannot infer if the rule change worked in the intended way or if app developers benefited from this on average. However, it is intuitive to assume that developers' actions reveal their preferences and that the upward shift in quality was a response designed to exploit the new environment. Similarly, Facebook's stated aim of increasing quality and user satisfaction is presumably (and plausibly) related to future (and current) opportunities for monetization for Facebook. Third, we study a particular episode in the evolution of a single, albeit an important, social media platform. We should be careful therefore in extrapolating the results to other platforms and rule changes. Nevertheless, our study offers an appealing look into the various ways in which platform owners can manage their platforms.

In summary, we study the changes in the success factors of Facebook apps following a change in the way apps could send out notifications. Our results suggest that app developers respond to "soft" quality incentives by launching better apps, in line with the goals that Facebook stated when announcing the change in notification rule. This study contributes to the emerging literature on the empirics of platform markets and we hope that it sheds light on the interaction between rules set by the platform owner and market dynamics in complementary goods markets.

References

- ABRAHAMSON, E. & ROSENKOPF, L. 1993. Institutional and competitive bandwagons: Using mathematical modeling as a tool to explore innovation diffusion. *Academy of Management Review*, 18, 487-517.
- ACQUISTI, A. & FONG, C. 2012. An Experiment in Hiring Discrimination via Online Social Networks. *Working Paper*.
- AGARWAL, R., ANIMESH, A. & PRASAD, K. 2009. Research Note—Social Interactions and the “Digital Divide”: Explaining Variations in Internet Use. *Information Systems Research*, 20, 277-294.
- ANDERSON, C. 2006. *The long tail: How endless choice is creating unlimited demand*, London, Random House.
- ANIMESH, A., VISWANATHAN, S. & AGARWAL, R. 2011. Competing “Creatively” in Sponsored Search Markets: The Effect of Rank, Differentiation Strategy, and Competition on Performance. *Information Systems Research*, 22, 153-169.
- ARAL, S. & WALKER, D. 2011. Creating Social Contagion Through Viral Product Design: A Randomized Trial of Peer Influence in Networks. *Management Science*, 57, 1623-1639.
- BAKSHY, E., ECKLES, D., YAN, R. & ROSENN, I. 2012. Social influence in social advertising: evidence from field experiments. *Proceedings of the 13th ACM Conference on Electronic Commerce*. Valencia, Spain: ACM.
- BAMPO, M., EWING, M. T., MATHER, D. R., STEWART, D. & WALLACE, M. 2008. The Effects of the Social Structure of Digital Networks on Viral Marketing Performance. *Information Systems Research*, 19, 273-290.
- BOUDREAU, K. 2008. Too Many Complementors? Evidence on Software Firms. Available at SSRN: <http://ssrn.com/abstract=943088>.
- BOUDREAU, K. & HAGIU, A. 2009. Platform Rules: Regulation of an Ecosystem by a Private Actor. In: GAWER, A. (ed.) *Platforms, Markets and Innovation*. Cheltenham, UK and Northampton, MA, US: Edward Elgar.
- BOUDREAU, K., LACETERA, N. & LAKHANI, K. 2008. Parallel search, incentives and problem type: Revisiting the competition and innovation link. *Harvard Business School Technology & Operations Mgt. Unit Working Papers*.
- BOYD, D. & ELLISON, N. 2007. Social network sites: Definition, history, and scholarship. *Journal of Computer-Mediated Communication*, 13, 210-230.
- BRAMOULLÉ, Y., DJEBBARI, H. & FORTIN, B. 2009. Identification of peer effects through social networks. *Journal of Econometrics*, 150, 41-55.
- CABRAL, L. 2006. Equilibrium, epidemic and catastrophe: Diffusion of innovations with network effects. *New Frontiers in the Economics of Innovation and New Technology: Essays in Honor of Paul David*, Edward Elgar, London, 427-437.
- CASADESUS-MASANELL, R. & HALABURDA, H. 2011. When Does a Platform Create Value by Limiting Choice? *Harvard Business School Working Paper 11-030*.
- CHURCH, J. & GANDAL, N. 2000. Systems Competition, Vertical Merger, and Foreclosure. *Journal of Economics and Management Strategy*, 9, 25-51.
- COUGHLAN, P. J. 2004. The Golden Age of Home Video Games: from the reign of Atari to the rise of Nintendo. *Harvard Business School Case Study 9-704-487*.
- DELLAROCAS, C. 2003. The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49, 1407-1424.

- DELLAROCAS, C., GAO, G. & NARAYAN, R. 2010. Are Consumers More Likely to Contribute Online Reviews for Hit or Niche Products? *Journal of Management Information Systems*, 27, 127-158.
- DUAN, W., GU, B. & WHINSTON, A. B. 2009. Informational cascades and software adoption on the internet: an empirical investigation. *MIS Quarterly*, 33, 23-48.
- EISENMANN, T. 2006. Internet companies growth strategies: Determinants of investment intensity and long-term performance. *Strategic Management Journal*, 27, 1183-1204.
- ERMECKE, R., MAYRHOFER, P. & WAGNER, S. Year. Agents of Diffusion - Insights from a Survey of Facebook Users. *In: 2nd Hawaii International Conference on System Sciences*, 2009. 1-10.
- GNYAWALI, D. R., FAN, W. & PENNER, J. 2010. Competitive Actions and Dynamics in the Digital Age: An Empirical Investigation of Social Networking Firms. *Information Systems Research*, 21, 594-613.
- GOH, K. Y., HENG, C. S. & LIN, Z. 2012. Social Media Brand Community and Consumer Behavior: Quantifying the Relative Impact of User-and Marketer-Generated Content. *Available at SSRN: <http://ssrn.com/abstract=2048614>*.
- GOLDFARB, A. & TUCKER, C. E. 2011. Privacy Regulation and Online Advertising. *Management Science*, 57, 57-71.
- GRAJEK, M. & KRETSCHMER, T. 2009. Usage and diffusion of cellular telephony, 1998–2004. *International Journal of Industrial Organization*, 27, 238-249.
- HAGIU, A. 2011. Quantity vs. Quality: Exclusion By Platforms With Network Effects. *Harvard Business School Working Paper 11-125*.
- HENDRICKS, K. & SORENSEN, A. 2009. Information and the skewness of music sales. *Journal of Political Economy*, 117, 324-369.
- HERVAS-DRANE, A. 2010. Word of mouth and taste matching: A theory of the long tail. *NET Institute Working Paper No. 07-41*
- KATZ, E. & LAZARFELD, P. 1955. *Personal influence: The part played by people in the flow of mass communications*, Glencoe, IL, Free Press.
- KATZ, M. & SHAPIRO, C. 1985. Network Externalities, Competition, and Compatibility. *American Economic Review*, 75, 424-440.
- KATZ, M. & SHAPIRO, C. 1986. Product compatibility choice in a market with technological progress. *Oxford Economic Papers*, 38, 146–65.
- LI, X. & HITT, L. M. 2008. Self-Selection and Information Role of Online Product Reviews. *Information Systems Research*, 19, 456-474.
- MANSKI, C. F. 1993. Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60, 531-542.
- OESTREICHER-SINGER, G. & SUNDARARAJAN, A. 2006. Network structure and the long tail of electronic commerce. *Proceedings of the Tenth INFORMS Conference on Information Systems and Technology (CIST 2006)*.
- PHELPS, J., LEWIS, R., MOBILIO, L., PERRY, D. & RAMAN, N. 2005. Viral marketing or electronic word-of-mouth advertising: Examining consumer responses and motivations to pass along email. *Journal of Advertising Research*, 44, 333-348.
- ROCHET, J. & TIROLE, J. 2003. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1, 990-1029.
- ROSEN, S. 1981. The economics of superstars. *American Economic Review*, 71, 845-858.
- SHAPIRO, C. & VARIAN, H. 1998. *Information rules: a strategic guide to the network economy*, Boston, MA, Harvard Business School Press.
- SUSARLA, A., OH, J. & TAN, Y. 2012a. Influentials or Susceptibles? Analyzing Cascades of Word-of-Mouth Conversations in Online Social Networks. *Available at SSRN: <http://ssrn.com/abstract=1984690>*.
- SUSARLA, A., OH, J. H. & TAN, Y. 2012b. Social networks and the diffusion of user-generated content: Evidence from YouTube. *Information Systems Research*, 23, 23-41.

TUCKER, C. 2011. Social networks, personalized advertising, and privacy controls.
TUCKER, C. & ZHANG, J. 2011. How Does Popularity Information Affect Choices? A Field Experiment.
Management Science, 57, 828-842.

Figures and tables

Figure 1: Announcement of rule change on Facebook’s developer blog²⁷

Feedback-based allocations for notifications

By Tom Whitnah - Wednesday, February 6, 2008 at 9:47am


To improve the Facebook Platform user experience and to reward compelling applications, we will be rolling out a feedback-based system that allots notifications in proportion to user response. Applications will no longer have a static upper limit of 40 notifications per user per day. Instead the number of notifications per application will be based on a range of factors including the rates that users ignore, hide, and report notifications as spam.

The new system aims to provide users with more compelling notifications and fewer notifications that they are likely to ignore or consider spam. We hope this change incentivizes developers to improve the quality of their notifications and encourage their users to send notifications to interested friends.

Before this change goes out, we will be providing two new Insights statistics tabs. These tabs provide developers with their current per user notification threshold as well as metrics on notifications that they can use to understand how users are responding. While there is a general correlation between good response rates and higher thresholds, other factors and metrics will be used to determine these scores as well and the allocations will adjust themselves accordingly.

The new Insights tabs will be available later this week and users will start seeing changes next week. Please send your feedback to developers-help@facebook.com with [notifications allocations] in the subject field.

Figure 2: Causal chain of app monetization



²⁷ Available at <http://developers.facebook.com/blog/post/77>, accessed July 10, 2012.

Figure 3: Average rating of apps launched before and after the rule change in 2008w7 (dotted line is the average of the average rating before/after the rule change)

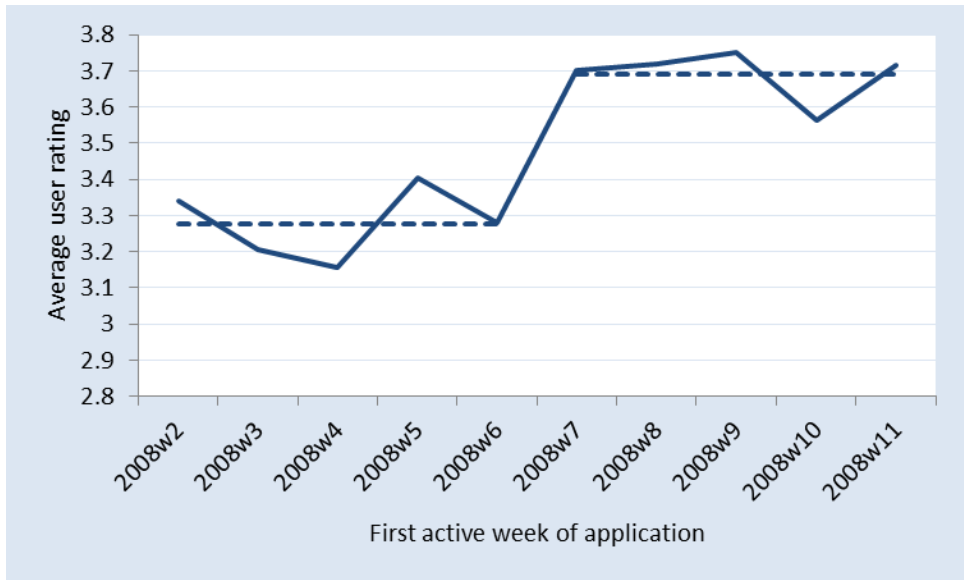


Figure 4: AppRating of apps launched in the five weeks before and after the rule change

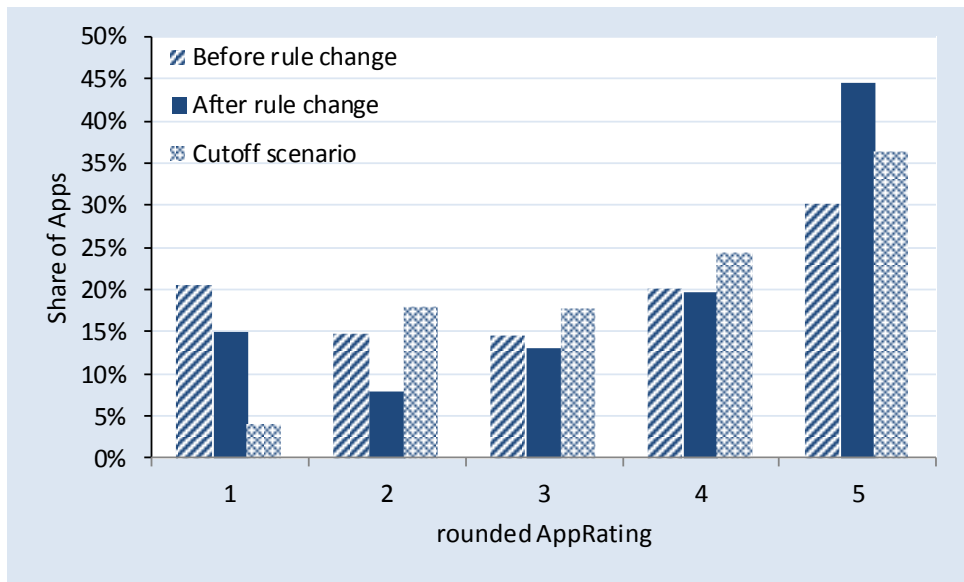


Table 1: Variable definitions

Variable	Definition
$AppRating_i$	Time-invariant average user rating of an app
$ActiveUsers_{it}$	Number of users that have interacted with the app in the last day (in thousand)
$RuleChange_t$	Dummy for the rule change (zero before sixth week of 2008 and one thereafter)
$UpdateFreq_{it}$	Total number of updates of an app divided by $WeeksActive_{it}$
$NumUsers_{it}$	Number of users that have installed an app (in million)
$WeeksActive_{it}$	Weeks since an app has first appeared in Facebook's app directory
$NumSisterApps_{it}$	Number of sister apps offered by the same developer

Table 2: Summary statistics

Variable	All observations				Before rule change				After rule change			
	Mean	SD	Min	Max	Mean	SD	Min	Max	Mean	SD	Min	Max
$AppRating_i$	3.59	1.38	0.65	5	3.58	1.38	1	5	3.59	1.39	1	5
$ActiveUsers_{it}$	5.74	67.87	0	4,510	7.26	83.70	0	4510	4.62	53.28	0.00	3010
$RuleChange_t$	0.58	0.49	0	1	0	0	0	0	1	0	1	1
$UpdateFreq_{it}$	0.15	0.24	0	5	0.13	0.23	0	5	0.18	0.24	0	5
$NumUsers_{it}$	0.16	1.08	0	69.86	0.15	0.97	0	24.69	0.16	1.15	0	69.86
$WeeksActive_{it}$	13.97	10.38	0	48	11.74	9.04	0	38	15.62	10.98	0	48
$NumSisterApps_{it}$	17.64	51.86	0	333	21.44	50.97	0	287	14.85	52.34	0	333

Notes: The number of observations for all variables is 109,233 (46,312 pre-change and 62,921 post-change). All observations are restricted to be within ten weeks of the rule change. Summary statistics are presented in linear form for all variables. In the regressions, the logarithm of $ActiveUsers_{it}$, $NumUsers_{it}$, $WeeksActive_{it}$, and $NumSisterApps_{it}$ is used.

Table 3: Pairwise correlations

Variable	[1]	[2]	[3]	[4]	[5]	[6]	[7]
[1] $AppRating_i$	1.000						
[2] $\ln(ActiveUsers_{it})$	-0.109	1.000					
[3] $RuleChange_t$	0.006	-0.088	1.000				
[4] $UpdateFreq_{it}$	0.027	0.126	0.105	1.000			
[5] $\ln(NumUsers_{it})$	-0.120	0.916	-0.002	0.058	1.000		
[6] $\ln(WeeksActive_{it})$	0.066	0.071	0.169	-0.138	0.341	1.000	
[7] $\ln(NumSisterApps_{it})$	0.054	0.002	-0.155	-0.098	0.018	-0.017	1.000

Note: All observations are restricted to be within ten weeks of the rule change.

Table 4: How did the rule change affect app ratings?

DEPENDENT VARIABLE: <i>AppRating_i</i>		
INDEPENDENT VARIABLES	(4-1)	(4-2)
<i>RuleChange_t</i>	0.346*** (0.0565)	
<i>2008w2 (before RuleChange)</i>		(omitted)
<i>2008w3 (before RuleChange)</i>		0.0202 (0.113)
<i>2008w4 (before RuleChange)</i>		-0.0725 (0.117)
<i>2008w5 (before RuleChange)</i>		0.149 (0.168)
<i>2008w6 (before RuleChange)</i>		-0.0339 (0.125)
<i>2008w7 (after RuleChange)</i>		0.304** (0.137)
<i>2008w8 (after RuleChange)</i>		0.385*** (0.119)
<i>2008w9 (after RuleChange)</i>		0.343*** (0.124)
<i>2008w10 (after RuleChange)</i>		0.166 (0.148)
<i>2008w11 (after RuleChange)</i>		0.388*** (0.136)
$\ln(\text{NumSisterApps}_{it})$	-0.0353* (0.0181)	-0.0387** (0.0188)
Category Dummies	YES	YES
Observations	2,690	2,690
R-squared	0.113	0.114

Notes: Random-effect OLS point estimates with robust standard errors in parentheses. Asterisks denote significance levels (***) $p < 0.01$, (**) $p < 0.05$, (*) $p < 0.1$). Each app is assigned to up to two out of 23 categories (such as photography, dating, or messaging). A constant is included but not reported. All observations are restricted to be within five weeks of the rule change.

Table 5: How are the drivers of active users affected by the rule change?

DEPENDENT VARIABLE: $\ln(\text{DailyActiveUsers}_{it})$		
INDEPENDENT VARIABLES	(5-1) Baseline Regression	(5-2) Rule Change
RuleChange_t	-0.117*** (0.00717)	-0.370*** (0.0313)
$\text{AppRating}_i * \text{RuleChange}_t$		0.0622*** (0.00461)
UpdateFreq_{it}	0.205*** (0.0267)	0.136*** (0.0285)
$\text{UpdateFreq}_{it} * \text{RuleChange}_t$		0.154*** (0.0350)
$\ln(\text{NumUsers}_{it})$	0.709*** (0.00832)	0.781*** (0.00867)
$\ln(\text{NumUsers}_{it}) * \text{RuleChange}_t$		-0.0597*** (0.00274)
$\ln(\text{WeeksActive}_{it})$	-0.774*** (0.00942)	-0.874*** (0.0108)
$\ln(\text{WeeksActive}_{it}) * \text{RuleChange}_t$		0.188*** (0.00892)
$\ln(\text{NumSisterApps}_{it})$	0.00238 (0.00407)	-0.0104*** (0.00404)
$\ln(\text{NumSisterApps}_{it}) * \text{RuleChange}_t$		0.00377 (0.00387)
Observations	109,233	109,233
Number of Apps	7,784	7,784
R ²	0.374	0.398

Notes: Fixed-effect OLS point estimates with standard errors clustered on the app-level in parentheses.

Asterisks denote significance levels (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). A constant is included but not reported. All observations are restricted to be within ten weeks of the rule change.

Appendix A

Table A. 1: Robustness checks for the results on app ratings

DEPENDENT VARIABLE: <i>AppRating_i</i>			
INDEPENDENT VARIABLES	(A1-1) Week control	(A1-2) +/- 10 weeks	(A1-3) Whole period
<i>RuleChange_t</i>	0.377*** (0.117)	0.295*** (0.0457)	0.232*** (0.0357)
<i>Week_t</i>	-0.00610 (0.0208)		
$\ln(\text{NumSisterApps}_{it})$	-0.0363** (0.0183)	-0.0410*** (0.0149)	-0.0381*** (0.0108)
Category Dummies	YES	YES	YES
Observations	2,690	4,340	8,338
R-squared	0.113	0.090	0.063

Notes: Random-effect OLS point estimates with robust standard errors in parentheses. Asterisks denote significance levels (***) $p < 0.01$, (**) $p < 0.05$, (*) $p < 0.1$). Each app is assigned to up to two out of 23 categories (such as photography, dating, or messaging). A constant is included but not reported. Observations are restricted to be within five weeks of the rule change in A1-1, within ten weeks in A1-2, and unrestricted in A1-3.

Table A. 2: Robustness checks for the results on active users

INDEPENDENT VARIABLES	DEPENDENT VARIABLE: $\ln(\text{DailyActiveUsers}_{it})$			
	(A2-1) Week Control	(A2-2) Only old apps	(A2-3) +/- 5 weeks	(A2-4) Whole period
$RuleChange_t$	-0.366*** (0.0321)	-0.201*** (0.0328)	-0.0874*** (0.0330)	-0.509*** (0.0315)
$AppRating_i * RuleChange_t$	0.0621*** (0.00461)	0.0617*** (0.00463)	0.0426*** (0.00442)	0.0640*** (0.00502)
$UpdateFreq_{it}$	0.136*** (0.0285)	0.191*** (0.0312)	0.151*** (0.0382)	0.129*** (0.0235)
$UpdateFreq_{it} * RuleChange_t$	0.154*** (0.0350)	0.264*** (0.0430)	0.221*** (0.0400)	0.163*** (0.0335)
$\ln(\text{NumUsers}_{it})$	0.781*** (0.00890)	0.764*** (0.00971)	0.645*** (0.0112)	0.865*** (0.00682)
$\ln(\text{NumUsers}_{it}) * RuleChange_t$	-0.0597*** (0.00274)	-0.0618*** (0.00277)	-0.0600*** (0.00260)	-0.0537*** (0.00293)
$\ln(\text{WeeksActive}_{it})$	-0.878*** (0.0134)	-0.890*** (0.0117)	-0.777*** (0.0150)	-0.974*** (0.00878)
$\ln(\text{WeeksActive}_{it}) * RuleChange_t$	0.184*** (0.00958)	0.139*** (0.00938)	0.101*** (0.00928)	0.227*** (0.00827)
$\ln(\text{NumSisterApps}_{it})$	-0.0101** (0.00409)	-0.0126*** (0.00409)	-0.0177*** (0.00421)	0.000582 (0.00396)
$\ln(\text{NumSisterApps}_{it}) * RuleChange_t$	0.00370 (0.00387)	0.00461 (0.00388)	0.0202*** (0.00362)	-0.0137*** (0.00419)
$Week_t$	0.000983 (0.00122)			
Observations	109,233	98,194	57,239	196,799
Number of Apps	7,784	6,027	7,158	8,300
R ²	0.398	0.393	0.307	0.493

Notes: Fixed-effect OLS point estimates with standard errors clustered on the app-level in parentheses.

Asterisks denote significance levels (***) $p < 0.01$, (**) $p < 0.05$, (*) $p < 0.1$). A constant is included but not reported. Observations are restricted to be within ten weeks of the rule change in A2-1 and A2-2, within five weeks in A2-3, and unrestricted in A2-4.