Peer Influence in Products with Network Externalities: Empirical Evidence from a Large Mobile Network

Rodrigo Belo[*], Pedro Ferreira[†]

February 28, 2018

Abstract

We study the effect of peer influence for products that exhibit positive network externalities to non-adopters. Users benefit from these products when their friends adopt even if they do not, which reduces their likelihood of adoption. However, using observational data to measure the effect of peer influence in these cases is likely to provide a positive estimate due to unobserved homophily. We suggest using the number of friends that end up adopting the product as a proxy for the unobserved user fixed effects and we show that this method recovers a negative effect of peer influence when expected. We use this method and data from a mobile service provider serving 5.7 million customers to estimate the effect of peer influence across a set of 5 different products with the feature. In all cases, we obtain negative estimates. Using randomization yields similar results providing robustness to our findings. Finally, we show that even for these products the effect of peer influence associated to the first friend that adopts the product is positive, which arises because these friends convey useful information about the product reducing uncertainty. The negative effect of peer influence arises only for subsequent friends that adopt the product. The latter are unlikely to convey new information about the product but each of them decreases the economic incentives for adoption resulting in a negative effect of peer influence.

1 Introduction

Peer influence has been shown to affect the diffusion of products and services across many different industries. Practitioners are interested in measuring its effect and in identifying conditions under which they can use it to productively increase sales. Scholars, in addition, are interested in finding the most appropriate empirical strategies to measure peer influence. Doing so has been facilitated in recent times by the pervasiveness of datasets on social interactions allowing to track consumers

[*]Rodrigo Belo, Erasmus University, rbelo@rsm.nl.
[†]Pedro Ferreira, Carnegie Mellon University, pedrof@cmu.edu.

with an unprecedented level of detail (Godinho de Matos *et al.*, 2012). However, it is still hard to econometrically identify the true effect of peer influence in social networks due to the confounding effects of unobserved covariates that may also explain adoption (Shalizi and Thomas, 2011). This paper contributes to both these goals. We analyze conditions under which, perhaps surprisingly, peer influence reduces the adoption of products that exhibit positive network externalities. We also suggest an empirical approach to measure peer influence that addresses some of the limitations associated to the traditional approaches used in the literature. Furthermore, we use this empirical approach over a very large dataset from a mobile provider and confirm that indeed the effect of peer influence on adoption may go both ways.

The seminal works by Rohlfs (1974) and Katz and Shapiro (1986) show how the number of consumers buying a product may generate positive network externalities leading to the well-known S-shaped curve. This has been empirically demonstrated in several instances such as automatic teller machines (Saloner and Shepard, 1992) and spreadsheet software (Brynjolfsson and Kemerer, 1996). A number of more recent studies on peer influence in social networks have focused on how the characteristics of the actors, e.g. influencers and influencees, affect diffusion. For example, using data from 1.3 million Facebook users, Aral and Walker (2012) show that younger users are more susceptible to influence than older users, men are more influential than women, women influence men more than they influence other women, and married individuals are the least susceptible to influence. Other studies have focused on how the topological structure of the social network and the specific placement of actors within it shapes the dynamics of diffusion. Again, Aral and Walker (2014), using the same dataset from Facebook, show that both embeddedness and tie strength increase peer influece influence. Relatively less effort has been put on understanding how specific characteristics of the products affect peer influence. In this paper, we study how the effect of peer influence changes when the products exhibit positive network externalities to non-adopters, that is, when consumers can benefit from them because their friends adopt and they do not. This allows us to contribute to address a significant gap in the literature given that all prior work to date focused only on products that do not change one's utility prior to adoption.

Products that exhibit these properties are pervasive in IT industries and include, for example, media subscriptions and communication services. For example, two or more people may use the same (shared) account to access services such as Netflix or Spotify. Even if the effect of word of mouth for these products is positive, that is, even if adopters convey a positive opinion about the product to their friends, the mere fact that they adopted the product and shared their account with non-adopters reduces the likelihood of adoption by the latter. This is an extreme case in which one single friend adopting the product is likely to dramatically reduce one's incentive to adopt. In general, some adoption by friends is still likely to increase one's likelihood of adoption, which is only likely to reduce when a significant number of friends adopt. Take the case, for example, of Skype. When first launched, premium Skype users could initiate multi-person video calls. Standard users could participate in these video calls but could not start them. Thus, they could already benefit from this feature even before adopting it. The first friend that purchases this feature is likely to essentially convey information about its usefulness potentially leading the ego to adopt. However, the more friends that adopt this feature the less likely the ego is to adopt because there are fewer new friends to video-call with whom she could not already video-call. Note that in this case peer influence may arise even if users do not explicitly exchange messages about who adopts what. The mere fact that adopters start video-calling non-adopters provides a (social) signal that may change the behavior of the latter, as discussed in detail in Peres *et al.* (2010). The core contribution of our paper is to study the role of both the informational and the economic mechanisms described above when they are likely to go in opposite directions. We do so in the context of the mobile industry by looking at tariff plans for voice calls with characteristics similar to the case illustrated above.

Our paper also studies ways to identify peer effects in social networks with observational data. The canonical model to do so is the peer effects model in which one regresses adoption on the number of friends that adopted Leenders (2002). However, this model tends to over estimate the effect of peer influence due to the confounding effect of homophily. Even an empirical strategy such as propensity score matching tends to do so given that, by definition, the effect of homophily on peer influence is positive. One way that is often used in practice to reduce this source of bias is to use fixed effects. However, this is not usually possible with panel data on adoption because users leave the panel once they adopt. In this paper, we show how using the number of friends that end up adopting the product can be used as a proxy for homophily that reduces this bias. In particular, we show that doing so yields a negative effect of peer influence in the case of products that exhibit positive network externalities for non-adopters, as one could expect given the discussion above. However, using OLS to estimate such a model would yield a positive estimate for the effect of peer influence in the case of these products, which would mislead both researchers as well as practitioners. This proxy brings new information to the analysis of peer influence that has not been used before in the literature. Therefore, we use results from multiple simulations to show that it works well across a wide range of parameters used to define the level of homophily in the social network. We also resort to randomization to estimate the effect of peer influence for this type of products Nyblom *et al.* (2003); Anagnostopoulos *et al.* (2008). Doing so confirms that indeed these products exhibit negative peer influence and thus the results that we obtain with the new proxy that we propose in this paper are not an artifact of method. We also show that using randomization with the peer effects model provides a lower bound, in magnitude, for the effect of peer influence, which can be seen as an advantage of this method vis-a-vis the ones that over estimate this effect.

Products that exhibit positive network externalities to non-adopters are often used by IT firms to attract new consumers, in particular in highly competitive settings. Showing that these products exhibit negative peer influence provides important information to practitioners, namely for production, inventory management and investment in marketing. For example, this result prevents them from over-provisioning markets that in the end are likely to register less adoption than anticipated. Showing how to measure how much peer influence reduces the adoption of this type of products provides valuable knowledge for managers that seek to optimize their operations, in particular the operations in IT firms shared by multiple products with positive network externalities that, in the end, do not benefit from peer influence in the same way.

The remainder of this paper is organized as follows. Section 2 discusses the relevant related work. Section 3 introduces our research question and hypotheses. Section 4 describes our empirical context and dataset. Section 5 introduces our empirical strategy and section 6 shows the results that we obtain. Finally, section 7 concludes.

2 Related Work

Our paper draws from the literature on product diffusion as well as from the literature on network theory. In the former, Mahajan (1985) characterizes diffusion as the dissemination of an innovation – for example, a trait, a product or a service – over a social system. The social system comprises the set of individuals, groups and institutions that might adopt the innovation. Mahajan (1985) discusses that diffusion may be triggered "by external influence, mechanisms internal to the social system, or by both". Leenders (2002) delves deeper into mechanisms internal to the social system and defines peer influence as the "dyadic process by which an individual shapes her behavior, beliefs, or attitudes according to what the other individuals in the social system think, express or how they behave". Therefore, peer influence can result from one individual observing the behavior of other individuals. An explicit message is not always exchanged between individuals for peer influence to arise. Leenders (2002) makes this case very clear: "social influence occurs when an actor adapts his behavior, attitude, or belief, to the behaviors, attitudes, or beliefs of other actors in the social system. It does not matter whether alter's influence on ego's behavior is intentional or unintentional and is not restricted to direct communication. The precondition for social influence to occur is the availability to ego of information about the attitudes or behavior of other actors." An ego observing the behavior of the alter is sufficient for the alter to provide information, or a (social) signal, to the ego and thus for peer influence to arise.

Along the same lines, Peres *et al.* (2010) provide a framework to classify the factors that drive adoption. The authors draw a clear distinction between factors that stem only from heterogeneity among individuals and factors that involve social interactions. The former includes only individual characteristics that determine whether and when adoption occurs, while the latter includes all forms of communication across individuals – peer influence, in short. In this context, diffusion is the process that leads new products and services to spread across markets, which is driven by "social influences". These influences "include all of the interdependencies among consumers that affect various market players with or without their explicit knowledge." (Peres *et al.*, 2010). We use this definition of peer influence in our paper. While other studies focus on communication among individuals over specific channels, such as messages on an online social network (Aral and Walker, 2012) or offline word-of-mouth (Mobius *et al.*, 2015). Instead, in this paper, we remain agnostic to the specific form of information transmission we focus on peer influence that may arise simply from individuals observing changes in the behavior of their friends.

Different streams of research have operationalized differently the mechanism by which individuals influence each other. For example, in epidemic models, such as the Susceptible-Infectious-Recovered (SIR) model (e.g., Kermack and McKendrick, 1927) and the Bass Model Bass (1969), new adopters are influenced by the proportion of previous adopters. In threshold models (e.g., Granovetter, 1978) adoption occurs when a given fraction of one's friends adopts the product. In hub models (e.g., Watts and Dodds, 2007), a number of well-informed central agents adopt the product leading others to adopt. All these models explore the idea that, under the right conditions, a small number of initial adopters may potentially lead to significant adoption. However, this is not always true. Watts and Dodds (2007) run a set of computer simulations to test this hypothesis finding that in most cases large cascades of peer influence were driven not by influential individuals (opinion leaders) but rather by a large number of easily influenced people. Other authors have also found that peer influence plays a limited role in adoption. For example, Goel *et al.* (2012) analyze diffusion patterns in seven online domains, such as Twitter and Yahoo. They find similarities across all domains, namely that most adoption is part of very simple cascades of only one hop, and that only a very small fraction of adoptions are associated to longer cascades. One reason why some studies provide large statistics for the effect of peer influence is that they failed to appropriately control for latent homophily, as discussed in (Shalizi and Thomas, 2011). 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 estimation strategy to distinguish homophily from influence. They conclude that the latter is responsible for at least 50% of the observed correlation. These results speak to the importance of using appropriate empirical strategies to avoid overestimating peer influence. One such strategy is to use randomization, as defined in Nyblom *et al.* (2003) and Anagnostopoulos *et al.* (2008). Randomization attempts to estimate, instead of eliminate, the effect of unobserved homophily to subtract it from the estimates that confound it with the effect of peer influence.

A parallel line of research has been looking at the effect of network structure on diffusion. The early works on diffusion considered peer influence and network effects as global phenomena that occurr in fully connected networks (e.g., Bass, 1969). More recent studies refine these assumptions and analyze diffusion in partially connected networks with local network effects (e.g., Goldenberg *et al.*, 2010). As a result, a vast literature emerged in recent times on "network games" Galeotti *et al.*

(2010); Jackson and Zenou (2012), in which outcomes depend directly and significantly on network structure. 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 in order to get more utility from when they eventually adopt. They use simulation to show how this "chilling effect" emerges. Bampo *et al.* (2008) use data from a real-world viral marketing campaign to calibrate a computer model and run simulations in different types of networks. The authors find that network structure has a critical role in spreading viral messages. In particular, scale-free networks seem to be more efficient for viral marketing campaigns than small-world or random networks because they tend to have more social hubs. In line with this finding, (Ugander *et al.*, 2012) show that contagion in the growth of Facebook is better explained by the number of components that individuals connect rather than their size.

Relatively less attention has been given to how product characteristics shape peer influence and thus affect adoption. Dou *et al.* (2013) study, from a theoretical point of view, how firms can optimize network effects by adjusting social media features, seeding and pricing strategies. They find a nuanced relationship among these variables in the cases of incomplete information. Aral and Walker (2011) offer one of the few empirical studies looking at how viral features affect diffusion. The authors explore results from a randomized field experiment run at Facebook and find that active-personalized viral messages are more effective at promoting adoption though passive-broadcast viral messages are used more often and thus responsible for a larger share of the observed adoption due to peer influence. Our paper contributes to the literature on product design by studying peer influence in a specific set of products. Namely, we study products that users benefit from if their friends adopt even if they do not. In these cases, friends borne the cost of adoption and provide a positive network externality to egos that lowers their incentive to adopt. This type of products trigger dynamics similar to those observed in "best-shot public games" (Hirshleifer, 1983), where everyone benefits from the effort of the individual that contributes the most to the public good. Hirshleifer (1983) shows that there is significant free-riding in these games leading to under-provision of the public good. To the best of our knowledge there have been no empirical studies at the individual level providing evidence of this behavior in real world settings, in particular with actual IT products. Our paper provides a first empirical example of such an effect and confirms that indeed peer influence reduces adoption in this case.

3 Research Question and Development of Hypotheses

Peer influence, as defined by Peres *et al.* (2010), is "the degree by which an action from an individual changes the behavior of someone else. This definition encompasses all types of interactions among individuals, including direct communication, i.e., word of mouth, as well indirect communication, such as social signals (adopters signal their peers about the quality of the product) and network externalities (adopters influence the utility of their peers with their own adoption)." These interactions among individuals can therefore trigger two broad types of mechanisms – informational and economic. Informational mechanisms pertain to the transmission of information about a new product, potentially leading to positive peer influence and increased adoption. Economic mechanisms pertain solely to the economic incentives associated to the decision to adopt given the information that the individual already has about who adopted and how this affects the utility that she can obtain from consuming the product.

These two types of mechanisms may entail direct communication, such as word of mouth, but also other forms of communication that do not require an explicit message to be exchanged. For example, if a given user adopts a product and because of that changes her behavior, this provides a signal to her friends, that may or may not, as a consequence, change their behavior. This is in line with the seminal definition of peer influence by Leenders (2002) who defines peer influence as the dyadic process by which people "...shape their behavior, beliefs and attitudes according to what other people in the social system think, express and do. Peer influence can be intentional or unintentional and is not limited to direct communication but, one way or another, information about the behavior and attitudes of friends needs to be available and shared."

The literature in peer influence has focused only on settings whether either only informational mechanisms are at play or both informational and economic mechanisms are present and reinforce each other, that is, both of them contribute positively for the effect of peer influence. For example, having friends that adopt a given product (e.g. a fax machine) increases both the likelihood that an individual knows about the product and the economic incentive to adopt it because the eventual adopter would benefit more from having more friends that adopted. Accordingly, the literature on peer influence focuses usually on the aggregate of these two effects instead of trying to separate them. However, the economic incentives to adopt a product not always increase with the number of friends that adopt the product, not even for all products that exhibit positive network externalities. For example, consider two types of tariff plans offered by a mobile telecommunications provider operating in a "bill & keep" regime, that is, callers pay for calls and callees do not. Both types of tariff plans allow users to initiate calls for free within the same provider, for a predefined period of time, in exchange for a fixed flat fee. One of these tariff plans, later called product P in our paper, allows calling for free users that also adopt the same tariff plan. The other tariff plan, latter called product N in our paper, allows calling for free any user (within the same provider) irrespective of whether the tariff plans that the latter adopted.

These two tariff plans offer different incentives for adoption. The more users that adopt product P the more the incentive to adopt because the number of users that one can call for free increases. However, product N offers the opposite incentives for adoption, namely the more users that adopt this product, the lower the incentive to adopt because adopters can already call their friends for free.[1] Note that both products P and N exhibit positive network externalities – users derive more

[1]We acknowledge a nuisance in the fact that users that did not adopt product N cannot initiate calls for free.

utility from the product when more friends adopt – the difference being that for product P these effects are realized only upon adoption while for product N they arise even before one adopts, which decreases the incentive for adoption. Therefore, very similar types of products, namely both with positive network externalities, can generate very different incentives for adoption, which in turn can translate into very different diffusion outcomes. In appendix A, we present a model that details the conditions under which users have lower incentives to adopt products that exhibit positive network externalities when more friends do.

This paper focuses on products that exhibit positive network externalities to non-adopters, that is, users benefit from these products when their friends adopt even if the former do not. If one uses $u_i(A_i, NFA_i)$ to represent the utility of user i as a function of whether she adopts the product (A_i) and of how many friends have adopted the product (NFA_i) then our paper studies products for which $u_i(0, x+1) \geq u_i(0, x)$ and $u_i(1, x+1) \geq u_i(1, x)$. The former condition embodies the fact that user i benefits from having more friends that adopt the product although she has not adopted. The latter condition indicates that the same is also true when she adopts. The prior literature has focused only on products with positive network externalities for adopters, that is, products for which the latter condition is true but the former is replaced by $u_i(0, x) = u_i(0, x+1)$ for all x.

Examples of the type of products that we study in this paper include IT products that can be used simultaneously by several people such as wi-fi networks. Adopters of wi-fi service may share their credentials with non-adopters. A non-adopter can therefore benefit from friends adoption even if she does not adopt. Furthermore, the more friends that adopt wi-fi service and share their credentials the more locations one can obtain service thus likely reducing one's incentive to adopt. This is a case in which an explicit message must have been exchanged between users in order to share the wi-fi credentials. Another example of a product with a similar feature was Skype Premium,

Instead, it is their friends that 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.

which allowed premium users to initiate multi-person video calls with non-premium users. Non-adopters of the premium service could still accrue benefits, that is, participate in multi-person video calls, without paying for it. Furthermore, this was so, irrespective of whether they explicitly knew that they were taking advantage of a premium feature. In other words, the economic mechanisms triggered by interactions among users may also arise in the absence of explicit messages about who had adopted what.

The informational and economic mechanisms described above are likely to contribute in different directions for the net effect of peer influence in the case of products of type N. The informational mechanisms are still likely to contribute positively for peer influence. However, the economic mechanisms are likely to contribute negatively. Consequently, the aggregate effect of peer influence may be negative for this type of products, namely when the effect of the economic mechanisms dominate. Therefore, assessing the net effect of peer influence for these products is ultimately an empirical question. Using the language introduced before, a product exhibits negative aggregate peer influence if there is x^* such that $u_i(0,x) \geq u_i(1,x)$ for all $x > x^*$, that is, the utility from not adopting the product supersedes the utility from doing so beyond a certain number of friends that adopt the product. This discussion leads to our research question. In this paper, we ask *how does accruing benefits from friends purchasing products affect one's willingness to purchase them?* Namely, *what happens to the diffusion of a product when one can benefit from friends' adoption even before one adopts?* The answers to these questions are likely to depend on the specific context studied. We expect informational mechanisms to dominate when products are complex or in the case of experience goods. Conversely, we expect informational mechanisms to become less important when products are relatively simple to understand and information about them abound. In the latter cases, the effect of the economic mechanisms is likely to dominate. This is likely to be the case for the products that we study in this paper. Therefore, our research hypothesis is:

H1: Exposure to adopters of a product that exhibits positive network externalities to non-adopters decreases the likelihood of one's adoption.

Even though the aggregate effect of peer influence may be negative for this type of products, the effect associated to the informational mechanisms may, at times, supersede the effect associated to the economic mechanism described above. For example, and in the context of the example provided before, every friend that adopts a tariff plan of type N provides information about it but also indicates, simultaneously, that she adopted this tariff plan. The former is likely to contribute positively to the effect of peer influence, as is usually the case in social networks. However, the latter is likely to contribute negatively to peer influence because it lowers the utility associated to adoption. The effect of the informational mechanism is likely large for the first friend that adopts the product (this friend conveys useful new information about the product), yet, the first friend that adopts the product is probably unlikely to reduce too much the utility to adopt because there are still too many friends to call for free with the new tariff plan. On the contrary, the more friends that adopt a tariff plan of type N the less useful the information they transmit about it and the lower the utility associated to adoption because every time a friend adopts this type of tariff plan there are fewer people that one can call upon adoption that one could not already talk to for free. Therefore, one would expect: 1) a positive estimate for the effect of peer influence associated to the first friend that adopts the product even for a product of type N, capturing the effect associated to the informational mechanism, and; 2) a negative estimate for the effect of peer influence associated to subsequent friends that adopt the product, capturing the economics associated to products of type N discussed above. Consequently, we split our research hypothesis above into two sub-hypotheses that highlight these heterogeneous effects:

H1a: Exposure to the first adopter of a product that exhibits positive network externalities to non-adopters increases the likelihood of adoption.

H1b: Subsequent exposure to adopters of a product that exhibits positive network externalities to non-adopters decreases the likelihood of adoption.

4 Context and Data

We use an anonymized panel of data with detailed information about all users of mobile service in a large network provider. There are roughly 5.7 million users active during our period of analysis. The data include Call Detailed Records (CDRs) for all calls placed by all individual users in this provider between August 2008 and June 2009. This dataset includes only individual users, that is, it does not include business clients nor call centers. CDRs include anonymized identifiers for the caller and the callee and the start time of each call. On an average day users generate about 4 million calls. Additionally, the data contain information about the users' pricing plans and supplementary services between January 2008 and June 2009. At every point in time, each user is associated with one pricing plan and possibly several supplementary services. Supplementary services are "a la carte" add-on services that users can subscribe, such as free calls on weekends or nights for a given period of time or simply voicemail. We analyze the full dataset comprising all the 5.7 million users and their calls.

During this 11-month period, the provider offered several products that allow users to place as many calls as they want to other users of the same provider. We consider two slightly different versions of these products. Product P allows calling for free users that also adopt this product. Products of type N allow calling for free any user irrespective of whether the latter adopted the same product. Using both types of products requires users to pay a fixed flat fee. We study the effect of peer influence on five products of type N, which we call $N(1), \ldots N(5)$, all of them corresponding to short-term offers that satisfy the following conditions: (1) the service was offered for a limited period time and within the 11-month period for which we have data; and (2) the service was adopted by at least 1% of the users. These criteria ensure that all adoption of these products occurred within a relatively short period of time (and therefore is completely captured in our dataset) and that there is a critical mass of adopters to analyze. For all these products of type N peer adoption generates a positive externality to the focal user even if she does not adopt the product, decreasing her own incentive to adopt.

We focus our analysis on the first of these products of type N, called $N(1)$. Later we show that the other 4 products of this type yield similar results. Product $N(1)$ became available in November 2008. Users that purchase this product pay a one time fixed flat fee and can call all users in the same provider for free during the month of December 2008 (independently of whether the latter adopted the same product). This product was adopted about 252K times between mid November and mid December 2008. Figure 1 shows the adoption of this product, which follows an S-shaped curve, as expected. Even though there is no explicit built-in mechanism to inform friends of who adopted this product, we believe that the short-term nature of its offering and the audiences that it was targeted to – namely students and young professionals – prompted adopters to mention it to their friends during their (free) calls or in person. In fact, it is very likely that adopters must have mentioned 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 their friends the reason behind the change in their behavior. In any case, we note that such an explicit message about who adopts the product is not needed for peer influence to arise in our setting. The simple fact that adopters start calling their friends more often after they adopt the product provides a (social) signal that egos may take into account when deciding whether to adopt.

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 (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

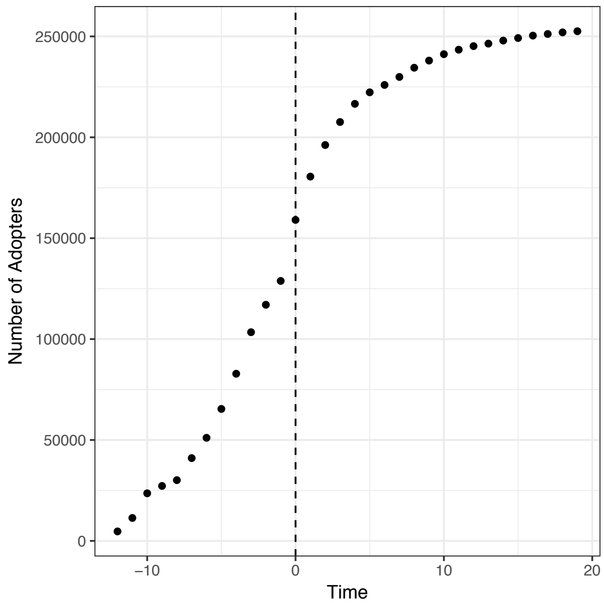

Figure 1: The adoption of product N(1) over time. The dashed line represents the moment when the product became active.

to the social network graph when two users exchange at least 3 calls during the month of November 2008 (the month preceding the activation of product N(1)). Figure 2(a) shows the distribution of the number of friends in our graph using this definition. As expected this distribution is highly skewed. The average degree in this graph is 2.90 with a standard deviation of 4.25. The most connected user in this graph connects to 94 other users. Figure 2(b) shows the that 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 that is involved in more calls is involved, on average, in 39.55 call per day during this month.

Table 1 describes what happens to the number of calls placed and received while product N(1) was active. This table shows results for the following regression:

$$y_{i,t} = \alpha + c_i + \gamma \, adopter_{i,t} + \delta \, exposure_{i,t} + c_i + \epsilon$$

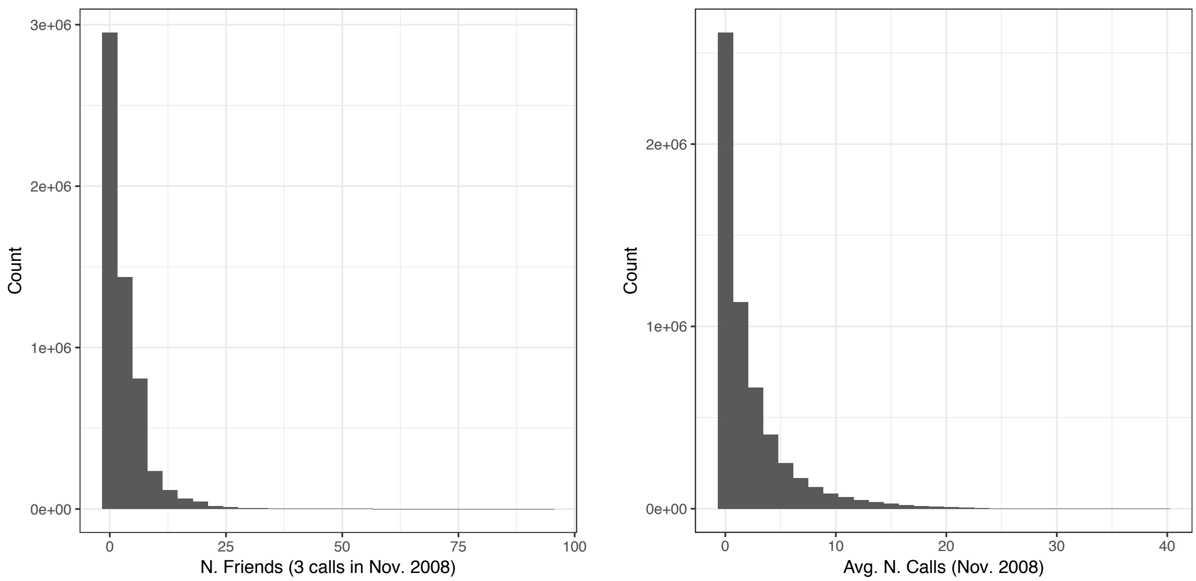

(a) Number of friends (b) Daily average number of calls

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

where i is a user, t is a day, $adopter_{it}$ indicates whether user i adopted the product by day t and $exposure_{it}$ indicates the number of friends of user i that have adopted the product by day t. The dependent variables are received calls and placed calls, labelled *calls_in* and *calls_out* below, from and to all other users, from and to all other users that adopted the product and from and to all other users that did not adopt the product in columns (1), (2) and (3), respectively. Likewise for columns (4), (5) and (6) with respect to *calls_out*. In this nomenclature, *calls_in* and *calls_out* indicate the number of calls received and the number of calls placed, respectively. These regressions are estimated using OLS with fixed effects and show that consumers behave as expected in our setting. We observe that adopters initiate more calls (column 4), and more so to non-adopters when they have no friends that adopted the product (row 1, column 6). They then initiate relatively more calls to adopters when they have friends that adopted the product (row 2, columns 5 and 6), who also call them creating a positive feedback loop. As expected, adopters also receive more calls (column 1), in particular from friends that adopted the product (columns 2 and 3). Being friends with users that adopted the product increases significantly the number of calls to other adopters compared to non-adopters (row 2, columns 2 and 3), again as part of a positive feedback loop created among adopters.

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

	Dependent variable:					
	Calls In			Calls Out		
	All	Adopted	Not Adopted	All	Adopted	Not Adopted
	(1)	(2)	(3)	(4)	(5)	(6)
adopter	0.102***	0.169***	-0.066***	0.417***	0.195***	0.221***
	(0.001)	(0.0002)	(0.001)	(0.002)	(0.0002)	(0.002)
exposure	0.042***	0.239***	-0.197***	0.018***	0.201***	-0.183***
	(0.0005)	(0.0001)	(0.0005)	(0.001)	(0.0001)	(0.001)
Observations	181,962,112	181,962,112	181,962,112	181,962,112	181,962,112	181,962,112
R^2	0.544	0.319	0.538	0.528	0.307	0.523

Note: *p<0.1; **p<0.05; ***p<0.01
Fixed Effects Estimator. Individual and Time Fixed Effects.

5 Empirical Strategy

We first study the aggregate effect of peer influence in products of type N, i.e., products that exhibit positive network externalities to non-adopters. To do so, we represent our data using an undirected graph, $G = (V, E)$, in which V is our set of users and E is a set of undirected edges. Edge e^{ij} connects user i and user 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. Additionally, we assume that the probability with which user i adopts the product at time t follows a distribution that depends on the number of friends that have already adopted the product – $exposure_{i,t} = \sum_{j:e^{ij} \in E} y_{j,t}$, and on her own observed and unobserved characteristics, $\mathbf{X}_i$ and c_i, respectively. Our interest is to measure the average effect of having one more friend that adopts the product on the likelihood of adoption. We use a Linear Probability Model (LPM) to do so because linear models provide good estimates for average marginal effects while, at the same time, rely on relatively weak assumptions (Wooldridge, 2001, pp. 453-454). This leads us to write a model of the form

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

where d_t are time dummies. We are interested in parameter ρ in this model, which measures the effect of peer influence. Equation (1) is the same formulation as the canonical peer effects model found in the prior literature Leenders (2002), which is $y_{i,t} = \alpha + \rho W y_{i,t} + \epsilon_{i,t}$ where W is an adjacency matrix coding the social network (each entry is binary and entry i, j is equal to one if i and j are friends) and thus $W y_{i,t}$ measures the number of friends of user i that adopted the product up to time t, which, above, we write more explicitly as $exposure_{i,t}$.

With homophily, similar users, or users of the same type, are likely to be connected in the social network, therefore c_i is likely correlated with the outcome variable and estimating equation 1 using OLS provides an upward-biased estimate for the effect of peer influence. If user types are observed then estimating this model including an indicator for user type removes the bias caused by homophily. Given that c_i is unobserved, the natural way to estimate equation (1) is to use fixed effects. However, fixed effects cannot be used 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 individual fixed effect captures adoption perfectly soaking up the effect of interest. Following (Wooldridge, 2001, pp. 63), 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. In settings like ours, in which one studies the diffusion of a product at the individual level over an observed network of users, one such proxy can be the number of friends that end up adopting the product, denoted by $final\_exposure_i$ in our paper. Final exposure is a good indicator for one's likelihood of adoption due to homophily. As shown in Appendix B, using such a proxy does indeed help reducing homophily bias. Finally, we note even an imperfect proxy for c_i might be very useful in our specific setting. As long as one believes that this proxy is positively correlated with one's likelihood of adoption — and this must be the case due to homophily in social networks — a negative estimate for the effect of friends' adoption is enough to show that the latter reduces diffusion. Therefore, we also estimate:

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

Furthermore, we separate the informational component from the economic component of the aggregate effect of peer influence by estimating the effect of peer influence associated to the first friend that adopts the product separately from the effect associated to the subsequent friends that adopt the product still controlling for final exposure. As discussed before in section 3, the rationale behind this model is that the informational component of peer influence must be strongest for the first friend that adopts the product, and then decrease for each subsequent friend that adopts, and vice-versa for the economic component of peer influence. Therefore, we also estimate the following model:

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

where $first\_exposure_{it}$ is an indicator of whether the focal user is exposed to at least one friend that adopted the product up to (and including) time t, and $additional\_exposure_{it}$ is the number of additional friends that have adopted up to (and including) time t. In what follows we also estimate equations (2) and (3) for product N(1).

6 Results

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

We start by estimating the aggregate effect of peer influence for product N(1). Table 2 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 build the social network to analyze the effect of peer influence in the adoption of product N(1) because this product was available in 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 before controlling for final exposure. The coefficient obtained for exposure is positive indicating a positive correlation between exposure and adoption. As discussed before, this positive correlation may be the result of unobserved homophily. Therefore, we run the same regression including final exposure as a proxy for the likelihood of adoption to control for such potential source 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 given the discussion in section 3. 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 an individual 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 2) per day results in about 7,500 fewer adoptions, or 3%, of the total adoption, that is, without the negative effect of peer influence the diffusion of this product would have been closer to 7%. This result comes in line with hypothesis *H1* that states that the likelihood of adoption decreases with exposure for this type of product. If one believes that our proxy for the likelihood of adoption is not perfect, and thus this estimate for the effect of peer influence on adoption is still confounded with some homophily, then the observed coefficient is an upper bound for the true effect of peer influence because such homophily biases this coefficient upwards. Therefore, finding a negative coefficient on exposure after using this proxy to control for the likelihood of adoption shows that the likelihood of adoption decreases with the number of friends that adopt this product.

Next we look at the results obtained estimating separately the informational and economic components of peer influence. Columns (3) and (4) show results from regressing adoption on $first\_exposure$ and $additional\_exposure$ using time dummies and controlling for the number of friends before and after controlling for final exposure, respectively. As with the aggregate effect, both coefficients are positive before controlling for final exposure because, just as before, this regression suffers severely from endogeneity. After controlling for final exposure, the coefficient associated with the informational component ($first\_exposure$) remains positive, while the coefficient associated to the economic component ($additional\_exposure$) becomes negative. This comes in line with what one would expect after our discussion in section 3. The first exposure has a positive effect on adoption, associated to learning about the product. Subsequent exposure has a negative effect on adoption, because additional exposure to adopters contributes less and less to gather information about the product but reduces more and more the economic incentive to adopt the product because there are fewer and fewer friends to call for free that have not yet adopted the product. The coefficient of 0.002 (row 2, column 4 in table 2) means that in our network, being connected to an adopter is positively associated with an increase in likelihood of adoption of 0.2 percentage points, but each additional adopter decreases the likelihood of adoption by 0.1 percentage points (row 3, column 4 in table 2).

Table 2: Peer Influence for product N(1) as a function of overall exposure, first exposure, additional exposure and final exposure.

	Dependent variable:			
	Adoption			
	(1)	(2)	(3)	(4)
Exposure	0.002^{***}	-0.0003^{***}		
	(0.0003)	(0.0001)		
First Exposure			0.004^{***}	0.002^{***}
			(0.00004)	(0.0002)
Additional Exposure			0.001^{***}	-0.001^{***}
			(0.0002)	(0.0001)
Final Exposure		0.003^{***}		0.003^{***}
		(0.0003)		(0.0003)
N. Friends	-0.0001^{***}	-0.0002^{***}	-0.00004^{***}	-0.0001^{***}
	(0.00001)	(0.00002)	(0.00000)	(0.00002)
Observations	179,174,082	179,174,082	179,174,082	179,174,082

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

6.2 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 for the other four products of type N in our dataset and we do so using four additional alternative definitions of friendship. Table 3 summarizes the results obtained from running regressions similar to those presented in columns (1) and (2) of table 2. For each product we show results using five network definitions. Two definitions require a minimum number of calls exchanged 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 5 calls. These definitions for friendship rely on recency of contact. They 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. These definitions may capture recent friendships better because they require two-way communication. They are called $3mutcall$ and $5mutcall$ below. Finally, we also use a threshold of 30 calls exchanged over the entire 11 month period. Using this definition for friend is likely to capture enduring friendships 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, which are more likely to rely on recency rather than on long-standing relationships. This fifth criteria to define friend is labelled $n\_calls30$ below.

The first line of Table 3 recovers the results represents incolumns (1) and (2) of Table 2. A total of 252,574 users adopted product N(1) (out of the 5.7 million in our dataset) 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 3. Namely, the coefficient on exposure is always positive before including our proxy for the likelihood of adoption and it becomes negative after doing so. This speaks to the danger associated to using plain peer effects models to measure peer influence and provides evidence that our findings are robust to a set of products of type N and to other definitions of friendship. This shows the powerful role that controlling for final exposure brings to improve the quality one's estimates of the effect of peer influence on product adoption. Finally, in Appendix C we show results using a Probit specification, which come all in line with the ones in table 3 obtained using LPM thus showing that our findings are not an artifact of functional form.

6.3 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 can place free calls to users that 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 (3) 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, which is different from products of type N. In contrast with what happens with the latter, we expect the coefficient of peer influence to be positive in this case P because the more friends that adopt the product the more utility one can derive from adopting it.

Table 3: Adoption of 5 products of type N as a function of exposure with and without controlling for final exposure using 5 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

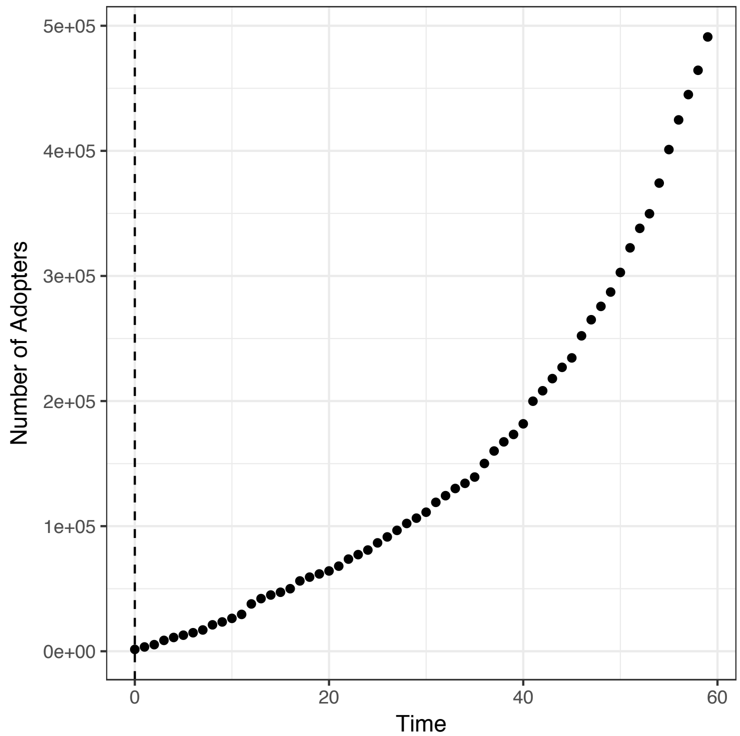

Figure 3: The adoption of a product of type P over time. The dashed line represents the moment when the product became active.

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-away communication. Likewise for the third and fourth rows with respect to 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 the case of this product, and in line with our expectation, the coefficient associated to exposure is 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 the case of a product for which such a coefficient should indeed be positive. Just as before for the case of product N(1),

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

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).

6.4 Using Randomization to Measure Peer Influence

An alternative approach to measure 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 on comparing the average of this distribution to the estimate obtained for this statistic with the original data. In appendix D we show that under mild conditions this procedure yields a lower bound, in absolute terms, for the magnitude of peer influence when the latter is measured using a peer effects model. This is an advantage vis-a-vis, for example, propensity score matching, which biases upwards the estimates of peer influence given that, by definition, homophily is positive.

Under the null hypothesis of no peer influence the probability of adoption of any user 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 a 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 (ρ), which captures the positive correlation between one's adoption and the number of friends that have already adopted the product arising from similarity in unobserved characteristics that friends and focal users share 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, it is desirable that these worlds 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 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[2]. This additional restriction further guards against potential concerns associated to 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

[2]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).

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 and every time period.

We compute the empirical distribution of ρ across our pseudo-worlds by running the model in equation (1) for each randomized version of the data. For each product of type N and for 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 sub-sample of 10,000 users[3]. The data from November 2008 and 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 with adoption dates shuffled among adopters as well as the parameter obtained using the original data. We observe that both the latter and the average of the empirical distribution are positive. The average of this distribution is 0.0043 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 0.0035,

[3] We use a subsample 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

outside the 95% confidence interval around the average of the empirical distribution. This means that without peer influence, the coefficient associated to the role of friends' adoption is higher than the coefficient obtained using the original data, and thus – after using randomization to account for unobserved effects such as homophily – we find that peer influence reduces adoption in the case of product N(1). In this case, peer influence reduced adoption by at least 6%. Appendix D shows that using randomization to measure 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 before without randomization (which was 31%).

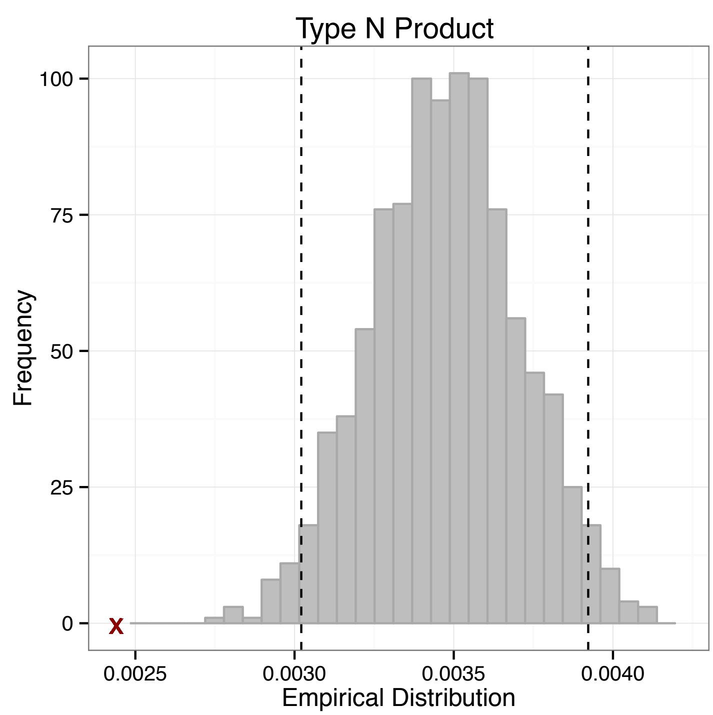

Figure 4: Distribution of the coefficient on exposure for product N(1) over 1,000 shuffles of the adoption dates.Dashed lines represent 95% confidence intervals. The '×' mark represents the coefficient obtained using the original data.

The last two columns in Table 5 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

Table 5: 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

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, providers 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 increase adoption in the case of products of type N and that randomization captures the expected direction of peer influence in these cases. Finally, and for sake of completeness, Appendix D shows that randomization in our setting provides the expected results for the case of product P.

6.5 Analyzing Alternative Hypotheses

The previous subsections provide robust empirical evidence that products of type N exhibit negative peer influence. One reason that may lead to this finding is budget coordination. For example, family members may coordinate on who pays the flat fixed fee to adopt a product of type N and then the designated buyer calld everyone else in the family for free. This is not an ideal scenario because two family members that do not purchase the product cannot still talk to each other directly for free. Yet, one can posit that this sort of behavior might arise with this type of product. Likewise, it may also arise among close friends. We investigate whether this behavior arises in our setting for the case of product N(1) by looking at what happens across highly connected clusters of users in our social graph (we use the social graph where we add edges if at least 3 calls are exchanged for this analysis). Specifically, we collect all maximal cliques with 3 to 6 members in this graph and define each clique as a "family". If a user belongs to more than one clique, we assign her to the clique they exchange most calls with. We get a total of 590,000 "families" and a total of 1.4 million users assigned to "families". We then use information on whether someone in the "family" adopted the product to a create a "family member adopter" dummy variable, represented below by $fam\_adopted$, which we include in our regressions. Table 6 shows the results for product N(1), which are qualitatively similar to the ones reported before. The exposure coefficient remains negative even after controlling for whether a family member adopted the product. This means that even if budget coordination within a family is at play, it does not explain the observed negative peer influence. Moreover, we see a positive coefficient in this new dummy, which means that, on average, budget coordination does not seem to happen. Finally, we show in appendix F that the coefficient on exposure remains negative for all products of type N even after controlling for whether a family member adopted, giving us confidence that the observed negative effect is not caused by within-family budget coordination.

Table 6: Adoption of product N(1) as a function of exposure, final exposure and of whether there is a “family” member that adopted the product.

	Dependent variable:			
	adopter			
	(1)	(2)	(3)	(4)
exposure	0.002^{***}	-0.0003^{***}	0.001^{***}	-0.001^{***}
	(0.0003)	(0.0001)	(0.0002)	(0.0001)
final_exposure		0.003^{***}		0.002^{***}
		(0.0003)		(0.0003)
fam_adopted			0.008^{***}	0.007^{***}
			(0.0003)	(0.0003)
n_friends	-0.0001^{***}	-0.0002^{***}	-0.00004^{***}	-0.0001^{***}
	(0.00001)	(0.00002)	(0.00000)	(0.00002)
Observations	179,174,082	179,174,082	179,174,082	179,174,082
Note:			*p<0.1; **p<0.05; ***p<0.01	

In any case, we believe, from anecdotal evidence, that call-back behavior was a common phenomenon in the country where our data come from. Mobile users could ring their friends and disconnect before the latter pick up or before getting to voicemail. They would then wait for a call back from their friends, which could be placed for free if the latter had previously adopted a product of type N. Unfortunately, the CDR data made available to us do not include such instances and, therefore, we cannot track this behavior in our setting. Still, we note that this call-back behavior requires the non-adopter who places a “ring” call to know beforehand that the her friend had already adopted the product in order for this strategy to be successfully used. In other words, the true mechanism by which peer influence arises would need to take place before ring calls are placed.

Finally, we note that products of type P may exhibit higher coordination costs than products of type N. An adopter of a product of type N can immediately start calling her friends. This behavior, just per si, gives the latter a signal that may lower their likelihood of adoption. However, an adopter of a product of type P may have a harder time to convey to her friends that she bought this product because she cannot call everyone for free. She will not pay to call other adopters of product P but she will need to pay for the call to users that did not adopt this product. This uncertainty may hamper the diffusion of signals for products of type P compared to products of type N. However, coordination costs cannot be responsible for the negative aggregate effect of peer influence that we find in our setting. If a product has low coordination costs then, by definition, they are not confounded in our estimates of peer influence. In the absence of coordination costs, users can easily estimate how many friends they have that adopted the product. They decide whether to adopt the product based on this (precise) information, diffusion develops and the appropriate market equilibrium arises. If a product has high coordination costs then users have only noisy signals about whether their friends purchased the product, and make decisions mostly based on their prior beliefs about whether their friends will adopt. That is, peer influence is unlikely to arise because the (social) signals are noisy.

More important, the lower coordination costs associated to products of type N cannot flip the sign of our estimate of ρ in a peer effects model of the form $y(i,t) = \alpha + \rho \, exposure(i,t) + \epsilon(i,t)$. Recall that in this model, $y(i,t)$ is an indicator for whether user i adopted the product up to time t and $exposure(i,t)$ is the number of friends of user i that did so. To see that coordination costs cannot flip the sign of our estimate of ρ in this model, start by positing that high coordination costs reduce the likelihood of adoption. Then $y(i,t)$ reduces as well as $exposure(i,t)$ (recall that $exposure(i,t)$ is the sum of $y(i,t)$ across friends). This leads to a biased estimate of ρ towards zero (Wooldridge, 2001, pp. 61-62). In addition, posit that high coordination costs increase the likelihood of adoption. Then $y(i,t)$ increases as well as $exposure(i,t)$. This, again, biases the estimate of ρ towards zero but cannot change its sign (Wooldridge, 2001, pp. 61-62). In sum, coordination costs may attenuate one's ability to estimate the effect of peer influence but cannot flip the sign of such an estimate. Hence, they cannot lead us to erroneously estimate a negative effect of peer influence for products of type N in lieu of a positive one.

7 Conclusion

Studying how products and services spread across social networks has been for a long time a subject at the top of research agendas across a number of adjacent disciplines, such as information systems, economics, management and marketing. Understanding diffusion allows for predicting market size over time, which has deep implications for many managerial tasks such as adjusting production, managing inventory or investing in marketing. Significant amounts of data now come from tracking interactions in social networks at an unprecedented level of detail. This provides fertile ground to measure peer influence and learn the extent to which people rely on signals from their friends to purchase products and services. However, measuring peer influence is a hard empirical task in observational studies because it is difficult to separate it from unobserved effects such as latent homophily. In fact, the first studies measuring peer influence offered upward biased estimates because they failed to partial out unobserved confounders. More recent studies found relatively smaller effects of peer influence, correcting for endogeneity using tools such as propensity score matching and instrumental variables. In other words, the large amounts of data on social interactions that are available for research today are not, per si, sufficient to correctly tease out the effect of peer influence. New tools to separate it from unobserved effects are required to appropriate benefit from such datasets. Our paper contributes for this line of research by studying the effect of peer influence for a set of products often overlooked in the prior literature and by providing a novel empirical way to approach the problems referred above.

Our paper studies peer influence for products that yield positive network externalities to non-adopters, i.e. products that people can benefit from when their friends adopt even though they did not. Although seldom studied in the literature, this class of products is pervasive throughout the economy, and in particular in IT related markets. Examples include communication services, such as skype and bluejeans, and Internet access services, such as wi-fi. The likelihood of adoption for these products reduces with the number of friends that adopt the product because one already benefits from them when friends adopt. This is in contrast with products that provide benefits only when one adopts the product, in which case, with positive network externalities, the more friends that adopt the more likely one is to adopt. We discuss two mechanisms that are usually present in social networks that may trigger peer influence. One such mechanism pertains to the transmission of information among users, namely the users that adopt a product are likely to convey information about it to their friends reducing the uncertainty of the latter thus potentially leading to increased adoption. The other mechanism pertains to the economic incentives for adoption, that is, users make decisions conditional on the information that they have about who adopted which products and factor in how such adoptions affect their utility.

We explore a large dataset with anonymized call detailed records that was shared with us by a large mobile provider. We use these records to infer a graph of friendships across users and we study how a set of add-on products diffuses across this social network. A number of products offered by this provider exhibit the properties described above. In particular, we study a set of tariff plans that allows users to call any other user in the same provider for free. Empirically, we find a positive correlation between one's adoption and friends' adoption for all these products when we use the traditional peer effects model to measure the effect of peer influence. This estimate comes against out intuition for this type of product and cannot be interpreted as the true effect of peer influence in this case. Our paper then offers a new empirical approach to estimate this type of model that corrects the sign of this estimate. The intuition behind the method that we propose is to control for unobserved heterogeneity across users using the number of friends that end up adopting the product as a proxy for homophily. This type of proxy for homophily has never been used before in the literature. However, and as we show in our results section, it carries new and useful information to estimate peer effects, namely, it leads researchers to obtain a negative effect of peer influence for products that should exhibit such an effect. Furthermore, this method still recovers a positive effect of peer influence for products that should exhibit such a positive estimate. In addition, we resort to multiple simulation results to show that this new method behaves consistently across a large set of unobserved parameters that hamper the identification of peer influence using only the plain peer effects model.

We also provide a number of robustness checks that increase our confidence in the results that we obtain. Namely, we use randomization to measure the effect of peer influence for the case of products that should exhibit negative estimates. In all cases, we find such negative estimates. In a few cases, these estimates are not statistically significant, which comes at no surprise once we also show (in appendix) that using randomization with peer effects models yields a lower bound for the effect of peer influence. In other words, randomization provides conservative estimates for the effect of peer influence, which may be seen as an advantage vis-a-vis other techniques used to measure peer influence in observational methods, such as propensity score matching, which typically over estimate the effect of peer influence leading researchers to claim peer effects when indeed they do not exist.

Our paper is the first to provide empirical evidence with real world data that indeed products with positive network externalities may also reduce adoption. Furthermore, we show how this result relates to whether one can benefit from these products without adopting them. In addition, our results are obtained in a truly large scale setting. We study the diffusion of these products in a network with 5.7 million users, which also increases our confidence in the results that we report. Still, 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 social signals. In particular, in our case, the adopters of tariff plans that allow users to call their friends for free may immediately start doing so, 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 calling friends for free would change 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 above were only available for short periods of time. This reduces significantly the chances that they might have actively changed the social network, in particular, the connections among friends that really matter for decision making. In fact, we provide empirical evidence that the network did not seem to change much after these tariff plans were available relative to how much it usually does.

References

Aris Anagnostopoulos, Ravi Kumar, and Mohammad Mahdian. 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, 2008. ACM.

S. Aral and D. Walker. Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Science*, 57(9):1623–1639, 2011.

Sinan Aral and Dylan Walker. Identifying influential and susceptible members of social networks. *Science*, 337(6092):337–341, 2012.

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

S. Aral, L. Muchnik, and A. Sundararajan. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*, 106(51):21544, 2009.

M. Bampo, M.T. Ewing, D.R. Mather, D. Stewart, and M. Wallace. The effects of the social structure of digital networks on viral marketing performance. *Information Systems Research*, 19(3):273, 2008.

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

Erik Brynjolfsson and Chris F Kemerer. Network externalities in microcomputer software: An econometric analysis of the spreadsheet market. *Management Science*, 42(12):1627–1647, 1996.

Yifan Dou, Marius F Niculescu, and DJ Wu. Engineering optimal network effects via social media features and seeding in markets for digital goods and services. *Information Systems Research*, 24(1):164–185, 2013.

Andrea Galeotti, Sanjeev Goyal, Matthew O Jackson, Fernando Vega-Redondo, and Leeat Yariv. Network games. *The review of economic studies*, 77(1):218–244, 2010.

Miguel Godinho de Matos, Pedro A Ferreira, and David Krackhardt. Peer influence and homophily in the diffusion of the iphone 3g in a very large social network. In *Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on and 2012 International Confernece on Social Computing (SocialCom)*, pages 134–143. IEEE, 2012.

S. Goel, D.J. Watts, and D.G. Goldstein. The structure of online diffusion networks. In *13th ACM Conference on Electronic Commerce*, 2012.

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

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

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

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

Matthew Jackson and Yves Zenou. Games on networks. *Handbook of Game Theory*, 4, 2012.

Michael L Katz and Carl Shapiro. Technology adoption in the presence of network externalities. *Journal of political economy*, 94(4):822–841, 1986.

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

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

Vijay Mahajan. *Innovation diffusion*. Wiley Online Library, 1985.

Markus Mobius, Tuan Phan, and Adam Szeidl. Treasure hunt: Social learning in the field. Technical report, National Bureau of Economic Research, 2015.

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

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

Jeffrey Rohlfs. A theory of interdependent demand for a communications service. *The Bell Journal of Economics and Management Science*, pages 16–37, 1974.

Garth Saloner and Andrea Shepard. Adoption of technologies with network effects: an empirical examination of the adoption of automated teller machines. Technical report, National Bureau of Economic Research, 1992.

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

Johan Ugander, Lars Backstrom, Cameron Marlow, and Jon Kleinberg. Structural diversity in social contagion. *Proceedings of the National Academy of Sciences*, 109(16):5962–5966, 2012.

D.J. Watts and P.S. Dodds. Influentials, networks, and public opinion formation. *Journal of Consumer Research*, 34(4):441–458, 2007.

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

Appendices

A Modeling Incentives to Adopt with Positive Network Externalities

We develop a model for how peer influence affects the adoption of products that exhibit positive network effects. We focus on the specific case of products that provide free-calls within the same carrier and look at the incentives for adoption of 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 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:

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

were $u_i|A_N$ 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 while that for product N such positive network effects are realized even before adoption, decreasing the incentive for adoption. Therefore, very similar types of products, namely both with positive network effects, can generate very different adoption incentives, which in turn can translate 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 in the fact that 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 come 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 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 profits of the firm with products of type N. If product N is unavailable then consumer i and consumer j each place $(a-p)/4b$ calls. The utility of each consumer is $a^2/4b-(a-p)p/4b$ and the profit of the firm is $(a-p)p/2b$. If product N is available and both consumer i and consumer j decide not to adopt it then each consumer places the same number of calls as if the product was unavailable, which results in the same utility for each of them as well as in the same profit for the firm. If one of the consumers adopts the product then she places $a/2b$ calls and her utility is given by $a^2/4b-f$. The other consumer places no calls and enjoys utility $a^2/4b$. The profit of the firm is f. If the consumer that adopts the product would not then her utility would change to $a^2/4b-(a-p)p/4b$. Likewise, the utility of the other consumer would also change to this amount. Therefore, the consumer that adopts the product would only do so if $f<(a-p)p/4b$ and consequently, the profit of the firm, f, is less than half of what it would be without product N in the market. This is the case when consumers coordinate on who adopts the product, which reduces significantly the profit of the firm.

Now, consider the case with over-adoption, that is, both consumers adopt the product. In this case, the profit of the firm would be $2f$. However, each consumer 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 consumers anticipate each others' actions and think alike, that is, when a consumer considers that she does not adopt the product then she needs to anticipate that other consumers 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 exactly the profit that the firm would enjoy had it not made product N available in the market. This is the case without coordination among consumers. 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 profit of the firm always reduces with products of type N. Intuitively, what happens is that consumers are free to choose whether to subscribe 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 otherwise, that is, if paying for the calls is cheaper than paying the one-time fee then consumers do not subscribe 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 consumer adopts product N only if the fee is sufficiently low, namely, smaller than the price paid for the share of calls that the consumer would like to place, and these two effects cancel out. In short, one only observes "too many consumers" adopting the product if the one-time fee is smaller than the price associated to the calls that each consumer would like to place but "too many consumers" adopting the product leads each of them to place a small number of calls, which, in turn, caps the one-time fee. If, on the other hand, consumers can coordinate on who adopts product N then they can still enjoy significant utility but, on aggregate, pay fewer one-time fees reducing significantly the profit of the firm.

A.1 Derivations for product P

We now derive the result for product P. Like for product N, a user that does not adopt product P pays for the calls placed. Assuming pairwise additive utility, the utility of user i when she does not adopt the product, represented by $u_i|\overline{A_P}$, is given by:

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

where F_i is her set of friends, $u_{ij}|\overline{A_P}$ 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 P then her utility, represented by $u_i|A_P$ becomes:

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

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

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

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

$$\partial u_i|A_P/\partial c_{ij} = a - 2b(c^*_{ij} + c_{ji}) - p1\{u_j|\overline{A_P} \geq u_j|A_P\} = 0$$

Consequently, the following cases arise:

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

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

Assume now that $d_i^{A_P}$ and $d_i^{\overline{A_P}}$ represent the number of friends of user i that adopt and do not adopt product N, respectively. Then, if user i adopts then $u^*_i|A_P = (a^2/4b)d_i^{A_P} + ((a^2 - ap)/4b)d_i^{\overline{A_P}} -$

f. If, however, user i does not adopt then $u_i^*|\overline{A_P} = ((a^2 - ap)/4b)(d_i^{A_P} + d_i^{\overline{A_P}})$. Therefore, user i adopts product N iff $d_i^{A_P} > 4bf/ap$. Thus user i adopts this product if a minimum number of friends also adopt. As expected, this threshold increases with the fee paid for product N and decreases with the price of calls.

B Comparing the Proxy Method to Spatial AutoRegressive models

In this appendix we show that using final exposure as a proxy for the type of user can help controlling for unobserved homophily. When homophily is present, OLS estimates for the coefficient of peer influence, ρ, are biased upwards in SAR model. This means that we may obtain a positive estimate for this coefficient even when peer influence is not present. Equation (8) in Appendix D provides the expression for the bias associated to 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 D, the estimate for ρ in 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 setup 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 that end up adopting the product is, most likely, a good proxy for the type of user, as long as one assumes some degree of homophily (provided by conditions 1-3 above). With homophily, adopters are likely to be connected to adopters. Thus, a user connected to many adopters is also likely to be an adopter herself. 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, but such a covariate is likely to be 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 to show the appropriateness of this empirical approach. 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. In addition, 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 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

Table 7: 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			

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 7 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 in 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 an extreme, worst-case scenario, in which all correlation between one's adoption and friends' adoption is due to homophily and we test whether the method 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 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 uaer 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 8 below shows the results obtained. For each combination of the probability of adoption 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 that for each specification 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,

Table 8: 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(adopt|A)	P(edge_ii)	(no controls) P(p-val<0.01)	type P(p-val<0.01)	final_exposure P(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

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 the type of user. The results in this column show clearly 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.

C Using Probit to Measure Peer Influence

Table 9 shows that the results obtained using Probit are qualitatively similar to the ones discussed in the main body of the paper obtained using LPM.

Table 9: Adoption of products of type N as a function of exposure with and without controlling for final exposure using Probit.

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	0.58*** (0.0028)	25,925 (10%)	-0.067*** (0.0025)	-7,622 (-3%)
N(1)	200811 - 3mutcall	252,574	0.107	32	0.8*** (0.0023)	43,374 (17%)	-0.005* (0.0029)	-589 (0%)
N(1)	200811 - 5call	252,574	0.094	32	0.92*** (0.0029)	33,909 (13%)	-0.00044 (0.0029)	
N(1)	200811 - 5mutcall	252,574	0.078	32	0.96*** (0.0026)	43,644 (17%)	0.04*** (0.0034)	3,925 (2%)
N(1)	year - n_calls30	252,574	0.147	32	0.71*** (0.0024)	37,289 (15%)	-0.02*** (0.0024)	-2,609 (-1%)
N(2)	200808 - 3call	159,619	0.113	31	0.69*** (0.0039)	16,629 (10%)	-0.035*** (0.0041)	-2,511 (-2%)
N(2)	200808 - 3mutcall	159,619	0.083	31	0.95*** (0.0031)	28,845 (18%)	-0.036*** (0.0039)	-2,634 (-2%)
N(2)	200808 - 5call	159,619	0.074	31	1*** (0.0038)	21,777 (14%)	-0.049*** (0.0044)	-2,857 (-2%)
N(2)	200808 - 5mutcall	159,619	0.062	31	1.1*** (0.0032)	29,064 (18%)	0.00012 (0.0044)	
N(2)	year - n_calls30	159,619	0.101	31	0.86*** (0.0035)	21,950 (14%)	-0.072*** (0.0037)	-4,810 (-3%)
N(3)	200808 - 3call	92,383	0.072	29	0.61*** (0.0024)	8,593 (9%)	-0.39*** (0.005)	-25,632 (-28%)
N(3)	200808 - 3mutcall	92,383	0.051	29	0.7*** (0.003)	8,436 (9%)	-0.44*** (0.0061)	-23,682 (-26%)
N(3)	200808 - 5call	92,383	0.047	29	0.76*** (0.0032)	7,746 (8%)	-0.44*** (0.0065)	-20,625 (-22%)
N(3)	200808 - 5mutcall	92,383	0.039	29	0.73*** (0.0037)	7,492 (8%)	-0.48*** (0.0071)	-20,308 (-22%)
N(3)	year - n_calls30	92,383	0.075	29	0.62*** (0.0024)	9,594 (10%)	-0.39*** (0.0049)	-28,770 (-31%)
N(4)	200808 - 3call	139,761	0.080	27	0.53*** (0.0016)	13,742 (10%)	-3.8e-05 (0.0025)	
N(4)	200808 - 3mutcall	139,761	0.057	27	0.59*** (0.002)	12,974 (9%)	0.009*** (0.003)	492 (0%)
N(4)	200808 - 5call	139,761	0.053	27	0.68*** (0.0021)	12,460 (9%)	0.011*** (0.0033)	552 (0%)
N(4)	200808 - 5mutcall	139,761	0.044	27	0.64*** (0.0025)	11,927 (9%)	0.019*** (0.0036)	841 (1%)
N(4)	year - n_calls30	139,761	0.084	27	0.54*** (0.0016)	15,509 (11%)	-4e-04 (0.0025)	
N(5)	200906 - 3call	54,147	0.035	36	1.2*** (0.0063)	5,254 (10%)	-0.23*** (0.0095)	-3,641 (-7%)
N(5)	200906 - 3mutcall	54,147	0.024	36	1.3*** (0.0081)	4,830 (9%)	-0.27*** (0.012)	-3,178 (-6%)
N(5)	200906 - 5call	54,147	0.022	36	1.4*** (0.0083)	4,270 (8%)	-0.25*** (0.013)	-2,673 (-5%)
N(5)	200906 - 5mutcall	54,147	0.018	36	1.5*** (0.0083)	4,132 (8%)	-0.27*** (0.014)	-2,486 (-5%)
N(5)	year - n_calls30	54,147	0.028	36	1.3*** (0.0071)	4,927 (9%)	-0.26*** (0.012)	-3,287 (-6%)

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

D 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:

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

Thus, 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.

E 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 5 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%.

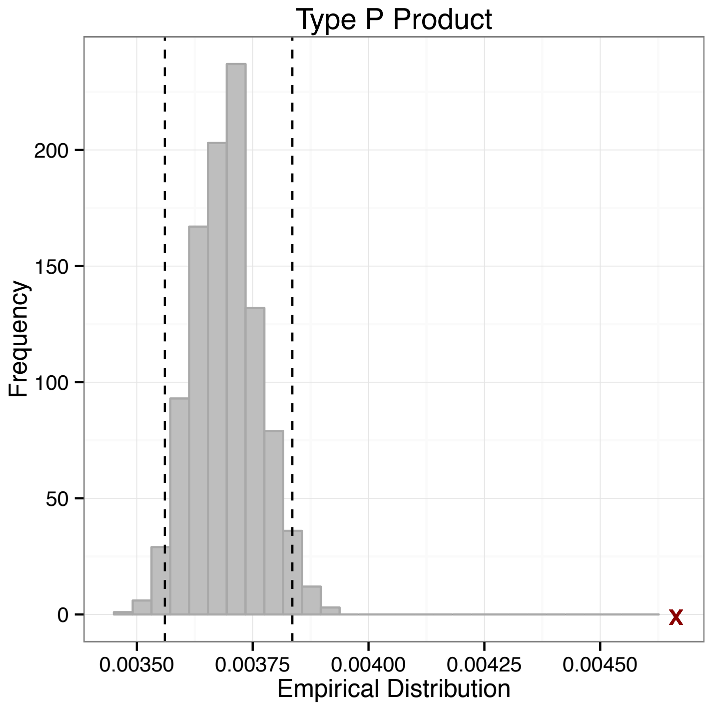

Figure 5: Distribution of coefficients on exposure for product P over 1,000 shuffles of the adoption dates.Dashed lines represent 95% confidence intervals. The '×' mark represents the coefficient obtained using the original data.

Table 10 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 the case of product P, as one could except 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 10 shows again that randomization yields smaller estimates for the effect of peer influence than those obtained without it.

Table 10: 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

F Adoption of products of type N controlling for family adoption

Table 11 shows the results obtained from regressing adoption on exposure, with and without controlling for final exposure and controlling for whether a user has a family member that adopted the product. The rationale for running this type of model is that the observed negative coefficient on exposure could be caused by family members coordinating on who adopts the product. Including a dummy for whether a family member adopted the product, allows us to control for such effects. The results in this table show that the coefficient on exposure remains negative for all the products of type N that we study in our paper even after controlling for whether a family member adopted the product. This provides additional confidence that the observed negative effect of peer influence is not caused by within-family budget coordination.

Table 11: Adoption of products of type N as a function of exposure with and without controlling for final exposure controlling and for whether a family member has adopted the product.

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	0.0012*** (2e-04)	31,821 (13%)	-0.00061*** (0.00012)	-16,534 (-7%)
N(1)	200811 - 5call	252, 574	0.094	32	0.0028*** (0.00019)	49,034 (19%)	-0.00077*** (5.9e-05)	-13,382 (-5%)
N(1)	year - n_calls30	252, 574	0.147	32	0.0022*** (0.00014)	58,911 (23%)	-0.00087*** (7e-05)	-23,683 (-9%)
N(2)	200808 - 3call	159, 619	0.113	31	0.0014*** (0.00022)	27,313 (17%)	-0.00071*** (0.00012)	-14,326 (-9%)
N(2)	200808 - 5call	159, 619	0.074	31	0.0027*** (0.00014)	36,327 (23%)	-0.00091*** (5.6e-05)	-12,095 (-8%)
N(2)	year - n_calls30	159, 619	0.101	31	0.002*** (8.1e-05)	35,211 (22%)	-0.001*** (6.9e-05)	-18,857 (-12%)
N(3)	200808 - 3call	92, 383	0.072	29	0.002*** (3.2e-05)	23,670 (26%)	-0.0042*** (0.00012)	-50,408 (-55%)
N(3)	200808 - 5call	92, 383	0.047	29	0.0024*** (2.8e-05)	18,752 (20%)	-0.0048*** (8.7e-05)	-37,752 (-41%)
N(3)	year - n_calls30	92, 383	0.075	29	0.0019*** (2.5e-05)	23,912 (26%)	-0.0043*** (6.4e-05)	-54,189 (-59%)
N(4)	200808 - 3call	139, 761	0.080	27	0.0033*** (6.5e-05)	41,164 (29%)	-0.00076*** (3.8e-05)	-9,417 (-7%)
N(4)	200808 - 5call	139, 761	0.053	27	0.0041*** (4e-05)	33,647 (24%)	-0.00083*** (4.6e-05)	-6,751 (-5%)
N(4)	year - n_calls30	139, 761	0.084	27	0.0034*** (2.8e-05)	43,770 (31%)	-0.00069*** (3.4e-05)	-9,009 (-6%)
N(5)	200906 - 3call	54, 147	0.035	36	0.0013*** (2.1e-05)	9,461 (17%)	-0.0014*** (4.8e-05)	-10,068 (-19%)
N(5)	200906 - 5call	54, 147	0.022	36	0.0016*** (2.2e-05)	7,096 (13%)	-0.0015*** (5.3e-05)	-6,776 (-13%)
N(5)	year - n_calls30	54, 147	0.028	36	0.0014*** (1.9e-05)	7,907 (15%)	-0.0015*** (4.8e-05)	-8,564 (-16%)

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