Identifying Influential and Susceptible Individuals in Social Networks: Evidence from a Randomized Experiment

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 2010, St. Louis, Missouri.

Summary: Perhaps the two most important and elusive questions at the heart of viral marketing today are when is peer-to-peer influence mobilized and who are the influentials that most effectively mobilize and propagate influence in networks and society? Answers to these questions are critical to the success of viral marketing campaigns and for understanding how behavioral contagions spread through social networks. Unfortunately, the econometric identification of peer influence and social contagion is non-trivial, making influence identification difficult in observational data. We therefore conduct a large scale randomized experiment involving millions of subjects on the popular social networking website Facebook.com to identify influence, and the characteristics of influential and susceptible individuals involved in the peer-to-peer spread of a commercial grade movie application. By activating a viral feature on the application that generates notifications of a user's activities to their social network peers and then randomizing receipt of the viral messages among peers, we create random assignment to treatment and control groups. We then estimate contagion models with random assignment of the receipt of potentially influential peer-to-peer viral messages. Our study demonstrates how randomized experiments can be used to identify influentials and susceptibles in large social networks of consumers.

Introduction: We examine how firms can take advantage of detailed consumer data generated by IT-enabled products and services to identify "influential" individuals in social networks – those who are likely to effectively encourage product adoption amongst their peers – and "susceptibles" – peers that are likely to respond positively to influence. The notion that influentials exist and that such individuals are catalysts for promoting diffusion of opinions, innovations and products is a popular one (Coleman, Katz, and Menzel 1957; Rogers 2003; Valente 1995; Van den Bulte and Joshi 2007). The complementary notion, however, that susceptibles may be even more critical to product diffusion and the formation of widespread contagions, has been systematically understudied in the literature (Watts and Dodds 2007). As Watts and Dodds note: "In the models that we have studied, in fact, it is generally the case that most social change is driven not by influentials but by easily influenced individuals influencing other easily influenced individuals." Despite receiving disproportionately less attention by researchers than the influentials hypothesis, the susceptibles hypothesis is well represented in a variety of theoretical threshold-based contagion models in which the adoption of an individual occurs when some number or proportion of his peers have adopted beyond his intrinsic adoption threshold.[1] Yet, empirical studies estimating the importance of influentials and susceptibles in the diffusion of products or behaviors in real-world networks significantly lag simulation models and theoretical models of influence-based contagion.

Early empirical studies of the influentials hypothesis have focused primarily on identifying structural network measures such as degree or betweenness that predict whether and to what extent an individual will be influential. However, in recent years there has been an explosion of detailed data on consumer characteristics relating to personal interests, preferences, offline and online behaviors, and product tastes. Despite this wealth of newly available data, there are almost no empirical studies of how individual attributes moderate peer influence.

This is perhaps not surprising as, in order to assess how individual attributes modulate peer influence, it is necessary to first econometrically identify peer influence in large-scale networks —a notoriously difficult task. Several sources of bias in both cross sectional and longitudinal data on interactions and outcomes among peers can confound assessments of peer

[1] In this context, a susceptible individual is one with a low intrinsic threshold.

influence and social contagion including simultaneity (Godes and Mayzlin 2004), unobserved heterogeneity (Van Den Bulte and Lilien 2001), homophily (Aral, Muchnik, and Sundararajan 2009), time-varying factors (Van Den Bulte and Lilien 2001), and other contextual and correlated effects (Manski 1993). If uncorrected, theses biases can lead researchers to incorrectly attribute observed correlations to the influence of individuals on their peers. Observational empirical studies must rely on one of the several existing methods of influence identification (Brock and Durlauf 2001; Bramoullé, Djebbari, and Fortin 2009; Manski 1993; Snijders, Steglich, and Schweinberger 2006; Tucker 2008; Sacerdote 2001; Aral, Muchnik, and Sundararajan 2009). 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. In contrast, randomized field experiments are believed to be the most effective way to obtain unbiased estimates of causal peer effects (Duflo, Glennerster, and Kremer 2007; Falk and Heckman 2009; Aral and Walker 2010). We therefore designed and conducted a randomized field experiment to identify peer influence and to assess the moderating effects of individual-level attributes on influenced adoption of a commercial Facebook application.

Experimental Design: The experiment employs a commercial Facebook application to deliver automated notifications on behalf of users to randomly selected subsets of their peers. Automated notifications are a good form of potentially influential communication for peer influence studies to examine because they do not allow influencing users to select which of their peers receive a communication, thereby circumventing potential confounding selection effects. We designate peers that receive notifications as treated and those that do not as untreated. The delivery of notifications to only a random subset of an individual's peers is ideal because it permits direct comparison of the response of treated peers to peers of the same application user that were not treated. This design decision was deliberately chosen to circumvent a common problem in empirical diffusion studies – even in the absence of influence, individual attributes may drive adoption decisions (Burt 1987) and, in conjunction with homophily (McPherson, Lovin, and Cook 2001)—the tendency for connected individuals to share similar attributes—this can lead to clustering of peer adoption decisions that are often spuriously attributed to peer influence effects. When the targets of potentially influential communications are randomized amongst peers of the same application user, any homophilous structure between an application user and his treated and untreated peers and the propensity to select a particular peer to notify are held constant and identical for treated and untreated groups. Other unrecorded factors that could potentially drive influenced adoption, such as offline or alternative online communications, can also be cleanly distinguished with this design, because treated and untreated peers in expectation share similar propensities to receive and be affected by such communications on average. Differences in adoption outcomes between treated and untreated peer groups can then be attributed solely to their treatment status, namely, whether or not they received a notification.

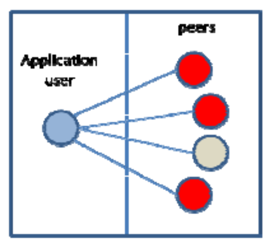

Treated peers comprise a randomly selected subset of an application user's peers that are recipients for automated notifications (treated peers are shaded red).

Data: The data were collected in collaboration with a commercial Facebook Application Developer. We first sampled the local ego network data (and mutual peer ties) of all users of one of several applications developed by this application developer. This sampled a social network of 12 million users and 360 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 12 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*).

Throughout the course of the experiment, we collected individual level profile data from 9,786 application users and their 1.4M distinct peers as well as time-stamped click stream data on notification delivery and subsequent peer responses. During this time, we recorded 69980 automated notifications randomly delivered to peers of application users, resulting in 666 peer adoptions.[2]

Empirical Methods: Our main statistical approach uses hazard modeling, which is the standard technique for assessing contagion in randomized trials in economics, marketing, and sociology literatures. We use several hazard modeling approaches to measure the moderating effect of individual-level attributes on *influence*, *susceptibility to influence* and *dyadic peer-to-peer* influence between pairs of users and particular peers. These analyses shed light on individual characteristics that drive influence, susceptibility and the particular dyad pathways through which influence is most likely to travel.

Influentials Models: Hazard rate models and binary choice models with duration dependence, which can be derived from utility theory and threshold based network models (Van den Bulte et al. 1999), are typically used to estimate social influence (Van den Bulte and Lilien 2001; Manchanda, Xie, and Youn 2008). However, our circumstances require a slightly different approach as we are interested in estimating the effect of individual-level attributes on the adoption of peers in the local networks of a focal application user, rather than the effects of focal users' social environments on their own adoption decisions. We therefore estimate the peer effects of the treatment 'outward' from an individual to their peers rather than estimating the effects of an individual's social environment 'inward' on their own adoption hazard, which is the standard approach. The extent of an application user's influence is measured by the adoption response within his local network by peers that received an automated notification. We employ continuous time multiple-failure proportional hazard models to assess how the individual-level attributes of an application user lead to adoptions (failures) in his local network. Failure times in our adoption data are ordered, meaning there is a natural sequential ordering of event times such that the time of the first adoption in a local network by definition precedes the time of the second adoption and so on. If t_{ik} is the adoption time for the k^{th} adoption in i's network, adoption times are sequential such that $t_{ik} > t_{ik-1}$. Because the social process of contagion can be affected by prior adoptions in a local network, for instance due to network externalities, we stratify hazard models

[2] This represents an adoption rate that outperforms click thorough rates for traditional banner advertizing and that is on par with click through rates for email campaigns (see Aral and Walker (2010): 25).

by allowing the baseline hazard function to vary over successive adoption events occurring within the application user's local peer network. We therefore estimate the following variance-corrected stratified proportional hazards model:

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

where stratification occurs over the *K* adoption events, $\lambda_{0k}(t)$ represents the baseline hazard of the k^{th} adoption event (*i*'s k^{th} friend adopting); X_{ki} represents a vector of individual-level attributes affecting the adoption of *i*'s peers.

Susceptibles Models: To estimate the effect of individual-level attributes on the susceptibility of a peer to influence, we employ continuous time single-failure proportional hazard models. To distinguish attribute-driven adoption from influence-based adoption, attribute covariates and interaction terms with treatment designation are both included in the models. This allows us to test, via randomization whether for instance women are more susceptible than men or whether individuals whoa re married are more susceptible than those who are single. We estimate the following proportional hazards model:

$$\lambda(t, X_j) = \lambda_0(t)e^{X_j\beta + T_j X_j \beta_T}$$

where $\lambda_0(t)$ represents the baseline hazard, X_j represents a vector of individual attributes affecting peer *j*'s adoption regardless of treatment status, and T_j represents the treatment status of peer *j* ($T_j = 1$ *iff* peer *j* received an automated notification). As peers may have received multiple notifications, we include modifications to the above by stratifying the model by number of notifications received and by including terms that cross the number of notifications received with individual covariates to examine the marginal effects of attributes on the influence of an additional (marginal) notification received on the adoption hazard.

Dyadic Peer-to-Peer Models: Models that consider sender-recipient pairs can be applied to our empirical data to shed light on several open questions in the study of peer-to-peer influence. This allows us to test, via randomization whether for instance women influence men, or rather whether men influence women, or alternatively whether the young influence the old or whether the old influence the young etc. One can imagine a series of interesting tests – are married individuals more susceptible than single individuals and if so is this due to a 'spouse effect' (married individuals respond at elevated rates to influence from their spouses)? We employ single-failure proportional hazard models to estimate the effect of influence moderating factors that depend on the dyadic combination of attributes of the sender and the recipient of the notification on the adoption hazard of peer *j* as follows:

$$\lambda(t, X_i, X_j) = \lambda_0(t)e^{S(X_i, X_j)(1+T_j)\beta_S}$$

where X_i represents a vector of the individual attributes of the sender, X_j represents a vector of the individual attributes of the recipient, and $S(X_{i,} X_j)$ represents a vector of covariates that characterize the joint the attributes of the sender-recipient pair. Examples of dyadic covariates include for example whether *sender and recipient are of the same gender, sender is older than recipient, come from the same hometown, attended the same high school or university,* or *work at the same company*. In addition, dyadic covariates can include similarity measures on individual attributes such as *difference in age*, *interests* and *product preference similarity*, as well as measures of total similarity across all measured attributes, enabling us to test whether homophily begets influence – considered to be a critical confounding factor in most influence studies.

Initial Results: In table 1 we present some of the more interesting initial results from influential and susceptible hazard models. Single failure hazard models for susceptibility to influence (models 1-6) include covariates for the individual attributes of peers of application users as well as covariates crossed with the number of notifications received by each recipient. The latter hazard coefficients reflect the modulating effect of the individual attribute on his hazard to adopt per automated notification received. Individuals that report their gender as male are significantly less susceptible to influence than individuals that do not report their gender (model 1). Susceptibility to influence increases non-linearly with the number of notifications received, as evidenced by the significance of the square of the number of notifications received (model 2). Models 3-7 indicate that an individual's susceptibility to influence increases with the commitment level of the reported relationship status until they are *married* (relative to individuals that do not report their relationship status), with susceptibility increasing from *single* to *in a relationship* to *engaged*. Individuals that report their relationship status as *married* do not seem to be significantly susceptible to influence. Finally, individual's that report their relationship status as *it's complicated* are the most susceptible to influence from their peers.

Models 7-13 are multiple failure hazard models for influence of an application user on his peers that include covariates for the individual attributes of application users as well as covariates interacted with the number of treated peers in his local network. The latter hazard coefficients reflect the moderating effect of the individual attribute on an application user's hazard to gain an additional adopter in his local network. Individuals that report their gender as male are significantly less influential than individuals that do not report their gender. Models 9-12 indicate that an individual's reported relationship status has little to no effect on how influential they are, with the exception of individuals who report their relationship status as *it's complicated*, who seem to exert negative influence on the adoption hazard of peers in their local network. The finding that the attributes considered do not significantly moderate the influence of individuals, while the same attributes do seem to significantly moderate susceptibility to influence is consistent with the finding by Watts and Dodds that most social change is driven not by influentials but by susceptibles. However, more robust studies are necessary to confirm this finding.

Plan of Work: We intend to run several models examining the moderating effect of a variety of other promising individual level attributes on influence, susceptibility to influence and dyadic peer to peer influence in the months leading up to WISE. We expect dyadic peer to peer model results to be particularly interesting. We have also begun analysis to dimensionally reduce a litany of high-dimensional, sparsely populated, heavy-tail distributed covariates (such as interests and product preferences) using multipartite community structure detection, to allow for viable inclusion of these individual attributes in hazard models.

Table 1	**Susceptiblity to Influence (Single Failure Proportional Hazards)**							**Influence (Multiple Failure Variance Corrected Proportional Hazards)**					
	1	2	3	4	5	6	7	8	9	10	11	12	13
	Hazard Ratio (SE)	*Hazard Ratio (SE)*	*Hazard Ratio (SE)*	*Hazard Ratio (SE)*	*Hazard Ratio (SE)*	*Hazard Ratio (SE)*	*Hazard Ratio (SE)*	*Hazard Ratio (SE)*	*Hazard Ratio (SE)*	*Hazard Ratio (SE)*	*Hazard Ratio (SE)*	*Hazard Ratio (SE)*	*Hazard Ratio (SE)*
Num. Notifications Rcvd (influentials: Num. Treated Peers)	5.10*** (0.158)	149*** (19.6)	3.68*** (0.146)	3.76*** (.140)	3.78*** (.138)	3.77*** (.142)	3.80*** (.138)	.0615*** (.00363)	.0729*** (.00687)	.0724*** (.00684)	.0720*** (.00687)	.0724*** (.00786)	.0765*** (.00614)
Num. Notifications Rcvd^2	.	.272*** (.0180)	.	.	.	.	.	.	.	.	.	.	.
Age X Num Not. Rcvd (Age X Num Treated Peers)	.		.	.	.	.	.	.	.	.	.	.	.
Gender (male)	2.43*** (0.189)		.	.	.	.	.	-0.217** (.0966)	.	.	.	.	.
Gender X Num Not. Rcvd (Gender X Num Treated Peers)	0.75*** (0.0400)		.	.	.	.	.	.0108* (.00629)	.	.	.	.	.
Single			.683** (.0891)	.756** (.0926)	.758** (.0928)	.756** (.0926)	.760** (.0931)		.667*** (.151)	.658*** (.150)	.657*** (.150)	.657*** (.150)	.665*** (.150)
In a Relationship			.526*** (.0905)	.476*** (.0885)	.535*** (.0921)	.533*** (.0918)	.536*** (.0923)		.418* (.223)	.418* (.228)	.416* (.223)	.418* (.223)	.420* (.224)
Engaged			.510** (0.173)	.516* (.176)	.394** (.157)	.516* (.176)	.519* (.176)		.272 (.351)	.273 (.351)	.219 (.405)	.274 (.350)	.257 (.351)
Married			.810 (.121)	.820 (.123)	.824 (.123)	.774 (.123)	.824 (.124)		.244 (.164)	.245 (.164)	.245 (.164)	.246 (.182)	.235 (.164)
It's Complicated			.434* (.219)	.441 (.222)	.442 (.223)	.441 (.222)	.283** (.180)		.0617 (.450)	.0669 (.449)	.0712 (.448)	.0665 (.452)	1.11** (.382)
Single X Num Not. Rcvd (Single X Num Treated Peers)			1.31** (.121)	.	.	.	.		-.00496 (.0197)	.	.	.	.
In a Relationship X Num Not Rcvd (In a Relationship X Num Treated Peers)				1.47** (.227)	.	.	.			-.000215 (.0241)	.	.	.
Engaged X Num Not. Rcvd (Engaged X Num Treated Peers)					1.76** (.414)	.	.				.00700 (.0209)	.	.
Married X Num Not. Rcvd (Married X Num Treated Peers)						1.20 (.149)	.					-.000180 (.00977)	.
It's Complicated X Num Not. Rcvd (It's Complicated X Num Treated Peers)							2.70** (1.05)						-1.38*** (.0476)
Log Likelihood	-11938	-11555	-5200	-5201	-5202	-5203	-5201	-7540	-2122	-2122	-2122	-2122	-2114
X^2 (d.f)	1327*** (3)	2095*** (2)	566.5*** (7)	564.0*** (7)	563.3*** (7)	561.1*** (7)	563.6*** (7)	358.2*** (3)	139.2*** (7)	139.6*** (7)	143.8*** (7)	163.2*** (7)	1359*** (7)
Observations	1388468	1388468	411569	411569	411569	411569	411569	28047	9748	9748	9748	9748	9748

Notes: ***p<.001; **p<.05; *p<.10; Variance Corrected Proportional Hazards Models are estimated with robust standard errors clustered around users' local network neighborhoods.

References

Aral, Sinan, Lev Muchnik, and Arun Sundararajan. 2009. "Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks." *Proceedings of the National Academy of Sciences* 106:21544-21549. http://www.pnas.org/content/106/51/21544.abstract.

Aral, Sinan, and Dylan Walker. 2010. "Creating Social Contagion Through Viral Product Design: A Randomized Trial of Peer Influence in Networks." *Social Science Research Network Working Paper Series*. http://ssrn.com/abstract=1564856.

Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin. 2009. "Identification of peer effects through social networks." *Journal of Econometrics* 150:41-55. http://linkinghub.elsevier.com/retrieve/pii/S0304407609000335.

Brock, W. A., and S. N. Durlauf. 2001. "Discrete Choice with Social Interactions." *Review of Economic Studies* 68:235-260.

Burt, Ronald S. 1987. "Social Contagion and Innovation: Cohesion Versus Structural Equivalence." *The American Journal of Sociology* 92:1287-1335. http://www.jstor.org/stable/2779839.

Coleman, James, Elihu Katz, and Herbert Menzel. 1957. "The Diffusion of an Innovation Among Physicians." *Sociometry* 20:253-270. http://www.jstor.org/stable/2785979.

Duflo, E, R Glennerster, and M Kremer. 2007. "Chapter 61 Using Randomization in Development Economics Research: A Toolkit★. "pp. 3895-3962 in, vol. 4, *Handbook of Development Economics*. Elsevier http://linkinghub.elsevier.com/retrieve/pii/S1573447107040612.

Falk, Armin, and James J. Heckman. 2009. "Lab Experiments Are a Major Source of Knowledge in the Social Sciences." *Science* 326:535-538. http://www.sciencemag.org/content/326/5952/535.abstract.

Godes, David, and Dina Mayzlin. 2004. "Using Online Conversations to Study Word-of-Mouth Communication." *MARKETING SCIENCE* 23:545-560. http://mktsci.journal.informs.org/cgi/content/abstract/23/4/545.

Manchanda, Puneet, Ying Xie, and Nara Youn. 2008. "The Role of Targeted Communication and Contagion in Product Adoption." *MARKETING SCIENCE* 27:961-976. http://mktsci.journal.informs.org/cgi/content/abstract/27/6/961.

Manski, Charles F. 1993. "Identification Problems in the Social Sciences." *Sociological Methodology* 23. http://www.jstor.org/stable/271005.

McPherson, Miller, Lynn Smith Lovin, and James M Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Annual Review of Sociology* 27:415-444. http://arjournals.annualreviews.org/doi/abs/10.1146/annurev.soc.27.1.415.

Rogers, Everett M. 2003. *Diffusion of Innovations, 5th Edition*. Original. Free Press http://www.amazon.ca/exec/obidos/redirect?tag=citeulike09-20&path=ASIN/0743222091.

Sacerdote, Bruce. 2001. "Peer Effects with Random Assignment: Results for Dartmouth Roommates*." *Quarterly Journal of Economics* 116:681-704. http://www.mitpressjournals.org/doi/abs/10.1162/00335530151144131.

Snijders, Tom A. B., Christian E. G. Steglich, and Michael Schweinberger. 2006. "Modeling the co-evolution of networks and behavior." http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.153.2361.

Tucker, Catherine. 2008. "Identifying Formal and Informal Influence in Technology Adoption with Network Externalities." *MANAGEMENT SCIENCE* 54:2024-2038.

Valente, Thomas W. 1995. *Network Models of the Diffusion of Innovations (Quantitative Methods in Communication Subseries)*. Hampton Press (NJ) http://www.amazon.ca/exec/obidos/redirect?tag=citeulike09-20&path=ASIN/1881303225.

Van Den Bulte, Christophe, and Gary L. Lilien. 2001. "Medical Innovation Revisited: Social Contagion versus Marketing Effort." *The American Journal of Sociology* 106:1409-1435. http://www.jstor.org/stable/3081115.

Van den Bulte, Christophe et al. 1999. *Integrating Models of Innovation Adoption: Social Network Thresholds, Utility Maximization, and Hazard Models*. http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.37.4215.

Van den Bulte, Christophe, and Yogesh V. Joshi. 2007. "New Product Diffusion with Influentials and Imitators." *MARKETING SCIENCE* 26:400-421. http://mktsci.journal.informs.org/cgi/content/abstract/26/3/400.

Watts, Duncan J., and Peter Sheridan Dodds. 2007. "Influentials, Networks, and Public Opinion Formation." *Journal of Consumer Research* 34:441-458. http://www.journals.uchicago.edu/doi/abs/10.1086/518527.