

Chapman University

Chapman University Digital Commons

Business Faculty Articles and Research

Business

10-6-2011

Identifying Social Influence in Networks Using Randomized Experiments

Sinan Aral

Dylan Walker

Follow this and additional works at: https://digitalcommons.chapman.edu/business_articles



Part of the [Advertising and Promotion Management Commons](#), [Communication Technology and New Media Commons](#), [Databases and Information Systems Commons](#), and the [Marketing Commons](#)

Identifying Social Influence in Networks Using Randomized Experiments

Comments

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in *IEEE Intelligent Systems*, volume 26, issue 5, in 2011 following peer review. This article may not exactly replicate the final published version. The definitive publisher-authenticated version is available online at <https://doi.org/10.1109/MIS.2011.89>.

Copyright

© 2011 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Identifying Social Influence in Networks Using Randomized Experiments

Sinan Aral
Stern School of Business, New York University
sinan@stern.nyu.edu

Dylan Walker
Stern School of Business, New York University
dwalker@stern.nyu.edu

Identifying causal estimates of peer-to-peer influence in networks is critical to marketing strategy, public policy and beyond. Unfortunately, separating correlation from causation in networked data is complicated. We argue that randomized experimentation in networks, made possible by the digitization of human interaction at population scale, can dramatically improve our understanding of the ebb and flow of market trends, product adoption and diffusion, the spread of health behaviors, the productivity of information workers and whether or not particular individuals in a social network have a disproportionate amount of influence on the system. We also discuss some of the complications that arise when conducting randomized experiments in networks by describing an experiment designed to test how different viral product design strategies affect peer influence and social contagion in new product diffusion.

Key words: Peer Influence, Social Contagion, Social Networks, Endogeneity, Causality, Randomized Experiments.

Forthcoming in IEEE Intelligent Systems

Introduction

The recent availability of massive amounts of networked data generated by email, instant messaging, mobile phone communications, micro-blogs and online social networks is enabling studies of population level human interaction on scales orders of magnitude greater than what was previously possible (Lazer et. al 2009). One important goal of applying statistical inference techniques to large networked datasets is to understand how behavioral contagions spread in human social networks. More precisely, understanding how people influence or are influenced by their peers can help us understand the ebb and flow of market trends, product adoption and diffusion, the spread of health behaviors such as smoking and exercise, the productivity of information workers and whether or not particular individuals in a social network have a disproportionate amount of influence on the system.

However, if we are to truly understand how social interaction and peer influence shape behavioral dynamics in large networked populations, we must be able to separate correlation from causation. By now, there is abundant empirical evidence that human behaviors tend to cluster in network space and in time. Recent studies have shown behavioral clustering for trends in obesity, smoking, product adoption, happiness, economic development and more. But, several alternative explanations besides peer influence and social contagion could also explain these patterns. For example, people tend to make friends with those that are like themselves, a social process called homophily. As a result, preferences and behaviors cluster in networks because we tend to choose friends who like the same things and behave in the same ways that we do. Peers are also more likely to be exposed to the same external stimuli. We tend to make friends with people we work with or who live nearby. As a result, our exposure to changes in health benefit plans at work or new restaurants opening in our neighborhoods is correlated with that of our friends. Our correlated exposure to such external stimuli can in turn drive patterns of correlation in our behaviors over time.

Separating correlation from causation in patterns of networked behavior over time is important for two reasons. The causal structure of the underlying dynamic process that governs the spread of a

behavior implies 1) different diffusion properties for the behavior – where it is likely to spread next and 2) different optimal containment and promotion policies. Consider two hypothetical scenarios in which marketing data show that the adoption of a new product is significantly correlated amongst linked consumers. In one scenario 90% of this correlation is explained by peer influence – friends convincing friends to adopt the product, and only 10% is explained by correlated preferences. In an alternative scenario, 10% of the correlation is explained by peer influence while 90% is explained by correlated preferences. In the first scenario, a peer-to-peer marketing strategy that creates incentives for adopters to spread positive word-of-mouth (WOM) about the product may be quite effective. In the second scenario, a traditional market segmentation strategy based on observable characteristics of consumers may be much more effective and a peer-to-peer strategy may not work at all. The same logic applies to whether the National Institutes of Health should allocate substantial funding to peer-to-peer obesity prevention programs, whether the American Lung Association should focus on peer smoking prevention strategies and whether advertisements such as the Department of Transportation’s ‘Friends Don’t Let Friends Drive Drunk’ campaign are likely to succeed. Developing methods that robustly identify causal estimates of peer-to-peer influence in social networks will therefore have dramatic implications for marketing strategy, public policy and organizations. Recent scientific debates about the veracity of a series of high profile studies of contagion in networks (e.g. Christakis and Fowler 2007, Lyons 2011) serve to highlight both the scientific stakes and the public policy implications of this important area of inquiry.

We discuss the difficulties of causal statistical estimation of peer influence in networks, review current approaches to establishing causality in observational data and propose methods based on randomized experimentation designed to control observable and unobservable confounding factors. We then present evidence from a randomized trial of peer influence in product adoption among 1.4 million friends on Facebook.com (Aral and Walker 2011a) which illustrates how randomized trials can be used to identify when and under what conditions influence propagates over network ties.

Endogeneity and the Identification of Peer Influence in Networks

Several sources of bias in both cross sectional and longitudinal data on interactions and outcomes among peers can confound assessments of peer influence and social contagion in networks including homophily (Aral et al. 2009, Shalizi and Thomas 2011), simultaneity (Godes and Mayzlin 2004), unobserved heterogeneity (Van den Bulte and Lilien 2001), truncation (Van den Bulte and Iyengar 2010) and other contextual and correlated effects (Manski 1993). If uncorrected, these biases can lead researchers to incorrectly attribute observed correlations to the influence of individuals on their peers resulting in misinterpretations of social contagion and the treatment effects of viral marketing campaigns, peer-to-peer behavioral interventions or viral product design strategies (Aral 2011).

Several approaches to the identification of peer effects have been proposed including peer effects models (e.g. Bramouille et al 2009), actor-oriented models (e.g. Snijders et al. 2006), instrumental variables methods based on natural experiments (e.g. Tucker 2008), dynamic matched sample estimation (Aral et al. 2009), structural models (e.g. Ghose and Han 2010) and ad hoc approaches (Christakis and Fowler 2007). However, randomized trials are believed to be the most effective way to obtain unbiased estimates of causal peer effects (Duflo et al 2006).

Randomized Trials of Peer Influence

Randomized trials in networks create unique estimation challenges that require particular attention to experimental design. Two primary experimental designs have been used in network studies – *network structure randomization* and *treatment randomization*. Network structure randomization tests how particular network configurations affect behavioral dynamics by randomly assigning experimental subjects into different network structures or different positions within a network structure (Kearns et al 2006, Centola 2010, Suri and Watts 2011). Treatment randomization designs on the other hand estimate the propagation of behaviors through peer influence by measuring the effect of a randomly assigned treatment on the behavior of peers of the treated (Duflo and Saez 2003, Leider et al 2009, Aral and Walker 2011 a, b). We focus here on the challenges inherent in creating successful treatment

randomization designs by considering a detailed example of an experiment that we conducted on the popular social networking website Facebook.com.

Treatment Randomization – An Experiment Measuring Peer Influence in Product Adoption

Experimental Design and Procedure

We designed a randomized field experiment on Facebook.com to test the effectiveness of different viral messaging capabilities in creating peer influence and social contagion in a product's adoption among the 1.4 million friends of 9,687 experimental users. We studied a commercial Facebook application which lets users share information and opinions about the film industry. The design enabled experimental group users to send broadcast notifications and personal invitations to their neighbors, while disabling these features for the baseline control group. The application then recorded data on the use of these viral features by experimental group users, as well as click stream data on recipient responses to viral messages, and their subsequent adoption and use of the application for all neighbors of experimental and control group users. This facilitated analysis of the average treatment effect of enabling notifications and invitations on peer adoption and network propagation as well as the mechanisms by which a particular viral channel influenced peer behavior. Below we outline several challenges and solutions to analyzing the results of field experiments that employ treatment randomization.

Inside-Out Experimental Designs

Randomized trials typically estimate the effect of a treatment on the treated. In contrast, measuring social influence requires experiments designed to measure the effect of a treatment on *peers* of the treated. The conventional approach to estimating peer influence and social contagion involves modeling individuals' likelihood or rate of adoption of a behavior as a function of their own personal characteristics and their exposure to social influence from their peers. This is typically accomplished by estimating the influence of an individual's social environment 'inward' on their own propensity to adopt the behavior or product. This conventional approach is infeasible in experimental settings because it is

difficult to comprehensively control the network environments of each user in the study. Fortunately, it is feasible to treat a user and observe the effect of treatment on the outcomes of their peers (illustrated in Figure 1). However, this approach creates two additional estimation challenges. First, it is unlikely that the baseline hazard of adoption remains constant from the first adopter to the second adopter to the third adopter in the local network around a recruited user. Second, peers of recruited users are not independent.

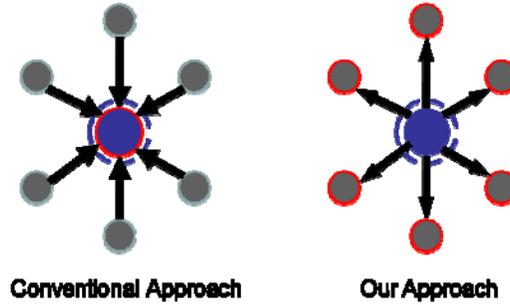


Figure 1. Arrows indicate the potential flow of influence that the experiment is designed to detect. The solid blue circle in the center represents the treated user and the red outlines indicate measured treatment effects.

We therefore measured the treatment effect of adding viral features to the applications of the treated on the adoption and usage behaviors of their peers by estimating the following variance corrected stratified proportional hazards model which accounts for the lack of independence among the multiple clustered failure times in the data and allows the baseline hazards to vary by adoption event across stages of the diffusion process:

$$\lambda_k(t, X_{ki}) = \lambda_{0k}(t) e^{X_{ki}\beta}$$

The model is stratified over the K adoption events. λ_{0k} represents the baseline hazard of the k th adoption event (i 's k th friend adopting) and X_{ki} represents a vector of covariates affecting the rate of adoption of i 's neighbors including i 's viral treatment status (whether they were assigned to active, passive, or baseline treatments).

Preventing Selection Effects

Selection effects can occur during both the initial recruitment of subjects into the experiment and when subjects enter through subsequent peer recruitment (for example, in our experiment, when peers adopt in response to a viral message). Steps must be taken to mitigate and measure both possibilities. The recruitment campaign we employed was designed to reach a representative audience of Facebook users. Observable attributes of recruited peers were measured and compared to population distributions to test for selection effects that occur through peer recruitment. However, peer-recruited populations may still differ on unobservable dimensions. For example, in our study, it could be that users and peers of users who use the viral features are systematically different from randomly selected Facebook users. To avoid these sources of selection bias, peers of recruited users only contributed to the local network peer adoption of originally recruited users and were not themselves used as test subjects.

Preventing Contamination and Leakage

In randomized experiments in networks, users assigned to different treatment groups may not be strictly isolated from one another and information leakage through indirect network pathways may contaminate results. In traditional studies, whether or not the network is measured, relationships may still exist between treatment and control populations that create leakage. One benefit of randomized experiments in networks is that researchers systematically observe how individuals in the study are connected, enabling us to mitigate leakage.

Since treatment is randomized, leakage will be uncorrelated with treatment assignment and cannot account for observed differences in responses to treatments. While it is possible that leakage will on average provide some common information to peers of treated users uniformly across treatment designations, this should only serve to make estimates across treatment groups more conservative by reducing differences between control and treatment groups. In addition, leakage is likely to decay with network distance (Aral et al 2007). Nonetheless, leakage could downward bias estimates of treatment effects and several measures can help prevent it.

First, in our social influence models, we only examined peers of initially recruited adopters. This helps avoid selection bias and also excludes individuals (and their potential adopter peers) that adopt in chains within a local neighborhood, lessening leakage effects.

Second, we accounted for users with multiple treated peers. Users may be peers of multiple treated users from different treatment groups making it impossible to link their observed outcomes to a single treatment. Furthermore, while peers of multiple treated users belonging to the same treatment group are clearly classified, measurements of their response may be incorrectly estimated as a consequence of being subject to influence from multiple treated friends.

Peer treatment randomization prevents us from guaranteeing that all treated friends of a peer will receive the same treatment. Consequently we labeled peers with multiple treated friends as contaminated as soon as they become so and excluded them from our analysis by censoring them. This procedure could underestimate the effect of clusters of adoption on the time to adoption or the number of adopters in a local network, but if this is the case, it will do so identically across all treatment types. Furthermore, as treatment groups are randomized, there can be no systematic correlation between the treatment of a user and that of her adopter peer.

The procedure we adopted for designating contaminated peers is detailed in Figure 2. In panel (a) for times $t > t_1$, peer 2 has multiple treated peers (R and 1) who may have been assigned different treatments. Peer 2 is therefore considered contaminated for times $t > t_1$. Panel (b) is identical, but no link exists between peers 1 and 2 and consequently neither user is considered contaminated. This procedure retains the maximal empirical support for hazard rate estimation and parameterizes our ignorance of what might happen subsequent to a user's contamination. These procedures enable a tightly controlled randomized trial of peer influence that addresses potential selection and leakage effects.

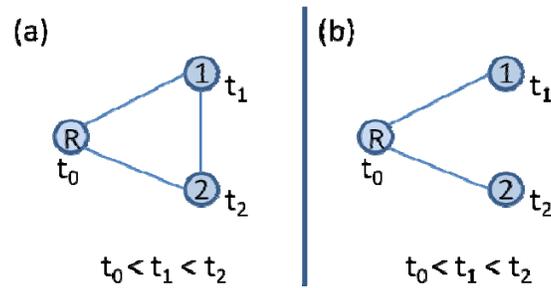


Figure 2. The initially recruited adopter, labeled R, adopts at time t_0 . Two peers of user R, labeled 1 and 2, adopt at subsequent times t_1 and t_2 respectively.

Results

Local Peer Influence

In our experiment, viral features that require more activity on the part of the user and are more personalized create greater marginal increases in the likelihood of adoption per message, but generate fewer total messages, creating countervailing effects on peer influence. On average, notifications, which are less personalized but also require less user effort, generated a 246% increase in local peer influence and social contagion while adding invitations, which are more personalized but require more user effort, only generated an additional 98% increase. Although invitations were more effective in encouraging adoption per message and were correlated with more sustained product use, they were used less often and therefore generated less total peer adoption in the network than passive-broadcast messaging.

Modeling Total Expected Adoption

Randomizing the treatment of peers of the treated enabled us to collect a greater number of subjects with randomized treatment status. On the other hand, maintaining consistency in the treatment amongst linked peers would have enabled us to directly estimate the total adoption and diffusion of a given version of the product in the population. Whether to randomize the treatment of peers of the treated or to maintain consistency in the treatment amongst linked peers is an important design choice in treatment randomization studies. This choice requires trading off the robustness of estimates of local peer effects with the ability to estimate the aggregate adoption of a behavior in the population. Randomizing

the treatment of recruited peers creates more randomized observations and mitigates the perpetuation of selection bias over successive stages of the diffusion process. On the other hand, maintaining consistent treatments for recruited peers enables direct estimation of total adoption and diffusion while reducing the chances of leakage from an adopter to a treated peer that receives a different treatment condition.

We randomized peer treatments to robustly estimate local treatment effects. However, one can model the total expected adoption of each type of application in the population by making a few reasonable assumptions. If we assume the average number of peer-recruited adopters of a given user depends only on that user's viral state v and his distance s from an originally recruited user,¹ we can model the total expected number of adopters of a given version of the product v as follows:

$$N_v^{total} = N_v^0 \left(1 + \sum_{s=1}^{\infty} \prod_{s=1}^s \langle n_{v,s} \rangle_v \right)$$

N_v^0 is the number of originally recruited users (adopter step $s = 0$) who have been randomly assigned to viral treatment state v . $\langle n_{v,s} \rangle_v$ is the average number of peers of a user with viral state v and distance s (the path length to an originally recruited user) who have adopted the application. One can then compare the total expected number of adopters under different treatment designations by assuming an equal number of originally recruited users for each treatment group, or comparing the ratios N_v^{total} / N_v^0 . We estimate the passive-broadcast treatment was associated with a 9.13% increase in total diffusion relative to the baseline treatment, while active-personalized treatment was associated with a 12.03% increase relative to the baseline.

Conclusion

Understanding how behaviors spread through social networks is essential to marketing strategy and public policy. But, separating correlation from causation in estimates of social contagion is critical to

¹ The reliance of average peer recruitment, $\langle n_{v,s} \rangle_v$, on the distance s from originally recruited subjects allows for the possibility that selection effects may exist and vary (e.g., be compounded) along chains of sequential user adoptions. This method further assumes that peer recruitment is Markovian in treatment status and thus depends only on a user's viral state and not on the viral states of users in the chain of adoptions that preceded and led to his own adoption.

designing successful policy interventions. Randomized trials in social networks are a promising strategy for understanding the dynamics of social influence at scale. But, such studies are not without their complications. Careful experiments, designed to maintain external validity while avoiding contamination, leakage, selection effects and the challenges of statistical estimation in datasets with complex interdependencies between observations is essential for our robust scientific understanding of social dynamics at scale and thus for informed policymaking. We therefore hope future research on human social dynamics in networks will make causal statistical inference a centerpiece of study design.

References

- Aral, S. (2011) "Identifying Social Influence: A Comment on Opinion Leadership and Social Contagion in New Product Diffusion." *Marketing Science*. 30(2) 217-223.
- Aral, S., E. Brynjolfsson, M. Van Alstyne. (2007). Productivity effects of information diffusion in networks. Proceedings of the 28th Annual International Conference on Information Systems. Montreal, CA.
- Aral, S., L. Muchnik, A. Sundararajan (2009) "Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks." *P. Natl. Acad. Sci. USA* 106 21544-21549.
- Aral, S. and Van Alstyne, M. W., (2011) "The Diversity-Bandwidth Tradeoff." *American Journal of Sociology*. In Press. <http://ssrn.com/abstract=958158>.
- Aral, S. & Walker, D. (2011a) "Creating Social Contagion through Viral Product Design: A Randomized Trial of Peer Influence in Networks." *Management Science*, In Press. <http://ssrn.com/abstract=1564856>.
- Aral, S. & Walker, D. (2011b) "An Experimental Method for Identifying Influential and Susceptible Members of Online Social Networks." *Working Paper*. NYU Stern School of Business.
- Bramouille Y, Djebbari H, & Fortin B (2009) "Identification of Peer Effects through Social Networks." *Journal of Econometrics* 150(1): 41-55.
- Centola, D. (2010) "The Spread of Behavior in an Online Social Network Experiment." *Science*, 329(5996): 1194-1197.
- Christakis NA & Fowler JH (2007) "The Spread of Obesity in a Large Social Network over 32 Years." *New England Journal of Medicine* 357(4): 370-379.
- Duflo, E. and E. Saez (2003) "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence From a Randomized Experiment" *Quart. Journal of Economics*, 118: 815-842.
- Duflo E, Glennerster R, & Kremer M (2006) "Using Randomization in Development Economics Research: A Toolkit." *Handbook of Developmental Economics*, Elsevier.
- Ghose, A. and Han, S-P. (2010). "An Empirical Analysis of User Content Generation and Usage Behavior in Mobile Media." *Working Paper*. NYU Stern School of Business.
- Godes D & Mayzlin D (2004) "Using Online Conversations to Study Word-of-Mouth Communication." *Marketing Science* 23(4): 545-560.
- Kearns, M., Suri, S and N. Montfort (2006) "An Experimental Study of the Coloring Problem on Human Subject Networks" *Science*, 313(5788): 824-827.
- Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabasi, A.L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., Jebara, A., King, G., Macy, M., Roy, D., Van Alstyne, M. (2009). "Computational Social Science," *Science*, February 6: 721-722.
- Leider, S., Mobius, M., Rosenblat, T., and QA Do (2009) "Directed Altruism and Enforced Reciprocity in Social Networks" *Quart. Journal of Economics*, 124(4): 1815-1851.

- Lyons, R (2011) "The Spread of Evidence-Poor Medicine via Flawed Social-Network Analysis," *Statistics, Politics, and Policy*: Vol. 2: Iss. 1, Article 2.
- Manski CF (1993) "Identification Problems in the Social Sciences." *Sociological Methodology* 23:1-56.
- Shalizi, C. and A. C. Thomas, (2011) "Homophily and Contagion Are Generically Confounded in Observational Social Network Studies", *Sociological Methods and Research* 40: 211-239.
- 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).
- Suri, S and D.J. Watts (2011) "Cooperation and Contagion in Web-Based, Networked Public Goods Experiments." *PLoS One*. 6(3): e16836.
- Tucker C (2008) "Identifying Formal and Informal Influence in Technology Adoption with Network Externalities." *Management Science* 54(12): 2024-2038.
- Van den Bulte, and R, Iyengar (2011), Tricked by Truncation: Spurious Duration Dependence and Social Contagion in Hazard Models, *Marketing Science*, 30 (2), 233-248.
- Van den Bulte C and Lilien GL (2001) "Medical Innovation Revisited: Social Contagion versus Marketing Effort." *The American Journal of Sociology* 106(5):1409-1435.