

Selection Effects in Online Sharing: Consequences for Peer Adoption

SEAN J. TAYLOR, NYU Stern, Facebook[†]
 EYTAN BAKSHY, Facebook
 SINAN ARAL, NYU Stern

Most models of social contagion take peer exposure to be a corollary of adoption, yet in many settings, the visibility of one’s adoption behavior happens through a separate decision process. In online systems, product designers can define how peer exposure mechanisms work: adoption behaviors can be shared in a passive, automatic fashion, or occur through explicit, active sharing. The consequences of these mechanisms are of substantial practical and theoretical interest: passive sharing may increase total peer exposure but active sharing may expose higher quality products to peers who are more likely to adopt.

We examine selection effects in online sharing through a large-scale field experiment on Facebook that randomizes whether or not adopters share Offers (coupons) in a passive manner. We derive and estimate a joint discrete choice model of adopters’ sharing decisions and their peers’ adoption decisions. Our results show that active sharing enables a selection effect that exposes peers who are more likely to adopt than the population exposed under passive sharing. We decompose the selection effect into two distinct mechanisms: active sharers expose peers to higher quality products, and the peers they share with are more likely to adopt independently of product quality. Simulation results show that the user-level mechanism comprises the bulk of the selection effect. The study’s findings are among the first to address downstream peer effects induced by online sharing mechanisms, and can inform design in settings where a surplus of sharing could be viewed as costly.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics

General Terms: Economics, Experimentation

Additional Key Words and Phrases: viral marketing; information diffusion; social advertising; econometrics

1. INTRODUCTION

Standard models of social contagion consider adoption decisions of agents in the presence of social signals, but often take peer exposure to be a consequence of adoption [Bass 1969; Granovetter 1978; Jackson and Yariv 2007; Schelling 1973]. This is natural for many situations where adoption creates a persistent signal that peers can observe; if an individual buys a car, she will find it difficult to prevent her peers from knowing about her adoption. While theory tends to conflate adoption and exposure, they reflect substantive design decisions in practice. In online settings, developers and marketers may seek to increase their “virality” by providing encouragements and incentives to spread their product or message to others.

The decision to adopt and peer exposure can range from perfectly correlated to completely independent. Applications which implement passive sharing automatically broadcast users’ behaviors to their peers [Aral and Walker 2011]. Similarly, Liking a Page on Facebook induces publicly visible connection that persists over time [Bakshy et al. 2012a]. However, in other settings (e.g. browsing the Web), individuals must actively share their behaviors.

[†] This research was conducted while the author was visiting Facebook.

Author’s addresses: Sean J. Taylor staylor@stern.nyu.edu, Eytan Bakshy eytan@fb.com, Sinan Aral sinan@stern.nyu.edu

Permission to make digital or hardcopies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credits permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
 EC’13, June 16–20, 2013, Philadelphia, USA. Copyright © 2013 ACM 978-1-4503-1962-1/13/06...\$15.00

In this paper, we examine the interaction between propagation and adoption. Exactly how an individual's adoption decision is linked to peer exposure can vary depending on the medium and product or message being spread. The relationship can be modulated by providing adopters with encouragements to share¹ that may impact the diffusion process. Developers can engineer features which encourage sharing or make peer exposure a more reliable consequence of product adoption or use², fulfilling the role that prestige and attractiveness might play in offline goods and services [Veblen 2005]. In order to predict the consequences of these strategies for product diffusion, it is important to understand how the nature of a user's decision to share may have an impact on her peers' decisions to adopt.

We implement a large-scale field experiment on Facebook to measure the effect of sharing interface on peer adoption, randomizing whether users of a coupon product share their redemption behavior in an active or passive manner. Our results show that when adoption is perfectly linked with sharing, exposed peers are less likely to adopt the product in turn. The adoption effect we observe could be explained by either selection of influential adopters and/or susceptible peers (*dyad selection*), or by selection of higher quality products. To disentangle these effects, we derive a discrete choice model of the adoption and sharing process and use Bayesian estimation techniques to fit the model to our experimental data. We find strong evidence for dyad selection – peers of users who share in the active sharing condition are more likely to adopt the products their peers share. We also find evidence for selection on product quality, but the effect on overall adoption outcomes is small. The major downstream effects of sharing regime are dominated by the selection of dyads.

We proceed as follows. In the first section we review relevant literature on information diffusion in social networks and show how our work contributes to it. In the following section, we derive an econometric model linking sharing with peer adoption in Section 3. We then describe our empirical context and experimental design in Section 4. We summarize our results and estimate the selection effects using the experimental data and model in Section 5. Finally, we conclude with a discussion of our results and their implications in Section 6.

2. RELATED WORK

Online social networks allow users to articulate their relationships to people, companies, and products, and enable studies of diffusion processes *in vivo* across consumer behaviors, including product recommendations [Leskovec et al. 2007], the adoption of social applications [Aral et al. 2009; Wei et al. 2010], and link re-sharing [Bakshy et al. 2011; Goel et al. 2012].

Recent studies are also beginning to analyze mechanisms of information transmission and their causal interpretations. Since individuals form relationships with similar others [McPherson et al. 2001], network autocorrelation does not necessarily imply that an individual influences their peers' behaviors [Hill et al. 2006; Aral et al. 2009]. This problem is exacerbated when the assumed exposure model omits backdoor paths which could plausibly account for the correlations [Shalizi and Thomas 2011]. Even given perfect observability of the network process and abundant behavioral data, latent homophily or confounding factors could drive the assortativity in peer outcomes.

One of the most promising approaches to address these confounds in diffusion studies is the use of randomized field experiments [Aral and Walker 2011; Aral and Walker 2012; Bakshy et al. 2012a; Bakshy et al. 2012b]. However, experiments thus far either compare

¹We use the term *share* because it broadly covers a range of diffusion phenomenon online, but we use it to address any choice that increases the visibility of an individual's adoption decision to others. In some settings, this action can be seen as an endorsement or a form of conspicuous consumption [Veblen 2005].

²Passive sharing, where an action is automatically made visible to an individual's peers, can be seen as the strongest possible encouragement to share, so that all users are equally as likely to share upon adoption.

overall effects via different channels of influence, or focus on the direct effects of social signals on individual behavior through a single mechanism. For example, Aral and Walker [2011] showed that active personalized messaging is more effective in encouraging adoption per message, while passive sharing generates greater total peer adoption in the network. However, it is not clear whether these differences in adoption rates are due to greater persuasiveness of the message format, differences in delivery³, or selection effects. Our work elucidates the mechanism for this selection effect (e.g. selecting peers who are more likely to adopt) by considering effects via single channel of communication and identical information content.

An intriguing aspect of word-of-mouth diffusion is the idea that social networks could be cleverly leveraged to increase the spread of desirable behaviors [Hill et al. 2006]. Much of the research in this area approaches the problem via mathematical models that are analyzed through proofs on a given graph structure or through simulation [Kempe et al. 2003; Aral et al. 2011; Chierichetti et al. 2012]. Influence maximization has thus far been studied through models that do not account network autocorrelation in susceptibility, or make a distinction between adoption rates and the decision to share. Part of the goal of our work is to shed light on how sharing decisions can affect downstream adoption to further the development of such models.

3. THEORY: SELECTION MECHANISMS IN PRODUCT DIFFUSION

Word-of-mouth (WOM) marketing or so-called “viral” diffusion is a repeated process of users adopting products and transmitting information about those products to their peers. While most studies focus on economic aspects of the *adoption decisions* or patterns of diffusion over the network, we examine *propagation decisions* and their selection effects on peer adoption.

3.1. Selection mechanisms

We posit that the conditions under which an adopter decides to share a product has a substantive impact the adoption rate of her peers. To see why, consider sharing to be a binary choice for all adopters and consider two extreme cases: passive (automatic) sharing and active (selective) sharing. In the former case, every adopter shares their adoption decision with all their peers, and sharing can perform no selective role. In the latter, perhaps only a small fraction of adopters share. This may happen, for instance, if sharing bears a cost – such the time to send an email or the initiative to bring up a product in conversation. To isolate selection, one can think of peers of sharers receiving identical signals indicating that their peer has adopted the product, e.g. a structured signal such a Like, a “+1”, or a check-in.

In a world where peers receive homogeneous signals – such as the aforementioned online social signals – differences in adoption rates among peers of adopters in the two sharing regimes must be caused by selection of some combination of sharer-peer-product groups from different sample populations. There are at least three possible selection effects that could alter peer adoption decisions when the sharing decision changes:

- (1) Adopter selection: Adopters who share are more influential than non-sharers.
- (2) Peer selection: Peers of sharers are more interested adopting the product than peers of non-sharers.
- (3) Product selection: Individuals share better products and decline to share worse ones.

All three selection mechanisms could be present if an individual has some pro-social or financial motivation to share products that her peers are more likely to adopt. In this case,

³In the context of the authors’ study, active messages were always received by the alter, whereas passive sharing could be aggregated or filtered, and therefore may not be as salient or seen by the alter.

she may use her private information about her peers' preferences to decide whether or not to share. Product selection may be more salient if users gain utility from portraying an association with prestigious brands or if users receive disutility from creating associations that embarrass them [Akerlof and Kranton 2000].

The difference between adopter and peer selection is subtle and worthy of further discussion. For example, a so-called "influential" may cause many peers to adopt merely because her peers are easily influenced [Watts and Dodds 2007]. It is clearly difficult to define a distinction between adopter and peer selection, and more difficult still to econometrically identify the difference between them.⁴ Accordingly, we will use the term *dyad selection* to refer to selection of adopter-peer pairs where peer adoption is more likely, regardless of whether it is due to characteristics of the adopter, their peer, or the relationship between them. *Product selection* will be used to refer to any selection effects which are observable between products.

3.2. Model

We now develop an economic model that formalizes our hypotheses and the tradeoffs we wish to describe.

3.2.1. Choice model. We index newly adopting individuals by $i = 1, \dots, N$ and products by $k = 1, \dots, K$. We model the probability an individual makes her adoption visible ($s_{ik} = 1$) or not ($s_{ik} = 0$) as a piecewise function depending on whether she is in the active sharing condition, indicated by $z_i = 1$.

$$s_{ik} = \begin{cases} 1 & : z_i = 0 \\ \mathbb{1}(\mu_k + \epsilon_{ik} \geq 0) & : z_i = 1 \end{cases}$$

Here $\mathbb{1}$ is the indicator function and $\mu_k \sim \text{Normal}(\alpha, \sigma_\mu^2)$ is the population average random utility for sharing product k . The unobserved term $\epsilon_{ik} \sim \text{Normal}(0, \sigma_\epsilon^2)$ represents individual-specific factors that increase the utility of i sharing k with her peers. This could include, for example, any utility she expects from sharing based on private information about her peers' preferences.

Individuals in the passive sharing condition live in a simple world: if they adopt, their behavior is visible to their peers. Subjects in the active sharing condition will share if the product itself is "shareable" enough – μ_k is sufficiently high – or if her idiosyncratic utility ϵ_{ik} from sharing is high enough. The first obvious result from this model is that sharing for those in the active sharing condition is always less than in the passive condition. We assume μ_k and ϵ_{ik} are independent random variables. In an empirical context, μ_k would be identified by variation in sharing rates between products.

Let $y_{ik} \in \{0, 1\}$ represent the decision of a peer of i , the subject's peer, to adopt product k after observing i 's behavior. We model the peer adoption decision as the following discrete choice:

$$y_{ik} = \mathbb{1}(\lambda_k + \nu_{ik} \geq 0).$$

In this equation $\lambda_k \sim \text{Normal}(\gamma, \sigma_\lambda^2)$ is the product utility from adoption. $\nu_{ik} \sim \text{Normal}(0, \sigma_\nu^2)$ is the unobserved utility from adoption that i 's peers receive for k . Note that the peer either receives a signal about the product or they do not, they have no information about the individual's treatment status z_i ⁵. Peers are only eligible to make this

⁴One strategy for identifying adopter-specific influence requires repeated observation of different adopters sending persuasive messages to members of the same population of peers. Since each adopter may have different peers, it is difficult to design an experiment where peers are sampled from the same population.

⁵It is important to note that our experimental design isolates a pure selection effect because the messages received by peers are the same regardless of the treatment status of sender.

decision if $s_{ik} = 1$. We will now differentiate between product and dyad selection mechanisms, correlations across different utility components which link an individual's sharing decision with her peers' adoption decisions. These mechanisms will only function if the individual's sharing decision occurs in the active sharing regime.

3.2.2. Selection on products. Sharing can act as a selection mechanism on product quality. Assume that μ_k and λ_k are correlated random variables with correlation coefficient ρ , which would be the case if the latent product quality, e.g. value, provides both sharing and adoption utilities:

$$\begin{bmatrix} \mu_k \\ \lambda_k \end{bmatrix} \sim \text{Normal} \left(\begin{bmatrix} \alpha \\ \gamma \end{bmatrix}, \begin{bmatrix} \sigma_\mu^2 & \rho\sigma_\mu\sigma_\lambda \\ \rho\sigma_\mu\sigma_\lambda & \sigma_\lambda^2 \end{bmatrix} \right).$$

The correlation coefficient ρ may be interpreted as a measure of product selection and we hypothesize that $\rho > 0$: products that are more likely to be shared are also more likely to be adopted, independently of dyadic preferences. The consequence of this hypothesis is that exposed peers decide whether to adopt products with a higher mean utility for adoption. To see how, let ϕ and Φ be the standard normal density and distribution functions respectively, and \mathbb{E} be the expectation operator. When adopters in the active sharing world share, the conditional distribution of λ_k for her peers' adoption decisions has a higher expected value:

$$\mathbb{E}[\lambda_k | \mu_k \geq -\epsilon_{ik}] = \gamma + \sigma_\lambda \rho \frac{\phi\left(-\frac{\epsilon_{ik}}{\sigma_\mu}\right)}{1 - \Phi\left(-\frac{\epsilon_{ik}}{\sigma_\mu}\right)} \geq \gamma. \quad (1)$$

In an empirical context, ρ can be identified by the correlation in sharing and adoption rates between products.

3.2.3. Selection on dyads. An individual who actively shares may be more likely to generate adoptions from her peers than those who do not. This could be because the individual is influential or her peers are susceptible to this particular product (e.g. the individual recommends a product to peers who are likely adopters).⁶

As in the Heckman [1979] model of sample selection bias, we will assume that our unobserved utility components ϵ_{ik} and ν_{ik} are distributed bivariate normal with correlation coefficient ψ :

$$\begin{bmatrix} \epsilon_{ik} \\ \nu_{ik} \end{bmatrix} \sim \text{Normal} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_\epsilon^2 & \psi\sigma_\epsilon\sigma_\nu \\ \psi\sigma_\epsilon\sigma_\nu & \sigma_\nu^2 \end{bmatrix} \right).$$

ψ is our measure of dyad selection and we hypothesize that $\psi > 0$: an individual is more likely to share when her peers are more likely to adopt the product, independent of the product quality. As in the product selection mechanism, the effect is driven by an increase in the mean of the distribution of the peer's idiosyncratic utility conditional on individual choosing to share:

$$\mathbb{E}[\nu_{ik} | \epsilon_{ik} \geq -\mu_k] = \sigma_\nu \psi \frac{\phi\left(-\frac{\mu_k}{\sigma_\nu}\right)}{1 - \Phi\left(-\frac{\mu_k}{\sigma_\nu}\right)} \geq 0. \quad (2)$$

⁶This type of selection mechanism through correlated unobservable variables is a contribution of Heckman [1979], which uses the example of researchers only observing the market wages of individuals who choose to enter the labor market. Here we only observe the adoption decisions of users whose peers chose to make their adoptions visible.

When $\psi > 0$, those who actively share will have peers who are more interested in adopting the product. In other words, passive sharing causes the individual to spread the product even when she knows her peers may not be likely to adopt.

If ρ is identified through repeated observations of sharing and adoption behavior of products, then the correlation in individual specific utilities ψ can be identified by exogenously assigning the individual to active and passive sharing interfaces. By assumption, peers of adopters in the passive condition have $\mathbb{E}[\nu_{ik}] = 0$, while peers of adopters in the active condition have idiosyncratic adoption utilities with conditional expectation, $\mathbb{E}[\nu_{ik} | \epsilon_{ik} > -\mu_k] > 0$. This difference, and therefore ψ , can be identified by variation in adoption rates within products and across peers of adopters in randomly assigned sharing interfaces.

If either of our two selection hypotheses are confirmed (i.e. $\rho > 0$ or $\psi > 0$), peers of adopters who passively share will have a lower probability of adoption than peers of adopters who actively share.

4. EXPERIMENT

We conducted a field experiment on Facebook to compare selection mechanisms present in active and passive sharing regimes using Facebook Offers, a marketing product that allows businesses to share discounts with customers by posting an Offers to their Facebook Pages. Offers are similar to coupons or discounts available through sites like Groupon or LivingSocial. When a user claims (adopts) an Offer, she receives an email which must either be shown at the businesses' physical location to get the discount, or can be used to receive a discount in an online store. Simultaneously, those who share passively⁷ share their claim activity with their peers (friends).

Offers are distributed via Facebook's News Feed. The News Feed is the primary means for users to consume *stories* about friends' activities, such as status updates from friends, or from *Pages*, which represent celebrities, businesses, and other organizations. Thus, there are two ways that a user may receive an Offer. First, the subscribers of a Page receive stories directly from the Page presenting the Offer (Figure 1a). Second, a user may be exposed to the Offer via a friend whose action was made visible after adopting it (Figure 1b,c). These two modes of diffusion correspond to having a single "big seed" (or broadcast node) [Watts et al. 2007] which initially spreads the Offer, after which point cascading effects may occur.

Our empirical context provides several advantages over other settings. First, we can observe the diffusion of a large sample of comparable units, so our analyses do not suffer from survivor bias (i.e. we observe even the unsuccessful cascades). Second, the behaviors we study (claiming Offers) provide users with valuable incentives, are low cost, and are expressly intended by marketers to achieve widespread distribution. Third, many Offers receive substantial distribution and many adoptions, so we can observe many distinct users interact with the same Offer, which is crucial to our identification strategy. Finally, we can plausibly observe almost all interactions between users and Offers because very little Offer transmission occurs outside of Facebook.

4.1. Experimental design

1.2 million users were randomly assigned to one of two experimental conditions – the active or passive sharing conditions – with equal probability at the time of adoption. That is, after subjects claimed an Offer (adopted) on a mobile device, they would either share their Offer redemption passively (Figure 2a), or were given a button that prompted the user to share their claim action with others (Figure 2b). For each Offer, we record an impression event each time a user sees the Offer in their News Feed (Figure 1) and if she claims (adopts) the Offer. We also record whether she shared the Offer after adopting. In the following analyses,

⁷This is the default behavior for many activities on Facebook such as Liking a Page.

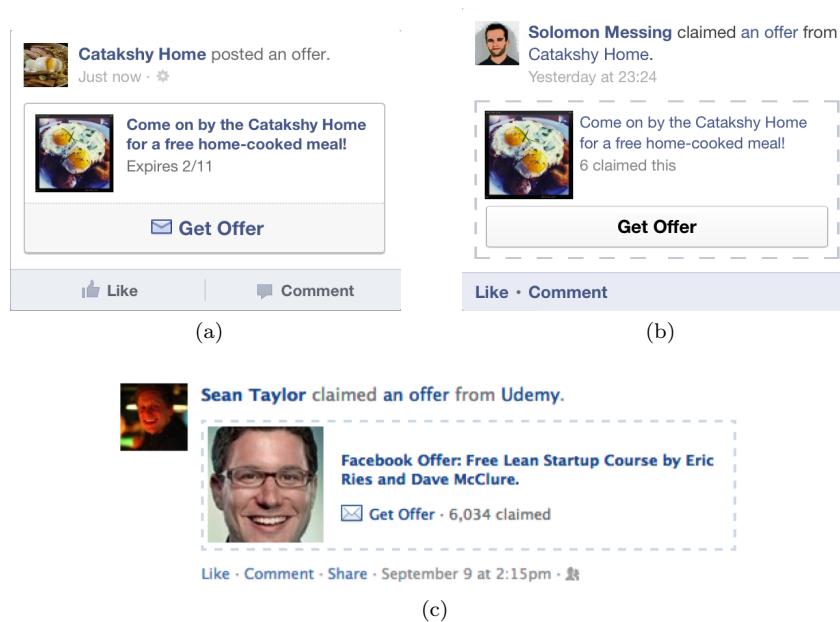


Fig. 1: A story for (a) a Page posting a new Offer on a mobile web browser and (b) a friend claiming (adopting) an Offer on the Facebook iPhone application (c) a friend claiming an Offer on the dekstop interface.

we use this data and consider Offers that were claimed by at least 25 users during a two month period in 2012.

We examine downstream effects of the sharing interface by measuring the subsequent behavior of peers who were exposed to subjects’ adoption activity. It is important to note that peers who see the activity of subjects who share under the passive sharing condition are different than those who share in the active condition. This introduces a selection effect that shapes the population of exposed peers, and this effect is what we intend to measure.

4.2. Interference

The experimental treatment – a change in the mobile sharing interface – is applied to adopters, but we measure the adoption outcomes of their peers. This approach can lead to interference if peers are exposed through multiple adopters (Figure 3), and is problematic for two reasons. First, the status of a peer is no longer well defined if she is exposed by subjects in different conditions. Second, even if a peer is exposed through multiple adopters with the same treatment, she may not be comparable to a user who is exposed through only one. Multiply-exposed individuals may have higher adoption rates due to increased homophily, multiple simultaneous social cues [Bakshy et al. 2012a] or multiple exposures over time [Centola 2010].

Because passive sharing may increase the number of multiply-exposed peers, interference can confound our ability to identify the selection effects we wish to estimate. Therefore our analysis only considers the peers who are exposed via a single adopter’s sharing action⁸.

⁸While we find effects from multiple exposures to be interesting, modeling these processes is beyond the scope of this paper.

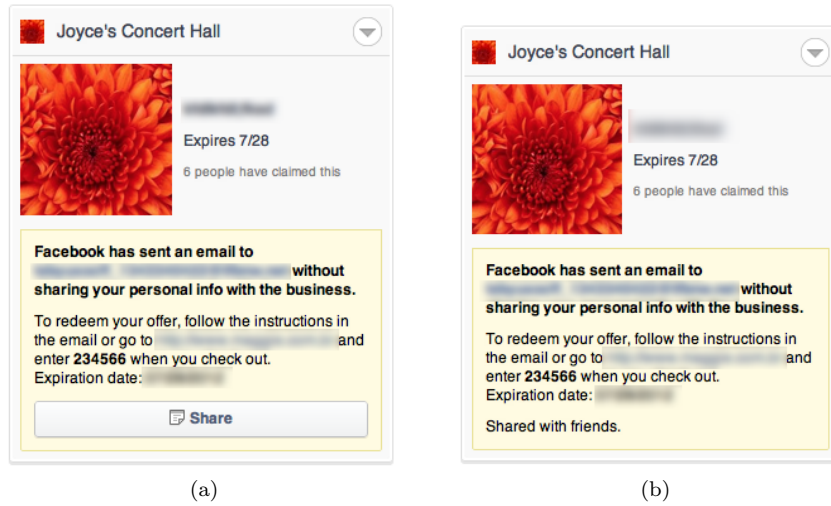


Fig. 2: Mobile interface presented to subjects after adoption for the (a) active and (b) passive sharing conditions.

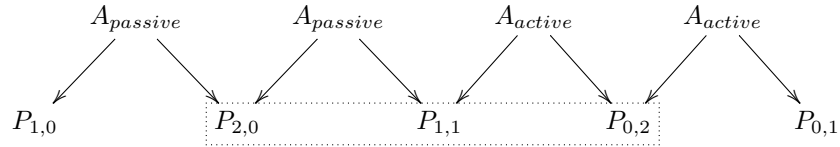


Fig. 3: An illustration of potential interference patterns in our experiment. The subscripts for each exposed peer (P) denotes the number of adopters (A) in the passive and active condition who had exposed them to an Offer. Our analysis only considers peers of type $P_{0,1}$ and $P_{1,0}$, omitting those exposed via more than one adopter (e.g. peers in the dotted box) in order to isolate selection effects. These two types constitute approximately 90% of exposed peers.

This preserves the vast majority (approximately 90%) of exposed peers while simplifying interpretation of the results.

5. RESULTS

We present our results through descriptive analysis and modeling. The first subsection provides a basic overview of the experimental data. In the following two subsections, we present results from reduced-form models which examine subjects' sharing decisions and peers' adoption decisions separately. We focus on separating variation in sharing and adoption outcomes into variation in Offer-specific effects and idiosyncratic user effects. In the fourth subsection, we present estimates from the joint decision model introduced in Section 3.2 to link the two models in a coherent system which can identify the correlation parameters we are interested in, allowing us to distinguish between product and dyad selection effects.

5.1. Descriptive statistics

Table I shows summary results from the direct effect of the experiment. Approximately the same number of subjects were exposed to each sharing interface, and subjects in each condition were exposed via approximately the same number of distinct Offers. While all

	Active	Passive
Subjects	577,933	573,113
Distinct Offers	23,102	23,251
Proportion shared	0.23	1.00

Table I: Summary of statistics for direct effects on subjects’ sharing behavior in the active and passive sharing conditions.

	Active	Passive
Mean friends exposed	59.17	66.62
Median friends exposed	41	46
Number of adoptions	20,591	87,686
Adoption rate	0.0050	0.0045
Adoptions per subject	0.036	0.153
Adoptions per sharer	0.157	0.154

Table II: Summary statistics for subjects’ exposed peers in the active and passive sharing conditions.

users in the passive condition shared, approximately one in five subjects in the active sharing condition shared the Offer with their peers.

After a subject shares an Offer, a story showing that the subject claimed the Offer was eligible to appear in her peers’ News Feeds. Table II provides descriptive statistics about how many peers were exposed to this story, as well as their subsequent adoption decisions.⁹ The mean and median number of exposed peers is slightly higher for sharing subjects in the passive sharing condition compared to those in the active condition. Figure 4 shows the distribution of the number of exposed peers by treatment condition. Here, we can see that active sharing shifts the distribution toward individuals who expose fewer friends to Offers. This effect is likely caused by selection on users who have fewer or less active peers. The result is fewer social exposures to the Offers from both less sharing as well as smaller number of exposures per sharing user.

Peers who are reached via active sharing are more responsive on average with about a 10% increase in the probability of adoption¹⁰ (95% confidence interval [1.063, 1.134]). However, the low sharing rate for subjects in the active condition means that it is about 4.3 times more effective to enable passive sharing as measured by aggregate peer adoptions.

Figure 5 provides an intuition for one of the mechanisms underlying selection effects in sharing decisions. Adoption rate of peers varies according to whether the alter was exposed as a member of a large group of exposed peers of the individual or a small group. Furthermore, the variability in adoption rate is greater for those who expose fewer peers.

5.2. Modeling variation in sharing behaviors

We first report the results for those in the active sharing condition. The share rate is approximately 23% and is very precisely estimated. We are interested in the extent to which Offer-level effects are driving sharing decisions by users. If there is little variability between Offers and most of the variation occurs at the dyad level, then it will not be possible

⁹Recall from Section 4.2 that we only consider the subpopulation of exposed peers who were exposed the Offer via a single friend.

¹⁰All confidence intervals reported in this section use the multiway bootstrap [Owen and Eckles 2012] with 500 replicates clustered by subjects and Offers. This bootstrap is expected to be accurate even in situations where treatment effects are expected to vary with both subjects and items [Bakshy and Eckles 2013].

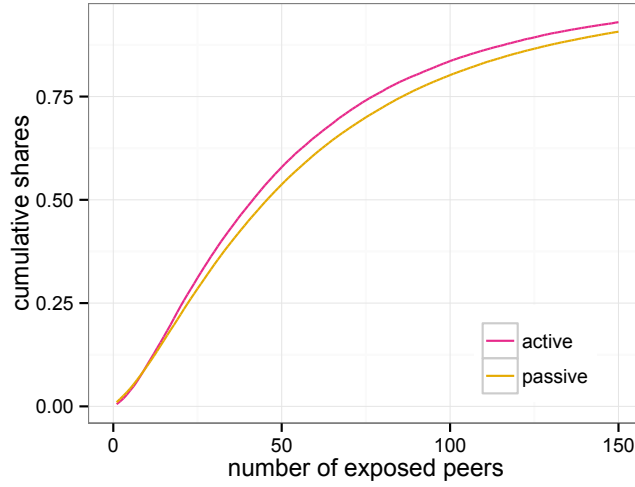


Fig. 4: The empirical cumulative distribution function for the number of exposed peers for each sharing subject by treatment condition. For clarity in comparison, the x-axis is truncated at the 90% percentile of the distribution. The empirical distributions show sharers in the passive condition usually expose more of their peers.

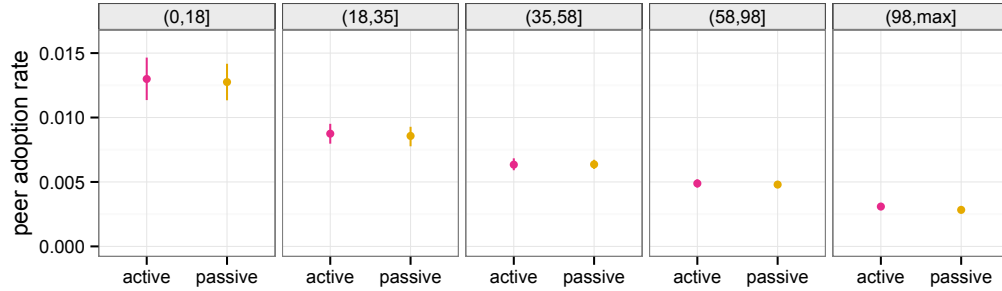


Fig. 5: Average adoption rate of sharing subjects' peers, broken down by quintiles of the total number of exposed peers. Error bars show the 95% confidence intervals.

for the product selection mechanism to function. To see why, assume there is no Offer-level variation in share rates. Then Offer characteristics do not affect sharing at all and the set of shared Offers will be sampled from the same distribution as the unshared Offers.

We estimate a random effects probit model,

$$Pr(s_{ik} = 1 | \mu_k) = \Phi(\mu_k + \epsilon_{ik}),$$

in order to estimate σ_μ^2 . To identify the parameters in the probit model, we let $\sigma_\epsilon^2 = 1$. Table III contains the model parameter estimates we obtain. The estimated intraclass correlation coefficient is $\frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\epsilon^2} = 0.064$, indicating that the Offer-level random effects do not explain much additional variance in the sharing model. This implies that the product selection mechanism is likely to be weak.

Parameter	Estimate
α	-0.812* (0.005)
σ_μ	0.262
Log-likelihood	-303,070
Groups	23,102
N	577,933

Table III: Maximum likelihood parameter estimates for probit regression predicting share rate with random effects at the Offer level. The estimated mean of μ_k is α , which is an estimate of average sharing utility. The variance of the random effects at the Offer level μ_k is small compared to the total variance. * denotes significance at the 0.001 level.

5.3. Effect of passive sharing on downstream adoption

In this section, we estimate an average treatment effect of active sharing on peer adoption rates. For each subject who shares – either because they were in the passive sharing condition or they chose to share in the active sharing condition – we measure two aggregate outcomes of the subjects’ peers: exposures and adoptions.

We define the number of peer exposures n_{ik} for user i and Offer k to be the number of unique peers who saw a story in News Feed about the subject claiming the Offer. We only count exposures which were unique, meaning that the alter must not have seen the Offer through any other user’s adoption. We count a peer as exposed just once regardless of how many impressions of claim story the user may have been served in her News Feed.

We define number of peer adoptions, $a_{ik} = \sum_{j=1}^{n_{ik}} y_{ijk}$, as a count of the number of peer exposures which generated an adoption. We assume that n_{ik} is exogenous, since it depends on the subject and her peers’ characteristics and Facebook usage behavior. Recall that z_i represents the exogenous (experimental) manipulation of the subject’s sharing interface and is equal to 1 in for users in the active sharing condition.

$$Pr(a_{ik} = L | n_{ik}) = \text{Binomial}(n_{ik}, \Phi(\beta z_i + \lambda_k + \nu_{ik})),$$

where β represents the average selection effect on the subject’s peers. As in the last section, we ignore the correlations between the unobserved parameters. We report parameter estimates for the regression model in Table IV. The coefficient β , measures all selection effects, is positive and significant, and therefore confirms our hypothesis that active sharing will increase the probability that an subject’s peers will adopt the product. The magnitude of β corresponds to about a 7% marginal increase in the relative risk of adoption for peers of users who share in the passive condition (95% confidence interval: [1.050, 1.089]).

As in the sharing model, we have assumed $\sigma_\nu^2 = 1$ in order to identify the other parameters. We can compute the intraclass correlation coefficient for adoption, $\frac{\sigma_\lambda^2}{\sigma_\lambda^2 + \sigma_\nu^2} = 0.029$. This is low, indicating that product quality does not explain a large amount of the variance in adoption outcomes.

5.4. Joint decision model

The structural model we specified in Section 3.2 unifies the regression models in the previous two sections by accounting for the correlations between the unobserved parameters. We assumed a correlation structure which accommodates two mechanisms for individual’s sharing decision to impact her peers’ adoption decisions, obviating the need for the β parameter in the adoption model. Estimating the joint model allows us to understand the relative contribution of each of the selection effects.

Parameter	Estimate
γ	-2.742* (0.004)
β	0.022* (0.003)
σ_λ	0.172
Log-likelihood	-151,521
Groups	25,726
N	702,090

Table IV: Maximum likelihood parameter estimates for binomial regression predicting the number of adopting alters with random effects at the Offer level. γ is the mean of the random effect λ_k , while β represents a reduced form measure of total selection effect. * denotes significance at the 0.001 level.

Our setup is similar to the simultaneous discrete-choice models with interdependent preferences considered in Yang et al. [2006], motivating a similar estimation procedure using Bayesian methods. Bayesian estimation is ideal for this setting because it allows us to flexibly perform inference on the correlation parameters ρ and ψ at the cost of parametric assumptions.

We use non-informative priors on all parameters and run Markov-chain Monte Carlo simulations to estimate their posterior distributions given the observed data. Due to the scale of our data, we used an efficient Hamiltonian Monte Carlo sampler [Hoffman and Gelman 2012] and computed our results using a state-of-the-art Bayesian model compiler [Stan Development Team 2013]. We simulated three Markov chains for 2,000 iterations, discarding the first 1,000 iterations for “burn-in.” We then used the last 1,000 draws for estimation. We evaluated convergence by computing a potential scale reduction factor for each estimated parameter in the model [Gelman and Rubin 1992].

Parameter		Mean	2.5%	Median	97.5%
Mean product sharing utility	α	-0.813	-0.821	-0.813	-0.804
Mean product adoption utility	γ	-2.739	-2.748	-2.739	-2.730
Std. dev. of product sharing utility	σ_μ	0.267	0.258	0.267	0.279
Std. dev. of product adoption utility	σ_λ	0.172	0.165	0.173	0.179
Product-level correlation coefficient	ρ	0.174	0.102	0.176	0.238
Dyad-selection correlation coefficient	ψ	0.025	0.019	0.025	0.032

Table V: HMC posterior mean, median, and 95% credible interval estimates for the parameters of the joint structural model of sharing and adoption specified in Section 3.2. Estimates for ρ and ψ are positive and significant, providing evidence for both product- and dyad-selection effects.

We estimate three main types of parameters in the model (Table V). First, there are mean sharing and adoption utilities, α and γ , which rationalize the average rates of sharing and adoption. Second, there are correlations between the unobserved utilities at the product level, ρ , and at the dyad level, ψ . Third, we estimate the standard deviations of the zero-mean product-level utilities, σ_μ and σ_λ .

Note that as in the regression models, σ_ϵ and σ_ν are fixed at 1 in order to identify the other parameters of the model. This is a typical assumption in this modeling situation where we have no absolute utility scale.

We find evidence for both types of selection effects that we hypothesized, $\rho > 0$ and $\psi > 0$, indicating that in the active sharing interface, users who shared Offers which were more likely to be adopted and seen by peers who were more likely to adopt them.¹¹

It is worth pointing out that our estimates of α , γ , and the variances of the random effects σ_μ^2 and σ_λ^2 are extremely close to the reduced-form models of the previous section. We have essentially replaced β with structural correlation parameters which allow us to distinguish between two mechanisms. But which of these mechanisms is more important?

Our estimate for ρ is substantially larger than our estimate for ψ , which could be interpreted as evidence for the relative importance of product selection effects. However, we must consider that the distributions of the Offer-level and dyad-level effects are different in scale. This warrants further analysis of the interaction between correlations and effect scales.

With posterior distributions for parameters in hand, we can use our model to decompose the treatment effect into product and peer selection through simulation. Recall the relative risk of peer adoption for active versus passive sharing had 95% confidence interval [1.063, 1.134], which is measured by

$$\text{RR} = \frac{\text{Pr}(y_{ik} = 1 | z_i = 1, s_{ik} = 1)}{\text{Pr}(y_{ik} = 1 | z_i = 0)}.$$

The effect of z_i works through two exhaustive mechanisms. First it changes ρ from 0 to our estimate $\hat{\rho} > 0$, enabling selection on product quality. Second, it changes ψ from 0 to our estimate $\hat{\psi} > 0$, enabling selection on dyads. We can simulate relative risk under counterfactual scenarios where only one of the two mechanisms is enabled:

$$\text{RR}_{\text{product}} = \frac{\text{Pr}(y_{ik} = 1 | s_{ik} = 1, \rho = \hat{\rho}, \psi = 0)}{\text{Pr}(y_{ik} = 1 | s_{ik} = 0)};$$

$$\text{RR}_{\text{dyad}} = \frac{\text{Pr}(y_{ik} = 1 | \rho = 0, \psi = \hat{\psi})}{\text{Pr}(y_{ik} = 1 | s_{ik} = 0)}.$$

To compute these counterfactual relative risks, we simulate sharing behavior and subsequent adoption rates by drawing from our posterior parameter distributions. For $\text{RR}_{\text{product}}$ we set $\psi = 0$ and then draw $(\epsilon_{ik}, \nu_{ik})$ pairs as independent random variables. For RR_{dyad} we set $\rho = 0$ and then draw independent (μ_k, λ_k) pairs. We then compute empirical relative risks over 500 generated sample populations and compute means and 95% confidence intervals for selection effects under each counterfactual scenario.

The results of this procedure are shown in Figure 6. We can see that disabling the product selection, leaving dyad selection only, retains most of the total selection effect in our simulations. In comparison, the product selection effect is weaker. Thus, despite the high correlation in Offer sharing and adoption utilities, their relatively low importance in the explaining the variance of these decisions limits the product selection effect.

6. DISCUSSION

We have presented a theoretical result and supporting evidence that encouraging so-called “virality” decreases the efficiency of marketing messages in social networks. Our study is the first to identify the interaction between adoption and propagation decisions. This relationship is important because peers of users who choose to share, and the products they share, are potentially different than the peers of users and products shared by the general

¹¹Recall from Section 3.2.3 that one cannot distinguish between adopter and peer selection mechanisms in our setting. Part of the positive correlation ψ may be explained by influential users’ higher propensity to share.

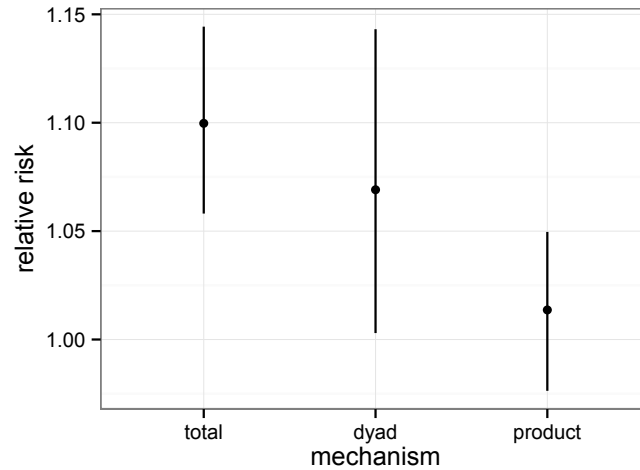


Fig. 6: Simulated relative risks of adoption for active versus passive sharing. 500 iterations were used, each with a set of parameters drawn from the estimated posterior distributions. Dyad selection comprises the bulk of the selection effect in most cases and the mechanisms are complementary. The confidence intervals on the total relative risk are larger than those we report earlier because they incorporate model uncertainty.

population of adopters. Our results suggest that the decision to share enhances efficiency of diffusion by increasing the probability of adoption for downstream users. Thus when users can choose to share, there are fewer wasted exposures generated in the diffusion process.

From a design perspective, our results show that while encouraging users to share their behaviors may increase the total number of adoptions, it can have negative consequences. There exists a tradeoff for platform providers for whom distribution is a scarce resource or brands using costly incentive strategies to improve rates of peer exposure.

In our experimental setup, either active or passive sharing distributed adoption stories through an automated content ranking system, exposing a potentially large audience to identical messages. In other settings, the audience and message resulting from an adopter’s sharing decision may be more variable. Adopters may decide how many peers they share with, with whom they share, and what exactly they choose to say when they share. It is possible that giving adopters tighter control over the outcome from sharing could yield stronger selection effects than we observed, resulting in higher adoption rates.

Our parameter estimates also seem to suggest a potential explanation for why campaigns rarely “go viral” [Bakshy et al. 2011; Goel et al. 2012]. In order to propagate through a network, a product must be adopted and shared at a high rate. In Facebook Offers, we found that product-level factors which predict adoption and sharing are only mildly correlated and explain only a small fraction of the variance in spreading behaviors. It may simply be rare to find examples of products which contribute high levels of sharing and adoption utility to all users.

6.1. Limitations and Future Work

While we are able to plausibly observe users’ interactions with respect to sharing and claiming Offers comprehensively, we are bound to investigate selection effects that occur via plausible changes to Facebook’s existing delivery mechanisms. For Facebook users, sharing means publishing content to a specified audience – often friends – so that the content appears in friends’ News Feeds. Like face-to-face situations, the likelihood of receiving information

via the Feed is determined by an individual's preferences and previous interactions. Thus, it is possible that Facebook's feed ranking algorithm automatically plays some selective role in the diffusion of Offers. It is possible that other platforms, especially those which do not use ranking, to exhibit even stronger sharing selection effects.

We used a randomized field experiment to estimate selection effects in the sharing process for single individuals and their peers. Since selection effects may compound over several steps of the diffusion process, it is possible that individual-level effects may differ for subjects had the experiment randomized over Offers. Furthermore, passive sharing is likely to increase the number of social signals an individual receives. As discussed in Section 4.2, multiply-exposed peers may behave differently, and we might also expect that overall macro effects would be different under other randomization schemes. While our experiment is designed specifically to estimate peer effects not caused by multiple reinforcing signals, examining how these different effects interact would be of substantial interest from a policy perspective.

There are also a number of other opportunities for further exploration. In our setting, adoption and sharing decisions were relatively costless for the subjects, requiring only a single touch on a mobile phone. It would be interesting to see if these results apply for more costly settings where adoption comes at some expense. Other types of encouragements to sharing could be explored, such as monetary incentives, which could generate smoother variation in the rate of sharing. Finally, other peer outcomes, such as using the Offer in brick-and-mortar stores, are also of great interest.

7. ACKNOWLEDGEMENTS

We would like to thank Dean Eckles, Rohit Dhawan, and Cameron Marlow for their feedback on this work.

REFERENCES

- AKERLOF, G. A. AND KRANTON, R. E. 2000. Economics and identity. *The Quarterly Journal of Economics* 115, 3, 715–753.
- ARAL, S., MUCHNIK, L., AND SUNDARARAJAN, A. 2009. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences* 106, 51, 21544.
- ARAL, S., MUCHNIK, L., AND SUNDARARAJAN, A. 2011. Engineering social contagions: Optimal network seeding and incentive strategies. In *available at SSRN: <http://ssrn.com/abstract>*. Vol. 1770982.
- ARAL, S. AND WALKER, D. 2011. Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Science* 57, 9, 1623–1639.
- ARAL, S. AND WALKER, D. 2012. Identifying influential and susceptible members of social networks. *Science* 337, 6092, 337–341.
- BAKSHY, E. AND ECKLES, D. 2013. Uncertainty in online experiments with dependent data: An evaluation of bootstrap methods. *Arxiv preprint*. <http://arxiv.org/abs/1304.7406>.
- BAKSHY, E., ECKLES, D., YAN, R., AND ROSENN, I. 2012a. Social influence in social advertising: Evidence from field experiments. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 146–161.
- BAKSHY, E., HOFMAN, J. M., MASON, W. A., AND WATTS, D. J. 2011. Everyone's an influencer: Quantifying influence on Twitter. In *Proceedings of the fourth ACM international conference on Web search and data mining*. WSDM '11. ACM, New York, NY, USA, 65–74.
- BAKSHY, E., ROSENN, I., MARLOW, C., AND ADAMIC, L. 2012b. The role of social networks in information diffusion. In *Proceedings of the 21st international conference on World Wide Web*. WWW '12. ACM, New York, NY, USA, 519–528.
- BASS, F. M. 1969. A new product growth for model consumer durables. *Management Science* 15, 5, 215–227.
- CENTOLA, D. 2010. The spread of behavior in an online social network experiment. *Science* 329, 5996, 1194–1197.
- CHIERICHETTI, F., KLEINBERG, J., AND PANCONESI, A. 2012. How to schedule a cascade in an arbitrary graph. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 355–368.

Proceedings Article

- GELMAN, A. AND RUBIN, D. 1992. Inference from iterative simulation using multiple sequences. *Statistical science* 7, 4, 457–472.
- GOEL, S., WATTS, D. J., AND GOLDSTEIN, D. G. 2012. The structure of online diffusion networks. In *Proceedings of the 13th ACM Conference on Electronic Commerce. EC '12*. ACM, New York, NY, USA, 623–638.
- GRANOVETTER, M. 1978. Threshold models of collective behavior. *The American Journal of Sociology* 83, 6, 1420–1443.
- HECKMAN, J. 1979. Sample selection bias as a specification error. *Econometrica: Journal of the econometric society*, 153–161.
- HILL, S., PROVOST, F., AND VOLINSKY, C. 2006. Network-based marketing: Identifying likely adopters via consumer networks. *Statistical Science* 21, 2, 256–276.
- HOFFMAN, M. D. AND GELMAN, A. 2012. The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*.
- JACKSON, M. AND YARIV, L. 2007. Diffusion of behavior and equilibrium properties in network games. *The American Economic Review* 97, 2, 92–98.
- KEMPE, D., KLEINBERG, J., AND TARDOS, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 137–146.
- LESKOVEC, J., ADAMIC, L., AND HUBERMAN, B. 2007. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)* 1, 1, 5.
- MCPHERSON, M., SMITH-LOVIN, L., AND COOK, J. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 415–444.
- OWEN, A. B. AND ECKLES, D. 2012. Bootstrapping data arrays of arbitrary order. *The Annals of Applied Statistics* 6, 3, 895–927.
- SCHELLING, T. 1973. Hockey helmets, concealed weapons, and daylight saving: A study of binary choices with externalities. *The Journal of Conflict Resolution* 17, 3, 381–428.
- SHALIZI, C. R. AND THOMAS, A. C. 2011. Homophily and contagion are generically confounded in observational social network studies. *Sociological Methods and Research* 27, 211–239.
- STAN DEVELOPMENT TEAM. 2013. Stan: A c++ library for probability and sampling, version 1.1.
- VEBLEN, T. 2005. *The theory of the leisure class; an economic study of institutions*. Aakar Books.
- WATTS, D. AND DODDS, P. 2007. Influentials, networks, and public opinion formation. *Journal of consumer research* 34, 4, 441–458.
- WATTS, D. J., PERETTI, J., AND FRUMIN, M. 2007. *Viral marketing for the real world*. Harvard Business School Pub.
- WEI, X., YANG, J., ADAMIC, L., DE ARAÚJO, R., AND REKHI, M. 2010. Diffusion dynamics of games on online social networks. In *Proceedings of the 3rd conference on Online social networks*. USENIX Association, 2–2.
- YANG, S., NARAYAN, V., AND ASSAEL, H. 2006. Estimating the Interdependence of Television Program Viewership Between Spouses: A Bayesian Simultaneous Equation Model. *Marketing Science* 25, 4, 336.