

Creating Social Contagion through Viral Product Design: A Randomized Trial of Peer Influence in Networks

Sinan Aral

NYU Stern School of Business & 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

We examine how firms can create word-of-mouth peer influence and social contagion by designing viral features into their products and marketing campaigns. Word-of-mouth (WOM) is generally considered to be more effective at promoting product contagion when it is personalized and active. Unfortunately, the relative effectiveness of different viral features has not been quantified, nor has their effectiveness been definitively established, largely because of difficulties surrounding econometric identification of endogenous peer effects. We therefore designed a randomized field experiment on a popular social networking website to test the effectiveness of a range of viral messaging capabilities in creating peer influence and social contagion among the 1.4 million friends of 9,687 experimental users. Overall, we find that viral product design features can indeed generate econometrically identifiable peer influence and social contagion effects. More surprisingly, we find that passive-broadcast viral messaging generates a 246% increase in local peer influence and social contagion effects, while adding active-personalized viral messaging only generates an additional 98% increase in contagion. Although active-personalized messaging is more effective in encouraging adoption per message and is correlated with more user engagement and sustained product use, passive-broadcast messaging is used more often enough to eclipse those benefits, generating more total peer adoption in the network. In addition to estimating the effects of viral product design on social contagion and product diffusion, our work also provides a model for how randomized trials can be used to identify peer influence effects in networks.

Key words: Peer Influence, Social Contagion, Social Networks, Viral Marketing, Information Systems, Randomized Experiment.

This Draft: November, 2010

* We are grateful to Erik Brynjolfsson, Haim Mendelson, Markus Mobius, Lev Muchnik, Foster Provost, Hal Varian and seminar participants at the National Bureau of Economic Research, the Workshop on Information Systems Economics, Laval, MIT Sloan, Northeastern, NYU Stern, and Wharton for invaluable discussions. Sinan Aral acknowledges generous financial support from Microsoft, the National Science Foundation (CAREER Award 0953832), IBM and the Institute for Innovation and Information Productivity.

1. Introduction

It is widely believed that social contagion and word-of-mouth (WOM) “buzz” about products drive product adoption and sales, and firms increasingly rely on “network” and “viral” marketing strategies (Hill et al 2006, Manchanda et al 2008, Nam et al 2010). Yet, two subjects central to the success of viral marketing efforts have been largely neglected in the WOM literature – the effectiveness of different *viral product design* strategies and *econometric identification of peer influence effects*. In order to address both topics we conducted a large-scale randomized field experiment to test the effectiveness of different viral product features in creating peer influence and social contagion in new product diffusion.

Viral product design – the process of explicitly engineering products so they are more likely to be shared amongst peers – has existed at least since the first chain letter was sent in 1888.¹ Today, products regularly use IT enabled features like automated broadcast notifications and personalized invitations to spread product awareness. Yet, although viral features have become more sophisticated and a central part of the design of products and marketing campaigns, there is almost no empirical evidence on the effectiveness of such features in generating social contagion and product adoption. We therefore investigate two basic questions: Can firms add viral features to products so they are more likely to be shared amongst peers? If so, which viral features are most effective in inducing WOM and peer-to-peer influence in product adoption?

Unfortunately, evaluating the effects of viral product design features is difficult because peer effects and WOM are typically endogenous (Manski 1993, Godes and Mayzlin 2004, 2009, Hartmann et al 2008, Aral et al 2009, Aral 2010). We therefore designed and conducted a randomized field experiment testing the effectiveness of two of the most widely used viral product features – active-personalized referrals and passive-broadcast notifications – in creating peer influence and social contagion among the 1.4 million friends of 9,687 experimental users of Facebook.com. The experiment uses a customized com-

¹ This earliest known example of a chain letter seems to have initiated by four women requesting donations for education efforts in New Hampshire: <http://www.silcom.com/~barnowl/chain-letter/evolution.html>.

mercial Facebook application to observe user behavior, communications traffic and the peer influence effects of randomly enabled viral messaging features on application diffusion and use in the local networks of experimental and control population users. By enabling and disabling viral features among randomly selected users, we were able to obtain relatively unbiased causal estimates of the impact of viral features on the adoption rates of peers in the local networks of adopters. Using detailed clickstream data on users' online behaviors we also explored whether positive network externalities generated by additional peer adopters inspired further product adoption and sustained product use.

WOM is generally considered to be more effective at promoting product contagion when it is personalized and active. Surprisingly, we find that designing products with passive-broadcast viral messaging capabilities generates a 246% increase in local peer influence and social contagion, while adding active-personalized viral messaging capabilities only generates an additional 98% increase. Although active-personalized messaging is more effective in encouraging adoption per message and is correlated with more user engagement and sustained product use, it is used less often and therefore generates less total peer adoption in the network. Overall, we find that viral product design features do generate econometrically identifiable peer influence and social contagion effects and provide a model for how randomized trials can identify peer influence in networks.

2. Viral Product Design

Since the early work of Katz and Lazarsfeld (1955) there has been great interest in how WOM drives consumer demand, public opinion and product diffusion (Brown and Reingen 1987, Godes and Mayzlin 2004, Aral et al 2009) and how firms can create broad, systematic propagation of WOM through consumer populations (Phelps et al 2004, Mayzlin 2006, Dellarocas 2006, Godes and Mayzlin 2009). Many campaigns target "influential" individuals who are likely to propagate organic WOM most broadly (Katz and Lazarsfeld 1955, Watts and Dodds 2007, Goldenberg et al 2009), using referral programs to create incentives for them to spread the word (Biyalogorsky et al 2001). Others use observational evidence on viral campaigns to inform viral branching models of WOM diffusion (Van der Lans et al 2010).

However, to this point, studies of viral product design have remained conspicuously absent from the literature on viral marketing.

Viral product design involves incorporating specific characteristics and features into a product's design to generate peer-to-peer influence that encourages adoption. A product's *viral characteristics* are fundamentally about its content and the psychological effects content can have on a user's desire to share the product with peers (Stephen and Berger 2009, Berger and Heath 2005, Heath, Bell and Sternberg 2001). A product's *viral features* on the other hand concern how the product is shared – how features enable and constrain a product's use in relation to other consumers. Viral features may enable communication, generate automated notifications of users' activities, facilitate personalized invitations or enable hypertext embedding of the product on publicly available websites and weblogs. Two of the most widely used viral product features are personalized referrals and automated broadcast notifications.

Personalized Referrals. Personalized referral features allow users to select their friends or contacts from a list and invite them to adopt the product or service, with the option of attaching a personalized message to the invitation.²

Automated Broadcast Notifications. Automated broadcast notifications are passively triggered by normal user activity. When a user engages the product in a certain way (e.g. sends a message, updates his or her status), those actions are broadcast as notifications to the user's list of contacts. Notifications build awareness among friends of new activities or products a user is adopting or engaging with, and can encourage those friends to eventually adopt the product themselves.³

Referrals are more personalized and targeted than broadcast notifications. Users actively select the subset of their social network to receive them (targeting) and can include personal messages in the

² Companies like Facebook enable users to 'invite their friends' to join the service through personalized referrals. When users send Gmail messages, an automated, pop-up hyperlink enables them to invite recipients to join Gmail.

³ When a user of LinkedIn.com joins a new group, changes their profile information, connects to a new contact or takes a new job, their contacts are informed via email about the activity. Facebook notifies friends when a user adopts a new application or achieves some application milestone.

referral (personalization).⁴ WOM is generally considered more effective at promoting product contagion when it is personalized and active. When individuals choose to share information about products and services with their friends, they tend to activate their strong-tie relationships (Frenzen and Nakamoto 1993, Aral & Van Alstyne 2009). Strong ties exhibit greater homophily (Jackson 2008), greater pressure for conformity (Coleman 1988) and deeper knowledge about one another. We tend to trust information from close “trusted” sources more and to respond more often to them due to reciprocity (Emerson 1962). In addition, the personalization of messages makes them more effective, especially in online environments in which we are bombarded with irrelevant information (Tam and Ho 2005, Tucker 2010).

For these reasons, one might suspect that personalized referrals are more effective than broadcast notifications. But, although each personalized referral may be more persuasive (more effective per message), the pervasiveness of broadcast messages may lead to greater overall peer adoption. The effort required by the user to actively select and invite peers to adopt the product could inhibit widespread use of the personalized referral and so limit its effectiveness in encouraging broad adoption.

Viral features can be broadly described using two dimensions: *activity* and *personalization*. Activity is higher for features that require active user engagement and lower for passive features that generate automated actions on behalf of users. Personalization ranges from broadcast features that are unselective in their audience to personalized features targeted and tailored toward specific peers (see Figure 1). Active-personalized viral features are generally considered more persuasive, but the additional effort they require may curtail their use. The relative overall effectiveness of these viral features is therefore ultimately an empirical question.

*** Figure 1 About Here ***

3. Empirical Methods

⁴ Targeting specifies whether the feature is directed at the broad population of potential consumers, a subset of consumers like a current user’s social network, or a specific person. Customization specifies whether the content of a feature’s engagement with the recipient can be tailored to a group of friends or a specific individual with a personalized message.

3.1. Identification of Peer Influence in Social Networks

The effects of viral product design features on social contagion are difficult to evaluate because peer effects and WOM are typically endogenous. Several sources of bias in analysis on interactions and outcomes among peers can confound assessments of peer influence and social contagion including simultaneity (Godes and Mayzlin 2004), unobserved heterogeneity (Van den Bulte and Lilien 2001), truncation (Van den Bulte and Iyengar 2010), homophily (Aral et al. 2009), time-varying factors (Van den Bulte and Lilien 2001), and other contextual and correlated effects (Manski 1993). If uncorrected, these biases can lead researchers to attribute observed correlations to peer influence, resulting in misinterpretations of the effectiveness of viral marketing or viral product design choices.

Several approaches for identifying peer effects have been proposed, including peer effects models and extended spatial autoregressive models (e.g. Kelejian and Prucha 1998, Oestreicher-Singer and Sundararajan 2008, Trusov et al 2009, Bramouille et al 2009), actor-oriented models (e.g. Snijders et al. 2006), instrumental variables methods based on natural experiments (e.g. Sacredote 2001, 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 considered to be one of the most effective ways to obtain unbiased estimates of causal peer effects (Duflo et al 2006).

3.2. Experimental Design and Procedures

We partnered with a firm that develops commercial applications hosted on the popular social networking website Facebook.com, and collected data on the peer influence effects of enabling viral features on the diffusion one of their applications. This application provides users the opportunity to share information and opinions about movies, actors, directors and the film industry in general. We designed multiple experimental versions of the application in which *personalized invitations* and *broadcast notifications* were enabled or disabled, and randomly assigned adopting users to various experimental and control conditions. As users adopted the application, each was randomly assigned to one of the two treatment conditions or the baseline control condition. The application collected personal attributes and preferences

from users' Facebook profiles, as well as data on their social networks and the personal attributes and preferences of their network neighbors.⁵

The experiment enabled experimental group users to use passive-broadcast and active-personalized viral messaging capabilities to exchange messages with their network neighbors, while disabling those 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 clickstream data on recipient responses to viral messages and their subsequent adoption and use of the application. When an individual adopted the application as a result of peer influence, their treatment status was also randomized to ensure that the Stable Unit Treatment Value Assumption held. This facilitated analysis of the relative effectiveness of different viral messaging channels in generating peer adoption and network propagation. Randomization also enabled exploration of the mechanisms by which a particular viral channel influenced recipient behavior. Two primary viral features were examined:

Automated Broadcast Notifications (Notifications). When enabled, notifications were generated automatically when an application user performed certain actions within the application, such as declaring a favorite movie or writing a movie review. When notifications were generated, they were distributed to a random subset of an application user's peers and displayed in a status bar at the bottom of the peers' Facebook environment. When a peer clicked on the notification, they were taken to an application canvas page where they were given the option to install the application. These notifications required no effort beyond normal use of the application, scoring low on the activity dimension. As they were randomly distributed to a Facebook user's peers and were not accompanied by a personalized message, they also exhibited low personalization.

Personalized Referrals or Invitations (Invites). When enabled, invites allowed application users to send their Facebook peers personalized invitations to install the application. Peers received the invitation in their Facebook inbox and could click on a referral link contained within the invitation. If they did

⁵ Facebook allows users to specify privacy settings that may restrict an application's access to some or part of their profile. This is unlikely to have a significant effect on the study, as it is estimated that less than 2% of Facebook users alter default privacy settings (Gross et al 2005).

so, they were taken to the application canvas page where they were given the opportunity to install the application. Each invite required a conscious and deliberate action from the user beyond typical application use, requiring more effort (activity) than notifications. As invites were targeted to specific peers and allowed the inclusion of a personalized message, they exhibited greater personalization.

The experimental design consisted of three treatment groups into which users were randomly assigned: *baseline*, *passive-broadcast*, and *active-personalized*. Users assigned to the baseline treatment group received a version of the application in which both notifications and invites were disabled. In the passive-broadcast treatment group (passive), only notifications were enabled. In the active-personalized treatment group (active), both notifications and invites were enabled. There were no other differences between baseline, passive and active applications. Throughout the experiment, each adopter of the application was randomly assigned to a treatment group according to the proportions displayed in Table 1. The proportion of users assigned to the baseline was chosen in agreement with the application developer to obtain a population size sufficient to establish a comparative baseline, while limiting potential adverse effects on the overall diffusion of the product.⁶

*** Table 1 About Here ***

Detailed logs of application user activity, adoption times, viral feature use, peer response, and application user and peer profile data were recorded, as were social network relationships for application adopters and mutual ties between peers of application users. Our experimental design allowed us to measure the effect of each of the viral features on the adoption response of peers as displayed in Figure 2. We took care to minimize contamination and leakage effects and describe our methods with regard to those and other considerations in detail in the Appendix.

*** Figure 2 About Here ***

⁶ The developer feared too many baseline users could stunt the viral diffusion of the application and therefore insisted that the number of baseline users be limited. Limiting baseline users should not bias results as the proportion of baseline users to either treatment group is constant across treatments and should only make our estimates more conservative in that analyses comparing a treatment group to the baseline group will have less power.

Recruitment. At the launch of the experiment, we designed an advertising campaign in collaboration with a second Facebook advertising firm to recruit a representative population of Facebook users. Advertisements were displayed to users through advertising space within Facebook and within existing Facebook applications. The campaign was conducted in three waves throughout the duration of the experiment and cost a total of \$6000 to recruit 9687 usable experimental subjects, or 62 cents per recruit.⁷ The number of impressions, clicks, and installation responses are displayed in Table 2. Summary statistics of the recruited study population are described in § 4. Comparisons to published demographic statistics indicate the sample is indeed representative of typical Facebook users (see Appendix).

*** Table 2 About Here ***

4. Analysis and Results

4.1. Data and Descriptive Statistics

The experiment was conducted over a 44-day period during which 9687 users adopted the application with 405 users randomly assigned to the baseline control group, 4600 users randomly assigned to the passive-broadcast treatment group, and 4682 users randomly assigned to the active-personalized treatment group. Users in these groups collectively had 1.4M distinct peers in their local social networks and sent a total of 70,140 viral messages to their peers, resulting in 992 peer adoptions – 682 of which were in direct response to viral messages. Three main observations arise from consideration of the summary statistics of the resultant data displayed in Table 3.

*** Table 3 About Here ***

First, assignment to control and treatment groups was clearly random with no significant mean or distributional differences between users in terms of their age, gender, network degree (number of Facebook friends), and level of Facebook activity (number of Facebook wall posts), confirming the integrity of the randomization procedure.

⁷ The cost per recruited user is several times smaller than the cost-per-user associated with recruitment for lab-based experiments. The low cost of recruitment makes online experiments an excellent source of experimental data.

Second, while their demographics and Facebook activity patterns were the same, measures of peer response in the network neighborhoods of treated users differed significantly across the treatment and control populations. T-tests show that the number and percentage of peer adopters in a user's local network are significantly higher for treated populations than for the baseline population. The number of peer adopters in a user's local network is roughly seven times greater for users in the passive-broadcast treatment group and ten times greater for users in the active-personalized group. Similarly, in comparison to the baseline group, the percentage of adopters in a user's local network is roughly 450% higher for users in the passive-broadcast group and 750% higher for users in the active-personalized group. Measures of the speed of adoption in a treated user's local network, as indicated by the time to the first, second, third, and fourth adoption events, reveal that the treatments increased the rate of adoption in a treated user's local network. The time to the first adopter is roughly 200% shorter for users in the passive-broadcast treatment group and roughly 300% shorter for users in the active-personalized group. The extent to which the effect of the treatment leads to adoption beyond a user's immediate local network can be measured by the maximal diffusion depth – the maximum network distance from a treated user to any peer adopter in a linked chain of adoptions. The average maximal diffusion depth is approximately 360% greater for the passive-broadcast treatment group and 450% greater for the active-personalized treatment group compared to baseline users. T-tests reveal these differences are highly significant.

Finally, the extent to which each treatment leads to increased application use is measured by users' average application activity. Average application activity is roughly 130% higher in the passive-broadcast treatment group and 140% higher in the active-personalized treatment group. We explore more formal models of peer influence and social contagion in the next three sections.

4.2. Model Specification

Our main statistical approach uses hazard modeling, which is the standard technique for assessing contagion in economics, marketing, and sociology (e.g. Van den Bulte and Lilien 2001, Iyengar et al

2010, Nam et al 2010). This approach typically represents the hazard of adoption of individual i at time t as a function of individual characteristics and social influence:

$$\lambda(t, x, w, y) = f(x_i(t)\gamma, \beta \sum_j w_{ij} y_j(t)),$$

where $\lambda(t)$ represents the baseline hazard of adoption; $x_i(t)$ is a vector of variables unrelated to social influence that affect i 's adoption decision; w_{ij} is the social exposure of i to peer j ; $y_j(t)$ is the adoption status of peer j at time t ; and γ and β are parameters to be estimated. However, our circumstances required a slightly different approach as we are interested in estimating the treatment effects of randomly assigned viral features on the adoption of peers in the local networks of focal experimental and control users, 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. Controlled "treatments" of each user's entire social environment are too complex and costly to be accomplished reliably in the field. Observation of the diffusion of the product also requires estimation of the adoption hazards of peers and the subsequent adoption hazards of peers of peers. An 'inside-out' strategy estimating the effects of treatment on adoption in a user's social environment is therefore the most appropriate modeling approach (see Appendix).

Our approach compares the hazards of adoption in the social environments of users treated with passive and active viral applications to the hazards of adoption in the social environments of users treated with the baseline application. The analysis therefore involves multiple failure time data in which multiple failures can occur for the same subject over time. In our case, we want to estimate the hazard of multiple occurrences of peer adoption in the local networks of treated and untreated users as a function of their exposure to different viral features. In multiple failure data, failure times are correlated within cluster (in our case within users' local networks), violating the independence of failure times assumption required in traditional survival analysis. The simplest way to analyze multiple failure data is to examine "time to first event." Several studies in the contagion literature take that approach (Iyengar et al 2010). Other studies

estimate the time to the first event and each subsequent event separately, which by construction assumes the baseline hazards of sequential adoption events are equal (Anderson and Gill 1982). However, those specifications overlook potentially relevant information and fail to consider the cascading effects of multiple adoption events in a network, such as the presence of non-linear network effects in product adoption. We therefore employ a variance-corrected stratified proportional hazards approach 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 to account for the possibility that adoption hazards vary across stages of a diffusion process.

Failure times in our data are sequential. The first adoption in a local network 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 $t_{ik} \geq t_{ik-1}$. As we observe time-stamped adoption of the application in minutes and seconds, no two events happen at the same time. As the social process of contagion can be affected by prior adoptions in a local network, for instance if network externalities are present, we assume that the baseline hazard function varies over adoption occurrences. 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 covariates affecting the adoption of i 's neighbors (including i 's viral treatment status (active, passive or baseline), a measure of i 's level of activity on the application (Application Activity), peer notifications sent (Notifications), and invites sent (Invites)); and β is a vector of unknown parameters to be estimated. We assume i 's k^{th} friend does not adopt until their $k-1$ friend adopts, as this is the case for all our data. Therefore the conditional risk set at time t for event k consists of all subjects under observation at time t who have experienced a $k-1$ adoption event. We estimate β using standard maximum likelihood estimation and adjust the covariance matrix to

account for non-independence across individuals i using the following robust covariance matrix (where G is a matrix of group efficient residuals):

$$V = I^{-1}G'GI^{-1}$$

Results are presented in Table 4. Robustness to different model specifications is shown in the Appendix.

*** Table 4 About Here ***

4.3. Effects of Viral Product Design on Peer Influence and Social Contagion

Table 4, Model 1 displays the average treatment effects of passive-broadcast and active-personalized viral treatments on peer influence and social contagion in the local networks of treated users above and beyond control group users who received the baseline application. Users of the passive-broadcast application experienced a 246% increase in the rate of application adoption by peers compared to the baseline group, while adding active-personalized viral messaging capabilities only generated an additional 98% increase (active-personalized users experienced a 344% increase over the baseline group). Models 2-4 decompose the variance in local network adoption rates explained by these treatments by estimating how intermediate variables such as overall application activity, notifications and invites explain the resultant increases in peer adoption. Model 3 shows that a significant amount of the treatment effects are explained by correlated increases in users' use of the application and the viral messages their use generates. Users assigned to passive-broadcast and active-personalized applications use their applications more and send more messages (invites and notifications) that generate greater peer adoption in their local networks. Model 4 reveals that invites have a greater marginal impact on the adoption rate of peers than notifications. One additional personal invite increases the rate of peer adoption by 6%, while one additional notification increases the rate of peer adoption by only 2% on average, confirming that more personalized active features have a greater marginal impact on the rate of peer adoption per message than passive broadcast features.

The clickstream data, which record each time-stamped viral message and any response to it by peers, corroborate these results. Table 5 displays the number of invitations and notifications sent, the responses to those messages that resulted in click-through installations of the application and the resultant adoption rate per message. Invitations are the least used but the most effective per message in creating peer influence and social contagion. Notifications, which require the least effort and are automatically sent to randomly selected peers, generate more messages, but are less effective per message in converting new users.

*** Table 5 About Here ***

These results together confirm the main findings of the study: viral product design features do in fact generate econometrically identifiable peer influence and social contagion effects. Features that require more activity on the part of the user and are more personalized to recipients create greater marginal increases in the likelihood of adoption per message, but also generate fewer messages resulting in less total peer adoption in the network.

*** Figure 3 About Here ***

Figures 3a) and 3b) plot the cumulative peer adoptions and the fractions of adopters in the local networks of baseline, passive and active treatment users, while 3d) plots the Kaplan-Meier survival estimates for baseline, passive and active treatments respectively.⁸ Susceptible peers of users in the passive-broadcast viral treatment group had an approximately seven-fold higher fraction of adopters in their local networks compared to baseline users. Susceptible peers of users in the active-personalized treatment group had over a ten-fold increase in adoption fraction compared to users in the baseline group, and an

⁸ Figure 3b) plots the fraction of susceptible peers that adopt the application t days after they become susceptible in active-personalized, passive-broadcast and baseline treatment and control groups, while Figure 3a) shows the cumulative adoption in each group. To assess the effect of the treatment group on the adoption of application user's peers through any influence-mediating channel, we identify the time of susceptibility to influence for all peers of buy-in users. To account for fixed-time effects, we look at the adoption response of all susceptible peers t days after they first became susceptible. We define the adoption fraction as, $A_f(t)$:

$$A_f(t) = \frac{\text{Number of susceptible peers that have adopted } t \text{ days after becoming susceptible}}{\text{Number of peers that are still susceptible } t \text{ days after becoming susceptible}}$$

and we plot the adoption fraction as a function of t for peers of buy-in users assigned to the baseline, passive, and active viral treatment groups.

additional 1.5-fold increase in adoption fraction over peers of users in the passive viral treatment group. These graphs confirm that viral feature design has an economically significant impact on the diffusion of product adoption.

4.4. Mechanisms Driving Social Contagion

Several social mechanisms could explain how viral features create product contagion. An unexpected result from the experiment enabled us to investigate these mechanisms. Interestingly, treated users not only had more peer adopters, but also used the application more than control group users (see Figure 3c). As Figure 3c shows, active-personalized users used the application more than passive-broadcast users, who in turn used the application more than baseline users. This result is surprising because users were randomly assigned to different applications and the versions were identical from users' point of view apart from the invitation option included in the active-personalized application. Understanding why use differed across treatment groups (despite randomization) provides insight into how viral features create contagion and sustained product use.

Viral feature inclusion, application use and peer adoption are correlated, as shown in Table 8, Model 1, and Figures 3a) and 3c). The randomized trial confirms that viral features cause peer adoption. We also know that because features are randomized and not controlled by the user, no other covariate can drive the existence of features.

While it may seem possible that omitted variables (such as unobserved user heterogeneity) could simultaneously drive application use and peer adoption, our data also rule out this possibility. Since feature inclusion is randomized, the distribution of any unobserved covariates must be the same across treatment groups and so, omitted variables cannot produce the discrepancy in peer adoption and application use across treatment groups shown in Figures 3a) and 3c). There could be an unobserved covariate which must first be activated by the existence of a feature to drive peer adoption and application use; however, that too is unsupported by the evidence. As we observe correlation between adoption and use beyond that which is explained by use of the invite and notifications features themselves (see Table 7), it

seems unlikely that a user characteristic that simultaneously drives peer adoption and use would be activated by a viral feature that the user does not use. Given these constraints, we depict the remaining possible causal relationships between feature inclusion, application use and peer adoption in Figure 4.

Figure 4 About Here

Cases (a) and (b) in Figure 4 cannot explain the observed correlation between viral feature inclusion and peer adoption. Specifically, they are inconsistent with the discrepancy in application use between users in different treatment groups (Figure 3c and Table 7).

Cases (c) and (d) represent a network externalities mechanism (Van den Bulte and Stremersch 2004), in which peer adoption drives increased application use by the original adopter. If peer adoption creates more sustained product use, generating more viral messages and a greater likelihood of peer influence, positive network externalities could create a virtuous cycle of engagement and social contagion. The number of peer adopters a user has is positively associated with their own sustained use of the application even when controlling for their treatment status, degree and overall Facebook activity (Table 7, Models 3). In addition, users of active-personalized and passive-broadcast applications exhibit more use (Table 7, Model 1) again controlling for observable differences in users' overall Facebook activity (Table 7, Model 2). These results are consistent with the existence of positive network externalities – as more of their peers adopted the application, users were more engaged and used the application more.⁹ The network effects explanation is supported by the evidence and seems plausible.

However, there could be alternative explanations for these results. For instance, cases (e) and (f) in Figure 4 represent a demand effect explanation, in which the correlation between features and application use is explained by an increased utility from the existence of viral features. In one variant of the demand effects explanation, the features themselves make the application more interesting and therefore simultaneously drive application use and peer adoption, creating a spurious correlation between peer adoption and product use. If the mere presence of the invite and notification features was correlated with

⁹ Although passive-broadcast features are associated with more product use than the baseline early on, this association disappears over time (see Figure 3c). That active-personalized features are associated with sustained product use while passive-broadcast features are not may suggest a *direct* network effect from interacting with *specific* peers.

both application use and peer adoption, and if peer adoption itself was not driving use, the correlation between the number of peer adopters and application use should disappear once we control for the use of invites and notifications. However, when we hold constant application use associated with both notifications and invites there is still a strong positive relationship between the number of peer adopters and application use (Table 7, Model 3), implying that increases in application use, beyond that explained by use of the viral features themselves, are correlated with more peer adoptions. This suggests demand effects do not fully explain the correlation between peer adoption and use.

It could still be that the viral state of the application itself makes the application more interesting - that the mere existence of features rather than their use increases users' utility. However, for that theory to be consistent with our data, users would have to derive utility from viral features they do not use. The data do not support this explanation either. When the viral states are entered into the regression they significantly predict application activity in the expected directions and magnitudes (Table 7, Model 1). When the number of peer adopters is controlled, those relationships disappear completely (Table 7, Models 3 and 4), indicating that the viral state of the application alone—or the utility from simply being able to notify or invite friends—does not predict application use.

An alternative explanation consistent with the causal relationships depicted in (e) and (f) is that there is a demand effects from the existence of viral features which inspires peer adoption – that peers' expected utility from adopting the application is higher because they expect to have access to viral features. For example, it could be that because a user received an invitation, they adopted the application because they valued the ability to invite others and expected to have this feature in the product they adopted, creating a demand effect from the expected utility of having viral features enabled. It seems unlikely that a significant portion of the expected utility from adopting the application comes from the existence of the viral features rather than the functions of the application itself. However, to address this alternative explanation we performed additional analysis. On average, application use by peer adopters is a reasonable proxy for their satisfaction with the product – the extent to which their expectations regarding the product conform to the product they actually received upon adoption. We therefore examined the ap-

plication use of peers that adopted through response to a viral message and divided those peers into two groups: those that received (through random assignment) a version of the product with the ability to send viral messages of the type they received from their influencing peer and those who received a version of the product without the ability to send viral messages of the type they received from their influencing peer. T-tests show that the use of the application by those who received applications with the features they would expect to receive and those that were “disappointed” (so to speak) by not receiving an application with the features they would expect to receive show no significant differences in application use (t -*statistic* = 0.9054; S.D. = 8.0377). Given this evidence it is unlikely that adoption and use are explained by the attractiveness of the applications with viral features. These results imply demand effects alone cannot explain increased use of applications with viral features. We must therefore reject cases (e) and (f).

The only remaining explanations are those depicted in (g) and (h), which capture both the demand effects arguments (features drive use) and network externalities (peer adoption drives use). Given these analyses, network externalities are at least in part responsible for driving application use. But, this is a conservative interpretation of the evidence. As we have ruled out most of the plausible demand effects explanations in (g) and (h) (that features drive use) it is likely that network effects are entirely responsible for the increased application use we observe among treated users.

Another piece of evidence corroborating network effects is that the hazard rate of adoption is increasing over adoption events, implying a reinforcement effect of prior adoptions on the likelihood of future adoption (Van den Bulte and Stremersch 2004). The hazard rate of adoption increases faster than exponentially for the first several adoption events, then more slowly, suggesting that reinforcement is approximately constant over peer adoptions. Although we interpret these results with caution because “one cannot distinguish between contagion and heterogeneity only on the basis of statistical properties of the distributional form” (Taibleson 1974: 878), the fact that the hazard rate of adoption is increasing in the k adoption events is consistent with a reinforcement effect of prior adoptions on future adoption.

Table 6, Table 8 and Figure 5 About Here

Finally, Table 8 presents correlates of application diffusion which corroborate results of the randomized trial. Models 4-6 confirm that peers of initial adopters also use the application because diffusion depth depends on peers' (and peers of peers) application use. Active-personalized and passive-broadcast treatments significantly increase average diffusion depth and these effects are again explained by application use and the viral features themselves (Model 6). Results in Model 3 also corroborate hazard model estimates, confirming that invitations are on average three times more effective per message in inspiring peer adoption than notifications. Taken together, evidence of a strong correlation between the number of adopter friends and application use and the distributional properties of the baseline hazards of adoption events suggest that network externalities accelerate contagion. As more of a user's friends adopt, they use they application more, creating a positive feedback loop.

5. Conclusion

We conducted a large-scale randomized experiment testing the effectiveness of viral product design features in creating social contagion. We found that viral product design has econometrically identifiable impacts on peer influence and social contagion in new product diffusion. Results of our randomized trial suggest that designing viral features into products can increase social contagion by up to 400%. Surprisingly, designing products with passive-broadcast viral messaging capabilities generates more total peer influence and social contagion than adding active-personalized viral messaging capabilities. Although active-personalized messaging is more effective in encouraging adoption per message and is correlated with more user engagement and sustained product use, it is used less often and therefore generates less total peer adoption in the network. Data on the distributional form of the diffusion process and on product use are consistent with the existence of positive network externalities that reinforce peer adoption and create a virtuous cycle of engagement and contagion. These results have broad implications for managers attempting to promote viral product diffusion and for theories of social contagion, opinion leadership and viral product design.

First, our estimates imply that viral product design may be more effective in encouraging new product adoption than traditional marketing strategies alone. Although conversion rates (CR) are always significantly smaller than click-through rates (CTR), the CR for notifications and invitations even outperformed published statistics on CTR for traditional banner advertising, paid search advertising and email marketing campaigns. The 1% CR on notifications outperforms the CTR for traditional banner advertising (which range from .10-.20% in publicly available statistics) and invitations are ten times as effective at generating conversions as traditional banner ads are at generating click-throughs. Compared to email campaign CTR (which range from 2% to 6% in publicly available statistics), invitations are again more effective at a 10% conversion rate.¹⁰ Notifications and invites also vastly outperform the ad campaign used in our recruitment phase on Facebook, which produced .07% and .01% click-through and conversion rates respectively. As Facebook currently has the largest market share of display advertising on the Web, these comparisons are even more impressive (Tucker 2010). Notifications and invites also outperform CTR and CR in paid search advertising, which have been estimated at .15% and .02% respectively (Ghose and Yang 2009). These comparisons show viral channels to be more effective at generating higher response rates than traditional digital advertising channels on the Web and in Facebook.

We also asked the directors of the firm with whom we partnered about their feature implementation and customer acquisition costs and learned that invites can be implemented for a total cost under \$600. Since implementing viral features incurs a low one-time fixed cost and the expected return is proportional to the increase in adopters the feature generates, viral product design may be a more cost-effective strategy than increasing spending on traditional digital advertising (which incurs costs proportional to impressions or clicks). It may be however that the success of viral product design efforts depends on traditional advertising to the extent that an initial base of users is needed to implement viral marketing. It is also important to consider the social cost of viral messages. Bombarding users with messages from

¹⁰ Click-through rates on banner ads have declined from 0.33% to 0.19% from 2004 through 2008 (Forrester, Go Big or Go Home Advertising, 2009); DoubleClick reports that in 2008 average CTR in the US was 0.10% for banner ads (DoubleClick, Benchmark Report, 2009). For email campaigns, estimated CTR in 2008 and 2009 remained stable at 5.9%. (Epsilon, October 2009), while Mailer Mailer reports average CTR on email campaigns at 2.80% (Mailer Mailer, June 2009) and Web Market Central reports a "2-3%" CTR (Web Market Central 2007).

peers may reduce the overall quality of the user experience. Future work should estimate the costs of viral product design more comprehensively and consider the implications of both marginal revenue and marginal cost on optimal product design.

Second, given that active-personalized features are more marginally effective but less globally effective than passive-broadcast features, a natural question is how managers can optimize the effectiveness of these viral features. As the main limitation of active-personalized features is that high effort costs curtail their use, one solution may be to couple active-personalized features with referral incentives that encourage their use (Biyalogorsky et al 2001). Optimally designed incentive strategies could encourage users to generate more personalized referrals and to target and personalize viral messages more effectively. It may also be possible to improve the low marginal effectiveness of passive-broadcast features by automatically targeting and personalizing broadcast messages algorithmically. If there is a social cost to viral messages, product and platform developers could seek to limit impersonal messages in adaptive ways that are tied to the effectiveness of the messages themselves.

Third, in the presence of viral features, network externalities drive a positive feedback loop in which product use drives peer adoption and peer adoption in turn drives product use. Managers should seek to enable this feedback loop by designing strong direct and indirect network externalities into their products. Maximizing engagement and minimizing churn may be obvious goals, but the effects of engagement and churn on product diffusion are less obvious. Interactions between network externalities, sustained use and customer churn may change over a product's lifecycle and may vary across products. More work on the relationships between social contagion, sustained use and customer churn over products and product lifecycles will help clarify when viral marketing is most effective.

Our work also has implications for platform developers that seek to benefit from social interactions taking place on their platforms. Platform developers can enable and constrain the viral features that operate in their ecosystem and engineer the user experience to increase sharing, interaction and the virality of products. If one considers the social cost of "spammy interactions," continuous redesign of social features for the purpose of optimizing the user experience is likely a rational, profit maximizing strategy.

From the platform developer's perspective optimization may take place over different variables and constraints, but improving the virality and use of applications available to users is likely an important goal.

Understanding optimal viral product design strategies, taking into account factors such as sustained product use, network externalities, social and economic costs, incentives and the marginal effectiveness of different viral features, could enable firms to optimally create and manage social contagion. The difficulty however is in determining what works and what does not. Numerous statistical challenges prevent clean causal estimation of the relationships between interventions and outcomes and the likely effects of changes in product design and platform policy. Fortunately, IT-based products and platforms provide natural vehicles for randomized experimentation. Given the low cost of conducting experiments, the rapid development and testing of viral design features, and the winner-take-all nature of markets with network externalities, this type of experimentation is likely to increase in the future and eventually to become commonplace in the development of many products and platforms. Our work sheds light on how viral products can be designed to generate social contagion, and offers a template for using randomized trials to identify peer influence in networks.

References

- Andersen, P. K. and Gill, R. D. (1982) "Cox's regression model for counting processes: a large sample Study." *The Annals of Statistics*, 10(4): 1100-1120.
- Aral, S. (Forthcoming) "Identifying Social Influence: A Comment on Opinion Leadership and Social Contagion in New Product Diffusion." *Marketing Science*.
- 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., (2009) "Networks, Information & Brokerage: The Diversity-Bandwidth Tradeoff." Available at SSRN: <http://ssrn.com/abstract=958158>
- Berger J and Heath C (2005) "Idea habitats: How the prevalence of environmental cues influences the success of ideas." *Cognitive Science* 29(2): 195-221.
- Biyalogorsky E, Gerstner E, & Libai B (2001) "Customer Referral Management: Optimal Reward Programs." *Marketing Science* 20(1): 82-95.
- Bramoulle Y, Djebbari H, & Fortin B (2009) "Identification of Peer Effects through Social Networks." *Journal of Econometrics* 150(1):41-55.
- Brown JJ & Reingen PH (1987) "Social Ties and Word-of-Mouth Referral Behavior." *The Journal of Consumer Research* 14(3):350-362.
- 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.
- Coleman JS (1988) "Social Capital in the Creation of Human Capital," *American Journal of Sociology*, 94: S95-S120.

- Dellarocas, C. (2006) "Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms" *Management Science*, 52(10):1577-1593.
- Duflo E, Glennerster R, & Kremer M (2006) "Using Randomization in Development Economics Research: A Toolkit." *Handbook of Developmental Economics*, Elsevier.
- Emerson RM (1962) "Power-Dependence Relations," *American Sociological Review* 27(1): 31-41.
- Frenzen J & Nakamoto K (1993) "Structure, cooperation, and the flow of market information," *Journal of Consumer Research* 20(3): 360-375.
- 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.*
- Ghose, A. and Yang, S. (2009). "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets." *Management Science*, 55(10): 1605-1622.
- Godes D & Mayzlin D (2004) "Using Online Conversations to Study Word-of-Mouth Communication." *Marketing Science* 23(4):545-560.
- Godes D & Mayzlin D (2009) "Firm-Created Word-of-Mouth Communication: Evidence from a Field Test." *Marketing Science* 28(4):721-739.
- Goldenberg J, Han S, Lehmann DR, & Hong JW (2009) "The Role of Hubs in the Adoption Process." *Journal of Marketing* 73(2):1-13.
- Gross R., Acquisti A, & Heinz J (2005) "Information Revelation and Privacy in Online Social Networks", *Workshop on Privacy in Proceedings of the 2005 ACM Workshop on Privacy.*
- Hartmann, W.R., Manchanda, P., Nair, H., Bothner, M., Dodds, P., Godes, D., Hosanagar, K., and C. Tucker (2008) *Marketing Letters*, 19: 287-304.
- Heath, C., Bell, C., and E. Sternberg (2001) "Emotional selection in Memes: The case of urban legends." *Journal of Personality and Social Psychology*. 81(6): 1028-1041.
- Hill S, Provost F, & Volinsky C (2006) "Network-Based Marketing: Identifying Likely Adopters via Consumer Networks." *Statist. Sci.* 21(2):256-276.
- Iyengar R, Van den Bulte C, & Valente TW. (Forthcoming) "Opinion Leadership and Social Contagion in New Product Diffusion." *Marketing Science.*
- Jackson MO (2008) "Average Distance, Diameter, and Clustering in Social Networks with Homophily." *Internet and Network Economics* 5385: 4-11.
- Katz, E., P.F. Lazarsfeld (1955) *Personal Influence*. Free Press, New York.
- Kelejian HH & Prucha IR (1998) "The method of elimination and substitution in the GMM estimation of mixed regressive, spatial autoregressive models." *Jrnl of Real Estate Fin. & Econ.* 17(1):99-121.
- Manchanda P, Xie Y, & Youn N (2008) "The Role of Targeted Communication and Contagion in Product Adoption." *Marketing Science* 27(6):961-976.
- Manski CF (1993) "Identification Problems in the Social Sciences." *Sociological Methodology* 23:1-56.
- Mayzlin, D. (2006) "Promotional Chat on the Internet." *Marketing Science* 25(2) 155-976.
- Nam, S, Manchanda, P and PK Chintagunta (2010) "The Effect of Signal Quality and Contiguous Word of Mouth on Customer Acquisition for a Video-On-Demand Service" *Marketing Science*: 1-11.
- Oestreicher-Singer G & Sundararajan A (2008) "The Visible Hand of Social Networks in Electronic Markets," *Working Paper*, 2008, Available at <http://ssrn.com/abstract=1268516>
- Phelps, JE, L Lewis, L, Mobilo, D Perry, N Raman (2004) "Viral Marketing or Electronic Word-of-Mouth Advertising: Examining Consumer Responses and Motivations to Pass Along Email." *Journal of Advertising Research* 44(4):333-348.
- 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).
- Stephen AT & Berger JA (2009) "Creating Contagious: How Social Networks and Item Characteristics Combine to Spur Ongoing Consumption and Reinforce Social Epidemics." *Working Paper, Wharton.*

- Taibleson, MH (1974) "Distinguishing between contagion, heterogeneity, and randomness in stochastic models." *American Soc. Review* 39(9): 877-880.
- Tam KY & Ho SY (2005) "Web Personalization as a Persuasion Strategy: An Elaboration Likelihood Model Perspective." *Information Systems Research* 16(3):271-291.
- Trusov, M, Bucklin, R.E., & Pauwels, K. (2009) "Effects of Word-of-Mouth Versus Traditional Marketing: Findings from an Internet Social Networking Site." *Journal of Marketing*, 73: 90-102.
- Tucker C (2008) "Identifying Formal and Informal Influence in Technology Adoption with Network Externalities." *Management Science* 54(12): 2024-2038.
- Tucker C (2010) "Social Networks, Personalized Advertising, and Privacy Controls" Net Institute Working Paper #10-07.
- Van den Bulte C & Iyengar R (2010) "Tricked by Truncation: Another Source of Spurious Social Contagion in New Product Adoption." *Working Paper, Wharton*.
- Van den Bulte C & Lilien GL (2001) "Medical Innovation Revisited: Social Contagion versus Marketing Effort." *The American Journal of Sociology* 106(5):1409-1435.
- Van den Bulte, C. and S. Stremersch (2004) "Social Contagion and Income Heterogeneity in New Product Diffusion: A Meta-Analytic Test." *Marketing Science*, 23(4):530-544.
- Van der Lans, R., van Bruggen, G, Eliashberg, J. and Wierenga, B. (2010) "A Viral Branching Model for Predicting the Spread of Electronic Word of Mouth." *Marketing Science* 29(2): 348-365.
- Watts DJ & Dodds PS (2007) "Influentials, Networks, and Public Opinion Formation," *Journal of Consumer Research* 34(4):441-458.

Tables and Figures

Table 1. Stratification Across Treatment Groups

Baseline Control	Passive-broadcast Treatment	Active-personalized Treatment
5%	47.5%	47.5%

Table 2. Recruitment Statistics Describing the Initial Advertising Campaign

Wave	Impressions	Clicks	Advertising Related Installs	Installs
1 (Day 0)	18,264,600	12,334	3,072	3,714
2 (Day 15)	20,912,880	25,709	2,619	3,474
3 (Day 20)	19,957,640	7,624	3,219	4,039
Total	59,135,120	45,667	8,910	11,227

Table 3. Summary Statistics and Mean Comparisons of Active, Passive and Baseline Users

	1	2	3	4	5	6
	Baseline (N = 405)	Passive (N = 4600)	Active (N = 4682)	t-statistic (B-P)	t-statistic (B-A)	t-statistic (P-A)
	Mean (SD)	Mean (SD)	Mean (SD)	t-statistic (SE)	t-statistic (SE)	t-statistic (SE)
Age	31.51 (13.80)	30.81 (13.31)	29.94 (13.27)	.46 (13.35)	1.03 (13.31)	1.45 (13.24)
Gender (1=Male)	.25 (.44)	.33 (.47)	.32 (.47)	-1.57 (.47)	-1.42 (.46)	.40 (.47)
Degree [†]	171.79 (223.88)	170.25 (278.64)	166.97 (248.77)	.09 (275.13)	.32 (247.15)	.55 (263.82)
Number of Facebook Wall Posts	40.52 (79.89)	36.45 (94.16)	37.07 (246.76)	.46 (93.11)	.15 (238.20)	-.09 (188.31)
Number of Adopters in User's Local Network	.01 (.12)	0.07 (.35)	0.10 (.44)	-2.84*** (.34)	-3.60*** (.43)	-3.64*** (.40)
Percentage of Adopters in User's Local Network	.02 (.002)	.09 (.01)	.15 (.01)	-1.92* (.01)	-2.35** (.01)	-2.83*** (.01)
Maximum Diffusion Depth	.01 (.11)	.04 (.22)	.05 (.24)	-2.53* (.21)	-3.01*** (.24)	-1.98*** (.23)
Time to 1st Adopter	9.40 (9.71)	4.77 (8.04)	3.17 (6.72)	1.27 (8.07)	2.04** (6.77)	2.45*** (7.30)
Time to 2nd Adopter	---	5.23 (8.17)	4.43 (6.97)	---	---	0.58 (7.45)
Time to 3rd Adopter	---	5.29 (8.07)	3.04 (5.25)	---	---	1.08 (6.33)
Time to 4th Adopter	---	6 (5.83)	1.17 (1.12)	---	---	2.84*** (3.58)
Application Use	3.17 (4.59)	4.17 (7.24)	4.56 (8.98)	-2.54** (7.08)	-2.89*** (8.73)	-2.20* (8.16)

Notes: This table reports means and standard deviations for demographic variables, peer adoption statistics and Facebook and application activity statistics of Baseline (Column 1), Passive (Column 2), and Active (Column 3) control and treatment group users, as well as results of T-tests of mean differences between Baseline and Passive Users (Column 4), Baseline and Active Users (Column 5) and Passive and Active Users (Column 6). Variables reported include: Age: Self-reported age on Facebook; Gender (1=Male): Self-reported gender on Facebook; Degree: Number of Facebook friends; Number of Facebook Wall Posts: Count of the number of 'wall posts' posted to an individual's Facebook profile recorded at the beginning of the study; Percentage of Adopters in User's Local Network: The percentage of an individual's Facebook friends who adopted the application calculated at the end of the observation period; Time to 1st, 2nd, 3rd, 4th Adopter: The time in days to the first, second, third and fourth adopters in the user's friend network; Application Use: A continuous measure of application calls from a user's account to the application server indicating the number of actions taken on the application. Statistical Significance of T-Tests: ***p<.001; **p<.05; *p<.10; † Kolmogorov-Smirnov Tests of Degree Distribution Differences: Baseline-Passive: .04, p = .80, N.S.; Baseline-Active: .04, p = .79, N.S.; Passive-Active: .01, p = .94, N.S.

Table 4: Variance-Corrected Stratified Proportional Hazards of Contagion in Networks of Baseline, Passive and Active Treatment Groups

	1	2	3	4
	<i>Hazard Ratio (SE)</i>	<i>Hazard Ratio (SE)</i>	<i>Hazard Ratio (SE)</i>	<i>Hazard Ratio (SE)</i>
Viral State = Passive	3.46*** (1.18)	3.35*** (1.15)	2.50** (.86)	2.51** (.86)
Viral State = Active	4.44*** (1.64)	4.21*** (1.56)	3.33*** (1.24)	3.31*** (1.24)
Application Use		1.02*** (.004)	1.02*** (.003)	1.02*** (.003)
Notifications			1.02*** (.002)	1.02*** (.002)
Invites				1.06** (.028)
Log Likelihood	-4694.359	-4631.795	-4544.845	-4542.577
X ² (d.f)	19.34*** (2)	57.41*** (3)	298.78*** (4)	307.47*** (5)
Observations	3929	3929	3929	3929

Notes: The table reports parameter estimates and standard errors from the Variance-Corrected Stratified Proportional Hazards Model specified on page 13 with robust standard errors clustered around users' local network neighborhoods. Variables reported include: Viral State = Passive: A dummy variable denoting Passive Viral Application users; Viral State = Active: A dummy variable denoting Active Viral Application users; Application Use: A continuous measure of application calls from a user's account to the application server indicating the number of actions taken on the application; Notifications: Integer count of the number of notifications sent; Invites: Integer count of the number of invites sent; Statistical Significance of parameters is reported as follows: ***p<.001; **p<.05; *p<.10.

Table 5: Clickstream Analysis of Responses to Viral Messages and Adoption

	1	2	3
	<i>Messages Sent</i>	<i>Adoptions via Click Through Installation</i>	<i>Adoption Rate (Marginal Impact)</i>
Invitations	160	16	.10
Notifications	69980	666	.01

Table 6: Baseline Hazards Over k Events λ_{0k} ($k = 1...6$)

	1	2	3	4
	<i>Mean (SD)</i>	<i>Min</i>	<i>Max</i>	<i>N</i>
λ_{01}	.0002 (.0001)	.0001	.001	523
λ_{02}	.002 (.001)	.001	.013	128
λ_{03}	.015 (.024)	.005	.14	42
λ_{04}	.034 (.010)	.021	.054	20
λ_{05}	.046 (.008)	.037	.067	15
λ_{06}	.099 (.044)	.053	.14	7

Notes: The table reports means, standard deviations, minimum and maximum values for baseline hazard lambda parameters of the k^{th} adoption events in users' networks, $k = 1...6$.

	1	2	3	4
	Application Use	Application Use	Application Use	Application Use
	<i>Beta</i> (<i>SE</i>)	<i>Beta</i> (<i>SE</i>)	<i>Beta</i> (<i>SE</i>)	<i>Beta</i> (<i>SE</i>)
Viral State = Passive	.129* (.074)	.112 (.079)	.062 (.076)	-.037 (.074)
Viral State = Active	.190*** (.074)	.171** (.079)	.091 (.076)	-.006 (.074)
Degree	-.0001 (.0001)	-.0001 (.0001)	-.0002** (.0001)	-.0002** (.0001)
Facebook Activity		.054*** (.016)	.042*** (.015)	.026* (.014)
Notifications				.022*** (.001)
Invites				.055** (.024)
Number of Adopters			.607*** (.030)	.360*** (.031)
F Value (d.f.)	3.51*** (3)	4.87*** (4)	83.54*** (5)	128.92*** (7)
R ²	.002	.003	.07	.14
Observations	6310	5766	5766	5766

Notes: This table reports OLS parameter estimates for a linear estimating equation regressing Application Use (defined in Table 3) on the variables listed, including: Facebook Activity: The normalized sum of integer counts of the number of wall posts, activities, affiliations, groups, interests, pages, notes, favorite books, movies, music, and TV shows, calculated at the beginning of the study. For all other variable definitions see Tables 3 and 4. Models are estimated with robust standard errors clustered around users' local network neighborhoods. Statistical significance of parameters is reported as follows: ***p<.001; **p<.05; *p<.10.

	1	2	3	4	5	6
	Number of Adopters	Number of Adopters	Number of Adopters	Diffusion Depth	Diffusion Depth	Diffusion Depth
	<i>Beta</i> (<i>SE</i>)	<i>Beta</i> (<i>SE</i>)	<i>Beta</i> (<i>SE</i>)	<i>Beta</i> (<i>SE</i>)	<i>Beta</i> (<i>SE</i>)	<i>Beta</i> (<i>SE</i>)
Viral State = Passive	.078** (.031)	.084** (.033)	.020 (.059)	.045** (.0178)	.048*** (.019)	.020 (.018)
Viral State = Active	.119*** (.031)	.131*** (.032)	.059* (.030)	.057*** (.018)	.063*** (.019)	.033* (.018)
Degree	.0001*** (.00002)	.0001** (.00003)	.0001** (.00002)	.0001*** (.00001)	.00004** (.00002)	.00003** (.00001)
Facebook Activity		.019*** (.006)	.006 (.006)		.013*** (.004)	.007** (.004)
Application Use			.061*** (.005)			.021*** (.003)
Notifications			.010*** (.0004)			.005*** (.0002)
Invites			.035*** (.010)			-.003 (.006)
F Value (d.f.)	12.20*** (3)	11.18*** (4)	157.94*** (7)	9.36*** (3)	10.11*** (4)	85.13*** (7)
R ²	.006	.007	.16	.004	.007	.09
Observations	8910	5766	5766	6310	5766	5766

Notes: This table reports OLS parameter estimates for linear estimating equations regressing the Number of Adopters (defined in Table 3) and Diffusion Depth: The maximum network distance from a treated user to any peer adopter in a linked chain of adoptions, on the variables listed. See Table 7 for all additional notes and variable definitions.

Figures.

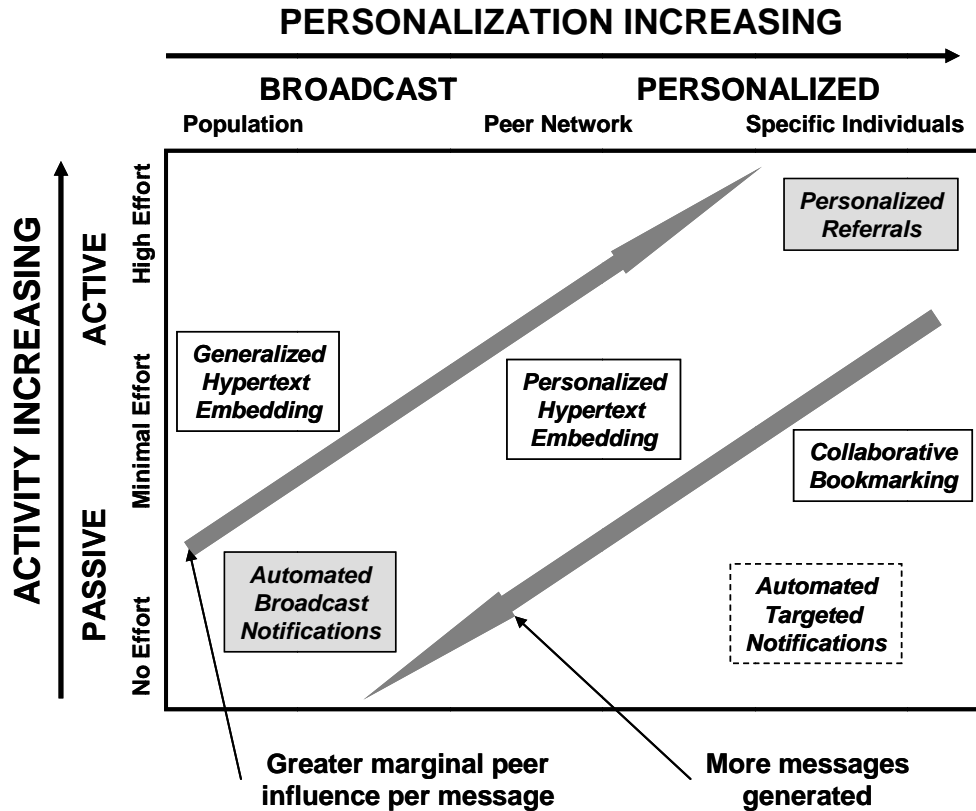


Figure 1. The Viral Product Feature Space. Describes the space of viral features along two dimensions - activity and personalization. We expect greater marginal peer influence per message to be generated by features near the upper right corner of the space (e.g. Personalized Referrals), and more messages to be generated by features near the bottom left corner of the space (e.g. Automated Broadcast Notifications). The space is not limited to personalized referrals and automated broadcast notifications. For example, Slide.com and RockYou.com use hypertext embedding to allow users to create and embed slideshows of pictures or other content on their websites, weblogs and social networking profile pages. As other users browse those items, hyperlinks allow them to download the products themselves. Personalized hypertext embedding, such as profile box installations on Facebook are more personalized than generalized hypertext embedding because they target a user’s personal social network. Collaborative bookmarking sites like Delicious.com are personalized but also include an element of algorithmic activity. The automated targeted notifications box represents a feature that could exist if notifications were targeted toward specific individuals using collaborative filtering.




Figure 2. Graphical Representation of the Experimental Comparison

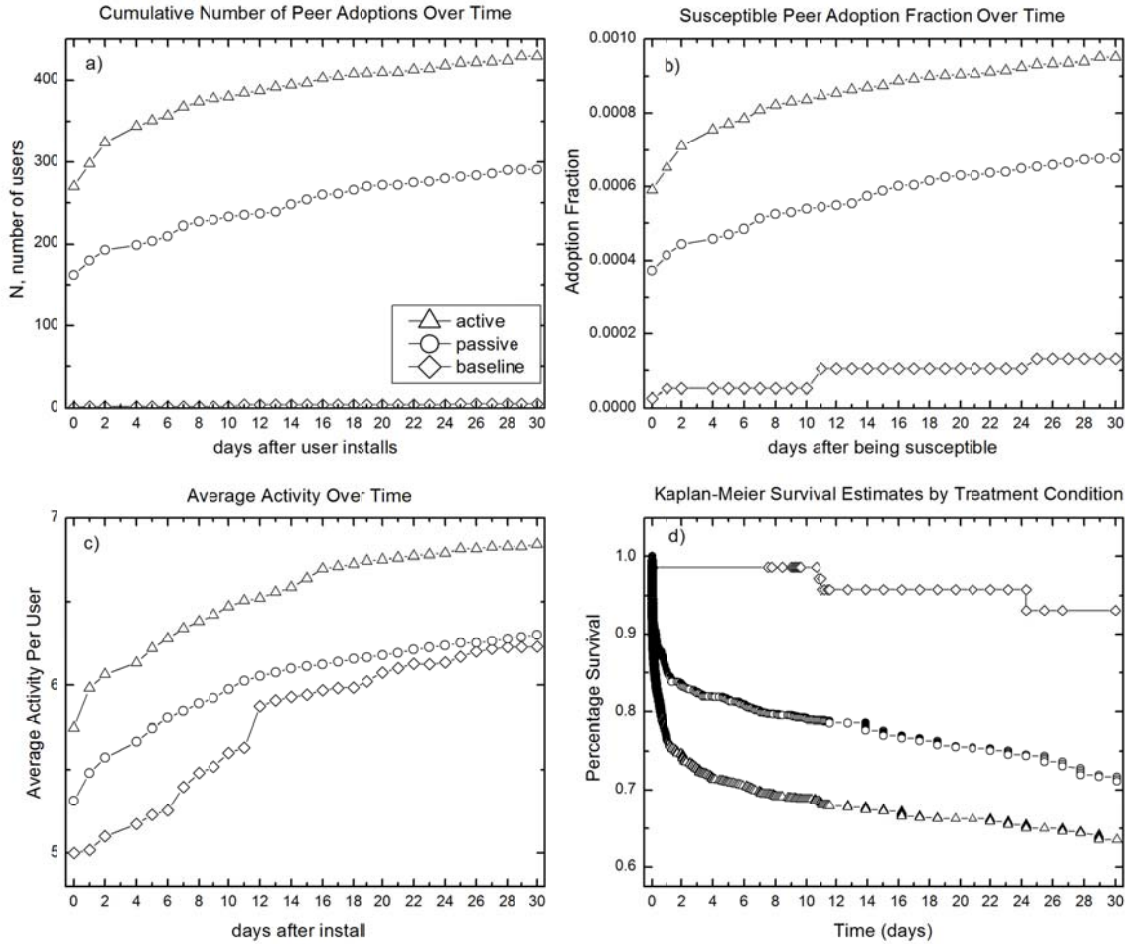


Figure 3. Plots a) the cumulative number of peer adoptions, b) the fraction of susceptible peer adopters, c) the average activity, and d) the Kaplan-Meier Survival Estimates over time for baseline, active and passive users.

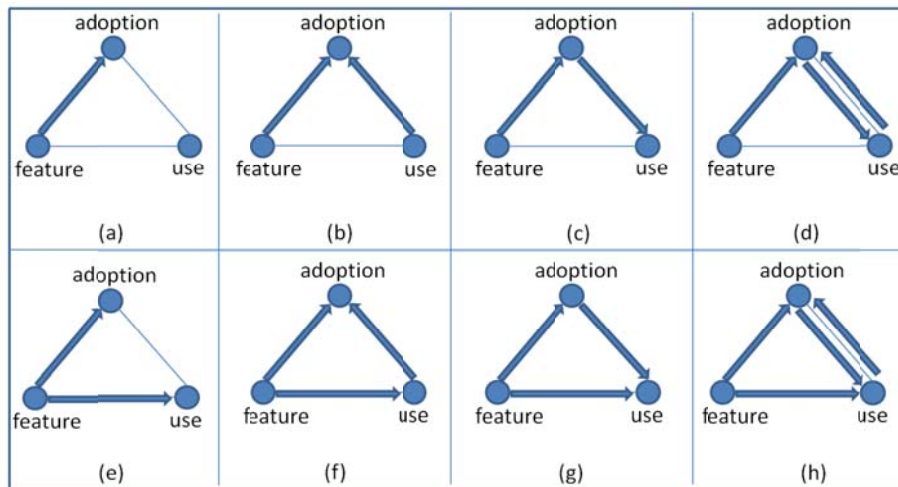


Figure 4: Possible causal relationships between the existence of application features, peer adoption, and application use. Arrows indicate causal direction.




Figure 5. Baseline Hazards (λ_{0k}) for $k = 1 \dots 6$ fitted to an Exponential and a Power Function

Appendix

1. Inside-Out Design

Randomized trials are traditionally used to estimate the effect of a treatment on the treated. To study the effect of viral feature incorporation on product adoption outcomes, we instead examine the effect of treatment on the peers of treated application users. The difference in these approaches is illustrated in Figure A1. 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 measurements of treatment effects. In social network environments, a conventional approach is infeasible because it is difficult to comprehensively control the network environments of each user in the study population. It is feasible however to treat a user and observe the effect of treatment on the outcomes of their peers.

The strength of our approach lies in its ability to capture effects of any form of influence-mediating communication channels between the treated user and her peers, including effects that arise through influence-mediating communication channels beyond those that can be explicitly recorded. For example, treated users could communicate with and influence their peers through offline interactions such as face-to-face communications or telephone conversations, as well as through unrecorded online communications such as email or external chat conversations. Because we measure the response of peers regardless of how they may or may not have been influenced by treated users, we are able to capture the effect of unrecorded influence-mediating communications on peer adoption.




Figure A1

2. Preventing Selection Effects

Selection effects could occur when a user chooses to adopt the experimental application through the recruitment campaign or when they adopt in response to a viral message. We took steps to mitigate and measure both possibilities. The recruitment campaign was designed to reach a representative audience of Facebook users and advertisements were displayed to users through advertising space within Facebook and within existing Facebook applications. Establishing to what extent the recruited population is representative of the general Facebook population is somewhat challenging because Facebook does not offi-

cially publish demographic statistics of their user base. However, through the use of a recently released social targeting advertisement service provided by Facebook, it is possible to obtain some official demographic statistics. Age and gender demographics sampled through this API and published online by *istrategylabs.com* are compared to the same demographic ranges for recruited study population users in Figure 2. Though our sample has a slightly higher percentage of women than the Facebook population and users in our sample have a slightly higher average degree (150 compared with 130 in Facebook statistics), the demographics of our study population are comparable to that of the broader Facebook population and the published Facebook demographics fall within one standard deviation of study population sample means.

In addition to issues of selection surrounding the population of recruited users, application users that adopt due to peer influence may be subject to selection effects and may be fundamentally different from application users that adopted via initial recruitment. It could be that users who use the viral features and peers of users who use viral features are systematically different from randomly selected Facebook users. We avoid these sources of selection bias in our analyses by only considering initially recruited users in the randomized treatment group to which they were assigned. Peers of recruited users only contribute to local network peer adoption of originally recruited users and are not themselves used as test subjects.




Figure A2

3. Preventing Leakage and Contamination

In randomized trials in network environments, users assigned to different treatment groups may not be strictly isolated from one another. This raises the concern that information leakage through indirect network pathways may contaminate the results of the study. It is important to note that in traditional studies, whether or not the network is measured, relationships may still exist between treatment and control populations that create leakage effects. One benefit of our design is that we systematically observe how individuals in the study are connected, enabling us to measure and prevent leakage.

Several factors reduce the likelihood that leakage is affecting our results. First, because treatment assignment is randomized, any leakage will be uncorrelated with treatment assignment and cannot account for the 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 the treatment designations, this effect should only serve to make our estimates across treatment groups more conservative as leakage should reduce differences between control and treatment groups. Second, information flows between individuals in a network typically decay rapidly with network distance (Wu et al 2004, Aral et al 2007). While all users may be connected through long friendship paths, leakage will diminish over successive hops in each path.

Nonetheless, leakage effects could downward bias our estimates of treatment effects toward zero and we therefore take several steps to prevent leakage. First, in hazard rate models, we only examine peers of initially recruited adopters. In addition to avoiding potential selection issues mentioned above, this also excludes individuals (and their potential adopter peers) that adopt in chains within a local neighborhood lessening leakage effects. Such individuals are likely to share more and shorter indirect paths with existing adopters than a randomly chosen peer, as a consequence of clustering and mutuality (Newman 2003).

Second, we account for users with multiple treated peers (of similar and/or different treatments). Existence of peers of multiple treated application users leads to two potential complications. First, users may be peers of multiple treated users from different treatment groups making it impossible to link their treatment effects to a single treatment. Second, peers of multiple treated users that belong to the same treatment group are clearly classified as peers of either baseline, active-personalized or passive-broadcast users; however measurements of their response may be incorrectly estimated as a consequence of being subject to influence from multiple treated friends. A peer with multiple treated friends in a given viral treatment group may exhibit an adoption outcome or time to adoption that is systematically different from those of peers with only one treated friend. These two scenarios are displayed in Figure A3.




Figure A3




Figure A4

The nature of treatment randomization does not allow us to simultaneously guarantee that all treated friends of a peer will receive the same treatment. Consequently we treat peers with multiple treated friends as contaminated as soon as they become so and exclude them from our analysis. This procedure could underestimate the effect of clusters of adoption on the time to adoption or number of adopters in a local network neighborhood, however if this is the case, it will do so in a manner that is the same for all treatment types. Furthermore, as treatment groups are randomized, there can be no systematic correlation between the type of treatment received by a user and that received by her subsequent adopter peer.

The procedure that we adopt for designating a peer as contaminated is detailed in Figure A4. 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. In panel (a) for times $t > t_1$, peer 2 has multiple treated peers (R and 1) that may have been assigned different treatments. Peer 2 is therefore considered contaminated for times $t > t_1$. In panel (b) a similar situation occurs, but no link exists between peers 1 and 2 and consequently neither user is considered contaminated. In our analysis, when a peer is designated as contaminated, she is removed from the hazard rate model for subsequent time periods. This procedure appropriately retains the maximal empirical support for hazard rate estimation and parameterizes our ignorance of what might happen subsequent to a user's contamination. The right-censoring of contaminated subjects has become standard practice in randomized clinical trials where a patient in a randomized treatment group undergoes some characteristic change that is beyond the researcher's control. Furthermore, by including right-censored observations in our data rather than truncating the data, we avoid problems caused by data truncation that could lead to spurious evidence of contagion (Van den Bulte and Iyengar, 2010).

We note that the exclusion of peers with multiple treated friends does not preclude measurement of network externalities. Peers of treated users that become adopters but are not connected are considered uncontaminated and are included in our analysis. For two peers of a treated user that are connected and eventually become adopters, the initial peer adopter is included in our analysis and only the peer that subsequently adopts is considered contaminated and excluded for all times subsequent to contamination. These procedures enable a tightly controlled randomized trial of peer influence that addresses the potential for selection and leakage effects.

4. Robustness Checks for Different Hazard Model Specifications

There are a limited number of survival models that apply to contexts with multiple failures. Among these models, the variance-corrected stratified proportional hazards model reported in the paper is the most appropriate specification given the structure of our data and the parameters we estimate. However, we also checked multiple other hazard model specifications in order to test the robustness of our results to changes in model specification and estimation strategy. For good reviews of appropriate specifications of survival models in multiple failure data we recommend Wei and Glidden (1997) and Ezell et al (2003).

Table A1 reports results of different hazard model specifications, all of which are similar to our own. We report the original variance-corrected stratified proportional hazards model specification detailed in the paper in Column 1. Column 2 reports an accelerated failure time model with a log-logistic survival distribution. Column 3 reports an exponential regression with log relative-hazard form. Column 4 reports results from a traditional Anderson-Gill model. Column 5 introduces a time-dependent covariate measuring the number of prior adopters to the traditional Anderson-Gill specification to capture the dependence structure among recurrence times which in our original model is captured by the adoption event strata *k*. Column 6 reports a Prentice William and Peterson proportional hazards specification with time dependent strata. Finally, Column 7 reports results of a Wei Lin and Weissfeld marginal risk set model. We note that all specifications produce similar results. However, we are most confident in our original specification which is best suited to our context and data.

	1	2	3	4	5	6	7
<i>Specification</i>	<i>VCSPHM</i>	<i>AFT</i>	<i>EXP</i>	<i>AG₁</i>	<i>AG₂</i>	<i>PWP</i>	<i>WLW</i>
Viral State =	2.51**	-2.41**	1.01***	2.60***	2.54***	2.51***	2.00*
Passive	(.86)	(1.16)	(.35)	(.91)	(.87)	(.865)	(.78)
Viral State =	3.31***	-3.66***	1.30***	3.51***	3.30***	3.31***	2.62**
Active	(1.24)	(1.22)	(.39)	(1.36)	(1.26)	(1.24)	(1.02)
Application	1.02***	-.119***	.015***	1.02***	1.02***	1.02***	1.00
Activity	(.003)	(.039)	(.003)	(.003)	(.003)	(.003)	(.002)
Notifications	1.02***	-.115***	.025***	1.02***	1.02***	1.02***	1.01***
	(.002)	(.010)	(.002)	(.002)	(.001)	(.002)	(.002)
Invites	1.06**	-.198	.090**	1.07*	1.06**	1.06**	1.02
	(.028)	(.259)	(.036)	(.037)	(.035)	(.027)	(.018)
Prior Adopters					1.50***		
					(.062)		
Time Dum- mies	NO	YES	YES	NO	NO	NO	NO
Log Likelihood	-4542.58	-2826.32	-4136.53	-5254.17	-5212.88	-4542.56	-4561.56
X ² (d.f)	307.47*** (5)	--	1656.60*** (11)	412.65*** (5)	435.88*** (6)	307.60*** (5)	109.17*** (5)
Observations	3929	3929	3929	3929	3929	3929	3929
Notes: ***p<.001; **p<.05; *p<.10; Standard errors are clusters around users' local network neighborhoods. " <i>VCSPHM</i> ": Variance Corrected Stratified Proportional Hazards Model as specified and reported in the paper; " <i>AFT</i> ": Accelerated Failure Time Model with Log-Logistic Survival Distribution; " <i>EXP</i> ": Exponential Regression with Log Relative-Hazard form; " <i>AG</i> ": Anderson-Gill Model; " <i>PWP</i> ": Prentice William and Peterson Proportional Hazards Model with Time Dependent Strata; " <i>WLW</i> ": Wei Lin and Weissfeld Marginal Risk Set Model.							

References

Andersen, P. K. and Gill, R. D. (1982) "Cox's regression model for counting processes: a large sample study," *The Annals of Statistics*, Vol. 10, No. 4 (1982) , pp. 1100-1120

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.

Ezell ME, Land KC, & Cohen LE (2003) Modeling Multiple Failure Time Data: A Survey of Variance-

- Corrected Proportional Hazards Models with Empirical Applications to Arrest Data. *Sociological Methodology* 33(1):111 - 167.
- Newman, M.E.J (2003), "Ego-Centered Networks and the Ripple Effect," *Social Networks* 25(1): 83-95.
- Prentice, RL, BJ Williams, and AV Peterson (1981). "On the regression analysis of multivariate failure time data." *Biometrika*. 79: 495-512.
- Van den Bulte, C. and Iyengar, R. (2010) "Tricked by Truncation: Another Source of Spurious Social Contagion in New Product Adoption," *Working Paper, Wharton School*.
- Wei, L.J and Glidden, D.V. (1997) "An overview of statistical methods for multiple failure time data in clinical trials," *Statistics in Medicine*, 16(8): 833-839.
- Wei, L.J., Lin, D.Y. and Weissfeld, L. (1989) "Regression analysis of multivariate incomplete failure time data by modeling marginal distributions," *Journal of American Statistical Association*, 84(408): 1065-1073.
- Wei, L.J., Ying, Z., and Lin, D.Y. (1990) "Linear regression analysis of censored survival data based on rank tests," *Biometrika* 77(4): 845-851.
- Wu, F. Huberman, B., Adamic, L., & J. Tyler. (2004). "Information flow in social groups." *Physica A: Statistical and Theoretical Physics*, 337(1-2): 327-335.