Supplemental Material (for online publication)

Journal: Experimental Economics

Title: Liking What Others "Like": Using Facebook to Identify Determinants of Conformity

Authors: Johan Egebark and Mathias Ekström

Affiliation and contact information corresponding author: NHH Norwegian School of Economics; mathias.ekstrom@nhh.no

The purpose of this document is to provide some additional details of the procedure that can be useful in order to replicate the experiment.

As mentioned in the article we needed to collaborate with a number of Facebook users to execute the experiment. In principle, all users involved in the execution of the experiment could have been recruited from a representative population. Our conjecture, however, was that very few would have accepted, leading to substantial and non-random attrition. More importantly, since it would be difficult to retain control of information leakage to subjects, we would risk losing the benefits that come with conducting a natural field experiment. We therefore adopted the strategy of approaching a small group of candidates, and selecting candidates on the basis that we were certain they would not reveal the experiment to anyone. Notably, all of the candidates that we approached accepted, and the fact that they gave us full access to their Facebook accounts highlights the degree of mutual trust. Confidentiality was our top priority and we repeatedly instructed the users never to reveal anything about our research. It should be noted that heretofore we have not received any indications that there was ever a breach of this confidentiality. In total, there were five accounts from which we posted updates. These five users were all between 27 and 29 years old, their number of Facebook friends ranged from 58 to 204, and two of the five accounts belonged to females.

As explained in the article, we used Facebook's privacy settings to construct treatment and control groups. This allowed us to expose both groups to identical content (except for the experimental manipulation). Importantly, the members of a group could only follow the

communication within the specific group, and this communication was displayed as normal to the selected members. Hence, we have no concern that the subjects perceived the updates that we posted differently from the ordinary stream of information on the News Feed. The updates in the experiment were thus authentic and appeared as a natural part of the ongoing communication on the website. Moreover, the updates were trivial in the sense that they were short, fairly easy to interpret and did not say anything which could be perceived as sensitive, such as political opinions or religious views.

Assignment to treatment followed a two stage randomization process. First, a random draw determined which treatment (T_{one}, T_{three} or T_{friend}) would be used on a given update. Second, we randomly assigned subjects to either a control group or a treatment group. After undertaking these two steps we published the same identical content to the treatment and the control group within a few minutes. The experimental manipulation, i.e., the Likes that the treated subjects observed, were added to the updates (by us) immediately after publication. The entire procedure was repeated every time we posted a new update.

The treatments alter the initial number of Likes that subjects face. Notably, if someone in the control group decided to Like an update, all the other control group members (who read the update after this event) observed one Like instead of none (the same holds for the other treatment conditions). Naturally, the more subjects a group consists of, the higher is the probability that at least one subject in the group Likes the update (and does so fast). Thus, by dividing the treatment and control groups into smaller entities it is possible to reduce the likelihood that the subjects were exposed to Likes posted by other subjects. This is also what we did. To illustrate: If we posted an update from a user with 120 friends, 60 subjects were randomly assigned to the treatment group and 60 to the control group. We then continued by splitting both the treatment and the control group in two. Hence, we ended up with two treatment groups with 30 subjects in each, and two control groups with 30 subjects in each. If one of the subjects in the control group Liked the update that would now only affect the remaining 29 subjects in his or her subgroup, instead of all the

59 subjects in the control group. On average these subgroups consisted of 30 subjects, and the strategy turned out to be successful: control group subjects were unexposed to Likes in more than 70 percent of all possible cases. Notably, the advent of endogenous Likes in the control group will, if anything, introduce a downward bias in our estimates. Note also that the unit of randomization is still the individual subject. That is, each subject is randomly assigned to be exposed to an update with initial Likes or not.