

IFN Working Paper No. 886, 2011

Like What You Like or Like What Others Like? Conformity and Peer Effects on Facebook

Johan Egebark and Mathias Ekström

Like What You Like or Like What Others Like?

-Conformity and Peer Effects on Facebook

Johan Egebark and Mathias Ekström*

October 14, 2011

Abstract

Users of the social networking service Facebook have the possibility to post status updates for their friends to read. In turn, friends may react to these short messages by writing comments or by pressing a *Like* button to show their appreciation. Making use of five Swedish accounts, we set up a natural field experiment to study whether users are more prone to *Like* an update if someone else has done so before. We distinguish between three different treatment conditions: (i) one unknown user *Likes* the update, (ii) three unknown users *Like* the update and (iii) one peer *Likes* the update. Whereas the first condition had no effect, both the second and the third increased the probability to express a positive opinion by a factor of two or more, suggesting that both number of predecessors and social proximity matters. We identify three reasonable explanations for the observed herding behavior and isolate conformity as the primary mechanism in our experiment.

Key words: Herding Behavior; Conformity; Peer Effects; Field Experiment

JEL classification: A14; C93; D03; D83

1 Introduction

Whenever a new trend arises—be it within fashion, on product markets or even in politics—it is relevant to ask if the popularity is explained by better quality or if it simply reflects a desire people have to do what everyone else does. The latter supposition, if true, has wide implications since it could explain, among other things, the formation of asset bubbles and dramatic shifts in voting

*Egebark: Department of Economics, Stockholm University and the Research Institute of Industrial Economics (IFN). Email: johan.egebark@ne.su.se. Ekström: Department of Economics, Stockholm University. Email: mathias.ekstrom@ne.su.se.

Acknowledgements: First we want to thank the Facebook users who made this experiment possible. We also want to express our gratitude to Pamela Campa, Stefano DellaVigna, Peter Fredriksson, Patricia Funk, Magnus Johannesson, Niklas Kaunitz, Erik Lindqvist, Martin Olsson and Robert Östling, as well as participants at the ESA Annual Meeting in Chicago 2011 and the National Conference of Swedish Economists in Uppsala 2011, for helpful discussions and valuable comments. Financial support from the Jan Wallander and Tom Hedelius Foundation is gratefully acknowledged. All remaining errors are our own.

behavior. Unfortunately, identifying herding behavior is by its nature difficult and hence we know little about the importance of this phenomenon.

In this paper, we use the world’s leading social networking service, Facebook, to study herding. Each Facebook user has a network of friends with whom he or she may easily interact through several different channels, e.g., by mailing, chatting or uploading photos or links. The most popular feature allows users to post status updates for friends to read; in turn, friends may react to these short messages by writing their own comments or by pressing a *Like* button to show they enjoyed reading it. We set up a natural field experiment to study whether users are more willing to *Like* an update if someone else has done so before. Making use of five Swedish users’ actual accounts, we create 44 updates in total during a seven month period.¹ For every new update, we randomly assign our user’s friends into either a treatment or a control group; hence, while both groups are exposed to identical status updates, treated individuals see the update after someone (controlled by us) has *Liked* it whereas individuals in the control group see it without anyone doing so. We separate between three different treatment conditions: (i) one unknown user *Likes* the update, (ii) three unknown users *Like* the update and (iii) one peer *Likes* the update. Our motivation for altering treatments is that it enables us to study whether the number of previous opinions as well as social proximity matters.² The result from this exercise is striking: whereas the first treatment condition left subjects unaffected, both the second and the third more than doubled the probability of *Liking* an update, and these effects are statistically significant.

We argue that conformity explains the behavior we observe in our experiment. Economists have defined conformity as an intrinsic taste to follow others (Goeree and Yariv, 2010), driven by factors such as popularity, esteem and respect (Bernheim, 1994). As Bernheim’s model suggests, actions that are publicly observable signal predispositions and therefore affect status. Hence, if status concerns are sufficiently important to individuals, they will deviate from self-serving preferences and conform.³ For many reasons, Facebook constitutes an environment where conformity potentially would occur. First it provides high visibility—at any given time, a large number of users observe each other’s actions—allowing signaling to occur. Second, much of the activity on the website revolves around expressing attitudes and beliefs which are likely to be important for projecting status. Third, since there is no obvious way for users to disagree once a norm is established (i.e., no dislike option exist), conformity pressure is unlikely to weaken over time.

Besides conformity, herding has been explained by: (i) correlated preferences, (ii) payoff externalities, (iii) limited attention and (iv) observational learning. We eliminate correlated effects due to the random assignment into treatment and control groups (see discussion in Cai et al., 2009).

¹The experiment took place between May and October 2010. The accounts we used were not created for the purpose of the experiment but rather borrowed from existing users.

²Social impact theory developed in Latané (1981) lists three important factors determining the size of social influence: strength, immediacy and number. Moreover, previous findings from social psychology show that the more unanimous predecessors are, the more likely it is that subsequent decision-makers follow suit (Asch, 1955). Finally, there is convincing evidence that peers can play an important role in determining behavior (see e.g., Bandiera et al., 2005; Mas and Moretti, 2009; Sacerdote, 2001; Kremer and Levy, 2008).

³Of course, people may also be inclined to express their independence by choosing a less popular option. Evidence of such behavior is found in Ariely and Levav (2000), Corazzini and Greiner (2007) and Weizsacker (2010).

Payoff externalities cause herding when each agent’s actions affect other agents’ payoffs in such a way that an equilibrium arises. One typical example is right-hand (or left-hand) traffic. This explanation can also be ruled out as there is no reason for supposing such payoffs arise in our setting (the absence of any equilibrium in type or number of responses in our experiment as well as on Facebook in general speaks against this explanation). Arguably the other two mechanisms are more relevant in our experiment and we therefore address them in more detail.

Limited attention is relevant in situations where agents make an optimal choice after delimiting the choice set (Huberman and Regev, 2001; Barber and Odean, 2008; Ariely and Simonsohn, 2008; DellaVigna and Pollet, 2009). We consider two cases—*saliency* and *searching*. Status updates on Facebook that someone has responded to may be more salient because of their altered physical appearance (a blue rectangular area is added beneath the update).⁴ Consequently, updates in treatment groups may have a higher probability of being read and this in turn could translate into more responses. However, the three treatment conditions we use affect the saliency of updates identically and since there is no treatment effect for one of the conditions, *saliency* cannot explain our results. *Searching*, on the other hand, would be an appropriate explanation if users want to save time or effort by looking for previous responses in order to find the best status updates quickly. Such screening will again increase the reading probability for updates in treatment groups, but since response behavior should be unaffected, we expect a treatment effect for both types of responses (*Likes* and comments). It turns out comments are unaffected in all of our treatment conditions and thus there is little support for the *searching* mechanism either.

Observational learning models fit situations where successors follow those who are believed to be better informed because this constitutes a best response.⁵ The key assumption is that agents obtain information by observing each other’s actions and this in turn helps them maximize intrinsic utility (for theoretical studies, see Banerjee, 1992 and Bikhchandani et al., 1992; for empirical evidence, see e.g., Anderson and Holt, 1997 and Alevy et al., 2007). We argue that in our setting, where subjects choose whether to *Like* a status update or not, observational learning is unlikely to exist. The obvious reason being that choices in our experiment are made after subjects have experienced the “product” and have been able to evaluate it against comparable alternatives. In essence, people have all the tools required to make a private quality assessment instantly without the need for information signals. Although we are unable to directly address this channel, we present findings which support this argument.

Our study is related to two different strands of literature within the economics discipline. On the one hand, we build on a growing body of experimental studies on herding behavior in different real-life settings, on the other, we tie in with the peer-effect literature.

Cialdini et al. (1990), Goldstein et al. (2008), Ayres et al. (2009) and Chen et al. (2010) all study

⁴The term saliency comes from neuroscience. The saliency of an item is the quality by which it stands out relative to its neighbors.

⁵Note the difference between limited attention and observational learning. The former refers to a situation where subjects are certain of the quality of an update after reading it but use some (conscious or subconscious) rule to shrink the choice set. The latter, on the other hand, is relevant if subjects, after reading an update, are unsure of its quality.

decisions related to public goods (littering, resource usage and contributions to an online community). Hence, the effects found could be explained by either conformity or conditional cooperation, or both. Salganik et al. (2006) set up an artificial online music market and show that previous downloads positively affect an individual's tendency to download a specific song. The effect is mainly driven by the fact that frequently downloaded songs are listed higher up on the website, i.e., are more salient. When position is independent of number of downloads, the effect is weaker, suggesting that limited attention is the prime explanation for their results. Cai et al. (2009) vary information on menus in a Chinese restaurant chain to separate observational learning from saliency. Since there is an effect on demand when the five most popular dishes are displayed but not when five randomly dishes are highlighted, the authors interpret the results in favor of observational learning. Martin and Randal (2008) varies the amount of money (and mixture of coins and bills) in a transparent box, used to collect voluntary visiting fees in an art gallery, to analyze how visitors' contributions are affected. However, no attempt is made to distinguish between different explanations. What separates our study is that we exploit a situation without payoff equilibria and where observational learning is negligible. Since status concerns are likely to be important, our focus is on conformity and how conformity pressures forms attitudes and beliefs.

The research on peer effects has found that peers may play an important role in affecting for example productivity at work (Bandiera et al., 2005 and Mas and Moretti, 2009) and savings decisions (Duflo and Saez, 2002, 2003). We show that social proximity is also important when it comes to expressing conforming preferences. Moreover, contrary to many previous studies, we offer a precise definition of what we mean by a peer; rather than just saying he or she is a colleague or a roommate, we use the degree centrality condition.⁶ This means we are able to explicitly study the role of what can be seen as a central person in a network of friends.

In a broader sense, this study is motivated by a growing interest in social interactions within economics. Starting with Becker's (1974) critique of traditional economic theory for neglecting the importance of social interactions, researchers have long tried to gain further insights into this aspect of human life. According to Manski (2000), theorist's have succeeded quite well whereas empirical research is lagging behind, mostly because identification generally has been too challenging. Manski's main conclusion is that more knowledge would be gained with well-designed experiments in controlled environments (Soetevent, 2006, concludes with a similar argument). Our hope is that this study will increase our understanding of herding behavior in general and conformity in particular, and encourage further research on these and other related topics.

The paper evolves as follows. Section 2 describes Facebook while section 3 presents the experimental design. Section 4 and 5 summarizes the data and presents the main findings. In Section 6 we discuss our results in a broader context and section 7 concludes.

⁶A seminal paper on different centrality conditions is Freeman (1979).

2 About Facebook

Facebook is the leading social networking site.⁷ The largest countries according to number of total users are the US (150 million), Indonesia (34 million) and the UK (28 million). Sweden has around 4 million users meaning penetration is about the same as in North America (somewhere between 40-50 percent).⁸ Currently, the website is the second most visited of all and it has attracted increasing attention as a marketing channel both within politics and the corporate environment. The company’s own statistics report that the average user has 130 friends, spends over one day per month on Facebook and creates 90 pieces of content each month (e.g., links, blog posts, notes and photo albums). Moreover, 50 percent of what Facebook defines as active users log on to the website in any given day.

Ultimately, Facebook is an arena for people who seek to interact with their network of friends. Other users are added to your network when they accept your *Friend Request*. Once you have become friends you may visit each other’s profiles and can easily interact through different channels, e.g. by mailing, chatting or uploading photos or links. The most popular feature, *Status*, allows users to inform their friends of their whereabouts and actions in status updates. These short messages are made visible to friends on the *News Feed* which displays updates as *Most Recent* or as *Top News*.⁹ Immediately after it has been posted, friends may react to a status update either by writing their own comments or by pressing a *Like* button to show their appreciation. Both types of responses show up together with the update and are thus clearly visible to the user who wrote the update and his or her network of friends. A status update is limited to 420 characters (including spaces), which means they are typically short, in most cases one sentence. Moreover, a majority of updates are current—they reveal for example what the user is doing right now or where he or she is—which means any reactions from friends are unlikely to show up after more than one or two days (the fact that none of the updates we used in the experiment generated any response after 20 hours confirms this).

The *Like* button was introduced in February 2009 and quickly became a widely popular way for users to express positive opinions about shared content. Facebook’s own description of how this feature works is as follows:

We’ve just introduced an easy way to tell friends that you like what they’re sharing

⁷comScore reports the website attracted 130 million unique visitors in May 2010 and Goldman Sachs estimates the website has more than 600 million users in total as of January 2011.

⁸The exact numbers varies depending on source. Figures reported here are from CheckFacebook.com which, although not affiliated with Facebook, claims to use data from its advertising tool.

⁹When we ran the experiment in 2010, users could choose between the alternative views themselves and easily change between them (see Figure 5 in the Appendix). Moreover, according to Facebook, the *Top News* algorithm was based on “how many friends are commenting on a certain piece of content, who posted the content, and what type of content it is (e.g. photo, video, or status update)” (see <http://www.facebook.com/help/?page=408>). This means the number of *Likes* did not determine if a status update appeared on *Top News* or not. During the experiment, we confirmed that this was the case. In light of this, it should be mentioned that Facebook continuously changes the interface of the website (arguably in order to develop and improve its functionality). A major change occurred in September 2011 which altered the way the *News Feed* presents information. This means there are some discrepancies between the presentation of Facebook in this paper (which is based on how things were in 2010) and its current format. Importantly, no major changes were made during the experimental period in 2010.

on Facebook with one easy click. Wherever you can add a comment on your friends' content, you'll also have the option to click "Like" to tell your friends exactly that: "I like this."

Leah Pearlman, The Facebook Blog

We include a print screen in the Appendix to show the way the *News Feed* presents recent status updates in reverse chronological order (Figure 5). As seen in this figure, the first update has no responses while the second has received a comment and the third a *Like*.

3 The Experiment

Posting a status update on Facebook usually means all of your friends can see it. However, each user has the possibility to control who sees a specific update through privacy settings. Thus, if users wish, they can create a subset of friends, e.g., family members or close friends, and write the message to this group only. We use this feature in our experiment since it allows us to post identical status updates, simultaneously, to different groups—in our case treatment and control groups. Importantly, members of a group can only follow the communication within the specific group and this communication is displayed as normal to the selected members. Hence, we do not worry that subjects perceive the status updates we post within the experiment differently from the ordinary stream of information on the *News Feed*.

We post 44 status updates in total during a seven month period using five Swedish Facebook accounts.¹⁰ Table 1 briefly describes the six steps we go through each time we post an update (every time we execute the process, we use one account which means the 44 updates are distributed over the five accounts). We do not want updates to stand out but rather be a natural part of the ongoing communication on the website. Therefore, in the first step, we ask one of our five users to text his or her status update to us whenever we decide it is time.¹¹ From the list of examples given in Table 7 in the Appendix, we see that updates are trivial in the sense that they are short, easy to interpret and do not say anything which could be perceived as “sensitive”, such as political opinions or religious views.¹² In the second step, after receiving the content for a status update, we randomly draw one of three types of treatment conditions: (i) one unknown user *Likes* the update, (ii) three unknown users *Like* the update, or (iii) one peer *Likes* the update. Our motivation for using three different treatments is to find out if the number of people expressing their opinion and social proximity is

¹⁰Note that we have access to more than five accounts in total (see discussion about the peer and unknown users below). All users were aware of the experiment and they personally gave us access to their accounts. We were very strict about anonymity, i.e., we instructed them under no circumstances to reveal anything about our research.

¹¹We explicitly instructed our users not to think of this as an update that would be used within an experiment. In fact, they themselves stressed that it was important that the updates we used expressed something they would have posted anyway, arguably because they did not want to gain a bad reputation by letting us post updates which they could not stand for. In some few cases, we came up with the content for a status update and then asked our users if we could use it.

¹²We want to study behavior in the simplest possible setting. It would be interesting in further research to see if conformity depends on the type of update but this question is outside the scope of this study.

important for the decision to *Like* an update. Randomization in this step means we can eliminate systematic differences in updates between treatments. The third step implies random assignment of subjects into either a control or a treatment group. In the fourth and the fifth step, we post the status update in treatment and control groups, and immediately expose treated subjects to the condition that was drawn in step 2.¹³ Finally, we collect data on responses.

An interesting feature of our experiment is that subjects have the possibility of responding in different ways, by *Liking*, commenting or both. Although we have three outcomes, our focus is on *Likes*. Compared to comments, which in some cases are quite complex, this response is easy to interpret and there is little doubt it captures what we want to measure: whether a person is more likely to express a positive opinion if someone else has done so before. We use comments in later analysis however to separate between different explanations for the observed behavior (see section 5.2).

Table 1: The experimental process

Step	Description
1	Ask for a status update from one of our five users
2	Random draw of one out of three treatment conditions
3	Random assignment of the user’s friends into treatment and control groups
4	Post identical status updates in treatment and control groups
5	Expose treated subjects to the condition drawn in step 2
6	Collect data on responses

To give a better sense of the three treatment conditions we use, Figure 1 provides a graphical illustration (note that one treatment is used for each update, i.e., updates vary between treatments). Treatment T_{one} , one unknown user *Liking* the update, is illustrated at the top of the figure. We randomly assign our user’s set of friends, F , into two equally sized groups, the control group, C , and the treatment group, T .¹⁴ We expose both groups to identical status updates but for treated individuals we add a *Like* made by an unknown user.¹⁵ The figure shows how the update is displayed in each of the two groups on Facebook’s *News Feed*. Since this treatment reveals only one person’s opinion and the person is a complete unknown, we think of this as the lowest possible trigger.

It is reasonable to assume that influence increases if predecessors are more unanimous (see Latané, 1981 and Asch, 1955). A natural extension, therefore, is to add more *Likes* to updates. The middle section of the figure illustrates treatment T_{three} where we increase the number of people liking the update to three. Importantly, we still want the users who *Like* the update to be unknown

¹³Note that updates are *Liked* by one (or three) of the users’ accounts available to us, i.e., we press the *Like* button using the accounts we have access to.

¹⁴To be more specific, we create four groups of equal size, two control and two treatment groups. The reason for clustering treatment and control groups into smaller entities is twofold. First, for practical reasons it was easier to handle smaller groups. Second, we wanted spontaneous communication in control groups to start as late as possible. Larger control groups increase the probability that someone *Likes* the update at any given time and if this happens, subjects in control groups are in some sense treated. A typical cluster contains around 30 subjects.

¹⁵The unknown users are people we added to our five users’ networks of friends. They were chosen in such a way that we are certain they are unknown to all subjects.

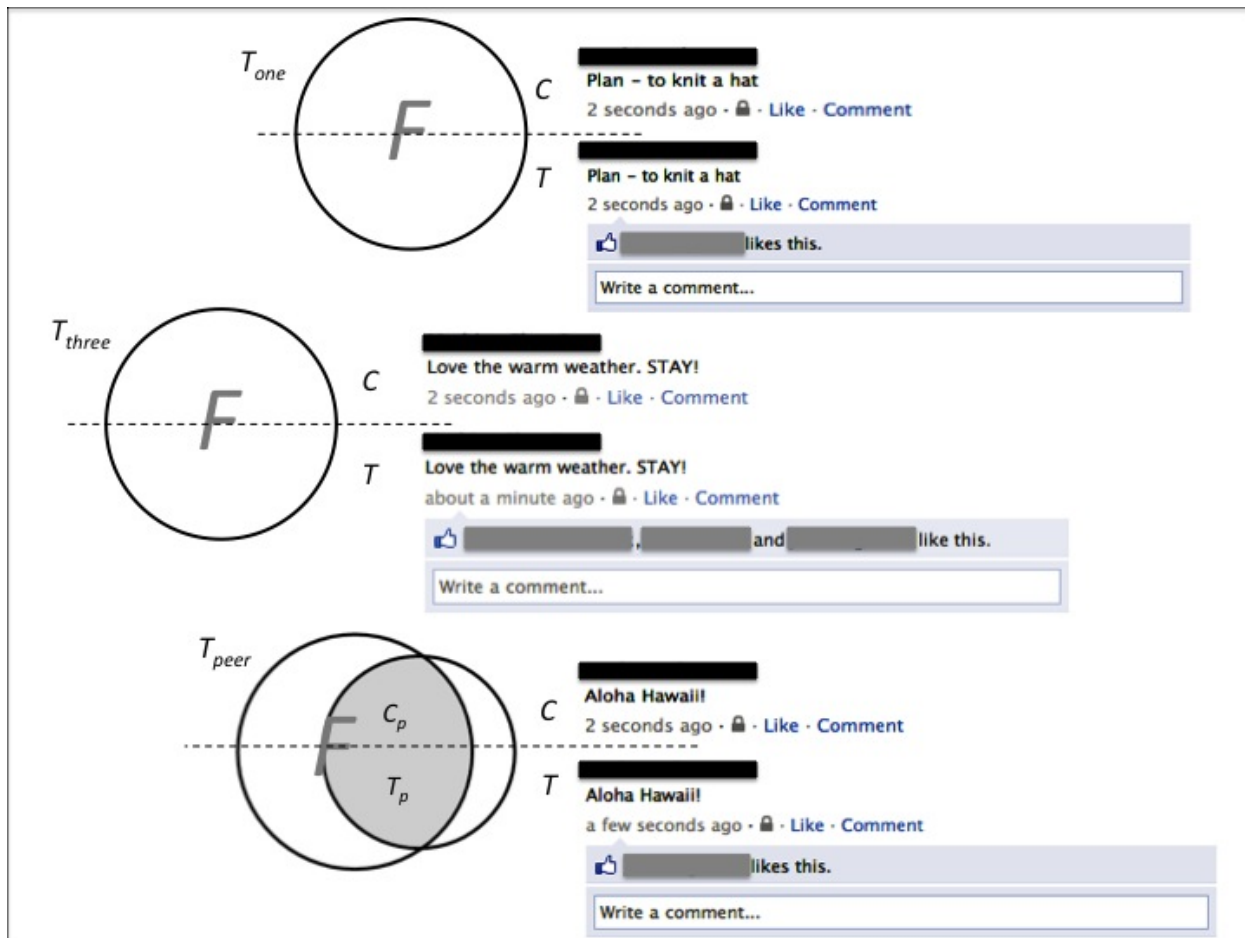
to subjects since this implies it is straightforward to compare results from the first two treatment conditions. If there turns out to be a difference in effects, we can learn something about whether the number of persons expressing their opinion matters. Again, we randomly assign friends into either a treatment or a control group and we expose both groups to identical status updates (note that this implies subjects’ treatment status varies between updates). Our decision to add exactly three *Likes* in this treatment condition has several reasons. First, increasing them one step at the time would have been too time-consuming. Further, we want to signal to some extent that the update in question is popular without making it stand out too much in the *News Feed*. Finally, the seminal studies by Solomon Asch (Asch, 1952, 1955, 1956) on how subjects change private answers to simple questions when exposed to group opinions show three confederates have the largest marginal influence on subjects’ decision to conform.¹⁶

Treatment T_{peer} measures whether subjects’ decision to follow depends on the strength of a relationship. The argument being that the better you know a person the more weight you put on his or her opinions. Facebook is built around the concept of friendship which means we can easily test for the existence of peer effects. Using four of the five accounts, we define what we call a peer group.¹⁷ Imagine again our user’s set of friends, F , illustrated by the bigger circle in the bottom of Figure 1. Each friend in this set has his or her own set of friends. For a majority of the friends in F , sets are overlapping but the number of friends in common varies. We identify the friend in F who has most friends in common with our user (i.e., the largest overlapping area). This friend is defined as the peer and the group of common friends, illustrated by the shaded area, is the peer group. In our last treatment, the peer is the one who *Likes* the update (random assignment of friends in F automatically splits the peer group into a treatment and a control group). Note that the peer group is always a subset of F —in our case the fraction of common friends is around 50 percent. Note also that for subjects who do not belong to the peer group, this third treatment is identical to the first treatment described above. The reason for defining our peer using the degree centrality condition is that we want to study what influence a central person has in a network of friends. Within social network theory, several centrality conditions exist. Only counting the number of nodes *directly* linked to a specific subject in the network means degree centrality is arguably the simplest condition. Other measures, such as the Bonacich centrality (Bonacich, 1987) or the intercentrality condition (Ballester et al., 2006), also take into account the centrality of the people you know and how important this indirect link is. At this stage we found it appropriate to focus on the simplest possible condition.

¹⁶Whether or not this result translates to our setting is an open question; nevertheless, we use this result as guidance.

¹⁷Unfortunately, for one user we were unable to use T_{peer} as treatment since we did not get access to the peer’s account.

Figure 1: Illustration of the experiment



4 Data

Table 2 describes the five users from whose accounts we post updates during the experiment. Ideally, we would have used accounts from more than one country and with more variation in the year of birth. However, the fact that we deal with peoples' private accounts made this difficult. At least, there is significant variation along variables such as gender and number of friends. Although the 44 updates are not evenly distributed across the five different users, they are so across gender and number of friends.¹⁸ Our unit of analysis is the subject, a unique friend-user combination, which means we have 960 observations for our first user (120 friends times 8 updates), 816 observations for our second user (204 friends times 4 updates) and so on. Consequently we have 710 subjects and 5660 observations in total.¹⁹ Columns 6–8 show that treatment condition T_{three} was drawn more frequently than the other two but the difference is quite small, especially when considering we could

¹⁸As much as we wanted to, we could not use every account as frequently.

¹⁹A small fraction of the 710 friends are friends with two out of five users. However, taking this into account does not affect any results and we therefore prefer the above mentioned, more precise, definition of a subject (it also makes sense since this is the level of randomization).

not use T_{peer} for one of the users. Finally, the last two columns give the distribution of responses. The two outcome variables are defined analogously: if the first response from the subject on a given update was to press the *Like* button (give a comment), we define this as a *Like* (comment). Some subjects who first *Liked* an update later also commented on the same update but the opposite never happened (our results are robust to different definitions of these outcome variables). We note that 90 *Likes* and 48 comments were made in total during the experiment.

Table 2: The five Facebook users

User	Gender	Born	Friends	n	Status updates				Responses	
					T_{one}	T_{three}	T_{peer}	Total	Like	Comment
1	Female	1981	120	960	3	5	-	8	19	8
2	Male	1983	204	816	2	1	1	4	10	12
3	Male	1983	152	608	1	1	2	4	6	4
4	Female	1983	176	2464	4	5	5	14	42	18
5	Male	1982	58	812	4	6	4	14	13	6
Total			710	5660	14	18	12	44	90	48

We are also able to observe some background characteristics for the 710 subjects. Table 3 summarizes these variables. The table is split in half and thus displays figures for the total sample as well as a subset of what we call responders. The smaller sample consists of subjects who at least once responded to our updates, irrespective of being in a treatment or in a control group when responding (as seen, responders constitute 16 percent of the sample). The reason for narrowing the sample in the right-hand panel is to highlight what kind of characteristics pertain to those who respond. For example, while females are only slightly overrepresented in the total sample they are in large majority among responders, implying females are more prone to respond to our updates.²⁰ We also note that there are less peer group subjects in the responder sample, and that there is at least some variation in age in both samples as indicated by a standard deviation of around five years in subjects' age.

Rows 5 and 6 of Table 3 show how active subjects are on Facebook (this data was collected somewhere midway through the experiment). The mean number of days since a subject's last response to any shared content on Facebook is approximately 19. However, as seen, there is much variation in activity and the real divider lies between those who have responded at least once within the last three days and those who have not (see Figure 6 included in the Appendix). We use this information to define what we denote active subjects, i.e., a subject is active if (and only if) he or she has responded within three days. Note the difference between responders and active subjects: while the first refers to those who respond to the updates we post, the latter are users who are active in a broader sense. As expected, our sample of responders is also more active in general. We use the distinction between active and less active subjects in later analysis when we investigate the possible role of observational learning (see Table 6 in section 5.2). Finally, we note that all background

²⁰This is in line with Facebook's own reporting saying women are behind 62 percent of all activity.

variables balance between treatment and control groups, which suggests the randomization process worked well (see Table 8 in the Appendix).

Table 3: Summary statistics

Variable	Total			Responders		
	n	Mean	S.d.	n	Mean	S.d.
Responder	5660	0.155	0.362	878	1	0
Female	5660	0.560	0.496	878	0.720	0.449
Peer group	4700	0.504	0.500	718	0.415	0.493
Year of birth	4018	1981	4.745	710	1980	5.749
Last response	5322	19.141	64.468	852	6.016	9.426
Active	5660	0.428	0.495	878	0.574	0.495

5 Results

5.1 Main Findings

Table 4 summarizes our main findings by presenting t-tests of differences in means between treatment and control groups for each treatment condition, together with the magnitude of the effect, where it is significant. In the first row, which presents results from the baseline treatment condition T_{one} , we see that there is no difference in point estimates between groups.²¹ Hence, we draw the conclusion that one unknown user *Liking* a status update does not increase the probability of further positive opinions. Moving to the second row, the result clearly alters: when three unknown people *Like* a status update, successors are more than two times as likely to press the *Like* button, and this effect is highly significant. Row 3–5 in Table 4 show what influence the peer has on subjects’ decision to *Like* an update. Beginning with the third row, we report the effect for the entire sample, i.e., we lump together subjects belonging to the peer group and those who do not. The treatment effect is significant at the 5 percent level and relatively large in magnitude. From the last two rows, where we split the sample, we draw the conclusion that this effect is entirely explained by reactions in the peer group—subjects are more than four times as likely to express a positive opinion when they observe their peer’s action. The fact that subjects who have no social relation to the peer are unaffected confirms our previous finding that one unknown user *Liking* an update is not enough to trigger herding behavior. As a complement to this presentation of the main results, Table 9 in the Appendix presents OLS regressions where we add standard errors clustered on the subject level and control for user or subject fixed-effects. As seen, results do not change.

²¹The mean values for responses may seem low at first sight but are not surprising considering how interaction on Facebook works. While some friends log in several times per day others have registered an account but do not use it. Moreover, because information runs quickly and is very current, active friends will often not observe our status updates and even if they do they will in most cases simply ignore it.


Table 4: Treatment effects by treatment condition

Treatment condition	Sample (n)	Dependent variable: Like			
		C (s.e.)	T (s.e.)	T-C (s.e.)	(T-C)/C*100 if significant
T_{one}	Full (1856)	0.015 (0.004)	0.012 (0.004)	-0.002 (0.005)	-
T_{three}	Full (2184)	0.011 (0.003)	0.027 (0.005)	0.016*** (0.006)	142 %
T_{peer}	Full (1620)	0.007 (0.003)	0.022 (0.005)	0.015** (0.006)	200 %
	Peer group (802)	0.007 (0.004)	0.031 (0.009)	0.024** (0.010)	333 %
	Not peer group (818)	0.008 (0.004)	0.014 (0.006)	0.007 (0.007)	-

Notes: The table reports sample means and the corresponding standard errors in parentheses. *, **, *** = significant at the 10, 5 or 1 percent level in a doubled sided t-test.

Figure 2 compares treatment effects across the five users. Although splitting the sample decreases the statistical power, there is a clear pattern showing that the probability of *Liking* a status update is consistently higher in treatment groups.²² This is reassuring since it indicates that the findings are not driven by one or a few of the users (note that this is also suggested by the fact that regression estimates are robust to the inclusion of user fixed effects in Table 9 in the Appendix).

Figure 2: Treatment effects across users



Note: Error bars represent standard errors of the mean.

²²Although we pool data over all three treatment conditions in Figure 2, results follow the same pattern if we analyze each condition separately or if we exclude T_{one} .

Table 5 looks at gender differences.²³ From the first two rows we draw the conclusion that neither males nor females are affected by T_{one} , a result which strengthens the conclusion of the zero-effect already discussed. For T_{three} and T_{peer} on the other hand, point estimates are considerably higher in treatment groups independent of gender, which suggests that both males and females are affected (however, the rather large difference for male peer group members in the last treatment does not reach statistical significance due to low power).

Table 5: Treatment effects by gender and treatment condition

Treatment condition	Gender (n)	Dependent variable: Like		
		C (s.e.)	T (s.e.)	T-C (s.e.)
T_{one}	Male (807)	0.010 (0.005)	0.008 (0.004)	0.002 (0.007)
	Female (1049)	0.019 (0.006)	0.015 (0.005)	0.004 (0.008)
T_{three}	Male (969)	0.002 (0.002)	0.011 (0.005)	0.009*
	Female (1215)	0.018 (0.006)	0.039 (0.008)	0.021**
T_{peer}	Male (436)	0.009 (0.006)	0.019 (0.010)	0.011 (0.011)
	Female (366)	0.005 (0.005)	0.045 (0.016)	0.040**

Notes: The table reports sample means and the corresponding standard errors in parentheses. *, **, *** = significant at the 10, 5 or 1 percent level in a doubled sided t-test. The last treatment condition (T_{peer}) includes peer group subjects only.

5.2 Explaining Observed Behavior

We argue that conformity explains the behavior we observe in our experiment. Theory predicts that conforming behavior occurs when status is signaled through publicly observed actions and individuals' concern about social status is sufficiently high (Bernheim, 1994). With this in mind, we would like to argue that Facebook constitutes a close to ideal environment for studying conformity. First, the fact that a large number of users observe each other's actions means signaling can occur. Second, much of the communication on the website revolves around attitudes and beliefs, values likely to be important for establishing and communicating status. Moreover, since there is no obvious way for users to disagree once a norm is established, conformity pressure is unlikely to weaken over time. Besides conformity, however, there are arguably two other explanations that are potentially relevant in our setting: limited attention and observational learning.

Limited attention refers to situations where agents make an optimal choice after delimiting the

²³Meta-studies of social influence (see e.g., Eagly and Carli, 1981) generally find that females are more likely than males to conform in public settings.

choice set (Huberman and Regev, 2001; Barber and Odean, 2008; Ariely and Simonsohn, 2008; DellaVigna and Pollet, 2009). Such behavior is typically found on financial markets where, due to binding time constraints, investors restrict their attention in some important way. If Facebook users act under similar conditions, they should use the same rationale, i.e. they limit the number of updates they actually read but after reading they make their preferred choice without taking previous responses into account. We consider two cases—*saliency* and *searching*—which both fall under the limited attention category.


Saliency refers to what we observe in Figure 1: status updates with at least one *Like* potentially stand out, relative to those without, because of their altered physical appearance (a blue rectangular area is added beneath the status update). Consequently, updates in treatment groups may have a higher probability of being read and this in turn could translate into more responses. Note, however, that the three treatment conditions affect the appearance of an update identically. Thus, if *saliency* is driving behavior, we would expect treatment effects in all three treatment conditions. From the first and the last row of Table 4, we know that there is one condition without any effect and we can thus rule out this channel as explanation for our results.

Searching, on the other hand, has to do with screening in another, more deliberate, sense. In order to save time or effort, users may look for previous responses to help them find the best updates quickly. Such a search rule will again increase the reading probability for updates in treatment groups and in the end the number of responses. The increased reading probability may vary with respect to either the number of persons liking the update or social proximity and hence we do not expect homogeneous effects across treatments as in the case with saliency. This means we cannot use the same logic as for *saliency* above. Instead we perform two tests to see if there is support for this explanation.

In the first test, we study differences in the type of response given to updates. Since *searching* only affects subjects’ choice set and not their response behavior, we expect comments to be affected in the same way as *Likes*. Figure 3 shows that while there is a treatment effect for the latter outcome there is no indication of an effect for comments.²⁴ Hence, this test gives little support for the *searching* mechanism (the same logic applies to *saliency* which means we have further evidence against that explanation). In control groups, by contrast, response probabilities for the two outcomes are similar, suggesting that the respective response modes are equally popular in absence of treatment.

²⁴Since T_{one} had no treatment effect, we focus on T_{three} and T_{peer} in Figure 3. Not surprisingly however, T_{one} leaves comments unaffected as well. Results are unchanged if we restrict attention to peer group subjects in condition T_{peer} (Figure 7 in the Appendix). The graphical evidence is confirmed by t-tests.

Figure 3: Response probability by treatment condition



Note: Error bars represent standard errors of the mean.


The *searching* explanation is further called into question by the findings presented in Figure 4. This figure shows the average number of *Likes* per update in treatment groups (on the y-axis) as a function of update quality (on the x-axis). We measure the quality of a status update by the number of control group subjects liking them (arguably, updates generating more than one *Like* in control groups should, in general, be of higher quality than those generating zero or one). We also plot the line which represents the average number of *Likes* under a null hypothesis. As seen, treatment effects are more prevalent for low quality updates.²⁵ But why does this contradict a model that assumes subjects use previous responses when searching for better updates? An example illustrates our main point. Imagine a user scrolling through the *News Feed* and finding an update that has been *Liked* by three unknown users or a peer. This is used as an indicator of high quality and the subject therefore reads the update. After reading, he or she knows the quality with certainty and therefore only responds to the update if the quality actually is high. Clearly, this means any effect would be exaggerated for updates that are of high quality and not, as in the figure, the other way around.

Given the success of limited attention as an explanation in other settings, it is perhaps surprising that this model does not fit our data. Nevertheless, status updates on Facebook are typically so short and easy to understand that the time or effort you can save by searching for previous responses (or noticing the blue rectangular area) is limited. Presumably, a more popular screening method is to limit your attention based on who posted the update, especially so since you will quickly learn who usually post updates that you like.²⁶

²⁵Our conclusion does not change if we study each treatment separately or if we focus on peer group subjects in the T_{peer} condition (Figure 8 in the Appendix).

²⁶As mentioned in section 2, Facebook simplifies this process through the *Top News* list view; importantly, such a search mechanism will only affect the number of potential responders in our experiment, not the results.

Figure 4: Treatment group *Likes* by update quality



Notes: Point estimates show the average number of *Likes* generated in treatment groups. Error bars represent the corresponding standard errors. Treatment condition T_{one} is excluded.

Setting limited attention aside, we now focus on observational learning. Within the economics literature, there has been a strong tradition towards using information based models to explain herding.²⁷ These models assume that agents obtain information by observing each other’s actions; as a result, successors are inclined to imitate those who are believed to be better informed in an attempt to maximize intrinsic utility. Cai et al. (2009) provides a good example of a setting where observational learning is at play. When diners in a Chinese restaurant chain are informed about the past weeks most popular dishes, the demand for these alternatives increases. Since choices about what dishes to order naturally involves uncertainty (customers cannot taste the food before eating), information about prior choices can serve as a quality signal and therefore help individuals make optimal choices. Contrary to the restaurant setting, choices in our experiment are made after subjects have experienced the “product” and have been able to evaluate it against comparable alternatives (as shown in Figure 5 in the Appendix, users can easily read and compare status updates on the *News Feed*). Consequently, since the decision we study is made after a private quality assessment has been made, observational learning should be less important.²⁸ We support this argument by looking at differences in behavior between active and inactive subjects in Table 6. If less active subjects change their behavior while those who are more active do not, this speaks

²⁷See Banerjee (1992), Bikhchandani et al. (1992), Anderson and Holt (1997) and Alevy et al. (2007).

²⁸Our main point is that it is unlikely that subjects cannot accurately estimate an update’s quality even in private. If we compare our study to Cai et al. (2009) again, what we do essentially is to test if customers, who have already tasted the food, are more likely to say they liked it if someone else said so before (a choice under perfect private information).

for the existence of observational learning. The reason is that less active users should benefit relatively more from any quality signal than those who are active (compare this to a restaurant environment: new diners typically need more guidance than regular customers who know what they prefer). Defining a subject’s activity as described in section 4, we see that both active and less active subjects change their behavior—in fact, for treatment condition T_{peer} it is only active subjects who are affected. In summary, there is little support for information based models such as those outlined in Banerjee (1992) or Bikhchandani et al. (1992).

Table 6: Treatment effects by activity

Treatment condition	Group (n)	Dependent variable: Like		
		C (s.e.)	T (s.e.)	T-C (s.e.)
T_{three}	Active (884)	0.018 (0.006)	0.037 (0.009)	0.019* (0.011)
	Less active (1300)	0.006 (0.003)	0.020 (0.006)	0.014** (0.006)
T_{peer}	Active (385)	0.011 (0.008)	0.056 (0.016)	0.045** (0.018)
	Less active (417)	0.004 (0.004)	0.005 (0.005)	0.001 (0.007)

Notes: The table reports sample means and the corresponding standard errors in parentheses. *, **, *** = significant at the 10, 5 or 1 percent level in a doubled sided t-test. T_{peer} includes peer group subjects only.

6 Discussion

When discussing our results, we first want to make clear that any evidence of conformity does not imply individuals are acting irrationally. In fact, Bernheim’s (1994) main point that modeling conformity does not require us to abandon the framework of consistent, self-interested optimization is directly applicable here. In addition to intrinsic preferences, Facebook users may care about status, which could be thought of as esteem, popularity or respect. If users, because of status concerns, are reluctant to be among the first to *Like* an update, three previous opinions could serve as social proof and diminish the predicted negative status effect that otherwise governs behavior. When a peer *Likes* an update, on the other hand, users can signal affiliation to that person and possibly reap “status points” by expressing similar preferences.

Another real life situation similar to the one we study is whether to applaud a performer or not. Even though you like the performance, there is a risk that you will be the only one who applaud; the presence of a large and unfamiliar crowd could therefore lead to self-censoring. Another person in the crowd is perhaps not very fond of the show but may still choose to applaud since his or her neighboring influential friend does so. Moreover, the fact that the number of people revealing their opinion matters in our experiment could potentially provide an alternative explanation for the rather

drastic shift in preferences in favor of extreme right parties, observed across Europe. The typical explanation for this phenomenon stresses the dissatisfaction against traditional political parties, but what if individuals had preferences matching those of the extreme right for a long time but were unwilling to express them (because of status concerns) until a large enough mass had done so as well? In this perspective, strategies to recover power would likely require different methods than the ones currently adopted. Being closely linked to status concerns, the findings from this experiment also resonate with the intriguing fact that soft drinks Coke and Pepsi attracts similar ratings in consumer blind tests, yet when labels are shown, Coke is preferred. Clearly, other factors than intrinsic utility must guide this switch in preferences, and they might originate from social status objectives.

Knowledge about the role of conformity is valuable for a number of reasons and could prove useful when designing decision environments or assessing the impact of different policies. The argument being that in some cases, we should care less about revealed preferences and focus on normative preferences (Beshears et al., 2008). One obvious example is board meetings with open voting procedures. Our results, generalized to this setting, suggest that decisions based on revealed preferences might be misleading and in the end possibly counterproductive. We also believe that a healthy suspicion to expressed opinions in general could be important, especially when taking into account the dynamic process under which specific opinions have evolved.

As a final note it is interesting to observe that conformity exists even though people act in front of a computer screen where they can easily hide (no one else actually knows whether they have seen the status update or not). This could suggest that the decision to conform may stem from a subconscious level, perhaps due to having been a successful strategy through our evolutionary past (for studies on the evolution of conformity see for example Henrich and Boyd, 2001 and Lachlan et al., 2004). Related to this topic, Bault et al. (2011) report findings which “suggest that the brain is equipped with the ability to detect and encode social signals, make social signals salient, and then, use these signals to optimize future behavior”.

7 Conclusions

In this paper, we set up a natural field experiment on the social network service Facebook to study whether people conform to previously stated opinions. We find that conforming behavior exists and does so to a significant degree. Our findings can be seen as empirical support for the theoretical model outlined in Bernheim (1994) where people care about social status in addition to intrinsic preferences. A key feature of the experimental design is the possibility to evaluate the effects along two important dimensions: the number of previous people expressing their opinion and social proximity. In accordance with social impact theory developed in Latané (1981), subjects only conform when there are sufficiently many people influencing them or when influence comes from someone they have a relation to. Importantly, we carefully address two other relevant explanations for herding behavior in this setting—limited attention and observational learning—and show that

these channels are unlikely to drive behavior in the experiment.

Hopefully, our study will stimulate further research on the importance of conformity, on Facebook and also outside the world of social media. For example, it would be interesting to collaborate with corporate entities and study if similar effects apply to messages written for commercial purposes. A more general question is whether monetary incentives affect the decision to conform. Since *Liking* a status update (unlike most other decisions) has no pecuniary cost, it may be an easy choice to follow others. A priori, however, it is difficult to say whether such incentives increase or decrease the probability to conform. Finally, we encourage attempts to test the generality and the relative strength of the peer effect found in this paper. While we used the degree centrality condition to define our peer, it should be fairly straightforward to evaluate if other conditions—such as the Bonacich network centrality (Bonacich, 1987)—developed in the literature are more, less or equally important. Another strategy is to follow the principles developed in Ballester et al. (2006) to locate the friend with maximum impact.

References

- Alevy, J., M. S. Haigh, and J. A. List (2007). Information cascades: Evidence from a field experiment with financial market professionals. *Journal of Finance* 62(1), 151–180.
- Anderson, L. R. and C. A. Holt (1997). Information cascades in the laboratory. *American Economic Review* 87(5), 847–862.
- Ariely, D. and J. Levav (2000). Sequential choice in group settings: Taking the road less traveled and less enjoyed. *Journal of Consumer Research* 27, 179–290.
- Ariely, D. and U. Simonsohn (2008). When rational sellers face nonrational buyers: Evidence from herding on ebay. *Management Science* 54(9), 1624–1637.
- Asch, S. (1952). *Social Psychology*. NJ: Prentice Hall.
- Asch, S. (1955). Opinions and social pressure. *Scientific American* 193(5), 31–35.
- Asch, S. (1956). Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs* 70(9).
- Ayres, I., S. Raseman, and A. Shih (2009). Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage. NBER Working Paper 15386.
- Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006). Who’s who in networks. wanted: The key player. *Econometrica* 74(5), 1403–1417.
- Bandiera, O., I. Barankay, and I. Rasul (2005). Social preferences and the response to incentives: Evidence from personnel data. *Quarterly Journal of Economics* 120, 917–62.

- Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics* 107(3), 797–817.
- Barber, B. M. and T. Odean (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies* 21(2), 785–818.
- Bault, N., M. Joffily, A. Rustichini, and G. Coricelli (2011). Medial prefrontal cortex and striatum mediate the influence of social comparison on the decision process. *Proceedings of the National Academy of Sciences* 108(38).
- Becker, G. S. (1974). A theory of social interactions. *Journal of Political Economy* 82(6), 1063–1093.
- Bernheim, B. D. (1994). A theory of conformity. *Journal of Political Economy* 102(5), 841–877.
- Beshears, J., J. J. Choi, D. Laibson, and B. C. Madrian (2008). How are preferences revealed? *Journal of Public Economics* 92, 1787–1794.
- Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100(5), 992–1026.
- Bonacich, P. (1987). Power and centrality: A family of measures. *American Journal of Sociology* 92(5), 1170–1182.
- Cai, H., Y. Chen, and H. Fang (2009). Observational learning: Evidence from a randomized natural field experiment. *American Economic Review* 99(3), 864–882.
- Chen, Y., F. M. Harper, J. Konstan, and S. X. Li (2010). Social comparisons and contributions to online communities: A field experiment on movielens. *American Economic Review* 100, 1358–1398.
- Cialdini, R. B., R. R. Reno, and C. A. Kallgren (1990). A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology* 58(6), 1015–1026.
- Corazzini, L. and B. Greiner (2007). Herding, social preferences and (non-)conformity. *Economics Letters* 97, 74–80.
- DellaVigna, S. and J. M. Pollet (2009). Investor inattention and friday earnings announcements. *Journal of Finance* 64, 709–749.
- Duflo, E. and E. Saez (2002). Participation and investment decisions in a retirement plan: the influence of colleagues’ choices. *Journal of Public Economics* 85, 121–148.
- Duflo, E. and E. Saez (2003). The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *Quarterly Journal of Economics* 118(3), 815–842.

- Eagly, A. H. and L. L. Carli (1981). Sex of researchers and sex-typed communications as determinants of sex differences in influenceability: A meta-analysis of social influence studies. *Psychological Bulletin* 90, 1–20.
- Freeman, L. C. (1979). Centrality in social networks conceptual clarification. *Social Networks* 1, 215–239.
- Goeree, J. K. and L. Yariv (2010). Conformity in the lab. Revise and resubmit Economic Journal.
- Goldstein, N. J., R. B. Cialdini, and V. Griskevicius (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research* 35, 472 – 482.
- Henrich, J. and R. Boyd (2001). Why people punish defectors: Weak conformist transmission can stabilize costly enforcement of norms in cooperative dilemmas. *Journal of Theoretical Biology* 208(1), 79 – 89.
- Huberman, G. and T. Regev (2001). Contagious speculation and a cure for cancer: A nonevent that made stock prices soar. *Journal of Finance* 56(1), 387–396.
- Kremer, M. and D. Levy (2008). Peer effects and alcohol use among college students. *Journal of Economic Perspectives* 22(3), 189–206.
- Lachlan, R., V. Janik, and P. Slater (2004). The evolution of conformity-enforcing behaviour in cultural communication systems. *Animal Behaviour* 68(3), 561 – 570.
- Latané, B. (1981). The psychology of social impact. *American Psychologist* 36, 343–356.
- Manski, C. F. (2000). Economic analysis of social interactions. *Journal of Economic Perspectives* 14(3), 115–136.
- Martin, R. and J. Randal (2008). How is donation behaviour affected by the donations of others? *Journal of Economic Behavior and Organization* 67, 228–238.
- Mas, A. and E. Moretti (2009). Peers at work. *American Economic Review* 99, 112–145.
- Sacerdote, B. (2001). Peer effects with random assignment: Results for dartmouth roommates. *Quarterly Journal of Economics* 116(2), 681–704.
- Salganik, M. J., P. S. Dodds, and D. J. Watts (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311, 854–856.
- Soetevent, A. R. (2006). Empirics of the identification of social interactions; an evaluation of the approaches and their results. *Journal of Economic Surveys* 20(2), 193–228.
- Weizsacker, G. (2010). Do we follow others when we should? a simple test of rational expectations. *American Economic Review* 100(5), 2340–2360.

Appendix

Figure 5: Print screen of a Facebook Homepage




Table 7: Examples of status updates from the experiment

Treatment condition	Content
T_{one}	I´m probably the only tourist who has visited Pisa but didn´t see the tower...
T_{one}	I don´t give a damn about your tax refund!
T_{one}	Party tonight. Prepare myself with intravenous drip and pain killers to be alive tomorrow...
T_{one}	Plan - to knit a hat
T_{three}	Love the warm weather. STAY!
T_{three}	A warm welcome to you, dishwasher!
T_{three}	I´ll be surprised if I don´t get an A on today´s exam
T_{three}	Towards the beach!
T_{peer}	Rhubarb desert before the running race. Hope the jogging tights still fits...
T_{peer}	Aloha Hawaii!
T_{peer}	Have the same posture as the Hunchback of Notre Dame. Lumbago please go away!
T_{peer}	Premiere on the PGA tour ;)

Note: Updates in the experiment were in Swedish, here we present translated versions.

Figure 6: Distribution of the variable Last response



Note: Observations above 30 are truncated.

Table 8: Background variables for treatment and control groups

Variable	C (s.e.)	T (s.e.)	T-C (s.e.)
Responder	0.152 (0.007)	0.158 (0.007)	0.006 (0.010)
Female	0.555 (0.009)	0.566 (0.009)	0.011 (0.013)
Peer group	0.513 (0.010)	0.495 (0.010)	-0.018 (0.015)
Year of birth	1980.894 (0.108)	1981.000 (0.104)	0.106 (0.150)
Last response	20.481 (1.272)	17.792 (1.227)	-2.690 (1.767)
Active	0.423 (0.009)	0.433 (0.009)	0.010 (0.013)


Notes: The table reports sample means and the corresponding standard errors in parentheses. *, **, *** = significant at the 10, 5 or 1 percent level in a doubled sided t-test.

Table 9: OLS regressions

Independent variable	Dependent variable: Like		
	(1)	(2)	(3)
<i>Constant</i>	0.015*** (0.004)	0.015*** (0.004)	0.015*** (0.004)
<i>T_{three}</i>	-0.004 (0.005)	-0.005 (0.005)	-0.005 (0.006)
<i>T_{peer}</i>	-0.008 (0.005)	-0.007 (0.005)	-0.007 (0.006)
<i>T_{one} × Treatment</i>	-0.003 (0.005)	-0.003 (0.005)	-0.004 (0.006)
<i>T_{three} × Treatment</i>	0.016*** (0.006)	0.016*** (0.006)	0.014** (0.006)
<i>T_{peer} × Treatment</i>	0.015** (0.006)	0.015** (0.006)	0.015** (0.006)
User FE	NO	YES	NO
Subject FE	NO	NO	YES
Observations	5,660	5,660	5,660
R-squared	0.003	0.003	0.181


Notes: The control group for *T_{one}* is used as constant. Standard errors clustered on the subject level in parenthesis. *, **, *** = significant at the 10, 5 or 1 percent level.

Figure 7: Response probability: peer group subjects only



Note: Error bars represent standard errors of the mean.

Figure 8: Average number of *Likes* per update in treatment groups by update quality: peer group subjects only



Notes: Point estimates show the average number of *Likes* generated in treatment groups. Error bars represent the corresponding standard errors. Treatment condition T_{one} is excluded.