Free Riding in Products with Positive Network Externalities: Empirical Evidence from a Large Mobile Network

Abstract

We study the effect of peer influence on products that exhibit positive network externalities to non-adopters, i.e., products that benefit adopters' friends even if they do not adopt. Contrary to products that exhibit positive network externalities upon adoption, this structure of incentives likely results in negative peer influence: the more friends that adopted the product, the smaller the incentives to adopt. We measure this effect empirically by using observational data from a large mobile carrier serving 5.7 million users. We estimate the effect of peer influence across five different products of this type. A naive approach to do so results in a positive estimate for peer influence, due to unobserved homophily. We follow two approaches to address this issue. First, we suggest using the number of friends that end up adopting the product as a proxy for unobserved user fixed effects. Second, we control for homophily by applying a shuffle test, i.e., we compare the effect of peer influence from the original data with the effect obtained from comparable randomly generated data without peer influence. We obtain negative estimates from both approaches, which provides robustness to our findings. Finally, we show that even for these products, the effect of peer influence associated with the first friends that adopt the product is positive, which arises because they still convey useful information about reducing uncertainty. The negative effect of peer influence arises only for the subsequent friends that adopt the product. These friends are unlikely to convey new information about the product, but each of them decreases the economic incentive to adopt, resulting in a negative aggregate effect of peer influence.

1 Introduction

Measuring and understanding peer influence has been a topic of significant research in information systems, economics, management, and marketing. Recent studies have focused on how the characteristics of individuals and products, as well as how the structure of the social network, shape the dynamics of diffusion (e.g. Aral and Walker, 2012, 2014, 2011). An essential aspect of the diffusion process is whether the product exhibits network externalities, i.e., whether adopting the product affects friends' utility.[1] Network externalities are common in Information Technology (IT) products, usually involving communication and collaboration among multiple users. In this paper, we study products that exhibit positive network externalities, i.e., products for which the focal user's utility increases when more of her friends adopt the product. We split these products into two categories: i) products of type P – those for which the utility of the focal user increases when more of her friends adopt *only after the focal user adopts*; ii) products of type N – those for which the utility of the focal user increases when more of her friends adopt *even if the focal user does not adopt*.

The canonical historical example of a product of type P is the fax machine. Having friends with a fax machine increases the likelihood that a user knows about it and thus increases the likelihood of purchasing a fax machine. Having an additional friend that buys a fax machine also increases the economic incentive to adopt a fax machine because, upon adoption, the user will have an additional person to whom communicate. The literature on network externalities and peer influence has mostly focused on products of type P, but products of type N also abound in the economy, particularly in IT settings. Overleaf — an online collaboration platform allowing users to work together on LaTeX documents — is an example of a product of type N. Overleaf offers a set of features for everyone but requires at least one of the document's collaborators to subscribe to the premium version so that all collaborators can use the premium features. The premium version offers access to features such as track changes and unlimited access to the document's

[1] To simplify the exposition, we use the term "friend" throughout the paper to refer to users connected in a network setting. Empirically, we use multiple definitions for friend, as detailed in section 3.

history. In this setting, a user who collaborates in multiple documents and can use the premium features in all of them because her co-authors have already subscribed Overleaf Premium, has little incentive to purchase the premium version herself because she already benefits from the Premium features without paying for them. Other examples of IT-related products with this characteristic include Skype Premium — which allowed non-premium users to benefit from free video calls with premium users — and LinkedIn — which allows (non-premium) job seekers to be contacted by (premium) recruiters not in their network vicinity, effectively reducing their incentive to purchase LinkedIn Premium.

In this paper, we measure the effect of peer influence in the diffusion of products of type P and type N in the context of mobile communications. More specifically, we study products that allow users unlimited calling for a specified period (e.g., a full month) in exchange for a small fixed upfront fee. Products of type N allow users unlimited calling to *any other user with service from the same telecom carrier*. In contrast, products of type P allow users unlimited calling *only to users with service from the same carrier that also adopt product P*. Despite their similarities, these products provide different incentives to potential adopters, likely to result in distinct diffusion processes. In particular, a user that adopts a product of type N is expected to call her friends more often, potentially leading the latter to call less and thus pay less for calls. Therefore, the friends of the user that adopts a product of type N benefit from it, although they did not buy the product themselves. In other words, the product of type N yields positive network externalities to non-adopters. In this paper, we ask *how does peer influence affect the diffusion of products that exhibit positive network externalities to non-adopters?* and more specifically, *which mechanisms contribute to the observed effect of peer influence in these cases?*

We use the peer effects model to measure the effect empirically of peer influence in our setting (Leenders, 2002). In this model, a user's adoption probability depends on how many of her friends have already adopted the product — a covariate usually called exposure– and user-level characteristics and time dummies. The coefficient on exposure is called the effect of peer influence and measures how an additional friend who adopts the product changes the focal user's likelihood of adoption. A naive approach to estimate peer influence with unobserved homophily results in a positive measurement for both products of type P and N. We follow two strategies to address this issue. First, we suggest using the number of friends that end up adopting the product as a proxy for unobserved user fixed effects. Second, we control for homophily by applying a shuffle test, i.e., we compare the effect of peer influence from the original data with the effect obtained from comparable randomly generated data without peer influence. We find that launching products as type N often contributes to decreasing adoption compared to products that exhibit no peer influence. More precisely, having friends who have adopted a product of type N contributes to reducing the likelihood of adopting the same product. When products exhibit positive network externalities to non-adopters, i.e., when users benefit from their friends' adoption even before they adopt, they do not need to adopt the product to already benefit from it. Empirically, this arises as a negative estimate for the effect of exposure in the peer effects model. However, no prior literature has discussed or documented a negative effect of peer influence. We set out to do so and thus measure to what extent launching products with positive network externalities to non-adopters reduce adoption.

We also show that this negative effect does not arise because of unfavorable word of mouth. We find evidence consistent with a positive effect of word of mouth for both types of products. Namely — and for both types of products — we observe that the first friends that adopt the product contribute positively to the likelihood of adoption by the focal user, which provides evidence of the *informational mechanism* at play, i.e., friends that adopt a product convey useful information about it. This information lowers the uncertainty of the focal user, increasing her likelihood to adopt the product. This effect is expected to be stronger for first friends that adopt the product, as they are more likely to convey new information about the product than later ones. When friends adopt the product, they may also change the economic incentives associated with the focal user's adoption. For a product of type P — just like the fax machine's case — an additional friend who adopts the product increases the economic incentive to adopt because, upon adoption, the focal user will have one more person to communicate. However, for a product of type N, an additional friend that adopts the product immediately lowers the economic incentive to adopt and more so the more friends that adopt because non-adopters can readily benefit from friends' adoption. This effect is akin to the idea of free-riding (Samuelson, 1954; Musgrave, 1959), in which one user benefits from the actions of other users to the point that she does not undertake the action herself. In our case, a user with many friends who have adopted a product of type N may not adopt altogether.

We find empirical evidence of this *economic mechanism* at play in our setting when we observe a positive effect of peer influence associated with the first friends that adopt a product of type N and a negative one associated with the subsequent friends that adopt the product. Thus, our empirical evidence is consistent with the idea that peer influence embodies informational and economic mechanisms that act in opposite directions as more friends adopt a product of type N. Hence, an overall negative effect of peer influence arises for products of type N when the economic mechanism dominates.

It is also well-known that identifying the effect of peer influence is empirically challenging, given that correlation in outcomes is not enough to claim causality. Namely, such correlation may arise due to the confounding effects of unobserved covariates, telling us little about whether friends' adoption causes one to adopt (Shalizi and Thomas, 2011). Our paper pursues a new approach to measure peer influence in observational studies with panel data. Namely, we propose to use final exposure (i.e., the number of friends that end up adopting the product) as a proxy for homophily, which helps researchers avoid overestimating the effect of peer influence. We explain in detail the rationale behind why this proxy works well for this purpose. This approach is an alternative to using the traditional fixed-effects model, which cannot be employed with panel data on adoption because users leave the panel once they adopt the product. We use this approach in our empirical context and provide results from multiple simulations showing that it works well across a wide range of homophily levels. As a robustness check, we also show that this approach behaves as expected in the case of products with positive peer influence, thus ruling out that it could only recover instances with negative effects. As a second robustness check, we follow Nyblom et al. (2003) and Anagnostopoulos et al. (2008) and apply randomization, which confirms the results obtained using the prior approach. Indeed, type-N products exhibit negative peer influence, and the results that we obtain using final exposure to proxy homophily are not an artifact of method or context. We also show that using randomization with the peer effects model provides a lower bound, in magnitude, for the effect of peer influence, which is an advantage of this method vis-a-vis the ones previously used in the literature that overestimated this effect (e.g., propensity score matching).

Prior literature in information systems, marketing, and management has shown at length that launching new products is likely to attract new consumers and retain existing ones. As a consequence, firms keep dynamically and strategically adjusting their portfolio of offers to compete. Our paper complements this research line by showing that the way products are launched has significant implications for how adoption propagates. Firms may increasingly make products available to consumers in ways that allow them to benefit from friends' adoption even if they do not adopt.

Such instances arise, for example, in collaboration platforms, communication ecosystems, and matching platforms (such as job markets, dating websites, and ride-hailing applications). Firms may strategically choose to do so to attract (or retain) even more consumers. Our paper shows that this also reduces adoption. Thus such strategies must be carefully considered to avoid backfiring because how many consumers end up adopting is crucial for brand awareness, which is, in turn, fundamental for business sustainability.

The remainder of this paper is organized as follows. Section 2 discusses the relevant related work, section 3 describes our empirical context and dataset, section 4 introduces our empirical strategy, and section 5 shows the results that we obtain. Finally, section 6 concludes.

2 Related Work

Our paper draws from the literature on peer influence and product diffusion. Peer influence is defined as the degree by which an action from a user changes the behavior of someone else (Peres et al., 2010), or as the dyadic process by which a user shapes her behavior, beliefs, or attitudes according to what the other users in the social system think, express, or how they behave (Leenders, 2002). These definitions encompass all types of interactions among users, including direct communication, i.e., word-of-mouth, and indirect communication, such as social signals (e.g., potential adopters observe how adopters behave) and network externalities (adopters affect the utility of their peers when they adopt).

Interactions among users in the same social system conceal two broad types of mechanisms that may lead to one's behavior to affect her friends' behavior. *Informational mechanisms* pertain to the transmission of information about a new product, which reduces uncertainty, leading to positive peer influence and increased adoption. *Economic mechanisms* relate only to the economic incentives associated with one's adoption decision given the information that the user has about who adopted and how such adoption affects the utility that she can obtain from adopting and consuming the product. Both types of mechanisms may be triggered by direct, explicit communication, such as word of mouth and other communication forms, where no explicit message is exchanged among friends. For example, a user may adopt a product and change her behavior because of such adoption. This change in behavior may provide a signal to her friends, who may, in turn, change their behavior. Take our empirical setting of mobile communications as an example. A user that purchases a tariff plan that allows unlimited calling to everyone with service from the same telecom carrier is likely to start calling her friends more often. This change in behavior may affect her friends' behavior regardless of whether she tells her friends about this tariff plan and about her decision to purchase it.

The existing models of peer influence have not operationalized the effect of the economic incentives described above. For example, classical peer influence models such as the Susceptible-Infectious-Recovered model (e.g., Kermack and McKendrick, 1927), the Bass model (e.g., Bass, 1969), threshold models (e.g., Granovetter, 1978; Centola and Macy, 2007) and hub models (e.g., Katz and Paul, 1955), offer predictions that explain well the observed aggregate patterns of diffusion over time. However, they are not user-level models and do not specify the exact mechanisms at play. These models focus essentially on the diffusion of information across the social system and disregard the effect of the economic incentives on adoption. Our work complements this research line by explicitly showing how both informational and economic mechanisms shape the diffusion curve.

Furthermore, all prior empirical studies that focus on the role of the informational mechanism on diffusion assume and find that peer influence never decreases adoption. Examples span many different fields, such as online communities (Aral and Walker, 2011), offline settings (Mobius et al., 2015), and also mobile handset diffusion (Godinho de Matos et al., 2014), which is closer to our setting. This practice arises from the fact that all these studies fail to acknowledge the effect of the economic incentives described above. However, and as we show in this paper, the effect of these incentives can be significantly negative, flipping the sign of the average effect of peer influence. To the best of our knowledge, our paper is the first in the literature to consider and measure a negative effect of peer influence, i.e., the fact that the likelihood of one's adoption may reduce when more friends adopt. These negative measurements arise in our empirical setting in the case of products of type N. Prior literature never obtained such negative estimates in part because it has never used the peer effects model to study this type of product in detail before.

There are a few theoretical studies on peer influence that have focused on how economic incentives affect diffusion. For example, the literature in network games has explored scenarios in which outcomes depend directly and significantly on user payoffs and network structure (Galeotti et al., 2010; Jackson and Zenou, 2012). Interestingly, Goldenberg et al. (2010) argue that some local network effects may slow down adoption because adopters tend to wait for their early-adopter friends to adopt to get more utility from when they eventually adopt. In our setting, this would be similar to having users postpone adopting a product of type P for when enough of their friends have adopted making adoption economically beneficial. These authors use simulation to show how this "chilling effect" emerges. Our paper looks at the diffusion of products of type N and thus explores further the idea that some economic incentive structures may indeed lead to a negative effect of peer influence, which hinders the diffusion process. In this context, free riding may arise as a limiting case of this "chilling effect", in which late adopters may end up not adopting altogether, decisively reducing aggregate diffusion. Yet, to date, free riding has been essentially studied in the context of public goods (such as national defense, public health, and transport infrastructures (Varian, 1984; Atkinson and Stiglitz, 1980; Nicholson, 1989). For example, Hirshleifer (1983) studies "best-shot public games", where everyone benefits from the effort of the user that contributes the most to the public good. Hirshleifer (1983) shows that there is significant free-riding in these games, leading to the under-provision of the public good. Consequently, these goods are better managed by centralized authorities that enforce everyone's participation through a system of rewards and punishments.

However, free riding is not a given even with public goods. For example, free riding seldom arises in single-shot games (McMillan, 1979a) because, in this case, individuals are usually unable to learn about payoffs correctly. As Schneider and Pommerehne (1981) put it "... the extent to which free-riding actually occurs, in reality, is a matter of empirical research". Yet, and to the best of our knowledge, there have been no empirical studies at the user level providing evidence of this behavior in real-world settings, particularly with actual IT products launched by private firms. Our paper provides the first empirical example of this effect at work. It confirms that, indeed, peer influence reduces adoption when one benefits from friends' adoption even if one does not adopt compared to worlds where peer influence does not arise. We highlight that our setting is significantly different from prior work in the economics of public goods. In our case, there is no common public good that several users contribute to. Instead, we have dyadic relationships among users that have to make decisions about adopting a private good. Our paper studies a context in which a private firm explicitly builds some element of free-riding into its product offerings. We complement the prior theoretical and simulation works by offering, for the first time, empirical evidence of negative peer influence.

Finally, we note that both peer influence and homophily are usually at play in social network contexts, contributing to co-adoption. For example, Ma et al. (2015) use call detailed records to analyze the role of peer influence and homophily in the diffusion of call ring-back tones. Using data from a large Indian telecom operator, they conclude that both play a significant role in the diffusion of this product. However, it is empirically hard to separate peer influence from homophily in observational studies. Indeed, several empirical studies provide large statistics for the effect of peer influence because they fail to appropriately control for latent homophily, as discussed in detail by Shalizi and Thomas (2011). Strategies to overcome the usually positive biases in peer influence include controlling for multiple observed variables, using matched samples (e.g., Aral et al., 2009), and applying randomization (e.g., Anagnostopoulos et al., 2008).

For example, Aral et al. (2009) show that failing to control for homophily can inflate the estimate of peer influence by 300-700%. The authors look at peer influence in an instant messaging network and use a matched-sample to distinguish homophily from peer influence. They conclude that the latter is responsible for at least 50% of the observed correlation in adoption. These results speak to the importance of using appropriate empirical strategies to avoid overestimating peer influence. In this paper, we apply two strategies to do so. The first strategy is to control for homophily by finding a proxy for how likely a user is to adopt the product in the absence of peer influence. We propose a new idea to accomplish this goal, namely using the number of friends that end up adopting the product as such a proxy. We provide extensive simulation results showing that this strategy works well over a wide range of homophily levels. The second strategy is to use randomization, as defined in Nyblom et al. (2003) and Anagnostopoulos et al. (2008), which consists of attempting to estimate, instead of eliminating, the effect of unobserved homophily, which can be subtracted from the estimates that confound it with the effect of peer influence. We detail each of these strategies in section 4.

3 Context and Data

We study add-ons to tariff plans in the mobile communications industry that allow unlimited calling for a certain period in exchange for a one-time upfront small fixed fee. We consider two different versions of these products. One such version allows for calling any user with service from the same telecom carrier, irrespective of whether she purchased the same add-on. This type of add-on is an example of a product of type N. The users that adopt one of these add-ons will increasingly call their friends, reducing their likelihood of purchasing such an add-on themselves. Another type of add-on allows for calling only users with service from the same telecom carrier who also buy the same add-on. This type of add-on is an example of a product of type P because users only benefit from it after they adopt it. Throughout this paper, and for the sake of space, we focus mostly on the add-ons of type N. Appendix A provides a theoretical model showing that users are likely to adopt an add-on of type N when a minimum number of their friends do not adopt. Derivations for a product of type P are similar and available on request by the authors. We also provide estimates for the effect of peer influence for an add-on of type P in Appendix B for the sake of comparison to the main results that we discuss in the paper.

We use an anonymized panel of data with detailed information about all users who subscribe to mobile service from a large telecommunications provider. There are roughly 5.7 million active users in this dataset during our period of analysis. The data include call detailed records for all calls placed by all users in this provider between August 2008 and June 2009. This dataset does not include business clients nor call centers. The record for each call includes anonymized identifiers for the caller and the callee and the call's start time. On an average day, users generate about 4 million calls in our dataset. Additionally, the data contain information about pricing plans and add-ons between January 2008 and June 2009. At every point in time, each user is associated with one tariff plan and possibly several add-ons. Add-ons are "a la carte" services that users can purchase, such as unlimited calling on weekends or at night for a given period of time (e.g., a full month).

During these 11 months, the provider offered five products of type N, which we call $N(1), \ldots N(5)$, all of which correspond to short-term offers that satisfy the following conditions: (1) the product was offered for a given period time and within the 11-month period for which we have data; and (2) the product was adopted by at least 1% of the users. These criteria ensure that all adoption of these products occurred within a relatively short period (and therefore is completely captured in our dataset) and that there is a critical mass of adopters to analyze.

We focus our analysis on the first of these products, called $N(1)$. Later, we show that the other four products of type N yield similar results. Product $N(1)$ allowed adopters unlimited calling to all users with service from the same telecom carrier during December 2008 (independently of whether the latter adopted the same product), in exchange for a small flat fee. This product became available in mid-November 2008 and was adopted by 252K users between then and mid-December 2008. Figure 1 shows the adoption curve for this product, which follows an S-shape, as expected. Even though there is no explicit built-in mechanism that users could use to inform their friends that they adopted this product, the short-term nature of its offering and the audiences that it was targeted to – namely students and young professionals – were likely to prompt adopters to mention it to their friends during their calls or in person. Adopters likely mention to their friends that they adopted this product because when they do so, they can call their friends more often and longer, and thus they likely need to explain to them the reason behind the (sudden) change in their behavior. In any case, such an explicit message about who adopts $N(1)$ is not required for peer influence to arise in our setting, and, similarly, for observing the adoption of the product itself. By calling their friends more often after adopting this product, adopters send non-adopters a signal that they may consider when deciding whether to adopt.

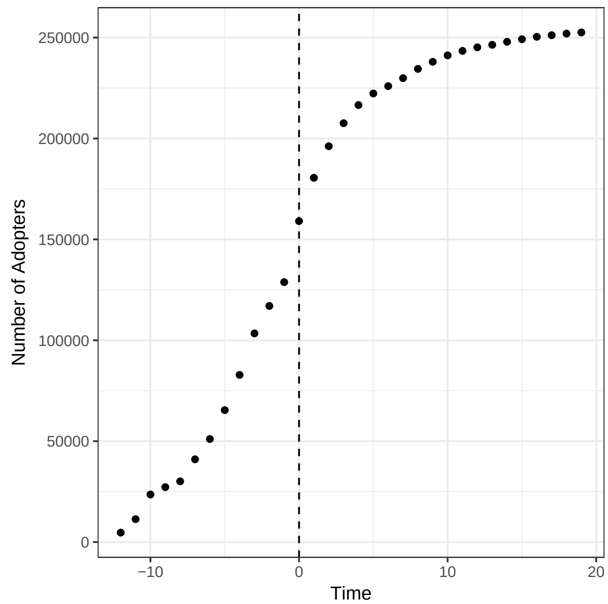

Figure 1: The adoption of product N(1) over time. The dashed line represents the moment when N(1) became active.

We define friends based on the number of calls that users exchange during a given period, as is customary when using data from call detailed records (e.g., Han and Ferreira, 2014). We use five different definitions of friend and show that our results are similar across them, thus providing evidence of robustness. For our main results, we use the social network generated by adding one edge to the social graph when two users exchange at least three calls during November 2008, which is the month preceding the activation of product N(1). Figure 2(a) shows the distribution of the number of friends in our social graph using this definition. As expected, this distribution is highly skewed. The average degree is 2.90, with a standard deviation of 4.25. The most connected user in this social graph has 94 friends. Figure 2(b) shows the distribution of the number of calls, placed plus received, during November 2008 (similar statistics are obtained for any other month in our data). As expected, this distribution is also very skewed. The average number of calls per day is

2.30, with a standard deviation of 3.52. The user involved in more calls is involved, on average, in 39.55 calls per day during this month.

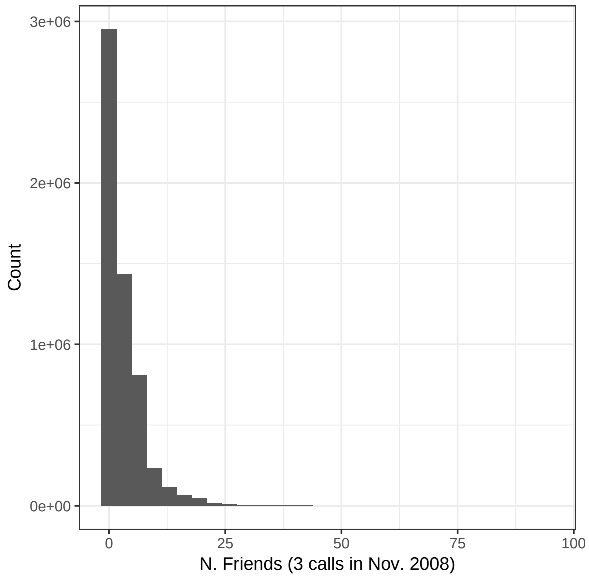

(a) Number of friends

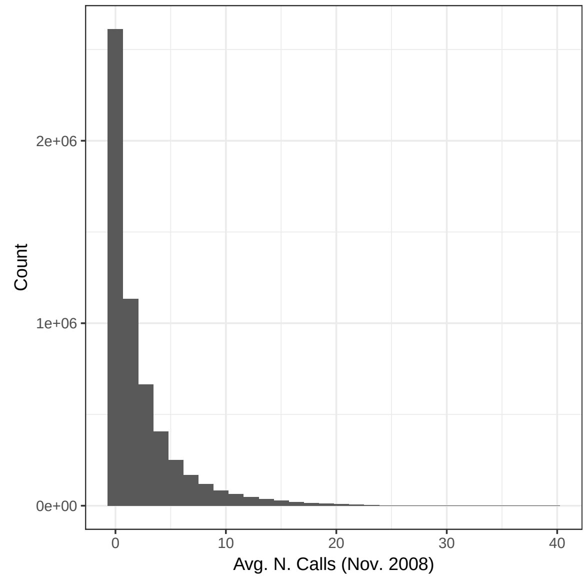

(b) Daily average number of calls

Figure 2: Histograms of number of friends and daily average number of calls (November 2008).

We also look at differences in calling behavior before and after the adoption of product N(1). Figure 3 shows that, after adoption, users place more calls to both adopters and non-adopters (in black) and get more calls from adopters, but fewer calls from non-adopters (in grey).[2] These results are in line with our expectations and with the prediction that, on average, non-adopters "take advantage" of adopters. The behavior of adopters changes as we would expect given the definition of product N(1), and thus adopters give their friends the expected signals for peer influence to potentially arise in our setting.

[2]Details on the regressions and the results that lead to this figure are available upon request from the authors.

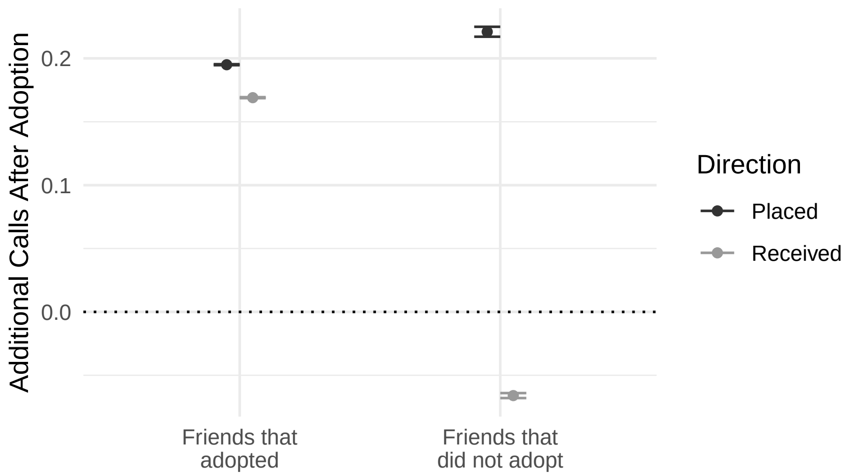

Figure 3: Calls placed and received to and from adopters and non-adopters of product N(1) during the time when this product was active.

4 Empirical Strategy

We start by studying the aggregate effect of peer influence on products of type N in our setting. To do so, we represent our data using an undirected graph, $G = (V, E)$, where V is our set of users and E is a set of undirected edges. Edge e^{ij} connects users i and j when they are friends. Each user $i \in V$ has a time-changing attribute $y_{i,t}$ indicating whether she adopted the product up to (and including) time t. We use the canonical peer effects model to measure peer influence (Leenders, 2002) and thus assume that the probability with which user i adopts the product at time t follows a distribution that depends on: (1) the number of friends that have already adopted that product, called $exposure_{i,t}$, and given by $\sum_{j:e^{ij} \in E} y_{j,t}$; (2) her own observed and unobserved characteristics, $\mathbf{X}_i$ and c_i, respectively. Our interest is to measure, on average, how having one more friend that adopts the product changes the likelihood of adoption. We use a linear probability model (LPM) to do so because linear models are computationally light, which is important in our case with 5.7 million users in our regressions, provide good estimates for average marginal effects while, at the same time, rely on relatively weak assumptions (Wooldridge, 2001, pp. 453-454). We thus write a model of the form

$$p(y_{i,t} = 1 | y_{i,t-1} = \cdots = y_{i,1} = 0, \mathbf{X}_i, c_i) = \alpha + \rho \, exposure_{i,t} + \mathbf{X}_i \beta + c_i + d_t, \tag{1}$$

where d_t are time dummies. We are interested in estimating this model to obtain parameter ρ, which is called the effect of peer influence. However, with homophily, similar users are likely to be friends in the social network. Therefore, c_i is likely positively correlated with the error term, and estimating equation 1 using OLS results in an upward-biased estimate for the effect of peer influence. If the user unobserved characteristics (c_i) were observed, including them in the model would remove the bias caused by homophily. Given that c_i is unobserved, the natural way to estimate this equation is to use fixed effects. However, one cannot use fixed effects in this case, given the nature of our panel. As is customary in panels measuring diffusion processes, users leave the panel when they adopt the product. Therefore, the user fixed effect captures adoption perfectly, soaking up the effect of interest.

One way to address this problem is to find a proxy for c_i that accounts, as best as we can, for unobserved factors contributing to the likelihood of adoption (Wooldridge, 2001). In settings like ours, one such proxy might be the number of friends that end up adopting the product, denoted by $final\_exposure$ in our paper. The overall rationale for using final exposure to disentangle homophily from peer influence is that it provides information about the focal user that would otherwise be discarded. This additional information comes not only from the friends who adopt before the focal user but also from friends who adopt after the focal user does. Information about friends' adoption after the focal user adopts is usually discarded because it is assumed to not play a role in the adoption decision of the focal user, as it happens after the adoption itself. Nevertheless, in a world without peer influence but with homophily, adopters will be more likely to be connected to adopters and non-adopters to non-adopters, leading to a positive correlation between being an adopter and having friends that also adopt. Thus, in a world with homophily, having information about whether the friends of the focal user adopt the product, either before or after the focal user does, provides information about whether the focal user is also a likely adopter. This information can be used to control for homophily and reduce bias in estimating peer influence. In Appendix C, we show mathematically that this is the case in our setting. We show simulation results illustrating that this proxy behaves well in practice independently of the network structure and over a wide range of values for the parameters of interest that configure settings similar to ours, such as the baseline level of adoption of the product and the level of homophily.

Finally, we note that even an imperfect proxy for c_i might be very useful when estimating peer influence for products of type N. As long as one believes that the proxy is positively correlated with one's likelihood of adoption—and this should be the case for final exposure due to homophily in social networks—a negative estimate for the effect of friends' adoption is enough to show that peer influence reduces diffusion in this case compared to a world without peer influence. With this in mind, we estimate:

$$p(y_{i,t} = 1 | y_{i,t-1} = \cdots = y_{i,1} = 0, \mathbf{X}_i, c_i) = \alpha + \rho\, exposure_{i,t} + \mathbf{X}_i \beta + \gamma\, final\_exposure_i + d_t, \quad (2)$$

where ρ is the effect of peer influence. A positive estimate for ρ indicates that users are more likely to adopt the product because of peer influence compared to a world without peer influence. Similarly, a negative estimate for ρ indicates that users are less likely to adopt the product because of peer influence compared to a world without peer influence. That is, estimating equation (2) leads us always to report the effect of peer influence vis-a-vis a counterfactual without peer influence. Note that this counterfactual still captures the traditional slow down of the diffusion curve triggered by reasons other than peer influence that may be unknown to us as researchers. In particular, the traditional S-shape of the diffusion curve is captured in our model by the time dummies, and thus the estimates that we report in our paper are net of that usual shape. For example, the diffusion curve for products of type N is likely still S-shaped but slows down more aggressively than it otherwise would in the absence of such negative peer influence. Our estimates in this paper for products of type N measure the loss in diffusion due to the negative effect of peer influence beyond the usual slow down of the diffusion curve.

Furthermore, we aim at separating the contributions of the informational and of the economic mechanisms to the aggregate effect of peer influence by estimating the effect associated to the first few friends that adopt the product separately from the effect associated to the subsequent friends that adopt the product still using the empirical strategy suggested above, i.e., still controlling for final exposure. As discussed before in section 3, our rationale here is that the effect of the informational mechanism must be strongest for the first friends that adopt the product, and then decrease for each subsequent friend that adopts, and vice-versa for the effect of the economic mechanism. Therefore, we also estimate the following model:

$$\begin{aligned} p(y_{it}=1|y_{i,t,\dots,t-1}=0,\mathbf{X}_i,c_i) = \quad & \alpha + \rho_1\, initial\_exposure\_x_{it} + \rho_2\, additional\_exposure\_x_{it} \\ & +\mathbf{X}_i\beta + \gamma\, final\_exposure_i + d_t, \end{aligned} \tag{3}$$

where $initial\_exposure\_x_{it}$ corresponds to the number of friends of user i that have adopted the product up to (and including) time t while exposure is below threshold x, and $additional\_exposure\_x_{it}$ is the number of friends of user i that have adopted the product up to (and including) time t when exposure is beyond threshold x. Mathematically, $initial\_exposure\_x_{it} = exposure_{it} * (exposure_{it} \leq x)$ and $additional\_exposure\_x_{it} = exposure_{it} * (exposure_{it} > x)$. Therefore, when regressing the likelihood of adoption on the adoption of (a) the first x friends that adopt and (b) the remaining friends (those that adopt after the first x friends do), we should observe the following: 1) a positive coefficient for (a) when x is small, capturing the importance of the informational mechanism for the very first friends that adopt; 2) a coefficient for (a) that reduces as x increases, capturing the idea that the more friends that adopt the product, the less new information they transmit about the product; 3) a negative coefficient for (b) for all x because additional adoption always reduces the incentive of the focal user to adopt.

5 Results

5.1 Peer Influence in the Adoption of Product N(1)

We start by estimating the aggregate effect of peer influence on product N(1). Table 1 shows the results obtained using the data from November 2008 to build the social network and a threshold of 3 calls to add edges to this network. We use these data to create the social network to analyze the effect of peer influence in the adoption of product N(1) because this product was available at the end of 2008 and became active on December 1st, 2008. In the next subsection, we show that our results do not change when we use other criteria to define the social network.

Column (1) shows results from regressing adoption on exposure using time dummies and controlling for the number of friends without controlling for final exposure. The coefficient obtained for exposure is positive, indicating a positive correlation between exposure and adoption. As discussed above, this positive correlation may be the result of unobserved homophily. Therefore, we run the same regression adding final exposure as a proxy for the likelihood of adoption to control for such potential sources of endogeneity. Column (2) shows the results obtained. In this case, the coefficient on exposure is negative, indicating that the more friends that adopt the product, the less likely one is to adopt it, as we would now expect for the case of this product. This result provides empirical evidence of the effect of using final exposure as a proxy for user type in our setting. Doing so corrects the sign of the effect of peer influence in the direction that we would expect. Out of a total of 5.7M users, about 252K (4%) adopted this product. Given that on average, a user was exposed to 0.147 adopters per day during the 32 days that the product was available for adoption, an effect of -0.0003 (row 1, column 2 in table 1) per day results in about 7,500 fewer adoptions, or 3%, of the total adoption.[3]

Next, we look at the results obtained by separately estimating the effect of the informational and the economic mechanisms of peer influence. Columns (3)-(6) show the results of using different thresholds for initial exposure. The coefficients of the form "$initial\_exposure\_x$" pertain to the adoption by the first x friends and "$additional\_exposure\_x$" pertain to the adoption of the remaining friends. Column (3), for example, shows results from regressing adoption on the first friend that adopts ($initial\_exposure\_1$) and on the remaining ones ($additional\_exposure\_1$) using time dummies and controlling for $final\_exposure$ (note that at any time t, $final\_exposure$ still differs from "$initial\_exposure\_x + additional\_exposure\_x$" by the number of friends that adopt the product after time t). The coefficient associated with the effect of the informational mechanism ($initial\_exposure\_1$) is positive, while the coefficient associated with the effect of the economic mechanism ($additional\_exposure\_1$) is negative. This result comes in line with what one would

[3] 5.7M total users $\times$ 0.147 average exposure $\times$ 32 periods $\times$ -0.0003 = -7,500 additional users. These 7,500 fewer users correspond to 3% of the 252K users that adopted the product.

expect after our discussion in section 3. Early exposure has a positive effect on adoption, associated with learning about the product. Subsequent exposure has a negative effect on adoption. Whenever a friend adopts the product, less new information about it is conveyed to the focal user. However, such adoption still reduces the economic incentive to adopt because there is one fewer friend to talk with upon adoption that could not already be reached before.

The coefficient of 0.001 (row 2, column 3) means that being friends with an adopter is positively associated with an increase in the likelihood of adoption of 0.1 percentage points, but each additional adopter decreases the likelihood of adoption also by 0.1 percentage points (row 3, column 3). Columns (4)-(6) show the results of using different thresholds, namely 3, 5, and 10, for x. The coefficient on additional exposure is always negative, as we would expect, capturing the effect of the economic mechanism. Also, as expected, the coefficient on initial exposure starts positive (for example, for $x = 3$), becomes statistically insignificant for a medium x (for example, $x = 5$), and becomes negative for a high x (for example, $x = 10$). This result shows that at some point, the effect of the economic mechanism dominates, as expected, given the negative aggregate effect of peer influence estimated in column (2). Finally, we note that as expected, the correlation between our measures of exposure, and in particular, between $exposure$ and $final\_exposure$ increases over time. However, the negative estimates that we obtain for the effect of peer influence for product N(1) do not arise because of such correlations. First, such correlations are also present for products of type P. Yet, adding $final\_exposure$ to estimate the effect of peer influence, in this case, retrieves positive estimates, as one would expect. Second, below in subsection 5.3, we find, again, negative estimates for the effect of peer influence for products of type N in our setting without adding $final\_exposure$ to our regressions. Finally, an alternative explanation for the negative effect of peer influence is budget coordination. For example, family members may coordinate on who pays the flat fixed fee to adopt a product of type N, having the designated buyer initiating free calls to the other family members. We have investigated this alternative and find no evidence of such behavior. The results of these analyses are available on request.

Table 1: Peer Influence for product N(1) as a function of initial exposure (first 1, 3, 5, and 10 friends that adopt), remaining additional exposure and final exposure.

	Dependent variable:					
	adopter					
	(1)	(2)	(3)	(4)	(5)	(6)
exposure	0.002***	−0.0003***				
	(0.00002)	(0.00002)				
initial_exposure_1			0.001***			
			(0.00003)			
additional_exposure_1			−0.001***			
			(0.00002)			
initial_exposure_3				0.0002***		
				(0.00002)		
additional_exposure_3				−0.002***		
				(0.00003)		
initial_exposure_5					0.00001	
					(0.00002)	
additional_exposure_5					−0.003***	
					(0.00003)	
initial_exposure_10						−0.0002***
						(0.00002)
additional_exposure_10						−0.003***
						(0.00004)
final_exposure		0.004***	0.004***	0.004***	0.004***	0.004***
		(0.00002)	(0.00002)	(0.00002)	(0.00002)	(0.00002)
n_friends	−0.00001***	−0.0001***	−0.0001***	−0.0001***	−0.0001***	−0.0001***
	(0.00000)	(0.00000)	(0.00000)	(0.00000)	(0.00000)	(0.00000)
Observations	179,172,418	179,172,418	179,172,418	179,172,418	179,172,418	179,172,418

Note: *p<0.1; **p<0.05; ***p<0.01
Fixed effects: Time

5.2 Using Additional Products of type N and Alternative Network Definitions

In this subsection, we present additional results that further corroborate our findings and show the robustness of our empirical method. Namely, we present estimates for the aggregate effect of peer influence on the other four products of type N in our dataset. We do so using additional alternative definitions of friendship. Table 2 summarizes the results obtained from running regressions similar to those presented in columns (1) and (2) of table 1. For each product, we show results using five network definitions. Two definitions require users to exchange a minimum number of calls in the month before the product was available. For example, the results in the preceding subsection for product N(1) used data from November 2008 and a threshold of 3 calls to add an edge to the social network. We now add another definition of friendship that requires users to exchange five calls. These definitions are called $3call$ and $5call$ below. Two other new definitions require again that users exchange 3 or 5 calls in the month before the product was available in the market but require at least one call in each direction. They are called $3mutcall$ and $5mutcall$ below. These four definitions rely on recency of contact. In a fifth definition, we require users to exchange at least 30 calls over the entire 11 months of data. Using this definition for friendship is likely to capture enduring relationships among users, who are potentially more likely to affect each other's decision to adopt products in our setting. Therefore, this approach to define the social network complements well the previous four. This last criterium to define friend is labeled $n\_calls30$ below.

The first line of Table 2 recovers the results in columns (1) and (2) of Table 1. A total of 252,574 users adopted product N(1), and the average exposure during the period that this product was available in the market for adoption is 0.147. The estimate for the coefficient of peer influence without controlling for final exposure is positive, indicating that peer influence would be responsible for 18% of the observed adoption. However, and as shown in the previous subsection, after controlling for final exposure, this coefficient becomes negative, indicating that peer influence is, in fact, responsible for a decrease of 3% in total adoption. The same pattern arises in all cases in Table 2. Namely, the coefficient on exposure is always positive before adding our proxy for the likelihood of adoption, and it becomes negative after doing so. This result speaks to the danger associated with using OLS to estimate plain peer effect models to measure peer influence and provides evidence that our findings are robust to a set of products of type N and other definitions of friendship. This result shows the powerful role that controlling for final exposure holds in improving one's estimates for the effect of peer influence on product adoption. Results using a Probit specification — available from the authors on request — come in line with the ones in Table 2 obtained using LPM, thus showing that our findings are not an artifact of functional form. Finally, Appendix D shows the diffusion curve for all products of type N in our dataset and product P, as well as their respective counterfactuals in the absence of peer influence. These figures show how adoption would have been higher for products of type N if there had been no peer influence and conversely for product P.

Table 2: Adoption of 5 products of type N as a function of exposure with and without controlling for final exposure using different definitions for the social network.

Product	Network	Adopters	Avg. Exposure	Periods	Exposure (no controls)	Extra Adoption	Exposure (controlling for final exposure)	Extra Adoption
N(1)	200811 - 3call	252,574	0.147	32	0.0017*** (0.00025)	46,063 (18%)	-0.00028*** (5e-05)	-7,512 (-3%)
N(1)	200811 - 3mutcall	252,574	0.107	32	0.0041*** (2.5e-05)	81,033 (32%)	-0.00017*** (2.9e-05)	-3,254 (-1%)
N(1)	200811 - 5call	252,574	0.094	32	0.0036*** (2e-04)	62,683 (25%)	-0.00017*** (2.8e-05)	-2,940 (-1%)
N(1)	200811 - 5mutcall	252,574	0.078	32	0.0052*** (2.6e-05)	75,399 (30%)	0.00016*** (3.6e-05)	2,299 (1%)
N(1)	year - n_calls30	252,574	0.147	32	0.0029*** (0.00015)	79,193 (31%)	-0.00028*** (2.8e-05)	-7,703 (-3%)
N(2)	200808 - 3call	159,619	0.113	31	0.0017*** (0.00025)	35,033 (22%)	-0.00044*** (7.5e-05)	-8,883 (-6%)
N(2)	200808 - 3mutcall	159,619	0.083	31	0.0035*** (3.9e-05)	51,283 (32%)	-6e-04*** (3.5e-05)	-8,841 (-6%)
N(2)	200808 - 5call	159,619	0.074	31	0.0032*** (0.00014)	42,789 (27%)	-5e-04*** (4.1e-05)	-6,646 (-4%)
N(2)	200808 - 5mutcall	159,619	0.062	31	0.0044*** (2.7e-05)	48,361 (30%)	-0.00044*** (4.1e-05)	-4,893 (-3%)
N(2)	year - n_calls30	159,619	0.101	31	0.0026*** (8.6e-05)	46,687 (29%)	-0.00052*** (3.8e-05)	-9,344 (-6%)
N(3)	200808 - 3call	92,383	0.072	29	0.0021*** (3e-05)	24,595 (27%)	-0.0041*** (0.00012)	-49,464 (-54%)
N(3)	200808 - 3mutcall	92,383	0.051	29	0.0024*** (2.3e-05)	20,458 (22%)	-0.0048*** (7.8e-05)	-41,068 (-44%)
N(3)	200808 - 5call	92,383	0.047	29	0.0025*** (2.7e-05)	19,556 (21%)	-0.0047*** (8.7e-05)	-36,945 (-40%)
N(3)	200808 - 5mutcall	92,383	0.039	29	0.0027*** (2.7e-05)	17,226 (19%)	-0.0051*** (9e-05)	-32,966 (-36%)
N(3)	year - n_calls30	92,383	0.075	29	0.002*** (2.4e-05)	24,506 (27%)	-0.0043*** (6.3e-05)	-53,414 (-58%)
N(4)	200808 - 3call	139,761	0.080	27	0.0034*** (6.1e-05)	42,667 (31%)	-0.00064*** (3.6e-05)	-8,016 (-6%)
N(4)	200808 - 3mutcall	139,761	0.057	27	0.0042*** (3.2e-05)	37,511 (27%)	-0.00071*** (4.4e-05)	-6,291 (-5%)
N(4)	200808 - 5call	139,761	0.053	27	0.0043*** (3.8e-05)	34,821 (25%)	-0.00069*** (4.5e-05)	-5,621 (-4%)
N(4)	200808 - 5mutcall	139,761	0.044	27	0.0047*** (3.8e-05)	31,919 (23%)	-0.00068*** (5.2e-05)	-4,634 (-3%)
N(4)	year - n_calls30	139,761	0.084	27	0.0034*** (2.6e-05)	44,449 (32%)	-0.00063*** (3.3e-05)	-8,178 (-6%)
N(5)	200906 - 3call	54,147	0.035	36	0.0014*** (2.1e-05)	9,760 (18%)	-0.0014*** (4.7e-05)	-9,789 (-18%)
N(5)	200906 - 3mutcall	54,147	0.024	36	0.0016*** (2.1e-05)	8,095 (15%)	-0.0016*** (5.2e-05)	-8,040 (-15%)
N(5)	200906 - 5call	54,147	0.022	36	0.0016*** (2.2e-05)	7,210 (13%)	-0.0015*** (5.3e-05)	-6,665 (-12%)
N(5)	200906 - 5mutcall	54,147	0.018	36	0.0017*** (2.5e-05)	6,398 (12%)	-0.0016*** (6.2e-05)	-5,873 (-11%)
N(5)	year - n_calls30	54,147	0.028	36	0.0014*** (1.8e-05)	8,278 (15%)	-0.0014*** (4.7e-05)	-8,196 (-15%)

*p<0.1; **p<0.05; ***p<0.01
Standard errors clustered at the subscriber level in parentheses

5.3 Using Randomization to Measure Peer Influence

An alternative approach to measuring the effect of peer influence in a network setting is to follow the strategy proposed in Nyblom et al. (2003) and Anagnostopoulos et al. (2008). This strategy relies on using randomized versions of the data to infer an empirical distribution for the effect of friends' adoption under the null hypothesis of no peer influence and then comparing the average of this distribution to the estimate obtained for this statistic with the original data. In appendix E, we show that under mild conditions, this procedure yields a lower bound, in absolute terms, for the magnitude of the effect of peer influence when the latter is measured using a peer effects model, which can be seen as an advantage vis-a-vis, for example, propensity score matching, which usually biases upwards the estimates of this effect given that, by definition, homophily is positive.

Under the null hypothesis of no peer influence, the probability of adoption is not determined by the number of friends that have already adopted the product. In our case, and again following Anagnostopoulos et al. (2008), we operationalize such a world by randomly shuffling the adoption dates across eventual adopters. In other words, we assume that the adoption dates of one's friends do not contribute to her adoption at any point in time under the null hypothesis. However, even if one estimates equation (1) using the pseudo dataset obtained by randomly shuffling the adoption dates of friends that adopt the product, one still obtains a positive estimate for the effect of peer influence, which captures the positive correlation between one's adoption and the number of friends that have already adopted the product arising from the similarity in unobserved characteristics between focal users and their friends that affect their propensity to adopt in the same way — homophily in short.

The randomized versions of the original data used to model pseudo-worlds need to be carefully generated. In particular, these worlds should exhibit the same aggregate descriptive statistics as the original one; otherwise, we could be modeling worlds that are hard to come by. Therefore, we randomly shuffle the adoption dates among the friends of focal users that eventually adopt the product. This approach ensures the same number of adopters at all points in time in all pseudo-worlds as in the original world. Therefore, this also guarantees that the aggregate diffusion curve for the product remains unchanged.

In addition to this procedure, we further restrict the random shuffles that we perform. We shuffle adoption dates only among friends of focal users whose adoption occurred within the same week[4]. This additional restriction further guards against potential concerns associated with the fact that adoption dates may still conceal unobserved effects that lead to adoption that we would erroneously interpret as peer influence. For example, by shuffling adoption dates among eventual adopters without restrictions, we are implicitly assuming that adoption dates are all drawn from the same distribution. However, if one's propensity to adopt is correlated not only with the number of friends that eventually adopt but also with whether one is an early or a late adopter (that is, if there is temporal clustering in adoption), then unrestricted shuffling may be unsatisfactory. If early adopters tend to be friends with each other, then shuffling adoption dates among all adopters would assign late adoption dates to early adopters' friends, changing the original data in undesired ways. Restricting shuffles to friends of focal users that adopt in the same week addresses this concern and still preserves aggregate network-level statistics, such as the total number of adoptions in each time period.

We compute the empirical distribution of the effect of peer influence – ρ – across our pseudo-worlds by running the model in equation (1) for each randomized version of the data. For each

[4]If a users' friend is also friends with someone that adopted in another week, she is assigned to the partition corresponding to the earlier week. This is not common in our dataset (10 cases only).

product of type N and each definition of the social network, we simulate 1,000 pseudo-worlds. Then, we reject our null hypothesis if the estimate of ρ obtained from estimating this same equation with the original data falls outside the 95% confidence interval of the parameter obtained from the empirical distribution. Figure 4 shows the results obtained for product N(1) using this method for a random sub-sample of 10,000 users[5]. The data from November 2008 with a threshold of 3 calls to add an edge to the graph were used to build the social network in this case. This figure shows the empirical distribution for parameter ρ obtained from estimating equation (1) across our pseudo-worlds and the estimate obtained using the original data. We observe that both the average of the empirical distribution and the estimate of ρ using the original data are positive. The average of this distribution is 0.0035, and the standard deviation is 1.8e-4. However, in this case, and in line with our expectation, the coefficient obtained with the original data is statistically lower than the average of the empirical distribution. With the original data, the coefficient on exposure is 0.0024, outside the 95% confidence interval around the average of the empirical distribution. This result means that without peer influence, the coefficient associated with the role of friends' adoption is higher than the coefficient obtained using the original data, and thus — after using randomization for accounting for unobserved effects such as homophily — we find that peer influence reduces adoption in the case of product N(1). Out of the 10,000 users included in this analysis, 534 adopted this product. Given that on average, a user was exposed to 0.122 adopters per day during the 32 days that the product was available for adoption, an effect of -0.0011 (row 1, column 10 in table 3) per day results in about 44 fewer adoptions, or 8%, of the total observed adoption.[6] Appendix

[5] We use a random sub-sample of 10,000 users due to computational restrictions. Applying randomization for the full dataset and 30 combinations of product-networks in our paper would require at least 30,000 hours, which corresponds to about three and a half years

[6] 10,000 total users $\times$ 0.122 average exposure $\times$ 32 periods $\times$ (0.0024 - 0.0035) = 44 fewer users. These 44 fewer users correspond to 8% of the 534 users who adopted the product in our subsample of 10,000 users.

E shows that using randomization to measure the effect of peer influence with peer effects models provides estimates that are biased towards zero. Therefore, is it expected that the estimate for the effect of peer influence obtained with randomization for product N(1) is lower than the one obtained without randomization (which was 18%).

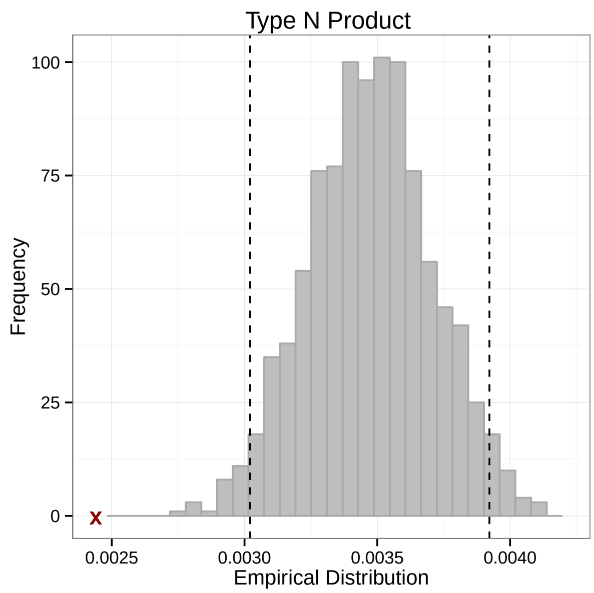

Figure 4: Distribution of the coefficient on exposure for product N(1) over 1,000 shuffles of the adoption dates.Dashed lines represent 95% confidence intervals. The '$\times$' mark represents the coefficient obtained using the original data.

The last two columns in Table 3 summarize the results obtained using randomization for all products of type N in our dataset. We observe a similar pattern across all these products. After randomization, the effect of peer influence on adoption is negative and smaller in magnitude than that obtained without randomization. The effect of peer influence for products N(3) and N(4) is negative but not statistically different from zero, which, again, may arise because randomization is a conservative empirical method that, as discussed above, provides only a lower bound for the true effect of peer influence. Overall, the results in this table provide further evidence that, indeed, peer influence does not contribute to increasing adoption in the case of products of type N and that randomization captures the expected direction of the effect of peer influence in these cases. Finally, and for the sake of completeness, Appendix F shows that randomization in our setting provides the expected results for the case of product P.

Table 3: Adoption of products of type N as function of exposure with and without controlling for final exposure and using randomization.

Product	Network	Adopters	Avg. Exposure	Periods	Exposure (no controls)	Extra Adoption	Exposure (controlling for final exposure)	Extra Adoption	Exposure (randomization)	Extra Adoption
N (1)	200811 - 3call	534	0.122	32	0.0024*** (0.00043)	96 (18%)	-0.0042*** (0.00071)	-165 (-31%)	-0.0011*** (0.00018)	-44 (-8%)
N (1)	200811 - 5call	534	0.082	32	0.0034*** (0.00058)	88 (17%)	-0.005*** (0.00092)	-132 (-25%)	-0.0014*** (0.00026)	-36 (-7%)
N (1)	year - n_calls30	534	0.134	32	0.0027*** (0.00042)	114 (21%)	-0.0036*** (0.00068)	-152 (-28%)	-0.001*** (0.00016)	-43 (-8%)
N (2)	200808 - 3call	357	0.099	31	0.0014*** (0.00038)	44 (12%)	-0.0037*** (0.00071)	-112 (-31%)	-0.00093*** (2e-04)	-29 (-8%)
N (2)	200808 - 5call	357	0.066	31	0.0024*** (0.00053)	48 (14%)	-0.0037*** (0.00095)	-76 (-21%)	-0.00089*** (0.00028)	-18 (-5%)
N (2)	year - n_calls30	357	0.094	31	0.0019*** (0.00041)	54 (15%)	-0.003*** (0.00076)	-88 (-25%)	-0.00062*** (2e-04)	-18 (-5%)
N (3)	200808 - 3call	205	0.074	29	0.00077** (0.00031)	17 (8%)	-0.0078*** (0.0017)	-168 (-82%)	3e-05 (0.00016)	
N (3)	200808 - 5call	205	0.048	29	0.00096** (0.00043)	13 (7%)	-0.0089*** (0.0021)	-124 (-60%)	2.5e-05 (0.00021)	
N (3)	year - n_calls30	205	0.081	29	0.00077** (0.00032)	18 (9%)	-0.0084*** (0.0015)	-197 (-96%)	-0.00012 (0.00016)	
N (4)	200808 - 3call	299	0.082	27	0.0023*** (0.00053)	54 (18%)	-0.0019** (0.00076)	-44 (-15%)	-1e-04 (0.00018)	
N (4)	200808 - 5call	299	0.053	27	0.0028*** (0.00068)	43 (14%)	-0.0024** (0.00098)	-36 (-12%)	-0.00017 (0.00027)	
N (4)	year - n_calls30	299	0.091	27	0.0022*** (0.00041)	59 (20%)	-0.0015** (0.00065)	-41 (-14%)	-5e-05 (0.00017)	
N (5)	200906 - 3call	110	0.037	36	7e-04** (3e-04)	10 (9%)	-0.0045*** (0.0012)	-62 (-57%)	-0.00053*** (0.00018)	-7 (-7%)
N (5)	200906 - 5call	110	0.024	36	0.0011** (0.00044)	9 (9%)	-0.0052*** (0.0015)	-46 (-42%)	-0.00065*** (0.00024)	-6 (-5%)
N (5)	year - n_calls30	110	0.030	36	0.001*** (0.00038)	12 (10%)	-0.0041*** (0.0013)	-46 (-42%)	-0.00045** (0.00019)	-5 (-5%)

*p<0.1; **p<0.05; ***p<0.01
Standard errors clustered at the subscriber level in parentheses

6 Conclusion

Studying how products and services spread across social networks has been a topic at the top of the research agenda in information systems, economics, management, and marketing. Understanding diffusion allows, for example, for predicting how market size is likely to evolve, which has profound implications for many managerial tasks such as adjusting production, managing inventory, or investing in brand awareness. In parallel, IT industries are rich in products with positive network externalities, i.e., products from which one user benefits when her friends adopt. Yet, little has been done to date to characterize diffusion when such benefits accrue to non-adopters, that is, for the case of products from which one user benefits when her friends adopt even if she does not — called products of type N in our paper. The incentives offered to users to purchase this type of product are significantly different from those that arise when the benefit from adoption still increases with the number of friends who adopt the product. Yet, such benefits are only realized upon adoption — called products of type P in our paper. We compare the diffusion of these two types of products in this paper and show that the former yield significantly less adoption.

We show that two mechanisms arise when users consider products of type N. On the one hand, and as is usually the case in social networks, users learn about the product from their friends, such as how to use it. This informational mechanism reduces uncertainty and increases the likelihood of adoption. This mechanism has been studied at length in the prior literature on the positive effects of peer influence. However, and on the other hand, when users benefit from friends' adoption even if they do not adopt, another important mechanism may come strongly into play, one of a purer economic nature, which reduces the likelihood of adoption. In the case of these products, non-adopters can already benefit from the product through their friends, which reduces their incentive to adopt the product themselves. We empirically estimate the effect of both these mechanisms on adoption using a peer effects model and confirm an average negative effect of peer influence for products of type N. However, prior literature on peer influence has not considered nor measured a negative effect of peer influence, mostly because this type of products has never been analyzed in detail before despite their pervasiveness across the economy. In this regard, our paper is the first to measure a negative effect of peer influence. We provide a simple economic model to show that, yet, this should be expected in the case of these products.

We estimate this negative effect of peer influence in the context of the mobile industry and confirm that, indeed, products of type N register less adoption than products without peer influence, and consequently, also less adoption than products with positive peer influence. We still find a positive effect of peer influence on the first friends that adopt these products; thus, when the information that they share about the product is most useful. In other words, the informational mechanism dominates for the first friends that adopt the product. However, the more friends that adopt the product, the less new information they convey about it, but they keep reducing the economic incentive of the focal user to adopt. That is, the economic mechanism dominates for the friends that adopt later, reducing the likelihood of adoption. Our empirical work provides evidence of these two mechanisms at work in a real-world setting.

Also, we develop a new empirical approach to measure peer influence in these settings. Specifically, we propose using the number of friends that end up adopting the product — final exposure — as a proxy for homophily. We show that estimating one's likelihood of adoption as a function of the contemporaneous number of friends that have adopted the product and of this proxy recovers the correct sign for the effect of peer influence, i.e., the effect of peer influence is negative for products of type N and positive otherwise. We also offer extensive simulation results showing that this new empirical strategy behaves well across a wide range of parameters that define the structure of the social network, the degree of homophily, and the baseline likelihood of adoption. We also use randomization techniques to estimate the effect of peer influence in our setting. These techniques do not use final exposure to proxy homophily. Instead, and appropriately shuffling the data, they attempt to estimate homophily to subtract it from the estimates of the effect of peer influence obtained using OLS to provide more accurate measurements. Using these techniques in our setting confirms that indeed the effect of peer influence is negative in the case of products of type

N. Unfortunately, these randomization techniques are much more computationally intensive than only controlling for final exposure as we propose in this paper, which is very relevant in practice when considering large social networks, such as the one in this paper (5.7 million users).

Unveiling the negative effect of peer influence on products of type N and explaining why it arises has significant implications for platform strategy and managers. We show that these products diffuse less than if they were launched as products of type P. Our paper provides a methodology to estimate the loss in adoption that arises from changing the structure of the economic incentives for adoption in the way described above. Firms may have other reasons that lead them to incorporate some element of free riding in the way they launch products of type N, for example, brand awareness. Whatever these reasons are, they need to make up for the loss in diffusion discussed in this paper. In this paper, the approach that we propose to estimate this loss provides a robust and easy-to-implement empirical strategy that managers can use in a timely fashion to obtain a benchmark for how much additional adoption they need to collect otherwise. Furthermore, note that the definitions of products of type P and of type N that we use in this paper are very general. While the examples that we discuss pertain only to IT-related industries, it is clear that other economically significant activities offer similar incentive structures for product and service adoption (e.g., any activity that involves belonging to a club and bringing friends along) and thus, our results are likely to apply much more broadly without change.

Finally, we acknowledge that our paper does not come without limitations. First, we acknowledge that, in our setting, we are unable to separate the contribution of explicit messages for peer influence (such as word-of-mouth) from the contribution arising from other social signals. In particular, in our case, the adopters of tariff plans that allow unlimited calling may immediately start calling their friends, which, just by itself, may reduce the likelihood of adoption among the latter.

While explicit messages are unnecessary for peer influence to arise, we acknowledge that it would be interesting, and useful from a managerial perspective, to measure whether different types of explicit messages trigger heterogeneous peer effects. Second, we do not study a general equilibrium model in which adopting products that allow unlimited calling changes the underlying social network. While this might occur in practice with long-standing offers for such products, in our setting, the tariff plans that we study were only available for short periods, which significantly reduces the chances that they might have actively and significantly changed the social network, in particular, the connections among friends that matter for decision making. Also, the products that we study in this paper are short-lived and do not require long term commitments from users. If this were the case, our user decision-making model would likely need to be much more complex and consider dynamic aspects.

References

Anagnostopoulos, A., Kumar, R., and Mahdian, M. (2008). Influence and correlation in social networks. In *KDD '08: Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 7–15, New York, NY, USA. ACM.

Anderson, S. P., Goeree, J. K., and Holt, C. A. (2004). Noisy directional learning and the logit equilibrium. *The Scandinavian Journal of Economics*, 106(3):581–602.

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. 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.

Aral, S. and Walker, D. (2014). Tie strength, embeddedness, and social influence: A large-scale networked experiment. *Management Science*, 60(6):1352–1370.

Arifovic, J. and Ledyard, J. (2012). Individual evolutionary learning, other-regarding preferences, and the voluntary contributions mechanism. *Journal of Public Economics*, 96(9-10):808–823.

Atkinson, A. B. and Stiglitz, J. (1980). *Public Goods and Publicly Provided Private Goods*. McGraw-Hill.

Bass, F. (1969). A New Product Growth for Model Consumer Durables. *Management Science*, 15(5):215.

Centola, D. and Macy, M. (2007). Complex contagions and the weakness of long ties. *American journal of Sociology*, 113(3):702–734.

Fischbacher, U. and Gachter, S. (2010). Social preferences, beliefs, and the dynamics of free riding in public goods experiments. *American Economic Review*, 100(1):541–56.

Fischbacher, U., Gächter, S., and Fehr, E. (2001). Are people conditionally cooperative? evidence from a public goods experiment. *Economics Letters*, 71(3):397–404.

Galeotti, A., Goyal, S., Jackson, M. O., Vega-Redondo, F., and Yariv, L. (2010). Network games. *The review of economic studies*, 77(1):218–244.

Godinho de Matos, M., Ferreira, P., and Krackhardt, D. (2014). Peer influence in the diffusion of the iphone 3g over a large social network. *MIS Quarterly*.

Goldenberg, J., Libai, B., and Muller, E. (2010). The chilling effects of network externalities. *International Journal of Research in Marketing*, 27(1):4–15.

Granovetter, M. (1978). Threshold models of collective behavior. *American journal of sociology*, pages 1420–1443.

Han, Q. and Ferreira, P. (2014). The role of peer influence in churn in wireless networks. In *Proceedings of the 2014 International Conference on Social Computing*, page 11. ACM.

Hirshleifer, J. (1983). From weakest-link to best-shot: The voluntary provision of public goods. *Public Choice*, 41(3):371–386.

Jackson, M. and Zenou, Y. (2012). Games on networks. *Handbook of Game Theory*, 4.

Janssen, M. and Ahn, T. (2006). Learning, signaling, and social preferences in public-good games. *Ecology and Society*, 11(2).

Katz, E. and Paul, F. (1955). Lazarsfeld (1955), personal influence. *The Part Played by People in the Flow of Mass Communication. New York*.

Kermack, W. and McKendrick, A. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A*, 115(772):700–721.

Leenders, R. T. A. (2002). Modeling social influence through network autocorrelation: constructing the weight matrix. *Social networks*, 24(1):21–47.

Ma, L., Krishnan, R., and Montgomery, A. L. (2015). Latent homophily or social influence? an empirical analysis of purchase within a social network. *Management Science*, 61(2):454–473.

McMillan, J. (1979a). Individual incentives in the supply of public inputs. *Journal of Public Economics*, 12(1):87–98.

McMillan, J. (1979b). Public Goods Equilibrium. Discussion Paper, University of Western Ontario.

Mobius, M., Phan, T., and Szeidl, A. (2015). Treasure hunt: Social learning in the field. Technical report, National Bureau of Economic Research.

Musgrave, R. (1959). *The Theory of Public Finance*. McGraw-Hill.

Nicholson, W. (1989). *Microeconomic Theory: Basic Principles and Extensions*. Dryden Press.

Nyblom, J., Borgatti, S., Roslakka, J., and Salo, M. A. (2003). Statistical analysis of network data—an application to diffusion of innovation. *Social Networks*, 25(2):175–195.

Peres, R., Muller, E., and Mahajan, V. (2010). Innovation diffusion and new product growth models: A critical review and research directions. *International Journal of Research in Marketing*, 27(2):91–106.

Samuelson, P. A. (1954). The pure theory of public expenditure. *The Review of Economics and Statistics*, 36(4):387–389.

Schneider, F. and Pommerehne, W. W. (1981). Free riding and collective action: An experiment in public microeconomics. *The Quarterly Journal of Economics*, 96(4):689–704.

Shalizi, C. and Thomas, A. (2011). Homophily and contagion are generically confounded in observational social network studies. *Sociological Methods & Research*, 40(2):211–239.

Varian, H. (1984). *Microeconomic Analysis*. W. W. Norton.

Wooldridge, J. M. (2001). *Econometric Analysis of Cross Section and Panel Data*, volume 1 of *MIT Press Books*. The MIT Press.

Appendices

A Modeling Incentives to Adopt with Positive Network Externalities

A.1 Derivations for product N

We develop a model for how peer influence affects the adoption of products that exhibit positive network effects to non-adopters. We focus on the specific case of products that provide free calls within the same carrier and look at the incentives for adopting two slightly different versions of these products. Consider two products. Product P allows for calling for free users that also adopt the same product. Product N allows for calling for free any user in the same carrier irrespective of whether the latter adopted this product. Both of them require users to pay a fixed flat fee to use them.

For sake of simplicity, assume that user i has a fixed number of friends, F_i and that the network structure does not change over time. Assume also that user i derives utility from communicating with user j, u_{ij}, independently of who has initiated the call and independently of communication with other users. Finally, and again for sake of simplicity, we assume quadratic pairwise utility in the number of calls between user i and user j, c_{ij}, and write:

$$u_i = \sum_{j \in F_i} u_{ij} = \sum_{j \in F_i} [a(c_{ij} + c_{ji}) - b(c_{ij} + c_{ji})^2 - pc_{ij}] \tag{4}$$

where p represents the price per call for the user that initiates the call — our setting is "bill & keep", in which callers pay for calls and callees do not. If user i adopts product P she will pay a fixed fee, f, but will not pay for calls she initiates to other users that have also adopted this product. Therefore, in this case we have:

$$u_i|A_P = \sum_{j \in F_i} [a(c_{ij} + c_{ji}) - b(c_{ij} + c_{ji})^2 - pc_{ij} 1\{u_j|\overline{A_P} \geq u_j|A_P\}] - f \tag{5}$$

where $u_i|A_P$ and $u_i|\overline{A_P}$ represent the utility derived by user i from adopting and not adopting product P, respectively, and $1\{u_j|\overline{A_P} \geq u_j|A_P\}$ takes the value of 1 if user j does not adopt product P. Below we show that maximizing utility leads user i to adopt product P iff:

$$d_i^{A_P} \geq \frac{4b}{ap} f$$

where $d_i^{A_P}$ represents the number of friends of user i who will eventually adopt product P. User i will adopt product P if there is a minimum number of friends that will also adopt it. This finding is consistent with a positive effect of peer influence on adoption: the more friends who adopt product P, the higher the likelihood of adoption.

The incentives to adopt product N are quite different. In this case, when user i adopts product N she can call all her friends within the same carrier for free. Therefore, in this case, we have:

$$u_i|A_N = \sum_{j \in F_i} [a(c_{ij} + c_{ji}) - b(c_{ij} + c_{ji})^2] - f$$

where $u_i|A_N$ represents the utility derived by user i from adopting product N. Below, we show that, in this case, user i adopts product N as long as there is at least a minimum number of friends that will not adopt product N:

$$d_i^{\overline{A_N}} \geq \frac{4b}{ap} f$$

where $d_i^{\overline{A_N}}$ represents the number of friends of user i who will not eventually adopt product N. Note that both types of products exhibit positive network effects — users derive more utility from the product when more friends adopt the product — the difference being that for product P these effects are realized only upon adoption. In contrast, for product N, such positive network effects are realized even before adoption, decreasing the adoption incentive. Therefore, very similar types of products, namely both with positive network effects, can generate very different adoption incentives, translating into very different diffusion outcomes. These two types of products offer distinct adoption incentives to users who have not adopted them. The more users that adopt product P, the higher the incentive to adopt it because the number of friends one can call for free increases. However, product N offers the opposite incentives for adoption: the more users that adopt it, the smaller the incentives to adopt it, given that adopters can already call their friends for free. Finally, we acknowledge a nuisance because users who did not purchase product N cannot initiate calls to their friends for free. Instead, it is their friends who have already adopted product N that can call them for free. In practice, what happens is that users ring their friends, the latter disconnect and call back for free.

We now detail the model for a product of type N. Derivations for a product of type P are included at the end of this appendix. Consider users i and j. Let c_{ij} and c_{ji} represent the number of calls placed by user i to user j and vice- versa, respectively. Assume that the utility of calls is quadratic in the number of calls and that calls received and placed contribute equally to utility. A user that does not adopt product N pays for the calls placed. Assuming pairwise additive utility, the utility of user i when she does not adopt the product, represented by $u_i|\overline{A_N}$, is given by:

$$u_i|\overline{A_N} = \sum_{j \in F_i} u_{ij}|\overline{A_N} = \sum_{j \in F_i} [a(c_{ij} + c_{ji}) - b(c_{ij} + c_{ji})^2 - pc_{ij}]$$

where F_i is her set of friends, $u_{ij}|\overline{A_N}$ is the utility from communicating with user j and p is the price per call. Both a and b are positive. If she adopts product N then her utility, represented by $u_i|A_N$ becomes:

$$u_i|A_N = \sum_{j \in F_i} u_{ij}|A_N - f = \sum_{j \in F_i} [a(c_{ij} + c_{ji}) - b(c_{ij} + c_{ji})^2] - f$$

where f represents the flat fee paid to acquire the product and $u_{ij}|A_N$ is the utility from communicating with user j. Consider now that user i does not adopt product N and maximizes her utility. The corresponding FOC yields,

$$\partial u_i|\overline{A_N} / \partial c_{ij} = a - 2b(c^*_{ij} + c_{ji}) - p = 0$$

Consider that instead user i adopts product N and maximizes her utility. In this case, the FOC yields:

$$\partial u_i | A_N / \partial c_{ij} = a - 2b(c^*_{ij} + c_{ji}) = 0$$

Consequently, the following cases arise:

i) if users i and j adopt then $\partial u_i | A_N / \partial c_{ij} = a - 2b(c^*_{ij} + c^*_{ji}) = 0$ and $\partial u_j | A_N / \partial c_{ji} = a - 2b(c^*_{ji} + c^*_{ij}) = 0$, which yields $c^*_{ij} + c^*_{ji} = a/2b$. By symmetry, we assume $c^*_{ij} = c^*_{ji} = a/4b$. This leads to $u^*_{ij} | A_N = u^*_{ji} | A_N = a^2/4b$;

ii) if user i adopts but user j does not then $c^*_{ij} = a/2b$ and $c^*_{ji} = 0$ which yields $u^*_{ij} | A_N = u^*_{ji} | \overline{A_N} = a^2/4b$. To see this note that if user i wants to deviate to $c^*_{ij} = a/2b + \epsilon$, with $\epsilon > -a/2b$ then $u^*_{ij} | A_N$ reduces to $a^2/4b - b\epsilon^2$; also, if user j wants to deviate to $c^*_{ji} = \epsilon > 0$ then $u^*_{ji} | \overline{A_N}$ reduces to $a^2/4b - b\epsilon^2 - p\epsilon$;

iii) if user i does not adopt but user j does then, by symmetry from case ii) above, $c^*_{ij} = 0$ and $c^*_{ji} = a/2b$, which yields $u^*_{ij} | \overline{A_N} = u^*_{ji} | A_N = a^2/4b$;

iv) if users i and j do not adopt then $\partial u_i | \overline{A_N} / \partial c_{ij} = a - 2b(c^*_{ij} + c^*_{ji}) - p = 0$ and $\partial u_j | \overline{A_N} / \partial c_{ji} = a - 2b(c^*_{ji} + c^*_{ij}) - p = 0$, which yields $c^*_{ij} + c^*_{ji} = (a-p)/2b$. Again, by symmetry, we assume $c^*_{ij} = c^*_{ji} = (a-p)/4b$. This leads to $u^*_{ij} | \overline{A_N} = u^*_{ji} | \overline{A_N} = (a^2 - ap)/4b$;

Note that cases ii) and iii) above embody the intuition referred to in the main text that when a friend adopts a product of type N, the ego will not call her. Instead, it is the friend who calls the ego (for free). Assume now that $d_i^{A_N}$ and $d_i^{\overline{A_N}}$ represent the number of friends of user i that adopt and do not adopt product N, respectively. Then, if user i adopts then $u_i^*|A_N = a^2/4b(d_i^{A_N} + d_i^{\overline{A_N}}) - f$. If, however, user i does not adopt then $u_i^*|\overline{A_N} = (a^2/4b)d_i^{A_N} + ((a^2 - ap)/4b)d_i^{\overline{A_N}}$. Therefore, user i adopts product N iff $d_i^{\overline{A_N}} > 4bf/ap$. Thus user i adopts this product if a minimum number of friends do not adopt. As expected, this threshold increases with the fee paid for product N and decreases with the price of calls.

Finally, we discuss what happens to the firm's profits with products of type N. If product N is unavailable, then user i and user j each place $(a - p)/4b$ calls. The utility of each user is $a^2/4b - (a-p)p/4b$, and the firm's profit is $(a-p)p/2b$. Suppose product N is available, and both user i and user j decide not to adopt it. In that case, each user places the same number of calls as if the product was unavailable, which results in the same utility for each of them and the same profit for the firm. If one of the users adopts the product, she places $a/2b$ calls, and her utility is given by $a^2/4b - f$. The other user places no calls and enjoys utility $a^2/4b$. The profit of the firm is f. If the user who adopts the product did not adopt, her utility would change to $a^2/4b - (a-p)p/4b$. Likewise, the utility of the other user would also change to this amount. Therefore, the user who adopts the product would only do so if $f < (a - p)p/4b$, and consequently, the firm's profit, f, is less than half of what it would be without product N in the market. This case arises when users coordinate on who adopts the product, which reduces the firm's profit significantly.

Now, consider the case with over-adoption, i.e., both users adopt the product. In this case, the profit of the firm would be $2f$. However, each user would only adopt the product if her utility, given by $a^2/4b - f$ in this case, was greater than that enjoyed by each of them when both of them choose not to adopt the product (here, users anticipate each others' actions and think alike, that is, when a user considers that she does not adopt the product, then she needs to anticipate that other users will do the same for the exact same reasons), which would be given by $a^2/4b - (a-p)p/4b$. That is, this case would only arise if $f < (a-p)p/4b$, and thus the profit of the firm would be less than $2f = (a-p)p/2b$, which is precisely the profit that the firm would enjoy had it not made product N available in the market. This is the case without coordination among users. The firm enjoys twice the fee, but, as discussed above, the fee is low, reducing the firms' profit compared to the case without product N in the market.

Therefore, and in sum, the firm's profit always reduces with products of type N. Intuitively, what happens is that users are free to choose whether to subscribe to such a product. If they do, then it must be that paying the one-time fee is cheaper than paying for the calls, which reduces the firm's profits. If paying for the calls is less expensive than paying the one-time fee, users do not subscribe to the product, and the firm's profits do not change. Also intuitively, the profit of the firm is only greater with product N if there is over adoption, allowing the firm to collect extra rents on the one-time fees. However, each user adopts product N only if the fee is sufficiently low, namely, smaller than the price paid for the share of calls that the user would like to place, and these two effects cancel out. In short, one only observes "too many users" adopting the product if the one-time fee is smaller than the price associated with the calls that each user would like to place, but "too many users" adopting the product leads each of them to place a small number of calls, which, in turn, caps the one-time fee. On the other hand, if users can coordinate on who adopts product N, they can still enjoy significant utility but, on aggregate, pay fewer one-time fees reducing the firm's profit significantly.

B Estimating Peer Influence for a Product of Type P

We find one product of type P in our dataset: a tariff plan in which users pay a fixed monthly fee and place free calls to users who have adopted the same tariff plan. Our dataset includes 526,651 adopters of this product between January 2008 and June 2009. This product is slightly different from the products of type N described in the previous subsections because it has been available in the market for a much longer period of time. Figure 5)shows its adoption over time. In this figure, each time period represents a week. We can observe that, in this case, our data covers the period in which the S-shaped curve is still ramping up. Below we study the effect of peer influence in the adoption of this product. Recall that users benefit from products of type P only if they adopt along with their friends. In contrast with what happens with products of type N, we expect the coefficient of peer influence to be positive for products of type P because the more friends that adopt, the more utility one can derive from adoption.

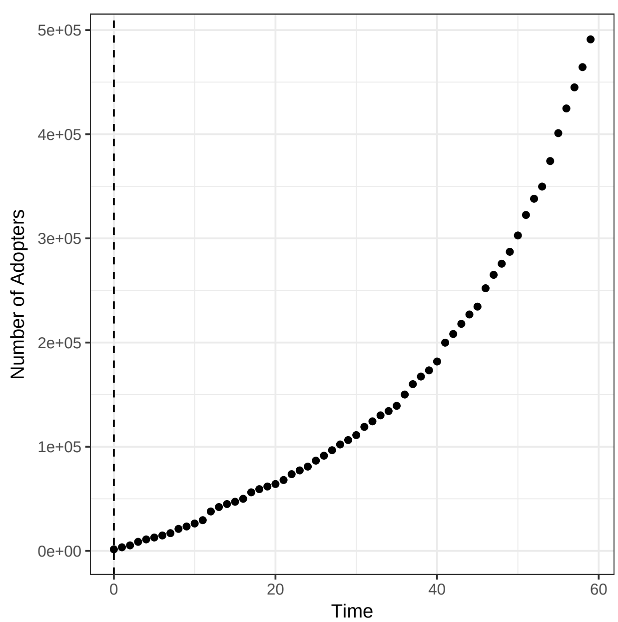

Figure 5: The adoption of a product of type P over time. The dashed line represents the moment when the product became active.

Table 4 shows the summary of the results obtained for this product. The first two rows use the data from August 2008 (the month before this product was available) to build the social network and a threshold of 3 calls to add an edge to this network. The second of these rows requires two-way communication, likewise, for the third and fourth rows concerning a threshold of 5 calls to add an edge to the social network. Finally, the fifth row in this table uses all the data available to us to define the social network and a threshold of 30 calls to add an edge between them. In line with our expectation, the coefficients associated with exposure are positive both before and after controlling for final exposure in all cases shown in this table. These results provide evidence that controlling for final exposure is not introducing a flawed downward bias on the coefficient of exposure that would always render it negative. Instead, we observe that controlling for final exposure maintains the positive coefficient on exposure in a product for which such a coefficient should indeed be positive. Just as before for the case of product N(1), the effect of peer influence that we measure with these regressions might still be biased away from zero because our proxy for the likelihood of adoption might not capture all sources of heterogeneity. Yet, we can see that the magnitude of the coefficient on exposure reduces after we control for final exposure, which provides some evidence that indeed this covariate is capturing some homophily. From our results, we conclude that peer influence increased the adoption of product P by at most 41% (the largest estimate in the last column of table 4).

C Using Final Exposure as Proxy for Unobserved Homophily

This appendix shows that using final exposure as a proxy for the type of user can help control for unobserved homophily. When homophily is present, OLS estimates for the coefficient of peer influence, ρ, are biased upwards, which means that we may obtain a positive estimate for this coef-

Table 4: Adoption of product P as function of exposure with and without controlling for final exposure.

Product	Network	Adopters	Avg. Exposure	Periods	Exposure (no controls)	Extra Adoption	Exposure (controlling for final exposure)	Extra Adoption
P	200808 - 3call	$526,651$	0.142	70	0.003*** (0.00026)	170,744 (32%)	0.0018*** (0.00023)	101,330 (19%)
P	200808 - 3mutcall	$526,651$	0.104	70	0.0048*** (2.4e-05)	203,722 (39%)	0.0034*** (2.5e-05)	141,854 (27%)
P	200808 - 5call	$526,651$	0.094	70	0.0047*** (0.00015)	179,471 (34%)	0.0031*** (0.00016)	117,299 (22%)
P	200808 - 5mutcall	$526,651$	0.079	70	0.0059*** (2.3e-05)	186,359 (35%)	0.0041*** (2.3e-05)	131,565 (25%)
P	year - n_calls30	$526,651$	0.164	70	0.0045*** (0.00014)	298,740 (57%)	0.0033*** (0.00018)	216,723 (41%)

*p<0.1; **p<0.05; ***p<0.01
Standard errors clustered at the subscriber level in parentheses

ficient even when peer influence is not present. Equation (8) in Appendix E provides the expression for the bias associated with this coefficient with homophily and no peer influence. In equation (1), c_i represents the type of user, which affects her likelihood of adoption (and d_t represent time dummies). As we show in Appendix E, the estimate for ρ obtained using this model is biased when all the following conditions are met (please refer to this appendix for additional information. Below $E_A[.]$ and $E_B[.]$ represent expected values over users of type A and B, respectively):

1. there are two types of users, A and B: the share of users of type A is $0 < \delta < 1$;
2. users of different types have different propensities to adopt: $E_A[y] \neq E_B[y]$;
3. each user type prefers to connect to users of the same type, which implies $E_A[a] \neq E_B[a]$;
4. user types are unobserved by the researcher.

The first three conditions set up a scenario with homophily, which the researcher is unable to control for when the fourth condition is also met, that is, when the researcher does not observe user types. If user types were observed, then estimating this model, including an indicator for user type, would remove the bias caused by homophily. We focus on what the researcher may instead do when the four conditions above are met, and thus user types are unobserved. The strategy that we follow here is to use a proxy for user type, c_i. The proxy that we suggest is the number of friends that adopt the product after all adoption has occurred, which we denote as "final exposure". More precisely, in a panel of length T the "final exposure" of user i is given by $a_{i,T}$. The number of friends who end up adopting the product is most likely a good proxy for the type of user, as long as one assumes some homophily (provided by conditions 1-3 above). We show this formally below.

Assume that heterogeneity is solely originated from having users of two different types, A and B, in which one type has a higher likelihood of adopting. Assume, without loss of generality, that $P(adoption|A) > P(adoption|B)$, i.e., users of type A, are more likely to adopt the product. Let $N_A|A$ represent the number of friends of type A for a user of type A and let $N_B|A$ represent the number of friends of type B or a user of type A. Likewise, define $N_A|B$ and $N_B|B$. Consider the average of final exposure for a user of type A, that is, number of friends (a below, in line with our empirical specification) of a user of type A who adopts the product by time T. This is given by

$$E[a|A] = E[N_A|A] \cdot P(adoption|A) + E[N_B|A] \cdot P(adoption|B)$$

Likewise, the average of final exposure for a user of type B, that is, the number of friends of a user of type B that adopt the product by time T, is given by

$$E[a|B] = E[N_A|B] \cdot P(adoption|A) + E[N_B|B] \cdot P(adoption|B)$$

Therefore,

$$E[a|A]-E[a|B] = (E[N_A|A]-E[N_A|B])P(adoption|A)+(E[N_B|A]-E[N_B|B])P(adoption|B)$$

By homophily $N_A|A > N_A|B$ and $N_B|B < N_B|A$, that is with homophily users are more likely to be connected to users of the same type. Then, $E[a|A] > E[a|B]$ when $(E[N_A|A] - E[N_A|B]) = -(E[N_B|A] - E[N_B|B])$, given that $P(Adoption|A) > P(adoption|B)$. Rewrite this condition as $(E[N_A|A] + E[N_B|A] = E[N_A|B]) + E[N_B|B])$ or $E[N_A + N_B|A] = E[N_A + N_B|B]$, and it is now clear that this amounts only saying that users of type A and users of type B have the same expected number of friends. This assumption is satisfied in all our empirical specifications, given that we always control for the total number of friends. Therefore, we expect users of type A to have higher final exposure, that is, $E[a|A] > E[a|B]$. Consequently, and with homophily, final exposure, i.e., the total number of friends who adopt the product provides information about the user type.

Coefficient ρ in equation (1) captures the correlation between a user's adoption and her friends' adoption only up to the moment when the former adopts the product. Information about how many friends end up adopting the product has not been used elsewhere in the literature. Still, such a covariate is likely correlated with each users' propensity to adopt in a scenario with homophily, even if there is no peer influence. Thus, we believe that this covariate is a good proxy for user type, i.e., for the propensity to adopt.

We create a set of simulations to illustrate our arguments above and show this empirical approach's appropriateness. For each simulation, we create a random graph with 10,000 users. Each user has a type, either A or B, and prefers to connect to users of the same type. Also, and without loss of generality, users of type A are more likely to adopt a hypothetical product. Each users' type is determined at random. Each user has a 50% chance to become either type. Connections among users in the social graph and whether a user adopts the product are also determined at random as a function of type. Specifically, the adoption date of users that adopt the product is randomly drawn

Table 5: Descriptive statistics by type of user for one particular simulation.

Variable	Statistic	B	A	Difference	t-Value	p-Value
N. Friends	Avg	2.97	3.03	-0.06	-1.66	0.10
	SD	1.74	1.74			
	N	4,969	5,031			
Adopter	Avg	0.01	0.11	-0.10	-21.77	0
	SD	0.09	0.31			
	N	4,969	5,031			
Adoption Period	Avg	10.10	9.64	0.46	0.80	0.43
	SD	3.63	3.11			
	N	42	542			
Final Exposure	Avg	0.08	0.28	-0.20	-23.52	0
	SD	0.28	0.52			
	N	4,969	5,031			

from a Poisson distribution with mean 10. We fixed the probability of adoption for type-B users at 1%, and the probability of a connection between two users of different types at 0.01%. The probability of adoption for type-A actors was set to different values across different sets of simulations (namely 5%, 10%, and 20%). Likewise, for the probability of a connection between two users of the same type (0.05%, 0.1%, 0.2%, and 1%). We run 1,000 simulations for each combination of these two parameters. As an example, Table 5 below shows descriptive statistics by type of user for one such simulation (all simulations provide similar descriptive statistics). For this particular simulation, we set the probability of adoption for type-A users at 20% and the probability of a connection between two users of the same type at 1%. As expected, about half of the users are from type A. Both user types have a similar number of friends (~ 3), and the users that adopt do so, on average, at a similar time period (~ 10). Also, as expected, the two types of users differ in the percentage of adopters (1% vs. 11%) and how many of their friends end up adopting the product (8% vs. 28%).

For each simulation, we estimate the peer influence coefficient ρ using three different empirical specifications. The first specification estimates this parameter without controlling for heterogeneity:

$$y_{it} = \beta_0 + \alpha a_{it} + \varepsilon$$

Given that our model was constructed without peer influence, i.e., adoption and adoption dates are by construction i.i.d, any statistically significant result obtained using this specification should be interpreted as a bias. In other words, we simulate a worst-case scenario, where all correlation between one's adoption and friends' adoption is due to homophily, and we test whether the method that we suggest in our paper is indeed capable of identifying all this effect as homophily, providing us with a null estimate for the effect of peer influence. Below we refer to this specification using the label "no controls". The second specification controls for the type of user:

$$y_{it} = \beta_0 + \alpha a_{it} + c_i + \varepsilon$$

where c_i is a dummy variable identifying the user type (A or B). Recall that this is the specification that the researcher would like to run but that she is precluded from running when the user type is unobserved. Certainly, adding type to the estimation eliminates the bias in the estimate of peer influence caused by homophily. Finally, the third specification we test uses the method that we suggest in our paper, that is, it uses $final\_exposure$ as a proxy for user type:

$$y_{it} = \beta_0 + \alpha a_{it} + final\_exposure_i + \varepsilon$$

The goal of our simulations is to show that the third specification performs well, that is, using $final\_exposure$ helps control for homophily in the absence of user type, which we assess by comparing the rate at which this specification and the specification using type directly provide null results for the effect of peer influence.

Table 6 below shows the results obtained. For each combination of the probability of adoption

Table 6: Fraction of simulations that yield a p-value lower than 0.01 for the coefficient of peer influence as function of whether type is known or final exposure is used as a proxy for user type.

$P(\text{adopt}\|A)$	$P(\text{edge}_{ii})$	(no controls) $P(\text{p-val} < 0.01)$	Type $P(\text{p-val} < 0.01)$	Final Exposure $P(\text{p-val} < 0.01)$
0.050	0.0005	0.124	0.025	0
0.050	0.001	0.137	0.030	0
0.050	0.002	0.119	0.024	0
0.050	0.010	0.144	0.033	0
0.100	0.0005	0.906	0.048	0.014
0.100	0.001	0.904	0.053	0.012
0.100	0.002	0.924	0.038	0.016
0.100	0.010	0.915	0.042	0.016
0.200	0.0005	1	0.081	0.092
0.200	0.001	1	0.076	0.111
0.200	0.002	1	0.072	0.103
0.200	0.010	1	0.071	0.105

of type-A users and of the probability of a connection between two users of the same type, reported in columns 1 and 2 of this table, respectively, we show the percentage of the estimates for each specification that are statistically significant at the 1% level (recall that each combination, that is, each row in this table was simulated 1,000 times, thus instantiating different network structures). The third column, which shows this percentage for the case of the specification without controls, provides a statistically significant estimate for ρ quite often, namely when there is significant homophily (high $P(edge_{ii})$) as expected, showing the inappropriateness of estimating SAR models using OLS to measure peer influence. The fourth column shows that when the researcher controls for the type of user, the number of instances that pick up a positive statistically significant effect of peer influence reduce dramatically (to less than 10% in call cases and less than 5% in most cases). Finally, the fifth column shows this same statistic for when final exposure is used instead of user type. The results in this column show that also, in this case, the number of instances that pick up a positive, statistically significant effect of peer influence, reduce dramatically, much like in the previous column. This means that $final\_exposure$ is a good proxy for $type$ and can be effectively used to reduce the bias induced by homophily when estimating peer influence.

D Adoption curves for products of type N and P

Figure 6 shows how the diffusion of each of our products looks like in our data (solid lines), and how it would have looked like if this coefficient would have been zero (dashed lines). Our counterfactuals, obtained using the results in Table 2 and the first social network, yield an s-shaped curve for the cases of N(1) and N(2). For other products, namely N(3), N(4), and N(5), adoption starts fast and then slows down. The exception is product P, for which we cannot observe all adoption (our data is right-censored for this product), and thus the diffusion is still in an ascending stage at the end of the observation period.

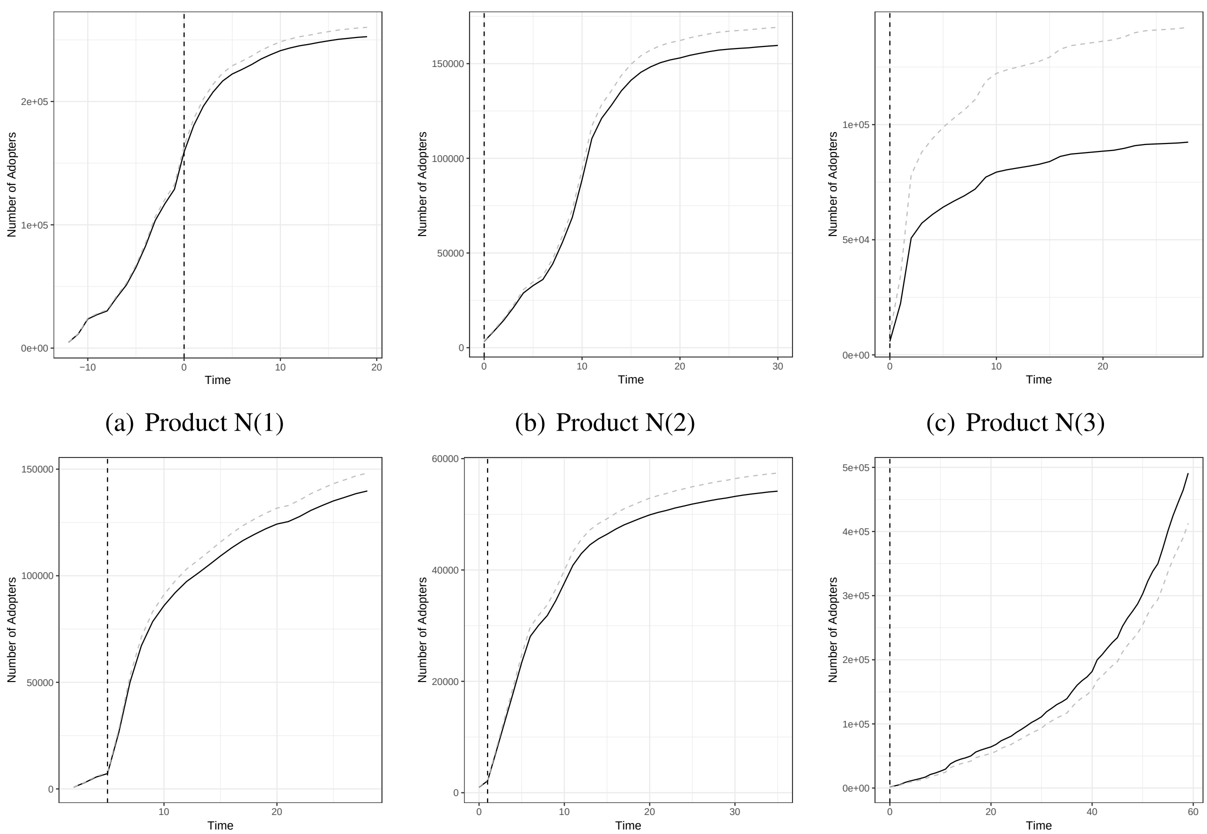

(a) Product N(1)

(b) Product N(2)

(c) Product N(3)

(d) Product N(4)

(e) Product N(5)

(f) Product P

Figure 6: Adoption curve of products N(1), ..., N(5), and P (solid line) and counterfactual in the absence of peer influence (dashed line)

E Randomization Provides a Lower Bound for the Effect of Peer Influence

This appendix shows that one can estimate a lower bound for the effect of peer influence using randomization. Assume that adoption, represented by y, depends on the number of friends that adopt, represented by a: $y = \alpha_0 + \alpha_1 a + \epsilon$. In this model, α_1 represents the effect of peer influence. With unobserved heterogeneity we have $Cov[a, \epsilon] \neq 0$. Thus, estimating this model with OLS provides a biased estimate for α_1. Assume that heterogeneity is solely originated from having users of two different types, A and B, in which one type has a higher likelihood of adopting. This means that there is no heterogeneity among users of the same type, and thus $Cov_A[a, \epsilon] = Cov_B[a, \epsilon] = 0$. Let δ represent the share of users of type A. The researcher's problem is that user types are unknown, otherwise one could correct for the unobserved homophily and obtain an unbiased measure for the effect of peer influence. Noting that $Cov_A[a, y] = Cov_A[a, \alpha_0 + \alpha_1 a + \epsilon] = \alpha_1 Var_A[a] + Cov_A[a, \epsilon] = \alpha_1 Var_A[a]$ and likewise $Cov_B[a, y] = \alpha_1 Var_B[a]$, it follows that:

$$
\begin{aligned}
\operatorname*{plim}_{N\to\infty} \hat{\alpha_1} &= \frac{Cov[a, y]}{Var[a]} \\
&= \frac{\delta Cov_A[a, y] + (1-\delta) Cov_B[a, y] + \delta(1-\delta)(E_A[a] - E_B[a])(E_A[y] - E_B[y])}{Var[a]} \\
&= \alpha_1 \frac{\delta Var_A[a] + (1-\delta) Var_B[a]}{Var[a]} + \delta(1-\delta)\frac{(E_A[a] - E_B[a])(E_A[y] - E_B[y])}{Var[a]} \qquad (6)
\end{aligned}
$$

Under the null hypothesis of no peer influence, that is, when $\alpha_1 = 0$, one obtains:

$$
\operatorname*{plim}_{N\to\infty} \hat{\alpha}_{1H,NI} = \delta(1-\delta)\frac{(E_A[a] - E_B[a])(E_A[y] - E_B[y])}{Var[a]} \qquad (7)
$$

where $\hat{\alpha}_{1H,NI}$ represents the estimator of the effect of the number of friends on adoption with homophily and without influence. The distribution of $\hat{\alpha}_{1H,NI}$ is obtained with randomization, i.e., by estimating α_1 on pseudo data sets with shuffled data. Therefore, one can consider $\hat{\alpha}_1 - \hat{\alpha}_{1H,NI}$

to estimate the effect of peer influence. One then has:

$$\underset{N\to\infty}{\text{plim}}\,(\hat{\alpha}_1 - \hat{\alpha}_{1H,NI}) \quad = \quad \alpha_1 \frac{\delta Var_A[a] + (1-\delta)Var_B[a]}{Var[a]} \tag{8}$$

Thus, $\text{plim}_{\,N\to\infty}(\hat{\alpha}_1 - \hat{\alpha}_{1H,NI})$ has always the same sign as α_1 and is always smaller in magnitude than the latter. Finally, note that this result is not specific to two types of users. The same reasoning can be applied to any number of user types, always yielding a lower bound for the effect of peer influence.

F Using Randomization to Estimate Peer Influence in the Case of Product P

In this appendix, we show results applying randomization in the case of product P. Figure 7 shows the results obtained for this product using all the data available to us to build the social network and a threshold of 30 calls to add an edge between users in the social network. This figure shows the empirical distribution of parameter ρ obtained, as before for the case of the products of type N, running the LPM model in equation (1) on 1,000 pseudo-samples with adoption dates shuffled among adopters. This distribution has a positive average, 0.0046, and a low standard deviation, 6.1e-5. Therefore, we reject the null hypothesis that $\rho = 0$, that is, peer influence and confounding factors such as homophily, result in a positive correlation between adoption and friends' adoption. The coefficient obtained using the original data is statistically higher than the average of the empirical distribution. With the original data, the coefficient obtained on exposure is 0.0052, outside the 95% confidence interval of the empirical distribution. Therefore, in this case, we conclude that peer influence increases the adoption of this product. Given that randomization provides a lower bound for the effect of peer influence, we find that, in this case, peer influence increases adoption by at least 10%.

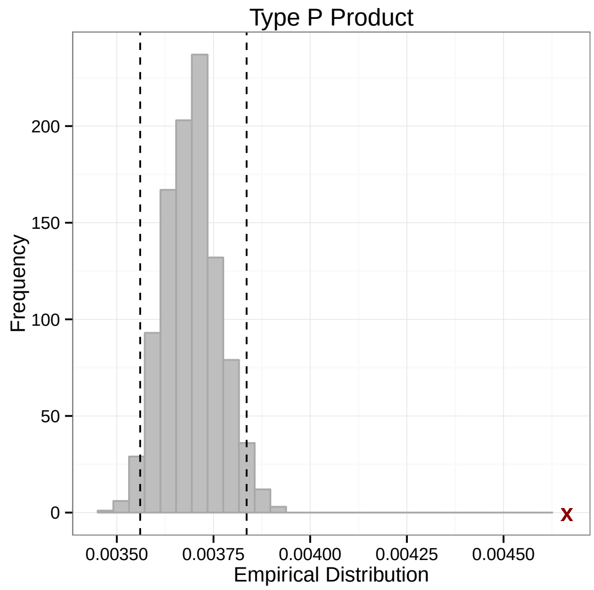

Figure 7: Distribution of coefficients on exposure for product P over 1,000 shuffles of the adoption dates.Dashed lines represent 95% confidence intervals. The '$\times$' mark represents the coefficient obtained using the original data.

Table 7 shows the results obtained for product P using other definitions for the social network, thus mimicking our analysis of the products of type N. We find similar results across these definitions. In particular, randomization recovers always a positive effect of peer influence in product P, as one could expect given the nature of the incentives for the adoption of this product. These results show that randomization is not always recovering a negative effect of peer influence. Instead, randomization recovers a negative effect of peer influence in the case of products of type N and a positive effect of peer influence in the case of product P, which comes in line with what our theoretical model predicts. Furthermore, and as before in the case of products of type N, Table 7 shows again that randomization yields smaller estimates for the effect of peer influence than those obtained without it.

Table 7: Adoption of product P as function of exposure with and without controlling for final exposure and using randomization.

Product	Network	Adopters	Avg. Exposure	Periods	Exposure (no controls)	Extra Adoption	Exposure (controlling for final exposure)	Extra Adoption	Exposure (randomization)	Extra Adoption
P	200808 - 3call	$1,126$	0.133	69	0.0039*** (0.00039)	360 (32%)	0.0024*** (4e-04)	220 (20%)	8.3e-05 (8.1e-05)	
P	200808 - 5call	$1,126$	0.089	70	0.005*** (0.00048)	308 (27%)	0.003*** (0.00051)	184 (16%)	0.0023*** (0.00013)	144 (13%)
P	year - n_calls30	$1,126$	0.164	70	0.0047*** (0.00032)	534 (47%)	0.0036*** (0.00034)	413 (37%)	0.00097*** (7.1e-05)	111 (10%)

*p<0.1; **p<0.05; ***p<0.01
Standard errors clustered at the subscriber level in parentheses