



Identifying Social Influence in Networks Using Randomized Experiments

Sinan Aral and Dylan Walker, *New York University*

The recent availability of massive amounts of networked data generated by email, instant messaging, mobile phone communications, microblogs, and online social networks is enabling studies of population-level human interaction on scales orders of magnitude greater than what was previously possible.^{1,2} 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 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.

Still, several alternative explanations besides peer influence and social contagion could also explain these patterns. For example, people tend to make friends with those who 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 preferences and 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 among linked consumers. In one scenario, 90 percent of this correlation is explained by peer influence—friends convincing friends to adopt the product—and only 10 percent is explained by correlated preferences. In an alternative scenario, 10 percent of the correlation is explained by peer influence while 90 percent 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 might be effective. In the second scenario, a traditional market segmentation strategy based on observable characteristics of consumers might be much more effective and a peer-to-peer strategy might 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, the American Lung Association should focus on peer smoking prevention strategies, and 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^{3,4} serve to highlight both the scientific stakes and the public-policy implications of this important area of inquiry.

In this article, 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.com⁵ that illustrates how randomized trials can be used to identify when and under what conditions influence propagates over network ties.

Endogeneity and Identifying 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,^{6,7} simultaneity,⁸ unobserved heterogeneity,⁹ truncation,¹⁰ and other contextual and correlated effects.¹¹ 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.¹²

Several approaches to identifying peer effects have been proposed including peer-effects models,¹³ actor-oriented models,¹⁴ instrumental variables methods based on natural experiments,¹⁵ dynamic-matched sample estimation,⁶ structural models,¹⁶ and ad hoc approaches.³ However, randomized trials are believed to be the most effective way to obtain unbiased estimates of causal peer effects.¹⁷

Randomized Trials of Peer Influence

Randomized trials in networks create unique estimation challenges that require particular attention to experimental design. Researchers have used two primary experimental designs 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.^{18–20} *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.^{5,21–23}

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.

Measuring Peer Influence in Product Adoption

We designed a randomized field experiment on Facebook 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 that 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.

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, however, 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 (see Figure 1), but this approach creates two additional estimation challenges. First, it is unlikely that the baseline hazard of adoption will remain 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.

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. That is, we estimated 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:

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

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 \mathbf{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).

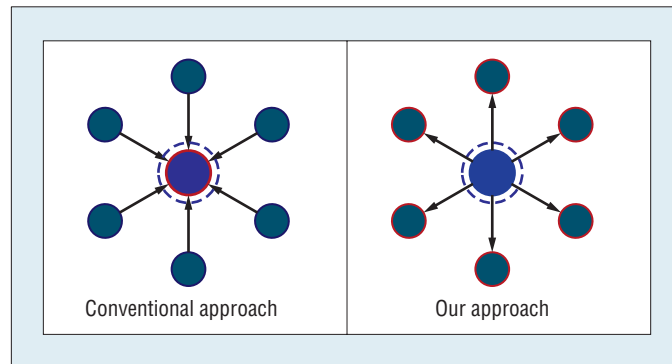


Figure 1. Estimating peer influence and social contagion in a conventional approach versus our approach. Arrows indicate the potential flow of influence that the experiment is designed to detect. The solid blue circle in the center represents the treated user, and the red outlines indicate measured treatment effects.

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. We measured observable attributes of recruited peers and compared it to population distributions to test for selection effects that occur through peer recruitment. However, peer-recruited populations might still differ on unobservable dimensions. For example, in our study, 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 might not be strictly

isolated from one another, and information leakage through indirect network pathways might contaminate results. In traditional studies, whether or not the network is measured, relationships might 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.

Because treatment is randomized, leakage will be uncorrelated with treatment assignment and cannot account for observed differences in responses to treatments. Although 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.²⁴ 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 might be peers of multiple treated users from different treatment groups, making it impossible to link their observed outcomes to a single treatment. Furthermore, although peers of

multiple treated users belonging to the same treatment group are clearly classified, measurements of their response might 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 adoption clusters 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 users and that of their adopter peers.

Figure 2 details the procedure we adopted for designating contaminated peers. In Figure 2a, for times $t > t_1$, peer 2 has multiple treated peers (R and 1), who might have been assigned different treatments. Peer 2 is therefore considered contaminated for times $t > t_1$. Figure 2b is identical except that 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. Designating contaminated peers. The initially recruited adopter R adopts at time t_0 . Peers 1 and 2 of user R adopt at subsequent times t_1 and t_2 , respectively. (a) In the first instance, Peer 2 is considered contaminated for times $t > t_1$, (b) but when no link exists between peers 1 and 2, neither user is considered contaminated.

Results

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 they 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 percent increase in local peer influence and social contagion. Invitations, which are more personalized but require more user effort, only generated an additional 98 percent 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.

Randomizing the treatment of peers of the treated let us collect a greater number of subjects with randomized treatment status. However, maintaining consistency in the treatment among 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 among 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, we could 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, we can model the total expected number of adopters of a given version of the product v as follows:

$$N_v^{\text{total}} = N_v^0 \left(1 + \sum_{\sigma=0}^{\infty} \prod_{s=0}^{\sigma} \langle n_a \rangle_v^s \right)$$

where N_v^0 is the number of originally recruited users (adopter step $s = 0$) who have been randomly assigned to viral treatment state v , and $\langle n_a \rangle_v^s$ is a user's average number of peers with viral state v and distance s (the path length to an originally recruited user) who have adopted the application. The reliance of average peer recruitment, $\langle n_a \rangle_v^s$, on the distance s

from originally recruited subjects allows for the possibility that selection effects might exist and vary (that is, might 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.

Thus, we can 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 by comparing the ratios N_v^{total}/N_v^0 . We estimate the passive-broadcast treatment was associated with a 9.13-percent increase in total diffusion relative to the baseline treatment, while active-personalized treatment was associated with a 12.03-percent increase relative to the baseline.

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 designing successful policy interventions. Randomized trials in social networks are a promising strategy for understanding the dynamics of social influence at scale. Such studies are not without their complications, however. 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, are 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

1. D. Lazer et al., "Computational Social Science," *Science*, vol. 323, no. 5915, 2009, pp. 721–722.
2. S. Aral and M.W. Van Alstyne, "The Diversity-Bandwidth Tradeoff," to be published in *Am. J. Sociology*, 2011; <http://ssrn.com/abstract=958158>.
3. N.A. Christakis and J.H. Fowler, "The Spread of Obesity in a Large Social Network over 32 Years," *New England J. Medicine*, vol. 357, no. 4, 2007, pp. 370–379.
4. R. Lyons, "The Spread of Evidence-Poor Medicine via Flawed Social-Network Analysis," *Statistics, Politics, and Policy*, vol. 2, no. 1, 2011, article 2.
5. S. Aral and D. Walker, "Creating Social Contagion through Viral Product Design: A Randomized Trial of Peer Influence in Networks," to be published in *Management Science*, 2011; <http://ssrn.com/abstract=1564856>.
6. S. Aral, L. Muchnik, and A. Sundararajan, "Distinguishing Influence-Based Contagion from Homophily-Driven Diffusion in Dynamic Networks," *Proc. Nat'l Academy Sciences*, vol. 106, 2009, pp. 21544–21549.
7. C. Shalizi and A.C. Thomas, "Homophily and Contagion Are Generically Confounded in Observational Social Network Studies," *Sociological Methods and Research*, vol. 40, no. 2, 2011, pp. 211–239.
8. D. Godes and D. Mayzlin, "Using Online Conversations to Study Word-of-Mouth Communication," *Marketing Science*, vol. 23, no. 4, 2004, pp. 545–560.
9. C. Van den Bulte and G.L. Lilien, "Medical Innovation Revisited: Social Contagion versus Marketing Effort," *Am. J. Sociology*, vol. 106, no. 5, 2001, pp. 1409–1435.
10. C. Van den Bulte and R. Iyengar, "Tricked by Truncation: Spurious Duration Dependence and Social Contagion in Hazard Models," *Marketing Science*, vol. 30, no. 2, 2011, pp. 233–248.
11. C.F. Manski, "Identification Problems in the Social Sciences," *Sociological Methodology*, vol. 23, 1993, pp. 1–56.
12. S. Aral, "Identifying Social Influence: A Comment on Opinion Leadership and Social Contagion in New Product Diffusion," *Marketing Science*, vol. 30, no. 2, 2001, pp. 217–223.
13. Y. Bramoulle, H. Djebbari, and B. Fortin, "Identification of Peer Effects through Social Networks," *J. Econometrics*, vol. 150, no. 1, 2009, pp. 41–55.
14. T. Snijders, C. Steglich, and M. Schweinberger, "Modeling the Co-evolution of Networks and Behavior," *Longitudinal Models in the Behavioral and Related Sciences*, K. van Montfort, J. Oud, and A. Satorra, eds., Lawrence Erlbaum, 2006.
15. C. Tucker, "Identifying Formal and Informal Influence in Technology Adoption with Network Externalities," *Management Science*, vol. 54, no. 12, 2008, pp. 2024–2038.
16. A. Ghose and S.-P. Han, "An Empirical Analysis of User Content Generation and Usage Behavior in Mobile Media," working paper, Stern School of Business, New York Univ., 2010.
17. E. Duflo, R. Glennerster, and M. Kremer, "Using Randomization in Development Economics Research: A Toolkit," *Handbook of Developmental Economics*, Elsevier, 2006.
18. D. Centola, "The Spread of Behavior in an Online Social Network Experiment," *Science*, vol. 329, no. 5996, 2010, pp. 1194–1197.
19. M. Kearns, S. Suri, and N. Montfort, "An Experimental Study of the Coloring Problem on Human Subject

- Networks,” *Science*, vol. 313, no. 5788, 2006, pp. 824–827.
20. S. Suri and D.J. Watts, “Cooperation and Contagion in Web-Based, Networked Public Goods Experiments,” *PLoS One*, vol. 6, no. 3, 2011, p. e16836.
 21. S. Aral and D. Walker, “An Experimental Method for Identifying Influential and Susceptible Members of Online Social Networks,” working paper, Stern School of Business, New York Univ., 2011.
 22. E. Duflo and E. Saez, “The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence From a Randomized Experiment,” *Quarterly J. Economics*, vol. 118, 2003, pp. 815–842.
 23. S. Leider et al., “Directed Altruism and Enforced Reciprocity in Social Networks,” *Quarterly J. Economics*, vol. 124, no. 4, 2009, pp. 1815–1851.
 24. S. Aral, E. Brynjolfsson, and M. Van Alstyne, “Productivity Effects of Information Diffusion in Networks,” *Proc. 28th Ann. Int’l Conf. Information Systems*, 2007, paper 17; <http://aisel.aisnet.org/icis2007/17>.

Sinan Aral is an assistant professor in the Stern School of Business at New York University and a research affiliate at the Massachusetts Institute of Technology’s Sloan School of Management. Contact him at sinan@stern.nyu.edu.

Dylan Walker is a senior research scientist in the Stern School of Business at New York University. Contact him at dwalker@stern.nyu.edu.

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.