Identifying Peer Influence in Massive Online Social Networks: A Platform for Randomized Experimentation on Facebook

Sinan Aral
NYU Stern School of Business and MIT, 44 West 4th Street Room: 8-81, New York, NY 10012
sinan@stern.nyu.edu

Dylan Walker
NYU Stern School of Business, 44 West 4th Street Room: 8-80, New York, NY 10012
dwalker@stern.nyu.edu

Extended Abstract of Research in Progress
Submitted to WISE 2009, Phoenix, Arizona.

Introduction

Identifying peer to peer influence in social networks is a difficult and fundamental problem in numerous disciplines including network theory, marketing, epidemiology and diffusion research. It is widely believed that the cleanest way to examine the magnitude of peer effects and the social, structural and personal circumstances under which influence is more or less pronounced is to observe large scale, real world randomized trials of person to person communications intended to influence particular economic decisions. In this research we explore this approach to the identification of peer effects using random experiments in a massive online social network.

Recent research demonstrates that individual economic behaviors (such as product adoption or work performance) tend to cluster in social networks, in both network space (assortative mixing) and in time (temporal clustering). Yet, while evidence of assortative mixing and temporal clustering of economic outcomes in networks may indicate peer to peer influence, these outcomes may also be explained by homophily - the demographic, technological, behavioral, and biological similarities of linked nodes (Jackson 2008). If ties are more likely between similar nodes, their outcomes could be correlated due to inherent similarities in their characteristics rather than as a consequence of their interactions. On one hand, linked nodes may directly influence one another to exhibit similar outcomes, creating viral contagions. On the other hand, linked nodes may simply have greater likelihoods of displaying correlated outcomes, in time and in network space, as a consequence of their similarities.

Several sophisticated econometric techniques have been developed in the hopes of untangling these explanations. Work on the identification of peer effects in networks (e.g. Brock and Durlauf 2001; Bramoulle 2007) has developed following (Manski 1993) and (Frank and Strauss 1986), or models of the co-evolution of networks and behaviors by (Snijders 2006), dynamic matched sampling techniques (Aral et al. 2008), methods based on exogenous shocks to peers (Tucker 2008), or examination of natural experiments (e.g. Sacerdote 2001). However, each of these methods suffers from its own limitations: identification conditions are strict, methods are not typically scalable to large networks, observation of naturally occurring random assignment is rare, and shocks to peers used as instruments are rarely truly exogenous because social relationships typically signal unobserved reasons why these shocks should be correlated amongst peers. We look to randomized trials involving millions of users of an online social network to overcome some of these limitations.

Over the last year we have built and tested a platform for the execution of randomized trials of social influence on Facebook.com. The platform utilizes multiple customized Facebook applications in concert to observe user behavior, communications traffic and outcomes related to social influence. We have designed several interrelated randomized trials and have collected data on social network relationships and online profile attributes from over 10 million Facebook users. The data comprise tens of thousands of direct users of our experimental applications and their social network neighbors as well as a rich set of covariates describing individual demographics, education and employment histories, interests, tastes and social behavior. Our initial results indicate that randomized trials can be quite effective in identifying peer influence and work we will complete by December 2009 should produce interesting tests of several well known hypotheses concerning when influence is more or less likely to be observed.

A Platform for Randomized Experimentation

The platform for randomized experimentation consists of two separate Facebook applications which we exhaustively monitor. As users adopt and use these applications we not only track all of their behaviors on the applications, but also call their personal profile information as well and their complete social networks and the profile information of their friends. In monitoring these applications we observe users sending viral messages to their friends notifying them of their use of a given application or inviting them to engage the application in some way. By randomizing application features (for example by turning viral messaging on and off for randomized control and experimental groups) we can test the conditions under which receipt of viral messages influences a user's friends to adopt an application or to engage in particular behaviors associated with an application. The experimentation platform has been designed to record several channels of communication through which application users may encourage application adoption or other behaviors by their peers. To investigate the efficacy of different types of viral messages that encourage product adoption, we define two broad categories of viral communications.

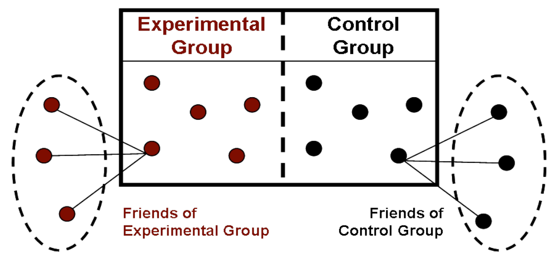

Fig.1. Illustrates the concept behind the design of our platform for random experimentation online. The key insight is that while the experimental group receives the 'treatment,' (e.g. an advertisement), the effect of peer influence or behavioral contagion is tested by measuring the uptake of the behavior by their friends compared to the friends of a untreated control group who do not receive the message.

Passive viral messages are those that are not directly addressed from one user to another but are impersonal messages automatically generated by the application to notify a user's friends of that user's activity. Examples of passive viral messages include wall posts, newsfeeds, and notifications of a user's activity sent to friends by the application (e.g. 'your friend has added a new survey… or …uploaded a new picture to application X'). Passive viral messages in online social networks are indicators of product use by peers that are visible to users in a given local network neighborhood, and are analogous to other forms of declaration in offline social interactions, such as wearing product-branded t-shirts or displaying product-related logos.

In contrast, *active viral messages* (such as requests and invitations) are direct solicitations by a user to his friends to adopt an application, join a group or engage in some other behavior on the application. These messages are personalized recommendations initiated and targeted by the user toward her friends and are particularly popular in cases where there is a perceived network value from using an application in conjunction with others.

We divide the population of application users into three categories corresponding to increasing levels of viral interactivity. As users download the applications they are randomly assigned to receive one of three versions with various viral features turned on or off. Users designated as "*non-viral*" cannot send (intentionally or otherwise) either active or passive viral communications. The behavior of these users and their friends serves as a comparative baseline. Users designated as "*passive viral*" may send passive viral messages triggered by actions taken within the application, but cannot send active viral messages such as personally directed requests or invites. Receipt of passive viral messages is also randomized in order to examine how a recipient's personal characteristics affect the influence of passive viral messages on their adoption and use. Passive viral messages are sent to a random subset of 20 of that user's friends. "*Active viral*" users receive, in addition to passive viral messaging capabilities, the option to send direct solicitations to friends in a manner that is integrated into the functionality of the application. Each of these time-stamped messages, their content, lists of recipients and recipients' subsequent behavioral responses are documented in our data.

In addition to investigation of influence through explicit channels of viral communication, we also investigate influence through latent or unrecorded channels of communication. To do so we randomly divide the population of application users into two categories irrespective of their viral status: "*ad*" and "*no-ad*" populations. Application users in the *ad* population are displayed an advertisement for a second application, while users in the *no-ad* population are displayed blank whitespace in the same location. This random assignment allows us to test whether friends of those who saw the ad are more likely to adopt the second application irrespective of the observed channels of viral communication. Application users are randomly assigned to one of these categories upon adoption and for the duration of application use.

Data

The data was collected in collaboration with a Facebook Application Developer. We first sampled the local ego network data of all users of one of several applications developed by this application developer. This sampled a social network of 10.5 million users and 280 million unique relationships. To obtain profile data from application users and friends of application users, in accordance with Facebook's data collection policy, ego and peer profile data was collected for each application user within a thirty minute window from the user's last access to the application. Using this procedure the profile data of 10.5 million Facebook users was obtained, including a rich set of detailed information regarding demographics (*age, gender, current location*), school and employment history, activities and interests, views (*political, religious*), product tastes (*movies, music, television shows, books*), and social participation (*communication activity, relationship status, online group membership, photo co-appearance*). Subsequent data collection throughout the course of our study will include weekly longitudinal changes to profile data over an extended period of time.

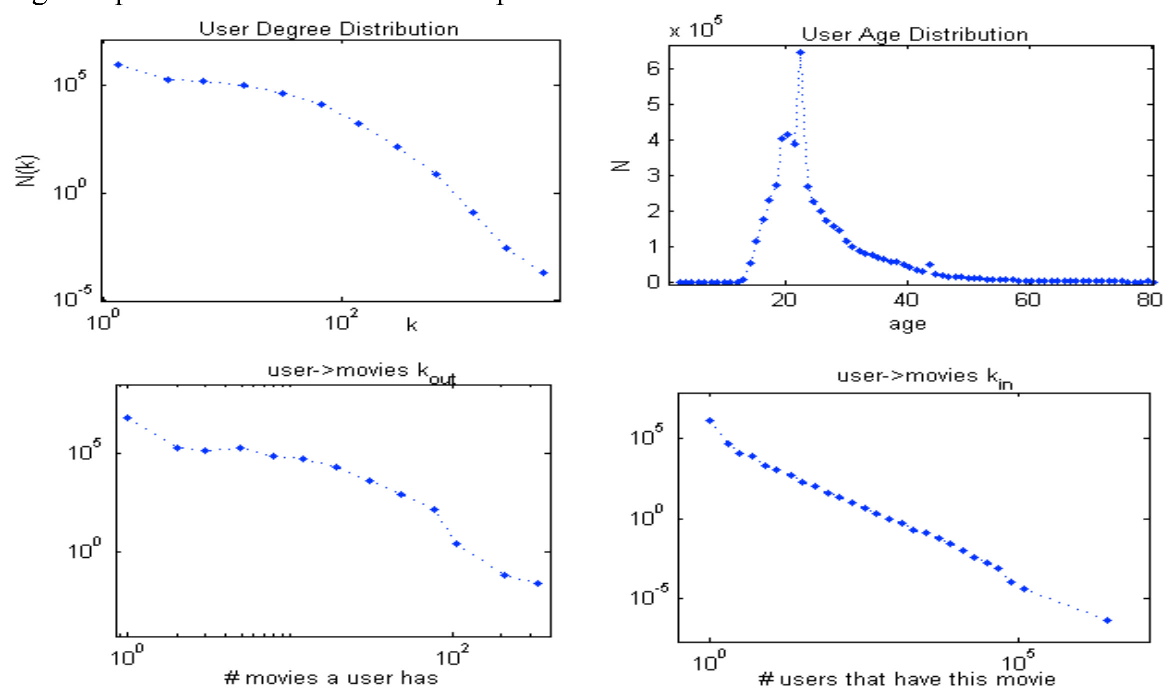

A small preliminary trial run (described below) was performed using the pre-launch campaign of a small Facebook application, during the period of July, 2, 2009 to Aug, 2, 2009. During this time user profile information for new users was collected and the viral communications between users of this application and their local network peers was recorded. The response of local network peers was also recording, including click through responses to viral communications, inbound traffic to the application, and application installation data. Sample descriptive statistics shown in Figure 2 describe the user degree distribution, age distribution, and the distribution of the number of movies a user likes and the number of users who like a given movie.

Influence Identification and Tests of Influence through Different Channels

Identifying Influence in Different Viral Channels (Active, Passive and Baseline)

We identify the influence effects of different viral messaging channels on the application adoption of users' friends by comparing the relative fraction of users' friends who adopt across treatment populations randomly assigned to *Active*, *Passive* and *Baseline* versions of Application 1. For example, to identify the effect of peer influence through *Active* viral channels we compare the relative adoption fraction by peers in the local networks of *Active* viral users $F(f\ riends\ of\ i | V_i = active, \{X_i\})$ to the relative adoption fraction by peers in the local networks of *Baseline* user $F(f\ riends\ of\ i | V_i = non-active, \{X_i\})$. Given random assignment, the difference in adoption rates across these two treatment populations represents the average treatment effect of turning on the *Active* viral messaging features of the application. We conduct similar experiments to test the effects of *Passive* viral messaging (by comparing *Passive* users to *Baseline* users) and the relative difference in the effectiveness of *Active* versus *Passive* messaging (by comparing *Active* users to *Passive* users).

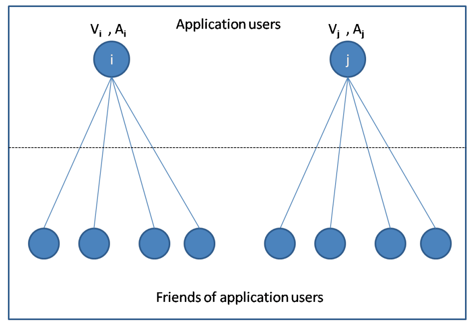

Given these baseline estimates, we test how message senders' individual characteristics moderate the influence effects of each viral channel. For example, we test whether men or women are more influential when sending *Active* or *Passive* viral messages, or how senders' age affects the influence of different viral channels. The result is a comprehensive randomized trial of the influence effects of different viral messaging channels evaluated across different segments of the user population.

Identifying Susceptibility to Influence by Randomizing the Recipients of Passive Viral Messages

We identify susceptibility of different types of peers to the influence of passive viral messages sent to them by their friends who are application users, by comparing the average click through response and subsequent application adoption of peers that received passive viral messages to those of peers of the same user who did not receive passive viral messages. As passive viral messages are sent to a random subset of an application user's peers, the difference in application adoption between recipient and non-recipient peers represents the average treatment effect of receiving a passive viral message. Given a baseline measure for this treatment effect, we test how the message recipients' individual characteristics moderate the likelihood of click through response and application adoption as described above. We also use the randomization of the receipt of passive viral messages to estimate the moderating effects of dyadic characteristics on the magnitude of peer influence. For example, we ask: are recipients of passive viral messages who are more similar in age or of the same gender as the sender more susceptible to peer influence than those that display less homophily to their respective senders?

Identifying Influence in Traditional Advertising through Non-Viral Peer to Peer Communication

We identify the influence effects of traditional advertising channels through (non-viral) peer to peer communications by randomly assigning users of Application 1 to *ad* and *no-ad* treatment populations, corresponding to versions of Application 1 that do and do not display advertisements for Application 2, and then comparing the relative fraction of users' friends that adopt Application 2. The average difference in the adoption fraction of an application users' local network between *ad* and *no-ad* populations represents the average treatment effect of an application user viewing an advertisement on the adoption likelihood of that user's friends. Given baseline comparisons between treatment populations, we assess the moderating effect of individual characteristics of an application user viewing an advertisement on the likelihood of adoption by that user's local network peers. For example, advertisement-viewing application users with higher levels of online participation or communication activity may be more influential in encouraging local network peers to adopt Application 2.

Initial results

We have performed a preliminary analysis of viral communication data, user response traffic and adoption status during the pre-launch phase of Application 1. During the pre-launch phase, 2,823 users sent 84,732 passive viral messages. These passive viral messages (notifications, newsfeed items, profile displays) generated a click through response from 2,496 individuals. A significant portion (47%) of these click throughs came from Facebook users who had not previously adopted the application. Of click throughs by non-application users, 62% resulted in adoption of the application, indicating a potentially significant response to passive viral messaging. More specifically in terms of the specific type of passive viral message, 65% of notifications induced an adoption event from a non-application user, 61% of newsfeed items induced an adoption event from a non-application user, and 53% of profile box items clicked by non-application users resulted in an installation. During the pre-launch phase, systems to record active viral invitations were not yet in place. These initial results demonstrate that the potential for influence is significant in this network. The real experiments are now running and we expect to have a full set of results on the analyses described above by the end of November 2009.

Plan of Work

We intend to complete the full set of experiments a month before WISE. We are currently collecting data on user click through and adoption behavior for Application 2. In addition to viral studies, this newly acquired data will allow us to perform the study to identify influence effects in traditional advertising scenarios. In the coming months, we will also leverage a larger amount of user profile data that has been collected since the initial pre-launch period to study the moderating effects of individual attributes on both viral and non-viral channels of influence. Beyond the work described above, we will implement a randomly served advertisement for Application 1 to users of Application 2. This will allow us to directly compare the influence effects of viral marketing strategies to those of more traditional advertising campaigns.

References

Aral, S., Muchnik, L., & Sundararajan, A. 2008. "Influence Dynamics in Large Complex Networks." Workshop on Information Systems Economics, Paris, France.

Bramoullé Y, Djebbari H, & Fortin B (2007) Identification of Peer Effects through Social Networks. IZA Discussion Papers 2652, Institute for the Study of Labor (IZA).

Brock WA & Durlauf SN (2001) Discrete Choice with Social Interactions. *The Review of Economic Studies* 68(2):235-260.

Frank O & Strauss D (1986) Markov Graphs. *Journal of the American Statistical Association* 81(395):832-842.

Jackson MO (2008) Average Distance, Diameter, and Clustering in Social Networks with Homophily. *Internet and Network Economics*, (Springer, Berlin / Heidelberg), Vol 5385/2008, pp 4-11.

Manski CF (1993) Identification of Endogenous Social Effects: The Reflection Problem. *Review of Economic Studies* 60(3):531-542.

Sacerdote B (2001) Peer Effects with Random Assignment: Results for Dartmouth Roommates. *Quarterly Journal of Economics* 116(2):681-704.

Snijders T, Steglich C, & Schweinberger M (2006) Modeling the co-evolution of networks and behavior. *Longitudinal models in the behavioral and related sciences*, eds Montfort Kv, Oud H, & Satorra A (Lawrence Erlbaum).

Tucker C (2008) Identifying Formal and Informal Influence in Technology Adoption with Network Externalities. *MANAGEMENT SCIENCE* 54(12): 2024-2038.