The Effect of Friends' Churn on Consumer Behavior in Mobile Networks

Pedro Ferreira
Associate Professor, Heinz College, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA, 15213, pedrof@cmu.edu; +1-412-2685526

Rahul Telang
Professor, Heinz College, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA, 15213, rtelang@cmu.edu; +1-412-2681185

Miguel Godinho de Matos
Assistant Professor, Catolica Lisbon
Palma de Cima, 1649-023 Lisboa, Portugal; miguel.godinhomatos@ucp.pt; +351-21-7270250

Acknowledgements: Miguel Godinho de Matos was supported by the Portuguese Foundation for Science and Technology [Grant UID/GES/00407/2013] and [Grant PTDC/EGE-OGE/27968/2017]. The authors thank António Moreira for excellent research assistantship and the industry partner for their support.

Pedro Ferreira

Associate Professor, Heinz College, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA, 15213, pedrof@cmu.edu; +1-412-2685526
Short bio: Pedro Ferreira's work focuses on how people use technology to consume experience goods and influence others to do so. These are inextricably linked to how firms behave and how public policies affect market structures. Ferreira's work focuses on the application of robust empirical identification methods to analyze large datasets obtained from organic in-vivo large- scale network-centric randomized experiments. His research spans two interrelated applied areas -- the impact of information and communication technologies on education and peer- influence and consumption in the media industry. The bulk of his work is on medialytics – using big data analytics to understand the future of the media industry. In addition, Ferreira has also been studying competition, consumer churn and switching costs in telecommunications.

Rahul Telang

Professor, Heinz College, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA, 15213, rtelang@cmu.edu; +1-412-2681185
Short bio: Professor Telang's research interests lie in Digital Media Industry with a particular focus on digitization of songs, movies, TV and books is affecting the incentives of content provider, content distributors as well public policy challenges in terms of innovation and copyright. In particular, he has examined the issue proliferation of distribution platforms including online piracy and its impact on traditional music, movies and books industry. Recently, he is investigating the role of social networks on music diffusion, technology adoption, and employee job search. Some of his prior work explored the challenges of interaction of multiple platforms (web portals vs telephony for customer service; SMS and voice for cellular phones). He was the recipient of Sloan Foundation Industry Study fellowship for his work in this domain and is a co-director of Digital Media Research Center at the Heinz College. His work is also funded extensively by industry participants including Google.

Miguel Godinho de Matos

Assistant Professor, Catolica Lisbon
Palma de Cima, 1649-023 Lisboa, Portugal; miguel.godinhomatos@ucp.pt; +351-21-7270250
Short bio: Miguel Godinho de Matos received a Ph.D. in Engineering and Public Policy with focus on Telecommunications Policy and Management and a M.Sc. in Engineering and Public Policy from Carnegie Mellon University. Miguel also holds a MSc and BSc in Computer Engineering from Instituto Superior Técnico in Lisbon. Miguel's research interests focus on the analysis of digitisation, social networks and peer influence on consumer behavior and consumer choice. Miguel's work has been published in top journals such as Management Science, Marketing Science, Management Information Systems Quarterly, in the Journal of Management Information Systems as well as in top peer-reviewed research conferences such as the International Conference of Information Systems and the Economics of Digitization Seminar Series of the National Bureau of Economic Research.

The Effect of Friends' Churn on Consumer Behavior in Mobile Networks

Abstract: We study how consumers decide which tariff plan to choose and whether to churn when their friends churn in the mobile industry. We develop a theoretical model showing conditions under which users remain with their carrier and conditions under which they churn when their friends do. We then use a large and rich anonymized longitudinal panel of call detailed records to characterize the consumers' path to death with unprecedented level of detail. We explore the structure of the network inferred from these data to derive instruments for friends' churn, which is typically endogenous in network settings. This allows us to econometrically identify the effect of peer influence in our setting. On average, we find that each additional friend that churns increases the monthly churn rate by 0.06%. The observed monthly churn rate across our dataset is 2.15%. We also find that firms introducing the pre-paid tariff plans that charge the same price to call users inside and outside the carrier helps retain consumers that would otherwise churn. In our setting, without this tariff plan the monthly churn rate could have been as high as 8.09%. We perform a number of robustness checks, in particular to how we define friends in the social graph, and show that our results remain unchanged. Our paper shows that the traditional definition of customer lifetime value underestimates the value of consumers, and in particular that of consumers with more friends due to the effect of contagious churn and, therefore, managers should actively take into account the structure of the social network when prioritizing whom to target during retention campaigns.

Key words and phrases: contagious churn, tariff plans, mobile industry

Introduction

Prices and switching costs have both decreased considerably in the mobile industry during the last couple of decades [17]. This trend has been propelled by the accelerated pace of innovation in wireless technologies as well as by regulatory changes. As a result, mobile markets became highly turbulent in recent times. Wireless carriers report churn rates as high as 2% per month both in the US and in Europe [43, 9]. New business opportunities arise when a quarter of the customer base changes every year. However, so do risks. As a result, churn management became a top priority for managers at wireless carriers who have now been investing significantly in churn prediction and customer retention.

One important element of user churn is the behavior of the network. In particular, industry reports and academic papers highlight how friends' behavior might influence one's behavior [4, 7, 3, 16, 24, 55, 64]. The argument goes that friends' churn encourages users to leave. One underlying mechanism for such influence is that when friends churn, they convey information about the quality of service at a competitor thus reducing uncertainty about it. This exchange of information through word of mouth may lead users to follow their friends. Another reason could be that many tariff plans have a differential pricing structure for calling inside vs. outside the network. This asymmetry in prices comes from termination charges among carriers to place calls across interconnected networks that pass along, in part, to users [50]. As friends churn, it may become economically sensible for users to also churn depending on their risk averseness. Besides churning, users may reduce their volume of calling. Of course, in response firms may offer tariff plans in which users pay similar prices to call inside and outside the carrier.

Thus, friends' churn may have a nuanced impact on a user's behavior. It can vary from reducing the amount of time spent calling inside vs. outside the carrier, to changing tariff plan, to eventually churning. No paper in the past has explored this question in this detail, in part, due to the lack of granular data. Our work closes this gap and contributes to the extant literature in several ways. First, we develop a structural model of consumer behavior in which consumers choose to keep their tariff plan, change tariff plan or churn. Second, we use Call Detailed Records (CDRs) from a large mobile operator to infer the social network across a random sample of consumers and their friends. Using these CDRs and the choices of tariff plans made by users, we then explore how changes in the structure of this network affect the behavior of users. Our analysis looks not only at eventual churn decisions but also at the changes in tariff plan choices and in the volume of calls placed. Thus, we are able to characterize the path to death (i.e. the behavior of consumers on their way to churn) with an unprecedented level of detail allowing us to develop new actionable knowledge to the firm on early warnings of churn. For example, [4] analyzes only consumption over time whereas we highlight the role played by tariff plans that charge a similar price to call inside and outside the carrier as an effective deterrent of churn. Our paper measures the value associated to offering this type of tariff plans, which has not been reported before in the literature.

It is well known that empirical identification on whether the behavior of friends affects one's behavior is hard to obtain with just observational data because homophily and reverse causality hinder our ability to do so [42, 54, 24]. Furthermore, and according to [35], homophily is a human characteristic that changes overtime rather than a static attribute. Exploring the variation in homophily over time may help separate it from peer influence. Yet, and ideally, one would run a randomized field experiment to identify peer effects [3, 7]. However, randomized experiments are not always available. They are expensive, hard to design, and difficult to carry on in-vivo in real world settings. To establish causation, researchers have used instrumental variables [59], shuffling tests [1] and propensity score matching [2] to try to identify peer effects. All these methods aim at reducing one's concern that the observed effect might be driven by unobserved covariates. In our paper, we use instrumental variables to try to identify the effect of peer influence. We use the churn of friends of friends who are not friends of an ego to instrument the churn of the friends of the ego (we use the term ego going forward to describe the user whose behavior we want to study). We provide multiple robustness checks showing that this instrumental variable is appropriate in our setup.

Using the identification strategy described above, we find that when only a few friends churn egos simply adjust their volume of calls taking into account the increased price of calling the friends that churn. When more friends churn egos change to a tariff plan that charges the same price to call inside and outside the carrier. Lastly, when a significant number of friends churn egos also churn. One contribution of our work is to study the effect of friends' churn on all these different outcome variables, thus allowing us to provide a more encompassing view of the consumers' path to churn. On average across subscribers in our dataset, every time an additional friend churns the ego's monthly likelihood of churn increases by 0.06%. Across our dataset, the monthly churn rate is 2.15%. We also find that a significant number of subscribers that would otherwise churn remain with the carrier and choose a tariff plan that charges the same price to call users inside and outside the carrier. Our results show that the introduction of this tariff plan avoided 5.94% in the monthly churn rate. Therefore, our results show clearly the value associated to tariff plans that charge the same price to call inside and outside the carrier, namely because of their effectiveness in reducing churn. Furthermore, a significant decrease in the amount of calling inside the network and a significant movement towards tariff plans that charge the same price to call inside and outside the network may provide early warnings of upcoming churn behavior (in a world without tariff plans that charge the same price to call inside and outside the carrier). These insights are likely to apply to other IT-related industries and, in general, to any industry with strong network effects, for example those with membership fees and where consumers are required to interact to complete the activity of interest (e.g. sports facilities).

Literature Review

Technical change and innovation in wireless technologies led to a significant decrease in the prices of handsets, voice communications and data transfer. The pervasiveness of Wireless Fidelity (Wi-Fi) and the massive deployment of Long-Term Evolution (LTE), both offering broadband speeds capable of supporting streaming services, allow for an unprecedented level of connectedness and mobility. In parallel, regulatory policy, such as Mobile Number Portability (MNP), which allows users to keep their phone number when changing carriers, also decreased switching costs considerably [17]. Mobile consumers can now churn more easily and do so more often. Mobile carriers in the US and Europe have reported annual churn rates as high as 25% [43, 9]. As a consequence, churn management became the number one priority for executives in mobile carriers. Managers now know that acquiring new customers is significantly more expensive than keeping current ones [8, 5, 29] and that long-term customers are more valuable [22, 25] because they generate more profits, are less sensitive to marketing campaigns from competitors, become less costly to serve and are more likely to provide positive word of mouth. Consequently, customer retention became the primary goal of most churn management strategies.

Customer retention has been shown to provide significant increases in profit [11, 20]. It typically involves two steps: 1) identify likely churners [45, 27]; and 2) offering retention incentives to top likely churners [62, 39]. The methods used to identify likely churners depend on the data available. With cross-sections, researchers have used logistic regression [44, 40, 32, 38, 46, 15], decision-trees [65, 3941, 49], neural networks [66, 57, 31, 47, 58, 6], support vector machines [33, 69, 18] and ensemble methods such as random forests [36, 67]. A comprehensive review of these models is provided in [63]. These methods aim at minimizing errors. They avoid misclassifying churners as non-churners and vice-versa. Unfortunately, they do not always agree on their findings, which depend largely on the datasets, parameters and configurations used [44, 31]. With longitudinal data researchers have used hazard models [8, 10, 52], Markov logic networks [21] and hidden Markov chains [4]. In these cases, researchers observe subscribers over time and characterize the dynamics of the "path to death" – how subscribers behave on their way to churn. In the context of pro-active churn management, top churners are then targeted with retentions incentives such as special offers, discounts and personalized messages [10, 25, 53].

The lack of integration between the two steps in customer retention referred above in a single framework to optimize profits leads to sub-optimal results [62]. For example, top churners might be easy to retain but may also yield little value [13, 14, 26]. Customer Lifetime Value (CLTV) was introduced to acknowledge and measure heterogeneity in value across customers [51, 61, 26, 23]. CLTV is the discounted value of the future stream of profits that the firm can make off the consumer. However, the most basic definition of CLTV does not reflect the fact that subscribers interact in social networks and that these interactions may yield further value to the company. Often, the CLTV of a subscriber is only associated with the profits generated from the services she uses and disregards the profits generated from increased usage of services by her friends who might have been led to use such services via peer influence, such as word of mouth. These ripple effects may be particularly important in the case of churn management because word of mouth may also increase one's propensity to churn when friends' churn. The goal of our paper is to study whether this is the case and, if so, to provide unbiased estimates for the effect of such peer influence.

However, obtaining unbiased estimates of peer influence in social networks is a challenging empirical question because reverse causality and latent homophily hinder our ability to do so using observational data [42, 54]. Several strategies have been pursued in the literature to try to identify peer influence. For example, [59] used exogenous shocks to the benefits of watching TV as instruments to identify peer influence in the adoption of video-messaging at an investment bank, [2] used dynamic propensity score matching to try to disentangle influence from homophily in the adoption of a mobile service application at Yahoo, [1] used randomization in distinct contexts to try to identify peer effects by comparing how their worlds were different from simulated worlds in which peer influence would have not played a role by design. More recently, [68] used a spatial autoregressive model to identify both the direct and indirect influence of peers in the adoption of caller ring back tones. Furthermore, a broad range of mechanisms can drive peer influence [56, 60, 37], such as information transmission, competition, conformity, network externalities and spatial proximity. These mechanisms can occur simultaneously and are difficult to disentangle empirically.

The early estimates of peer influence in social networks overestimated the effect because they failed to separate it from homophily. More recent work found smaller, yet positive, effects of peer influence. For example, [28] analyze diffusion patterns in several online domains, including 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. Also, [2] look at peer influence in Yahoo messaging and conclude that homophily is responsible for at least 50% of the observed correlation. These works speak to the importance of using appropriate empirical strategies to avoid overestimating peer influence. In our paper we are interested in measuring peer influence, namely how a user changes her behavior when her friends churn. We use the structure of the social network to derive instruments to avoid overestimating the effect of peer influence in our setting. Furthermore, we use an ordered probit model allowing us to characterize the stages that users go through before churning. This allows to present evidence that a significant number of users tend to first adjust calling volumes, then they tend to change tariff plan inside the carrier and only latter consider churning to a competitor. The understanding that consumers exhibit this nuanced behavior before they churn may provide some lead time for the firm to act in order to retain consumers.

Structural Model

Consider a set of users indexed by i. Let F_i represent the set of friends of user i. Let F_i^{in} and F_i^{out} represent the set of friends of user i inside and outside the carrier used by user i, respectively. We have $F_i^{in} \cap F_i^{out} = \emptyset$ and $F_i^{in} \cup F_i^{out} = F_i$. Friends' churn moves them from F_i^{in} to F_i^{out} and vice-versa for friends' join. Let $\theta_{i,j} \geq 0$ represent the intrinsic utility from calling user j. Let $\theta_i^{in} = \sum_{j \in F_i^{in}} \theta_{i,j}$ and $\theta_i^{out} = \sum_{j \in F_i^{out}} \theta_{i,j}$ be associated with the aggregated utility from calling friends inside and outside the carrier, respectively.

Users subscribe to a tariff plan. There are two types of tariff plans in the market. One type charges price p^{in} to call inside the carrier and price p^{out} to call outside the carrier. Typically, $p^{in} < p^{out}$ due to interconnection charges across carriers that pass, in part, to consumers. The other type of tariff plan charges the same price $\tilde{p}$ to call both inside and outside the carrier. Typically, $p^{in} < \tilde{p} < p^{out}$. Let $sp_i \in \{0,1\}$ indicate whether user i subscribers a tariff plan that charges the same price to calls inside and outside the carrier.

Let $\sum_{j \in F_i} \theta_{i,j} log(q_{i,j}) + \sum_{j \in F_i} \beta q_{ij} sp_i$ represent the utility of user i, where q_{ij} represents the number of calls placed by user i to user j, $\theta_{i,j} > 0$ and β is a (positive or negative) real number. This functional form captures well the utility derived from placing calls. The higher the $\theta_{i,j}$ the more value derived from calling user j. The more calls made to user j the higher the utility and the lower the marginal utility from calling her. Note that with this functional form $\theta_{i,j}$ is readily interpreted as the share of calls placed to user j. In this setting, parameter β measures the additional utility, potentially negative, that users derive from subscribing a tariff plan that charges the same price to call inside and outside the carrier. This parameter measures, for example, the user's aversion to risk and is related to her beliefs about what her friends will do. For example, a user who believes that a significant number of her friends will churn (in the same time period or in the future) will derive more utility from subscribing in advance a tariff plan that charges the same price to call inside and outside the carrier to mitigate the risk. In this case, we expect β to be positive and large. On the contrary, users who are more risk loving are unlikely to act upon this risk and thus subscribing in advance a tariff plan that charges the same price to call inside and outside the carrier would reduce their utility. In this case, β would be negative. Alternatively, the second term in the utility function could also be written as $(sp_i)\left[\beta_{in}\sum_{j\in F_i^{in}} q_{ij} + \beta_{out}\sum_{j\in F_i^{out}} q_{ij}\right]$. In this specification, β_{in} would be negative and capture the penalty from adopting a tariff plan that charges the same price to call inside and outside the network, thus indexed to the amount of calls placed inside the carrier, and β_{out} would be positive and capture the benefit from using such type of tariff plan, thus indexed to the amount of calls placed to outside the carrier. Yet, the former specification with only one parameter β readily allows for measuring the aggregate effect of offering a tariff plan that charges the same price to call inside and outside the carrier, and thus easier to use in order to provide a measure of the value of introducing this type of tariff plan in the market.

Let w_i represent the fixed budget allocated by user i to calls. This has been shown to change little from month to month. Users maximize their utility subject to the budget constraint $\sum_{j\in F_i} q_{ij}\left(sp_i\tilde{p} + (1 - sp_i)p_{ij}^{in/out}\right) \leq w_i$, where $p_{ij}^{in/out} = p^{in}\{j \in F_i^{in}\} + p^{out}\{j \in F_i^{out}\}$. The budget constraint limits the number of calls that users can place by taking their price into account, and in particular by taking into account that asymmetric tariff plans charge different prices for call inside and outside the carrier. This yields three distinct regimes depending on the risk profile of the user, as depicted in Figure 1:

1) risk loving: if $\beta > \beta_h$ then users will always prefer to subscribe the tariff plan that charges the same price to call inside and outside the carrier;

2) risk averse: if $\beta < \beta_m$ users will always prefer to subscriber a tariff plan that changes a lower price to call inside the carrier and will churn iff $\theta^{out}/\theta^{in} > 1$;

3) otherwise: if $\beta_m < \beta < \beta_h$ then users will switch between tariff plan types depending on θ^{out}/θ^{in}. If $\theta^{out}/\theta^{in} < t_{o/i}$ they will prefer a tariff plan that charges a lower price to call inside the carrier. If $\theta^{out}/\theta^{in} > 1/t_{o/i}$ they will prefer to churn and subscribe a tariff plan that charges a lower price to call inside the carrier. In between, they will prefer a tariff plan that charges the same price to call inside and outside the carrier.

In the above $\beta_m = (\theta_i\tilde{p}/w)log\left(\tilde{p}/\sqrt{p^{in}p^{out}}\right)$ and $\beta_h = (\theta_i\tilde{p}/w)log(\tilde{p}/p^{in})$. In addition, $t_{o/i} = \left[log(\tilde{p}/p^{in}) - (\beta w)/(\theta_i\tilde{p})\right]/[log(p^{out}/\tilde{p}) + (\beta w)/(\theta_i\tilde{p})]$. Appendix 8.1 includes the derivation of the above conditions that Figure 9 illustrates. The intuition behind this figure is the following. Users who are very risk averse will subscribe a tariff plan that charges the same price to call inside and outside the carrier at all times. On the other hand, users who are very risk loving will subscribe a tariff plan that charges a lower price to call inside the carrier at all times. The former also prefer to churn when they have many friends outside their carrier. Finally, users in between switch between tariff plan types depending on which carrier their friends are. If a significant number of friends are inside their carrier then users prefer a tariff plan that charges a lower price to call inside the carrier. If a significant number of friends are outside their carrier then users prefer to churn and subscribe this type of tariff plan with the competitor. In between, users prefer a tariff plan that charges the same price to call inside and outside the carrier. In this case they are indifferent with respect to which carrier they subscribe. Still, it is still reasonable to assume that when changing from a tariff plan that charges a lower price to call inside the carrier to a tariff plan that charges the same price to call inside and outside the carrier users will remain with their current carrier. Although pre-paid users face always little switching costs it is still harder to switch carrier than switching tariff plans within the same carrier.

XXX FIGURE 1 XXX

Context, Dataset and Descriptive Statistics

Dataset, Communications Graph and Friends

Our dataset covers 11 months of records between August 2008 and June 2009. We have the anonymized phone numbers of the caller and the callee, a timestamp for when the call is placed and the duration of the call for all calls originated or received by a subscriber of OurNet. We also have the anonymized phone numbers of the sender and the receiver for every text message sent or received by a subscriber of OurNet. We have a set of characteristics such as date of birth and gender but only for a small fraction of OurNet subscribers. We also know their tariff plan and when they changed it. Our dataset comprises roughly 3.7 billion calls and 13 billion text messages.

We use this dataset to define an undirected graph of communications across OurNet subscribers over this period of 11 months. Two subscribers are connected in this graph if one of them called or sent a text message to the other and the latter answered back with a call or a text message within the same calendar month. This procedure allows us to eliminate sporadic communications that are unlikely to proxy social proximity such as those with message bots, short numbers and call centers. Our graph has 4,986,131 subscribers and 57,069,798 undirected edges. The density of the graph, defined by $2E/\left(V(V-1)\right)$, where E is the number of edges and V is the number of subscribers, is $4.6 * 10^{-6}$. We say that two subscribers are friends inside OurNet if they are connected in this graph.

Definition of Churn

Prepaid subscribers can easily obtain SIM cards, for example from a supermarket shelf, without providing much information to OurNet. The only indication that a prepaid subscriber had churned from OurNet is prolonged absence of activity. OurNet follows the industry standard and considers that a subscriber churns when she does not place calls or sends text messages for at least 3 months in a row. We followed this definition of churn in this paper, which has also been used in several prior works, namely [34, 30]. The number of prepaid subscribers who place calls or send text messages after 3 months of inactivity is less than 1% in our dataset, therefore using a longer definition for churn would have a negligible impact on our results. If a subscriber does not place calls in months $x + 2$, $x + 3$ and $x + 4$ we consider that she decided to churn in month x. We assume that month $x + 1$ is used to wind down any remaining balance. Our panel of data covers 11 months but gets 4 months shorter due to our definition of churn. Therefore, throughout this paper, our period of analysis is 7 months, between September 2008 and March 2009.

Sampling Strategy and Number of Friends

We random sampled 10,000 subscribers from our graph. From these subscribers, 280 were discarded because their out degree was 3 standard deviations above the mean. Trimming these subscribers allows us to eliminate PBXs machines as well as to keep the size of the social graph computationally manageable for analysis. Other 120 subscribers were discarded because they joined OurNet after June 2009 and thus we have no information about who they called or text messaged. 700 subscribers were discarded because they had no friends. An additional 555 subscribers were discarded because we cannot know for sure whether they churned and if so when: 239 of them were last active before October 2008, the other 316 joined OurNet after March 2009. Finally, we drop an additional 162 subscribers who never place calls. The remaining customers place an average of 32.9 calls per month and receive an average of 40.0 calls per month (in the top quartile these statistics become 42 and 53, respectively).

After these adjustments, we were left with 8,183 subscribers in our sample, which we used throughout the analysis in this paper. We use 50,888 observations of which 58% are for 5,579 unique subscribers using a tariff plan that charges the same price to call inside and outside the carrier. The remainder of the observations pertain to 3,891 unique subscribers observed using a tariff plan that charges a lower price to call inside the carrier. The average number of friends inside the carrier is 19.23 and 32.35 for users with and without a tariff plan that charges the same price to call inside the carrier, respectively (the standard deviation is 22.05 and 27.70, respectively). The median is 11 and 24 friends, respectively. The size of the social network of a subscriber is her number of friends inside and outside the carrier. We count the number of friends outside the carrier using the same definition of friend as for inside the carrier. The average size of the social network is 25.63 and 50.54 for subscribers with and without a tariff plan that charges the same price to call inside the carrier, respectively (the standard deviation is 26.48 and 32.52 friends, respectively). The median is 16 and 32 friends, respectively. We note that according to this measure, the size of the social network does not change much over time. On average, it increases 0.10 and 0.15 friends per month, respectively. Therefore, subscribers with tariff plans that charge a lower price to call inside the carrier have significantly larger social networks.

The 8,183 subscribers in our sample are called egos. In this paper we need to characterize the churn of the friends of egos and the churn of the friends of the friends of egos. We do so by using our communications graph to snowball two waves out from the egos. The average number of friends of friends of our egos is 974 with a standard deviation of 1,152. The 8,183 subscribers in our sample have 191,996 unique friends and 2,566,283 unique friends of friends and thus, in fact, we are working with millions of subscribers in this analysis.

Calling Behavior

Figure 2 shows the average number of calling minutes per month to friends inside and outside the carrier. The average number of minutes per month is 139.95 and 8.05 inside and outside the carrier, respectively for users with a tariff plan that changes a lower price to call inside the carrier. These statistics are 20.07 and 8.23, respectively, for users with tariff plans that charge the same price to call inside and outside the carrier. Therefore, the latter exhibit significantly lower calling activity. As Figure 3 shows the distribution of airtime inside the carrier is highly skewed. In March 2009, 40% of subscribers with a tariff plan that charges a lower price to call inside the carrier call at least 1 hour per month and 10% of them call more than 3 hours per month. These statistics are 8% and 2%, respectively, for users with a tariff plan that charges the same price to call inside and outside the carrier.

XXX FIGURE 2 XXX

XXX FIGURE 3 XXX

This staggering difference between calling time inside and outside the carrier, which also characterizes the distribution of sms inside and outside the network, is related to the fact that OurNet operates under a sender-pays-all model. Under this regime, the caller pays for all calls placed. The callee does not pay for calls received. As a consequence, carriers compensate each other for calls terminated at other carriers. The price charged to the caller covers the costs to provision the call end-to-end. At the interconnection point to another network the carrier where the call originates pays a termination fee to the carrier where the call terminates so that the latter can provision the appropriate network infrastructure to terminate the call. Carriers pass part of this cost to consumers under tariff plans that charge a higher price to call outside the career. Naturally, there are no termination charges for calls within the same network.

Calling behavior is intimately related to the choice of tariff plan. 53 different tariff plans are used by subscribers in our sample, of which 39 charge the same price to call inside and outside the network. The announced average price per minute across these tariff plans is 21.4 cents to call inside the carrier and 34.6 cents to call outside the carrier. Subscribers in our sample spend, on average, 13.42 euros per month. Eventual churners and non-churners spend, on average, 9.52 and 13.70 euros per month, respectively. Subscribers with and without a tariff plan that charges the same price to call inside and outside the carrier spend 9.37 and 18.51 euros per month on average, respectively. Therefore, subscribers with a tariff plan that charges the same price to call inside and outside the carrier spend more per month though exhibit less calling activity. Summing up over all observations, users with a tariff plan that charges the same price to call inside and outside the carrier generate 251k (39%) and 393k (61%) euros of revenue over the period of analysis (7 months). The amount spent per month changes only slightly over time. Eventual churners spend, on average, 36 cents less every month. Non-churners spend 9 cents more every month, on average.

Ego's Churn and Join

Churn is significant in our dataset and seems to have increased in the second half of our panel. 1,095 subscribers in our sample churned during the period of analysis (13.38%): 830 with a tariff plan that charges the same price to call inside and outside the carrier (14.8%) and 265 with a tariff plan that charges a lower price to call inside the carrier (6.8%). Figure 4 shows that among subscribers with a tariff plan that charges a lower price to call inside the carrier, those that spend more time calling other networks churn at a higher rate (1.93% vs. 1.19% ($p - value < 0.01$) per month on average) than those who spend more time calling inside the carrier. Intuitively, subscribers that call more outside the carrier are "closer" to the other networks and thus might be among the first ones to consider churn when calling friends in other networks costs more. Figure 4 also shows that the reverse is true for subscribers with a tariff plan that charges the same price to call inside and outside the carrier, for whom, these statistics are 2.48% and 2.89%, respectively ($p - value < 0.04$).

XXX FIGURE 4 XXX

The churn rate is significant in our dataset (13.8% in 7 months extrapolates linearly to roughly 24% per year) but so is the join rate as Figure 4 shows. This allows the carrier to keep a steady size for its customer base. However, newer consumers are known to be less loyal than older ones. The monthly churn rate as a function of age peaks at 6% for subscribers that have been 1 month with the company, declines to roughly 3% for subscribers that are half a year with the company and is, on average, 0.9% for subscribers that have been with the company for more than a year and a half.

Finally, the top of Table 1 shows the number of subscribers and observations per type of tariff plan and as function of whether there are more or less friends inside the carrier. We note that the majority of observations have more friends inside the carrier (roughly 95% and 85% for users with ($sp = 0$) and without ($sp = 1$) a tariff plan that charges a lower price to call inside the carrier, respectively). The bottom part of this table shows averages and standard deviations for churn rates. The churn rate is always higher for users with tariff plans that charge the same price to call inside and outside the carrier (0.013 vs 0.028 ($p - value < 0.001$), 0.012 vs 0.027 ($p - value < 0.001$), for subscribers with and without more friends inside the carrier, respectively, and 0.012 vs 0.028 ($p - value < 0.001$) for all users) but statistically similar for users with and without more friends inside the carrier irrespective of their tariff plan type (0.024 vs. 0.021 ($p - value = 0.19$), 0.013 vs 0.012 ($p - value =$

0.81) and 0.028 vs. 0.027 ($p - value = 0.70$) for total, $sp = 0$ and $sp = 1$, respectively). These descriptive statistics show that churn is significant across users with a tariff plan that charges the same price to call inside and outside the carrier, which may mean that this type of tariff plan is ineffective at reducing churn.

XXX TABLE 1 XXX

Friends' Churn and Join

Figure 5 shows the average number of friends that churn for eventual churners and non-churners both per month and cumulative since the beginning of our panel for users with and without a tariff plan that charges the same price to call inside the carrier. The average number of friends that churn, either per month or cumulatively, is always higher for users with a tariff plan that charges a lower price to call inside the carrier (0.148 vs. 0.237 ($p - value < 0.001$) and 0.464 vs. 0.748 ($p - value < 0.001$), respectively). 10% and 25% of the subscribers with and without a tariff plan that charges the same price to call inside and outside the carrier have at least one friend that churned over the period of analysis, respectively. In both cases, 10% of the subscribers have at least one friend that churned per month. However, note that the number of friends that churn, either per month or cumulative, is significantly larger for eventual churners among users with a tariff plan that charges a lower price to call inside the carrier (0.23 vs. 0.49 ($p - value < 0.001$) and 0.73 vs. 1.20 ($p - value < 0.001$), respectively), which is not true for the case of cumulative number of friends that churn for users with a tariff plan that charges the same price to call inside and outside the carrier (0.15 vs. 0.17 ($p - value < 0.02$) and 0.46 vs. 0.46 ($p - value = 0.83$), respectively).

XXX FIGURE 5 XXX

Figure 6 shows the average number of friends that join for eventual churners and non-churners both per month and cumulative since the beginning of our panel for users with and without a tariff plan that charges the same price to call inside and outside the carrier. In this case, this statistic is lower for eventual churners for the former type of user and higher for the latter (0.43 vs. 0.69 ($p - value < 0.001$) and 1.89 vs. 2.15 ($p - value < 0.02$) per month and cumulative, respectively, for the former and 0.30 vs. 0.23 ($p - value < 0.001$) and 1.25 vs. 0.73 ($p - value < 0.01$) per month and cumulative, respectively, for the latter). All these analysis remains unchanged if we consider the ratio of the number of friends that churn and join to the logarithm of the size of the social network, which we will use later in our regressions to control for the number of friends.

XXX FIGURE 6 XXX

Empirical Strategy

The theoretical model developed before shows that, in our setting, choosing a tariff plan and churning is a function of θ^{out}/θ^{in}. Therefore, our empirical strategy to study the effect of this covariate on churn is to regress churn on it. For this purpose, let the ordinal covariate $state_{i,t}$ indicate whether user i in month t i) has a tariff plan that charges a higher price to call outside the carrier; ii) has a tariff plan that charges the same price to call inside and outside the carrier; iii) churns. Later in this paper, we provide empirical evidence showing that this is the appropriate ordering for these outcomes. Accordingly, we estimate

$$state_{i,t} = \gamma_1 + \gamma_2 \theta_{i,t}^{out}/\theta_{i,t}^{in} + d_t + \epsilon_{i,t}$$

where d_t are month dummies and $\epsilon_{i,t}$ is the idiosyncratic error term. Assuming that the error term is normally distributed yields the traditional ordered Probit regression, which we will use to estimate two cut-off points: t_1, beyond which users are no longer better off with a tariff plan that charges a lower price to call inside the carrier; and t_2, beyond which users are better off churning. In addition, we are interested in parameter γ_2, which measures how marginal changes in friends churning and joining affects the decision to change tariff plan and/or carrier. However, recall that $\theta_{i,t}^{out} = \theta_{i,t-1}^{out} + frd\_churn_{i,t}$ (where $frd\_churn_{i,t}$ indicates the cumulative number of friends of user i than churned up to time t), and therefore this parameter is unidentified in this regression to the extent that $\epsilon_{i,t}$ includes time-varying effects that affect the likelihood of churn of both user i and her friends. For example, user i and her friends can be exposed to marketing campaigns that trigger all of them to churn irrespective of whatever influence there might have been at play. The month dummies in d_t capture some dynamic effects that may explain heterogeneity in churn. For example, time dummies capture the effect of a reduction in price by a competitor, which may lead some consumers to churn from OurNet. There is no price discrimination across new consumers in the market that we analyze, that is, price reductions by competitors are similar for all users that would like to switch carriers and, therefore, captured by time dummies. However, time dummies are not enough to provide identification because they do not control for the heterogeneity associated to time varying effects unknown to us.

We use an instrumental variable to alleviate these endogeneity concerns. Consider three users labeled i, j and k. User i is a friend of user j and user j is a friend of user k. Our goal is to use the churn decision of user k to instrument the churn decision of user j. We argue that because users j and k are friends their churn decisions correlate. Significant empirical evidence in many network contexts has shown that this correlation, termed 1-hop homophily in the literature, is generally strong [2]. Our first stage results will show that this is also strong in our empirical case. However, the churn decision of user k may correlate to the churn decision of user i through unobserved mechanisms. In particular, this is likely to be the case with 1-hop homophily when user i and user k are friends. To lessen this concern, we consider only the churn decisions of users k that are not friends of user i to instrument user j's decision to churn. We call these users friends of friends not friends of the ego, "ffnf" for short, and use $churn\_ffnf$ to refer to their churn decisions. Therefore, we assume absence of 2-hop homophily in order to use the churn decision of such a user k to instrument the churn decision of user j. To further support this assumption, we show that the churn decision of user k has no bearing on the churn decision of user i conditional on the churn decision of user j. For example, when user j does not churn the churn decision of user k is irrelevant for the churn decision of user i. We use the following regression for this purpose:

$$churn_{i,t} = \alpha_a + \alpha_b fncfnf\_churn_{i,t} + X_{i,t}\alpha_c + u_i + d_t + \nu_{i,t}$$

where $fncfnf\_churn$ represents the churn decisions of users k when the users j that connect them to user i do not churn, that is the friends that do not churn of friends of the ego, not (direct) friends of the ego – "fncfnf" for short. As Figure 7 shows, user i and such a user k belong to different cliques in the social network, which reduces largely the likelihood of unobserved effects that act upon both of them.

XXX FIGURE 7 XXX

As a robustness check we rely only on the churn decisions of users k that share only one friend, user j, with user i to instrument user j's decision to churn. We denote these decisions by $churn\_ffnf\_1path$. As Figure 8 shows, this ensures not only that user i and user k belong in different cliques but also that user j is the only user in common between these cliques, which further reduces the likelihood of unobserved effects that act upon both user i and user k. In particular, in this case, it is likely that the correlation between the churn decisions of user k and user i flows through the only user connecting their cliques or, in other words, our instrument correlates with our outcome only through the endogenous variable (user j's decision to churn). We also show that the churn decision of such a user k has no bearing on the churn decision of user i conditional on the churn decision of user j by regressing $churn_{i,t}$ on $churn\_fncfnf\_1path_{i,t}$, where the latter represents the churn decisions of such users k when user j, the only user in between a user k and user i, does not churn. An additional restriction that could potentially be introduced would be to consider only the churn of friends of friends of the ego (not friends of the ego) that live in other cities than the ego. This would introduce separation not only in the call graph but also in terms of geographical proximity. Unfortunately, in our case, we are unable to generate this type of instrument because we do not have geographical information. However, future research may be able to do so and introduce even more separation between egos and their instruments.

XXX FIGURE 8 XXX

With these instruments in hands, which are inspired in the work of [12] and [19], we develop the corresponding instruments for $\theta_{i,t}^{out}/\theta_{i,t}^{in}$. Given that $\theta_{i,t}^{out} = \theta_{i,t-1}^{out} + frd\_churn_{i,t}$ we use $\left(\theta_{i,t-1}^{out} + churn\_ffnf_{i,t}\right)/\theta_{i,t}^{in}$ as an instrument for the former. For the abovementioned robustness check we use $\left(\theta_{i,t-1}^{out} + churn\_ffnf\_1path_{i,t}\right)/\theta_{i,t}^{in}$ as instrument.

Empirical Results

Shift in Calls When Friends Churn

When friends churn the likely reaction from users who subscribe tariff plans that charge a higher price to call outside the carrier is to adjust the calling pattern. These users are likely to call less often the friends that churn because it becomes more expensive to do so. On the other hand, users who subscribe to tariff plans that charge the same price to call inside and outside the carrier are less likely to make such an adjustment. Table 2 shows precisely this. Column (1) show results for users with a tariff plan that charges a lower price to call inside the carrier and columns (2) show results for users with a tariff plan that charges the same price to call inside and outside the carrier. In both cases the dependent variable is the percentage of monthly costs with calls inside the carrier. This decreases only for the former users when their friends churn, about 0.20% per each 1% increase in the cumulative number of friends that churn. This change in the volume of calls inside vs. outside the network might be seen as the firm as an early warning for churn.

XXX TABLE 2 XXX

Irrelevance of 2-Hop Churn Decisions

Table 3 shows the results from regressing the churn decisions of egos in our dataset (subscribers i in our model) on the churn decisions of the friends of their friends that are not friends of the egos (subscribers k in our model) when the subscribers that connect them do not churn. The latter are used as instruments for the former. This table shows that when the friends of the ego do not churn the churn decisions of their friends have no bearing on the churn decision of the ego. In other words, in this case, whether the subscribers we use as instruments churn is irrelevant for the ego's decision to churn and thus this test provides some evidence that our instrumental variable is exogenous in our setting.

XXX TABLE 3 XXX

Choice of Tariff Plan and Churn

Table 4 shows our empirical results. Column (1) shows ordered Probit results while columns (2) and (3) show IV ordered Probit results using the instruments discussed in previous sections. In all specifications increases in θ^{out}/θ^{in} reduce the probability of subscribing a tariff plan that charges a lower price to place calls inside the carrier ($sp_i = 0$) and increase the probability of subscribing a tariff plan that charges the same price to place calls inside and outside the carrier ($sp_i = 1$) and of churning. Using the results in column (3) yields that a marginal increase in θ^{out}/θ^{in} reduces the probability of $sp_i = 0$ by 11.8 percentage points, increases the probability of $sp_i = 1$ by 10.3 percentage points and increases the probability of churn by 1.5 percentage points. At the average, one more friend churning increases θ^{out}/θ^{in} from 0.46 to 0.50, which leads to an increase of 0.06% in the monthly churn rate. The average monthly churn rate in our dataset is 2.15%.

XXX TABLE 4 XXX

The estimate of 2.195 for the cut-off point from $sp_i = 1$ (last column of table 4) to churn implies that, on average, users churn if $\theta^{out}/\theta^{in} > (cut_2 - d_t)/\gamma_2 = (2.195 - 0.024)/0.304 = 7.14$, and thus $t_{o/i} = 1/7.14 = 0.14$ on average (0.024 is the average effect of the time dummies). This result, in turn, implies $\beta = \left(log(\tilde{p}/p^{in}) - 1/t_{o/i} log(p^{out}/\tilde{p})\right)(\theta_i\tilde{p})/\left(w(1/t_{o/i} + 1)\right) = 0.0028$ on average. As expected, this estimate is between the average $b_m = 0.00145$ and the average $b_h = 0.00324$ in our data. Therefore, offering a tariff plan that charges the same price to call inside and outside the carrier leads this provider to retain a number of users who would otherwise churn. Column (3) in Table 4 shows that the predicted penetration of the tariff plan that charges the same price to call inside and outside the carrier is 0.568, that is $Pr(sp = 1) = 0.568$. The predicted penetration of this tariff plan increases to 0.696 when computed only over users with more friends outside the carrier, that is $Pr(\, sp = 1 \mid \theta^{out}/\theta^{in} > 1\,) = 0.696$. We want to predict how many subscribers have this type of tariff plan and more friends outside the carrier, that is $Pr(sp = 1 \wedge \theta^{out}/\theta^{in} > 1)$ to measure how the introduction of this tariff plan helped retained consumers that would have otherwise churned. Using the Bayes rule we have $Pr(sp = 1 \wedge \theta^{out}/\theta^{in} > 1) = Pr(\, sp = 1 \mid \theta^{out}/\theta^{in} > 1\,)Pr(\theta^{out}/\theta^{in} > 1)$. In our dataset we have 4346 observations out of 50888 with $\theta^{out} > \theta^{in}$. Therefore, the introduction of the tariff plan that charges the same price to call inside and outside the carrier avoided 5.94% ($0.696 * 4346/50888$) in the monthly churn rate. For comparison, the monthly churn rate across our dataset is 2.15%. Hence, the introduction of this tariff plan had a significant effect on customer retention.

Finally, our ordered probit model provides evidence of a path to death that firms can potentially identify in a timely fashion as early warnings of churn. Recall that we coded our dependent variable, $state_{i,t}$ in order from i) a tariff plan that charges a lower price to call inside the carrier; ii) a tariff plan that charges the same price to call inside and outside the carrier; iii) and churn. Our empirical analysis shows the adjustment in the volume of call under i) and positive probabilities of transition from i) to ii) and from ii) to iii). Therefore, using the granular data available to us we are able to characterize the users' behavior towards churn with a level of detail that the current literature has not yet been able to provide.

Robustness Checks

We perform a number of robustness checks to our IV results to show that our findings are not an artifact of the definition of friend nor of the structure imposed by our theoretical model. For the former, we start by noting that friends are likely to call each other more often. Therefore, consider user i and let us define that user j is a friend of user i only if she is in the top quartile of the distribution of calls exchanged (that is, placed and received) by user i. This definition ensures that user i and user j talk significantly and, at the same time, corrects for the calling behavior of user i, that is, not all users j that talk a lot with user i are friends of user i but only those that do so proportionally more. We redefine the social graph using this definition of friend and redo the analysis reported in the previous subsection. Columns (1) and (2) of Table 5 show the results obtained, which are qualitatively similar to the ones obtained before. Namely, increases in θ^{out}/θ^{in} reduce the probability of subscribing a tariff plan that charges a lower price to place calls inside the carrier ($sp_i = 0$) and increase the probability of subscribing a tariff plan that charges the same price to place calls inside and outside the carrier ($sp_i = 1$) and of churning. In another specification, let user j be a strongfriend of user i if she is in the top quartile of the distribution of calls exchanged by user i and let user k be a weakfriend of user i is she is in the bottom quartile of the distribution of calls exchanged by user i. We redo the analysis reported in the previous subsection considering both (and simultaneously) the effect of strong and weak friends. Columns (3) and (4) of Table 5 show the results obtained, which are in line with our intuition for our setting. The likelihood of moving to a tariff plan that charges the same price to call inside and outside the carrier and of churning increases more when a strong friend churns compared to when a weak friend does. Finally, consider the case of a fully connected social graph in which the edge between users i and j is labeled with the airtime between them. In this graph, we can defined weighted friends using the share of airtime, that is, the weight associated to user j for user i is given by $airtime_{ij} / \sum_{k \in F_i} airtime_{ik}$. We can then redo the analysis reported in the previous subsection using these weights to weight friends' churn (likewise for the churn of friends of friends for our IVs). Columns (5) and (6) of Table 5 show the results obtain. Again, we still find that increases in θ^{out}/θ^{in} increase the likelihood of moving to a tariff plan that charges the same price to call inside and outside the carrier and of churn. In sum, all these results provide evidence that our findings do not seem to arise because of the way friends are defined. We also note that all these models include the number of calls placed and received as additional controls.

XXX TABLE 5 XXX

We also test for the appropriateness of our ordered model. To this end, we start by running separate models for the decision to choose tariff plan type and for the decision to churn. Columns (1) and (2) in Table 6 show the results obtained for the former decision. These results show that increases in θ^{out}/θ^{in} increase the likelihood of moving to a tariff plan that charges the same price to call inside and outside the carrier. Columns (3) and (4) in this table show the results obtained for the latter decision. We also observe here that increases in θ^{out}/θ^{in} increase the likelihood of churn. It is therefore clear that even with separate models increases in θ^{out}/θ^{in} behave as expected in both cases. We then model our setting using a multinomial logit model, which does not impose a specific ordering of outcomes as our ordered probit model does. In our multinomial logit model, consumers choose one of three alternatives: i) a tariff plan that charges a lower price to call inside and outside the carrier; ii) a tariff plan that charges the same price to call inside and outside the carrier; iii) churn, and they make a decision as a function of θ^{out}/θ^{in}. Columns (5) and (6) in table 6 show our results, which are obtained using the control function approach to address endogeneity, that is, we project our endogenous variable θ^{out}/θ^{in} onto our instruments (first stage) and bootstrap 1000 times controlling for the first stage residuals to adjust the standard errors (see [48] for more information on this empirical strategy). The results shown in these columns are against choosing the tariff plan that charges the same price to call inside and outside the carrier, which allows us to show clear evidence that the ordering of outcomes is indeed as one could expect. The estimates in row 1 of these columns show that increases in θ^{out}/θ^{in} reduce the likelihood of choosing a tariff plan that charges a lower price to call inside the carrier (compared to the baseline alternative, i.e. against choosing a tariff plan that charges the same price to call inside and outside the carrier). The estimates in row 4 show that increases in θ^{out}/θ^{in} increase the likelihood of churn (again compared to the baseline alternative, i.e. choosing a tariff plan that charges the same price to call inside and outside the carrier). In sum, this exercise provides clear evidence in favor of the ordered outcomes as used in our theoretical model and in our ordered probit regressions.

XXX TABLE 6 XXX

Finally, we study the nature of the effect of friends' churn. If this effect indeed arises from word of mouth among friends, then one may expect some sort of compounding or immediateness to arise. In the analysis in the previous subsection, friends' churn is a cumulative measure, that is, $friends\_churn_{i,t}$ is the cumulative number of friends of user i that churned up to time t. Mathematically, $friends\_churn_{i,t} = \sum_{\tau \leq t} \sum_{j \in F_i} churn_{j,\tau}$. Let us know affect friends' churn with a "remembrance" discount factor. Call this factor $x \in [0,1]$ and redefine our endogenous covariate as $friends\_churn_{i,t} = \sum_{\tau \leq t} \sum_{j \in F_i} x^{t-\tau}\, churn_{j,\tau}$.The cumulative number of friends that churn up to time t is obtained with $x = 1$. The lower the x in this specification the more the consumer discounts the past. In the limit, when $x = 0$, only contemporaneous churn matters. The same approach is used to compute our instruments, that is, the number of friends of friends (not friends of the ego) that churn is also affected by the discount factor x. Tables 7 and 8 show the results obtained for the two instruments used throughout our paper. We observe that the effect of friends' churn is still statistically significant in all specifications but, and as expected, its magnitude reduces with x. The more consumers forget the past the fewer signals from friends they have (or retain) to inform their decisions. In addition, we observe that there is still contagious churn when only contemporaneous signals are used by consumers (the case of $x = 0$). Therefore, we find evidence that contagious churn may have both a compounding nature as well as an immediate nature. The compounding nature comes from the fact that contagious churn is stronger when consumers remember the past more and compound, at some rate, all prior signals from their friends. The immediate nature comes from the fact that there is still contagious churn even when only contemporaneous churn from friends is taken into account.

XXX TABLE 7 XXX

XXX TABLE 8 XXX

Conclusions

Churn management has become a nuclear task in customer relationship management. One way to prevent churn is to sign consumers into long-term contracts, which require consumers to pay significant fees to terminate them before expiry. For this reason, and in particular in mobile markets, a number of consumers prefers pre-paid SIM cards, which they top up as they go. Placing calls with pre-paid tariff plans is usually more expensive than with long-term contracts but consumers can terminate their relationship with the carrier by simply winding down their remaining balance. Two major types of tariff plans are available in the market today for pre-paid consumers. Under one type of tariff plan placing calls to subscribers in the same carrier is cheaper than placing calls to subscribers in other carriers. Under another type of tariff plan all calls are priced equally. Some consumers choose one type of tariff plan over the other. For example, consumers that are risk averse adopt a tariff plan that charges the same price to call inside and outside the carrier and thus they do not need to consider which carrier their friend belongs to before calling them. On the other hand, users that are risk loving will choose a tariff plan that charges a lower price to call subscribers inside the carrier if most of their friends are in this carrier. Otherwise, if most of their friends are in another carrier then they are likely to move to that carrier and choose this type of tariff plan there. More interestingly, users in between will not always choose the same type of tariff plan and will move from a tariff plan that charges a lower price to call users inside the same carrier to a tariff plan that charges the same price for all calls if a significant number of their friends churn. Our paper identifies this nuanced consumer behavior as an early warning of churn. Another nuanced behavior that we find is that even before changing the type of tariff plan consumers are likely to adjust their volume of calls inside vs. outside the carrier when their friends start churning, which provides yet another early warning to the firm that consumers may be closer to churn.

Our paper looks at how the decisions of friends about which carrier to choose affect one's decision to change tariff plan or to churn. Given the strong positive network externalities in mobile markets, consumers are likely to start by subscribing service from the carrier which most of their friends subscribe to. A subscriber whose friends churn and move to a competitor is then likely to adjust her behavior. If only a few friends churn, then the likely behavior might be to simply adjust calls taking into account the increased price of calling the friends that churned. If more friends churn, then the likely behavior might be to change to a tariff plan that charges the same price to call inside and outside the carrier. If a significant number of friends churn, then the likely behavior might be to just churn altogether. In this paper, we develop a theoretical model showing how consumers make these decisions thus explaining in detail the behavioral process towards churn. As a result of using this approach to model consumer behavior, we show that 1) in a world without a tariff plan that charges the same price to call inside and outside the carrier, subscribers choose the carrier where most of their friends are; 2) the introduction of this type of tariff plan helps retain some consumers with more friends outside the carrier that would otherwise churn.

We use a 11-month long panel of data from a large mobile provider to empirically test our hypotheses. We random sample 10,000 consumers and analyze their friends and friends of their friends, using data from more than 2.5 million subscribers. We show how friends' churn affects the propensity of subscribers to choose a tariff plan that charges the same price to call inside and outside the carrier and the propensity to churn. A contribution of our paper is the methodology we use to identify this effect that is typically endogenous in network settings. We use the churn decisions from friends of friends not friends of the ego to instrument the churn decision of friends of the ego. With respect to the exogeneity of this instrumental variable we show that the churn decision of friends of friends is irrelevant for the churn decision of the ego when the friend that connects them does not churn. Armed with this instrument, we estimate an ordered probit model to account for the order in which consumers may move from i) adjusting their volume of calls inside vs. outside the carrier; to

ii) choosing a tariff plan that charges the same price to call inside and outside the carrier; to iii) churning. Our empirical analyses show that, on average across subscribers in our dataset, every time an additional friend churns the ego's monthly likelihood of churn increases by 0.06%. Across our dataset, the monthly churn rate is 2.15%. This model also shows that a significant number of subscribers that would otherwise churn remain with the carrier and choose a tariff plan that charges the same price to call users inside and outside the carrier. We show that the introduction of this tariff plan avoided 5.94% in the monthly churn rate.

Our results are significant in three ways. First, we provide empirical evidence that consumers that churn in mobile markets exhibit a path to death. When only a few friends churn the immediate behavioral response is to adjust the volume of calls. If more friends churn, then users are likely to switch to a tariff plan that charges the same price to call inside and outside the carrier. Finally, if many friends churn then users are likely to churn too. This ordered stream of user behaviors that telecommunication firms can observe may therefore provide them with early warnings of churn. The prior literature connected the path to death in telecommunications services to usage. Here, we show that this concept can be enriched by looking, in addition, at the type of tariff plan that consumers subscribe and whether they change it. Second, we show empirical evidence of contagious churn in mobile markets. This implies that a definition of Customer Lifetime Value (CLTV) that includes only the present value of the future stream of profits on each specific consumer underestimates the value of consumers and particularly so for consumers that have many friends. Instead, CLTV should also take into account the role of peer influence in social networks and the fact that some consumers are worthier than others because, if nothing else, they have more friends to influence. Third, we show how introducing tariff plans with particular features can help reduce churn. In the context of mobile, we provide empirical evidence that tariff plans that charge the same price to call inside and outside the network are effective at reducing churn and we measure their economic value.

A concept similar to tariff plans that charge the same price to call inside and outside the network can be applied in other markets and across industries. Take, for example, the case of sports facilities. Consider a city with two of such clubs where people go play tennis. Each person is a member of one club (and pays an annual fee) and pays to rent the tennis court on an hourly basis. One type of membership allows for renting the court for a small price (p_{in} in our model) if you are playing another member of the club but requires a higher price (p_{out}) for the court if you want to bring in a guest from outside. Another type of membership charges always the same price ($\tilde{p}$) for the court no matter whether your opponent is a member of the same club or a guest from outside. In this market, the same dynamics that we study in our paper arise. A member that usually plays with other members of the same club should opt for the former type of membership. A member who mixes up and plays players from her club but also from other clubs should opt for the latter type of membership. Finally, a player that essentially plays players from another club should join that club. The latter type of membership retains the player that brings in players from other clubs to play with. The same dynamics arise, for example, with a video content distributor that charges a flat monthly fee to provide access to a library with content from multiple providers. This should reduce the likelihood of losing consumers to the specific video stores that own each piece of content.

Appendices

Derivation of Optimal Tariff Plans and Churn Decisions

Let $U_i(q_i) = \sum_{j \in F_i} \theta_{i,j} log(q_{i,j}) + \sum_{j \in F_i} \beta q_{ij} sp_i$ represent the utility of user i with tariff plan of type sp_i. This user maximizes her utility subject to the budget constraint $\sum_{j \in F_i} q_{ij} (sp_i \tilde{p} +$

$(1 - sp_i)p_{ij}^{in/out}) \leq w$, where $p_{ij}^{in/out} = p^{in}\{j \in F_i^{in}\} + p^{out}\{j \in F_i^{out}\}$. Let $p_{ij}^{out/in} = p^{out}\{j \in F_i^{in}\} + p^{in}\{j \in F_i^{out}\}$. Therefore, the Lagrangian and its derivatives become:

$$\mathcal{L}(q_i; \lambda) = \sum_{j \in F_i} \theta_{i,j} log(q_{i,j}) + \sum_{j \in F_i} \beta q_{ij} sp_i + \lambda \left(w - \sum_{j \in F_i} q_{ij}\left(sp_i \tilde{p} + (1 - sp_i) p_{ij}^{in/out}\right) \right)$$

$$\partial \mathcal{L}(q_i; \lambda) / \partial q_{ij} = \theta_{ij}/q_{ij} + \beta sp_i - \lambda\left(sp_i \tilde{p} + (1 - sp_i) p_{ij}^{in/out}\right)$$

$$\partial \mathcal{L}(q_i; \lambda) / \partial \lambda = w - \sum_{j \in F_i} q_{ij}\left(sp_i \tilde{p} + (1 - sp_i) p_{ij}^{in/out}\right)$$

Setting these derivatives to zero and solving for q_{ij} and λ yields:

$$\left(q_{ij}, \lambda\right) = \left(\theta_{ij} w\left(sp_i/\tilde{p} + (1 - sp_i)/p_{ij}^{in/out}\right)/\theta_i, \theta_i/w + sp_i \beta/\tilde{p}\right)$$

Substituting into the user's utility function and assuming a symmetric competitor in the market (that is, a competitor that offers the same type of tariff plans with the same prices), yields:

$$U(\boldsymbol{q_i}) = \begin{cases} \sum_{j \in F_i} \theta_{ij} log\left((w\theta_{ij})/(\theta_i p_{ij}^{in/out})\right), & if\ sp_i = 0 \\ \sum_{j \in F_i} \theta_{ij} log\left((w\theta_{ij})/(\theta_i \tilde{p})\right) + \beta w/\tilde{p}, & if\ sp_i = 1 \vee sp_i = 1\ and\ churn \\ \sum_{j \in F_i} \theta_{ij} log\left((w\theta_{ij})/(\theta_i p_{ij}^{out/in})\right), & if\ sp_i = 0\ and\ churn \end{cases}$$

Let $\beta_l = (\theta_i \tilde{p}/w) log(\tilde{p}/p^{out})$, $\beta_m = (\theta_i \tilde{p}/w) log\left(\tilde{p}/\sqrt{\{p^{in} p^{out}\}}\right\}$ and $\beta_h = (\theta_i \tilde{p}/w) log\left(\tilde{p}/p^{in}\right)$. In addition, let $t_{o/i} = \left[log\left(\tilde{p}/p^{in}\right) - (\beta w)/(\theta_i \tilde{p})\right]/[log(p^{out}/\tilde{p}) + (\beta w)/(\theta_i \tilde{p})]$. The above formulation is particularly useful to derive the following propositions, which combined generate Figure 9:

1) $sp_i = 0$ is preferred to $sp_i = 0$ and churn for any β if $\theta^{out}/\theta^{in} < 1$ and otherwise if $\theta^{out}/\theta^{in} > 1$.

Proof: $\sum_{j\in F_i} \theta_{ij} log\left((w\theta_{ij})/(\theta_i p_{ij}^{in/out})\right) > \sum_{j\in F_i} \theta_{ij} log\left((w\theta_{ij})/(\theta_i p_{ij}^{out/in})\right) \Leftrightarrow$

$\sum_{j\in F_i} \theta_{ij} log\left(p_{ij}^{out/in}/p_{ij}^{in/out}\right) > 0 \Leftrightarrow \theta^{in} log\left(p^{out}/p^{in}\right) + \theta^{out} log\left(p^{in}/p^{out}\right) > 0 \Leftrightarrow$

$\theta^{out}/\theta^{in} < 1$.

2) $sp_i = 0$ is preferred to $sp_i = 1 \vee sp_i = 1$ and churn for any $\left(\theta^{out}, \theta^{in}\right)$ if $\beta < \beta_l$, for $0 \leq \theta^{out}/\theta^{in} < t_{o/i}$ if $\beta_l < \beta < \beta_h$ and for no $\left(\theta^{out}, \theta^{in}\right)$ when $\beta > \beta_h$.

Proof: we want to show that $\sum_{j\in F_i} \theta_{ij} log\left((w\theta_{ij})/(\theta_i p_{ij}^{in/out})\right) >$

$\sum_{j\in F_i} \theta_{ij} log\left((w\theta_{ij})/(\theta_i \tilde{p})\right) + \beta w/\tilde{p}$, which is the same as $\theta^{in} log\left(\tilde{p}/p^{in}\right) +$

$\theta^{out} log(\tilde{p}/p^{out}) > \beta w/\tilde{p}$. It is immediate that this is true for all $\left(\theta^{out}, \theta^{in}\right)$ when

$\beta < (\theta_i \tilde{p}/w) log(\tilde{p}/p^{out})$ because $p^{in} < p^{out}$ implies $\theta^{in} log\left(\tilde{p}/p^{in}\right) +$

$\theta^{out} log(\tilde{p}/p^{out}) > \theta_i log(\tilde{p}/p^{out})$. Likewise, this is never true when $\beta >$

$(\theta_i \tilde{p}/w) log\left(\tilde{p}/p^{in}\right)$ because $p^{in} < p^{out}$ implies $\theta^{in} log\left(\tilde{p}/p^{in}\right) + \theta^{out} log(\tilde{p}/$

$p^{out}) > \theta_i log\left(\tilde{p}/p^{in}\right)$. Finally, for $(\theta_i \tilde{p}/w) log(\tilde{p}/p^{out}) < \beta < (\theta_i \tilde{p}/w) log\left(\tilde{p}/\right.$

$\left.p^{in}\right)$ this proposition amounts to show that $\theta^{out}\left(log(\tilde{p}/p^{out}) + (\beta w)/(\tilde{p}\theta_i)\right) <$

$\theta^{in}\left(log\left(\tilde{p}/p^{in}\right) - (\beta w)/(\tilde{p}\theta_i)\right)$ and therefore $\theta^{out}/\theta^{in} < t_{o/i}$.

3) $sp_i = 0$ *and churn* is preferred to $sp_i = 1 \vee sp_i = 1$ *and churn* for any $\left(\theta^{out}, \theta^{in}\right)$ if $\beta > \beta_h$, for $\theta^{out}/\theta^{in} > 1/t_{o/i}$ if $\beta_l < \beta < \beta_h$ and for no $\left(\theta^{out}, \theta^{in}\right)$ when $\beta < \beta_l$.

Proof: we want to show that $\sum_{j\in F_i} \theta_{ij} log\left((w\theta_{ij})/(\theta_i p_{ij}^{out/in})\right) >$

$\sum_{j\in F_i} \theta_{ij} log\left((w\theta_{ij})/(\theta_i \tilde{p})\right) + \beta w/\tilde{p}$, which is the same as $\theta^{in} log(\tilde{p}/p^{out}) +$

$\theta^{out} log\left(\tilde{p}/p^{in}\right) > \beta w/\tilde{p}$. It is immediate that this is true for all $\left(\theta^{out}, \theta^{in}\right)$ when

$\beta < (\theta_i \tilde{p}/w) log(\tilde{p}/p^{out})$ because $p^{in} < p^{out}$ implies $\theta^{in} log(\tilde{p}/p^{out}) +$

$\theta^{out} log(\tilde{p}/p^{in}) > \theta_i log(\tilde{p}/p^{out})$. Likewise, this is never true when $\beta > (\theta_i \tilde{p}/w) log(\tilde{p}/p^{in})$ because $p^{in} < p^{out}$ implies $\theta^{in} log(\tilde{p}/p^{out}) + \theta^{out} log(\tilde{p}/p^{in}) < \theta_i log(\tilde{p}/p^{in})$. Finally, for $(\theta_i \tilde{p}/w) log(\tilde{p}/p^{out}) < \beta < (\theta_i \tilde{p}/w) log(\tilde{p}/p^{in})$ this proposition amounts to show that $\theta^{out}\left(log(p^{in}/\tilde{p}) + (\beta w)/(\tilde{p}\theta_i)\right) < \theta^{in}\left(log(\tilde{p}/p^{out}) - (\beta w)/(\tilde{p}\theta_i)\right)$ and therefore $\theta^{out}/\theta^{in} > 1/t_{o/i}$.

4) $t_{o/i} < 1 \Leftrightarrow 1/t_{o/i} > 1 \Leftrightarrow \tilde{p} < \sqrt{p^{in}p^{out}}\, \exp\{(\beta w)/(\theta_i \tilde{p})\}$ when $log(\tilde{p}/p^{out}) < (\beta w)/(\theta_i \tilde{p}) < log(\tilde{p}/p^{in})$.

Proof: the first part follows from noting that $t_{o/i} > 0$ for these values of β. Also for this values of β, $t_{o/i} < 1 \Leftrightarrow log(\tilde{p}/p^{in}) - (\beta w)/(\theta_i \tilde{p}) < log(p^{out}/\tilde{p}) + (\beta w)/(\theta_i \tilde{p}) \Leftrightarrow log(\widetilde{p^2}/p^{in}p^{out}) < 2(\beta w)/(\theta_i \tilde{p})$ from where the second part follows immediately

XXX FIGURE 9 XXX

References

1. Anagnostopoulos, A.; Kumar, R.; and Mahdian, M. Influence and correlation in social networks. In Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (New York, NY, USA, 2008), KDD ’08, ACM, pp. 7-15.

2. Aral, S.; Muchnik, L.; and Sundararajan, A. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. Proceedings of the National Academy of Sciences 106, 51 (2009), 21544-21549.

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

4. Ascarza, E.; and Hardie, B. G. S. A joint model of usage and churn in contractual settings. Marketing Science 32, 4 (2013), 570-590.

5. Athanassopoulos, A. D. Customer satisfaction cues to support market segmentation and explain switching behavior. Journal of Business Research 47, 3 (2000), 191-207.

6. Au, W.-H.; Chan, K. C. C.; and Yao, X. A novel evolutionary data mining algorithm with applications to churn prediction. IEEE Transactions on Evolutionary Computation 7, 6 (Dec 2003), 532-545.

7. Bapna, R.; and Umyarov, A. Do your online friends make you pay? a randomized field experiment on peer influence in online social networks. Management Science 61, 8 (2015), 1902-1920.

8. Bhattacharya, C. B. When customers are members: Customer retention in paid membership contexts. Journal of the Academy of Marketing Science 26, 1 (1998), 31-44.

9. Bobbier, T. Keeping the customer satisfied: The dynamics of customer defection, and the changing role of the loss adjuster. CILA Report, http://www.cila.co.uk/publication/articles/keeping-customer-satisfied-tony-boobier, accessed July (2013).

10. Bolton, R. N. A dynamic model of the duration of the customer's relationship with a continuous service provider: The role of satisfaction. Marketing Science 17, 1 (1998), 45-65.

11. Braff, A.; Passmore, W. J.; and Simpson, M. Going the distance with telecom customers. McKinsey Quarterly 13 (2003), 1-16.

12. Bramoulle, Y.; Djebbari, H.; and Fortin, B. Identification of peer effects through social networks. Journal of Econometrics 150, 1 (2009), 41-55.

13. Bult, J. R. Semiparametric versus parametric classification models: An application to direct marketing. Journal of Marketing Research 30, 3 (1993), 380-390.

14. Bult, J. R.; and Wittink, D. R. Estimating and validating asymmetric heterogeneous loss functions applied to health care fund raising. International Journal of Research in Marketing 13, 3 (1996), 215-226.

15. Burez, J.; and den Poel, D. V. Handling class imbalance in customer churn prediction, Expert Systems with Applications 36, 3, Part 1 (2009), 4626-4636.

16. Cavusoglu, H.; Hu, N.; Li, Y.; and Ma, D. Information technology diffusion with influentials, imitators, and opponents. Journal of Management Information Systems, 27, 2 (Fall 2010), 305-334

17. Cho, D.; Ferreira, P.; and Telang, R. Empirical analysis of mobile number portability and exclusive handset deal on switching costs and market competition. In 40th Telecommunications Policy Research Conference (TPRC) (2012).

18. Coussement, K.; and den Poel, D. V. Churn prediction in subscription services: An application of support vector machines while comparing two parameter-selection techniques. Expert Systems with Applications 34, 1 (2008), 313-327.

19. de Matos, M. G.; Ferreira, P. A.; and Krackhardt, D. Peer influence in the diffusion of iphone 3g over a large social network. MIS Quarterly 38, 4 (2014), 1103-1133.

20. den Poel, D. V.; and Lariviere, B. Customer attrition analysis for financial services using proportional hazard models. European Journal of Operational Research 157, 1 (2004), 196-217. Smooth and Nonsmooth Optimization.

21. Dierkes, T.; Bichler, M.; and Krishnan, R. Estimating the effect of word of mouth on churn and cross-buying in the mobile phone market with markov logic networks. Decision Support Systems 51, 3 (2011), 361-371.

22. F., R. The Loyalty Effect. Harvard Business School Press, Boston, Massachusetts, 1996.

23. Fader, P. S.; and Hardie, B. G. S. Customer-base valuation in a contractual setting: The perils of ignoring heterogeneity. Marketing Science 29, 1 (2010), 85-93.

24. Garg, R.; Smith, M. D.; and Telang, R. Measuring information diffusion in an online community. Journal of Management Information Systems, 28, 2 (2011), 11-38

25. Ganesh, J.; Arnold, M. J.; and Reynolds, K. E. Understanding the customer base of service providers: An examination of the differences between switchers and stayers. Journal of Marketing 64, 3 (2000), 65-87.

26. Glady, N.; Baesens, B.; and Croux, C. A modified pareto/nbd approach for predicting customer lifetime value. Expert Systems with Applications 36, 2, Part 1 (2009), 2062-2071.

27. Godinho de Matos, M.; Pedro, F.; and Rodrigo, B. Target the ego or target the group: Evidence from a randomized experiment in proactive churn management. Marketing Science 37, 5 (2018), 793-811.

28. Goel, S.; Watts, D. J.; and Goldstein, D. G. The structure of online diffusion networks. In Proceedings of the 13th ACM Conference on Electronic Commerce (New York, NY, USA, 2012), EC ’12, ACM, pp. 623-638.

29. Hadden, J.; Tiwari, A.; Roy, R.; and Ruta, D. Computer assisted customer churn management: State-of-the-art and future trends. Computers and Operations Research 34, 10 (2007), 2902-2917.

30. Han, Q.; and Pedro Ferreira, P. Role of peer influence in churn in wireless networks. In the International Conference on Social Computing (SocialCom) (2014).

31. Hung, S. Y.; Yen, D. C.; and Wang, H. Y. Applying data mining to telecom churn management. Expert Systems with Applications 31, 3 (2006), 515-524.

32. Hwang, H.; Jung, T.; and Suh, E. An LVT model and customer segmentation based on customer value: a case study on the wireless telecommunication industry. Expert Systems with Applications 26, 2 (2004), 181-188.

33. Kim, S.; Shin, K.; and Park, K. An application of support vector machines for customer churn analysis: Credit card case. In Advances in Natural Computation (Berlin, Heidelberg, 2005), L. Wang, K. Chen, and Y. S. Ong, Eds., Springer Berlin Heidelberg, pp. 636-647.

34. Krivitsky, P. N.; Ferreira, P.; and Telang, R. Network neighbor effects on customer churn in cell phone networks. In Symposium on Statistical Challenges in Electronic Commerce Research (2011).

35. Kwon, H. E.; Oh, W.; and Kim, T. Platform structures, homing preferences, and homophilous propensities in online social networks. Journal of Management Information Systems 34, 3 (2017), 768-802.

36. Lariviere, B.; and den Poel, D. V. Predicting customer retention and profitability by using random forests and regression forests techniques. Expert Systems with Applications 29, 2 (2005), 472-484.

37. Leenders, R. T. Modeling social influence through network autocorrelation: constructing the weight matrix. Social Networks 24, 1 (2002), 21-47.

38. Lemmens, A.; and Croux, C. Bagging and boosting classification trees to predict churn. Journal of Marketing Research 43, 2 (2006), 276-286.

39. Lemmens, A.; and Gupta, S. Managing churn to maximize profits. 2017.

40. Lemon, K. N.; White, T. B.; and Winer, R. S. Dynamic customer relationship management: Incorporating future considerations into the service retention decision. Journal of Marketing 66, 1 (1 2002), 1-14.

41. Lima, E. Domain Knowledge Integration in Data Mining for Churn and Customer Lifetime Value Modeling. PhD thesis, University of Southampton, Faculty of Law, Arts and Social Sciences, Southampton, UK, 2009.

42. Manski, C. F. Identification of Endogenous Social Effects: The Reflection Problem. The Review of Economic Studies 60, 3 (07 1993), 531-542.

43. Mattison, R. The Telco Churn Management Handbook. XiT Press, Illinois, USA, 2005.

44. Mozer, M. C.; Wolniewicz, R.; Grimes, D. B.; Johnson, E.; and Kaushansky, H. Predicting subscriber dissatisfaction and improving retention in the wireless telecommunications industry. IEEE Transactions on Neural Networks 11, 3 (May 2000), 690-696.

45. Neslin, S. A. Cell2Cell Case Notes. Tuck School of Business, Dartmouth College, 2002.

46. Neslin, S. A.; Gupta, S.; Kamakura, W.; Lu, J.; and Mason, C. H. Defection detection: Measuring and understanding the predictive accuracy of customer churn models. Journal of Marketing Research 43, 2 (2006), 204-211.

47. Pendharkar, P. C. Genetic algorithm based neural network approaches for predicting churn in cellular wireless network services. Expert Systems with Applications 36, 3, Part 2 (2009), 6714-6720.

48. Petrin, A.; and Train, K. A control function approach to endogeneity in consumer choice models. Journal of Marketing Research 47, 1 (2010), 3-13.

49. Risselada, H.; Verhoef, P. C.; and Bijmolt, T. H. Staying power of churn prediction models. Journal of Interactive Marketing 24, 3 (2010), 198-208.

50. Rochet, J. C.; and Tirole, J. Two-sided markets: A progress report. The RAND Journal of Economics 37, 3 (2006), 645-667.

51. Rosset, S.; Neumann, E.; Eick, U.; Vatnik, N.; and Idan, Y. Customer lifetime value modeling and its use for customer retention planning. In Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (New York, NY, USA, 2002), KDD ’02, ACM, pp. 332-340.

52. Schweidel, D. A.; Fader, P. S.; and Bradlow, E. T. Understanding service retention within and across cohorts using limited information. Journal of Marketing 72, 1 (2008), 82-94.

53. Shaffer, G.; and Zhang, Z. J. Competitive one-to-one promotions. Management Science 48, 9 (2002), 1143-1160.

54. Shalizi, C. R.; and Thomas, A. C. Homophily and contagion are generically confounded in observational social network studies. Sociological Methods and Research 40, 2 (2011), 211-239. PMID: 22523436.

55. Stieglitz, S.; and Dang-Xuan, L. Emotions and Information Diffusion in Social Media-Sentiment of Microblogs and Sharing Behavior. Journal of Management Information Systems, 29(4), (2013), 217-248

56. Strang, D.; and Soule, S. A. Diffusion in organizations and social movements: From hybrid corn to poison pills. Annual Review of Sociology 24, 1 (1998), 265-290.

57. Thieme, R. J.; Song, M.; and Calantone, R. J. Artificial neural network decision support systems for new product development project selection. Journal of Marketing Research 37, 4 (2000), 499-507.

58. Tsai, C. F.; and Lu, Y. H. Customer churn prediction by hybrid neural networks. Expert Systems with Applications 36, 10 (2009), 12547-12553.

59. Tucker, C. Identifying formal and informal influence in technology adoption with network externalities. Management Science 54, 12 (2008), 2024-2038.

60. VandenBulte, C.; and Lilien, G. Medical innovation revisited social contagion versus marketing effort. American Journal of Sociology 106, 5 (2001), 1409-1435.

61. Venkatesan, R.; and Kumar, V. A customer lifetime value framework for customer selection and resource allocation strategy. Journal of Marketing 68, 4 (2004), 106-125.

62. Verbeke, W.; Dejaeger, K.; Martens, D.; Hur, J.; and Baesens, B. New insights into churn prediction in the telecommunication sector: A profit driven data mining approach. European Journal of Operational Research 218, 1 (2012), 211-229.

63. Verbeke, W.; Martens, D.; Mues, C.; and Baesens, B. Building comprehensible customer churn prediction models with advanced rule induction techniques. Expert Systems with Applications 38, 3 (2011), 2354-2364.

64. Wattal, S.; Racherla, P.; and Mandviwalla, M. Network externalities and technology use: A quantitative analysis of intraorganizational blogs. Journal of Management Information Systems, 27, 1 (Summer 2010), 145-174

65. Wei, C. P.; and Chiu, I. T. Turning telecommunications call details to churn prediction: a data mining approach. Expert Systems with Applications 23, 2 (2002), 103-112.

66. West, P. M.; Brockett, P. L.; and Golden, L. L. A comparative analysis of neural networks and statistical methods for predicting consumer choice. Marketing Science 16, 4 (1997), 370-391.

67. Xie, Y.; Li, X.; Ngai, E.; and Ying, W. Customer churn prediction using improved balanced random forests. Expert Systems with Applications 36, 3, Part 1 (2009), 5445-5449.

68. Zhang, B.; Pavlou, P. A.; and Krishnan, R. On direct vs. indirect peer influence in large social networks. Information Systems Research 29, 2 (2018), 292-314.

69. Zhao, Y.; Li, B.; Li, X.; Liu, W.; and Ren, S. Customer churn prediction using improved one-class support vector machine. In Advanced Data Mining and Applications (Berlin, Heidelberg, 2005), X. Li, S. Wang, and Z. Y. Dong, Eds., Springer Berlin Heidelberg, pp. 300-306.

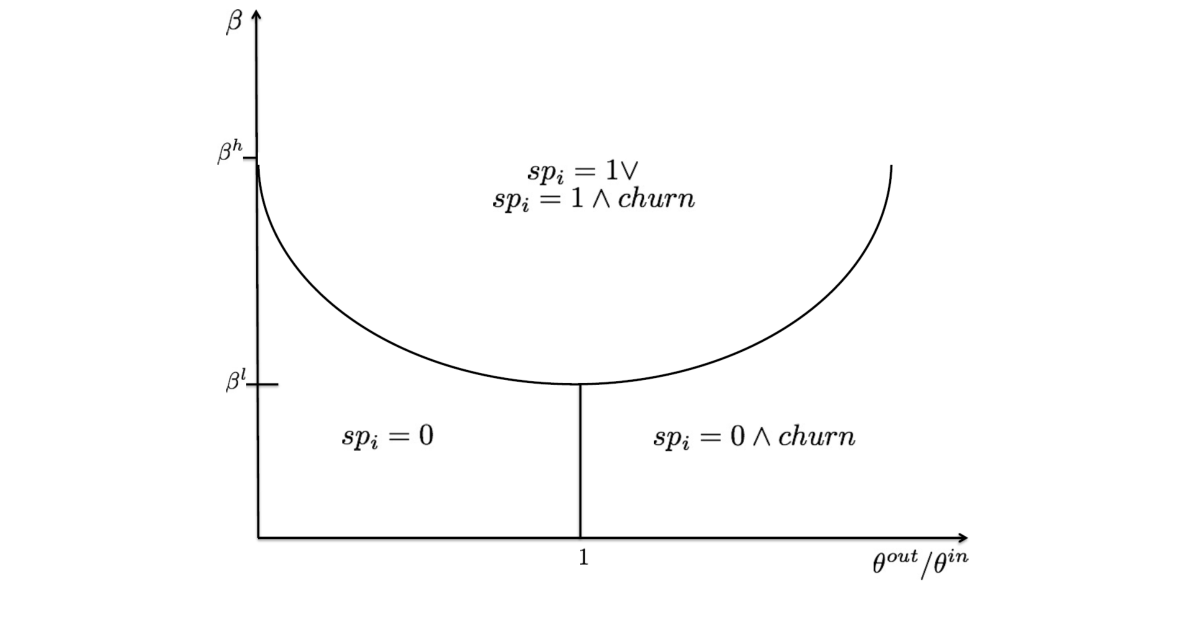

Figure 1 - Optimal tariff plan type ($sp = 0/1$) and churn decision a function of the risk profile (β) and of the ratio of aggregated utility inside (θ^{in}) and outside (θ^{out}) the carrier.

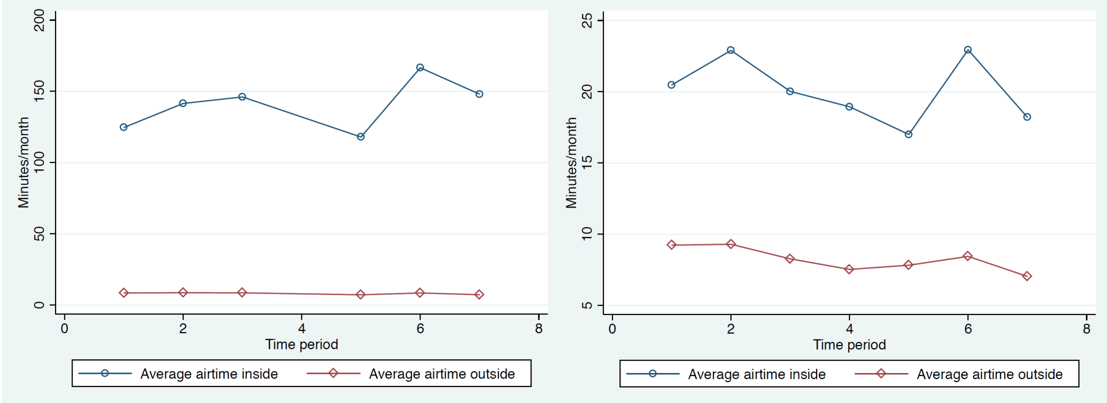

(a) $sp = 0$	(b) $sp = 1$

Figure 2 - Average airtime for calls placed inside and outside the carrier over time.

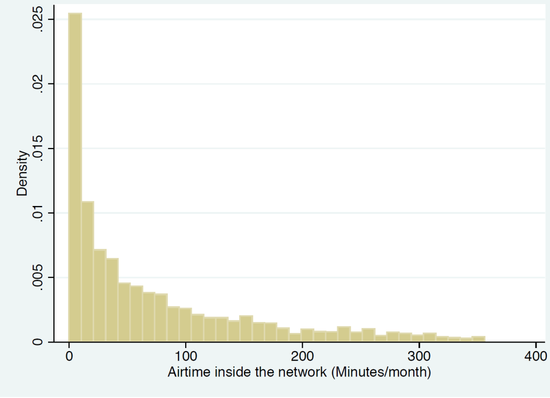

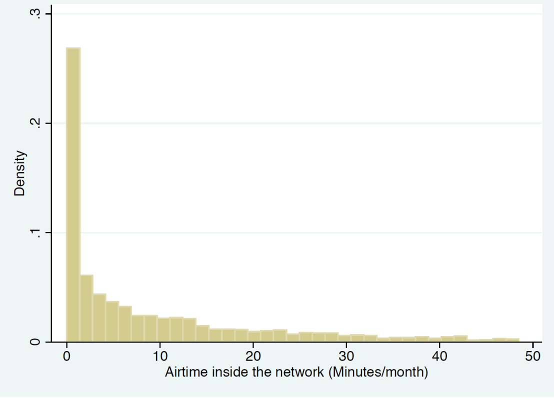

(a) $sp = 0$	(b) $sp = 1$

Figure 3 - Histogram for airtime inside the carrier in March 2009 (trimmed at 90% percentile).

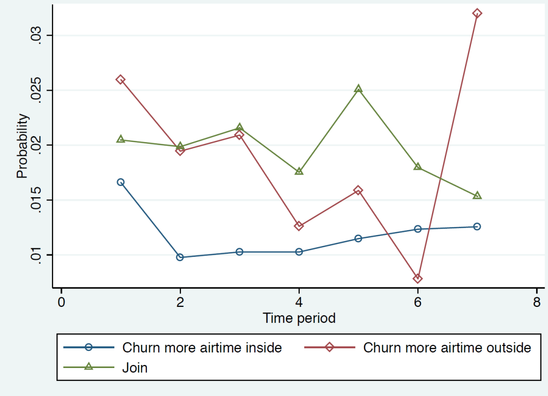

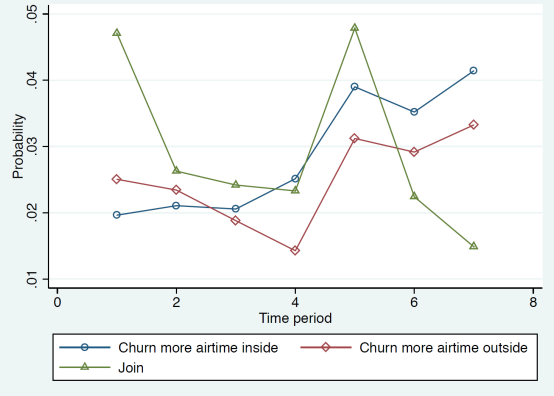

(a) $sp = 0$	(b) $sp = 1$

Figure 4 - Join rate and churn rate over time broken-down by whether subscribers spent more or less time calling inside the network.

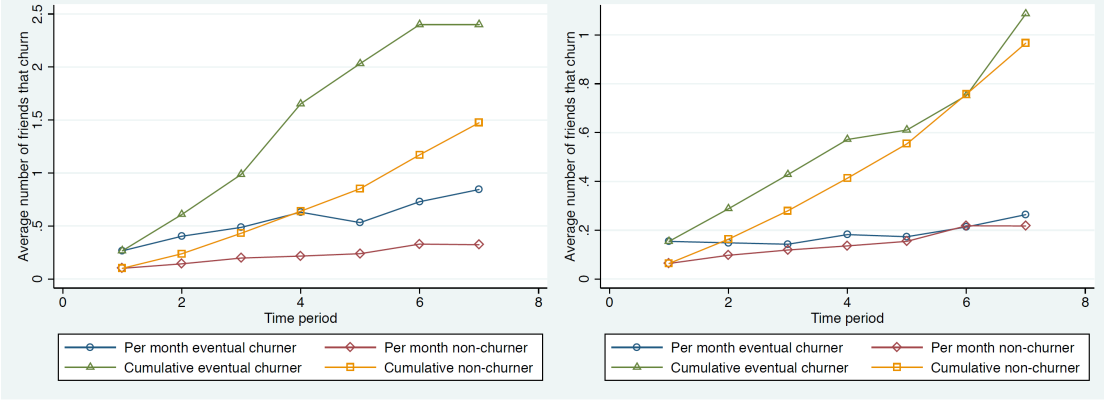

(a) $sp = 0$ (b) $sp = 1$

Figure 5 - Average number of friends that churn for eventual churners and non-churners over time and cumulative ($sp = 0$ indicates using a tariff plan that charges a lower price to call inside the carrier, $sp = 1$ indicates using a tariff plan that charges the same price to call inside and outside the carrier).

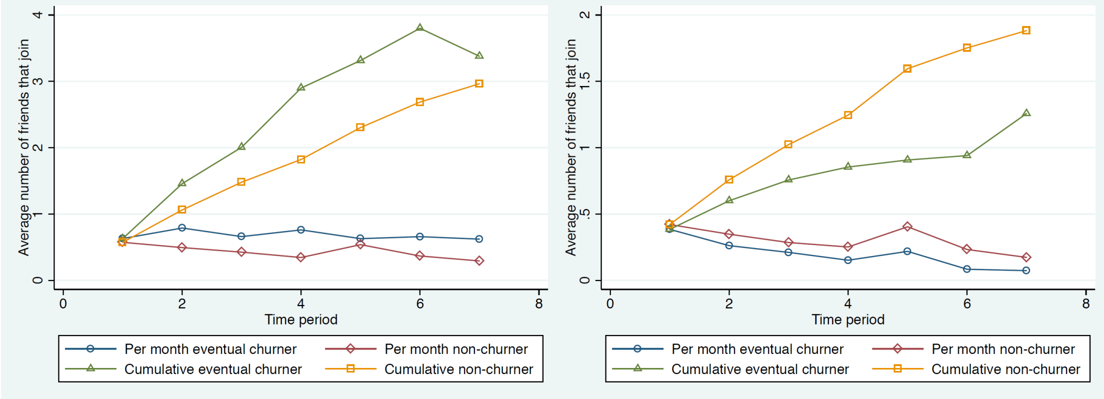

(a) $sp = 0$ (b) $sp = 1$

Figure 6 - Average number of friends that join for eventual churners and non-churners over time and cumulative ($sp = 0$ indicates using a tariff plan that charges a lower price to call inside the carrier, $sp = 1$ indicates using a tariff plan that charges the same price to call inside and outside the carrier).

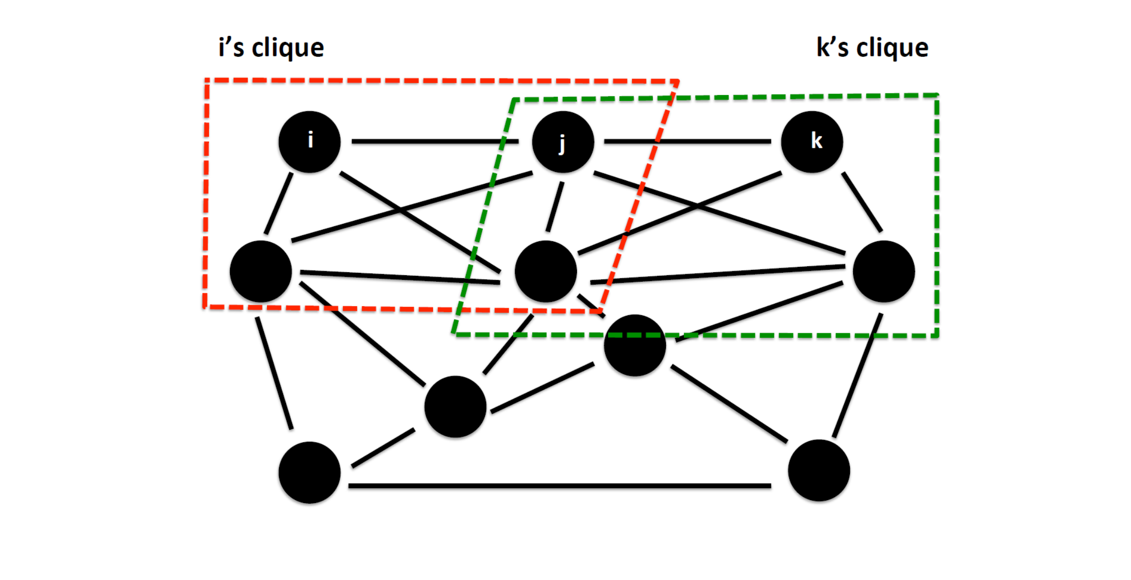

Figure 7 - Friends of friends (k) not friends of the ego (i) lie in different network cliques. The churn decision of the former is used to instrument the churn decision of the common friends (j).

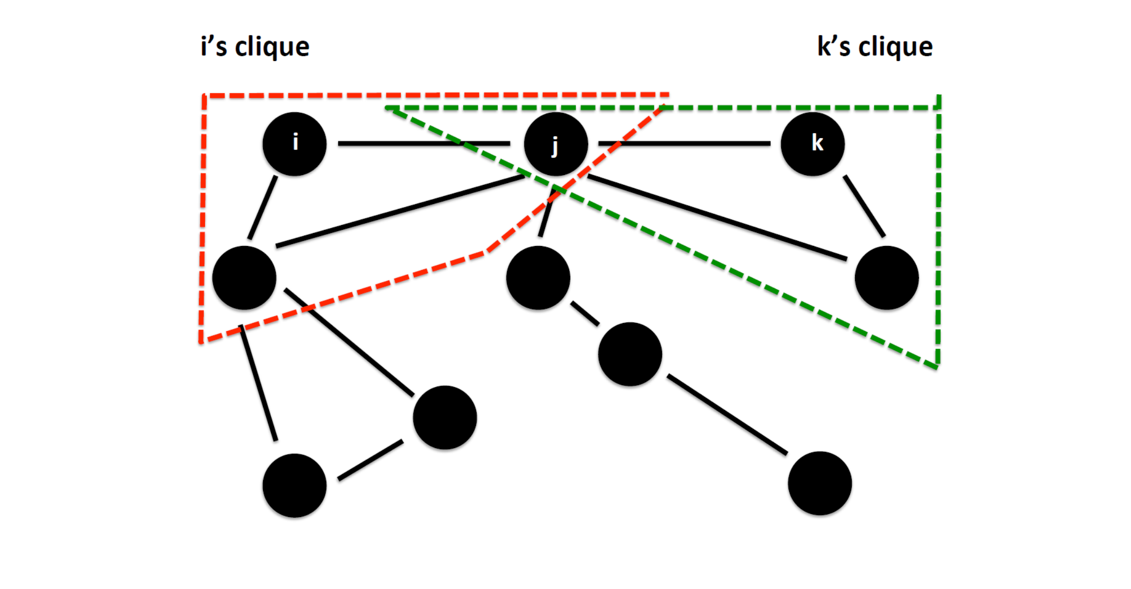

Figure 8 - Friends of friends not friends (k) of the ego (i) with only one common friend (j) lie in different network cliques that intersect only at that friend. The churn decision of the former is used to instrument the churn decision of the only common friend (j).

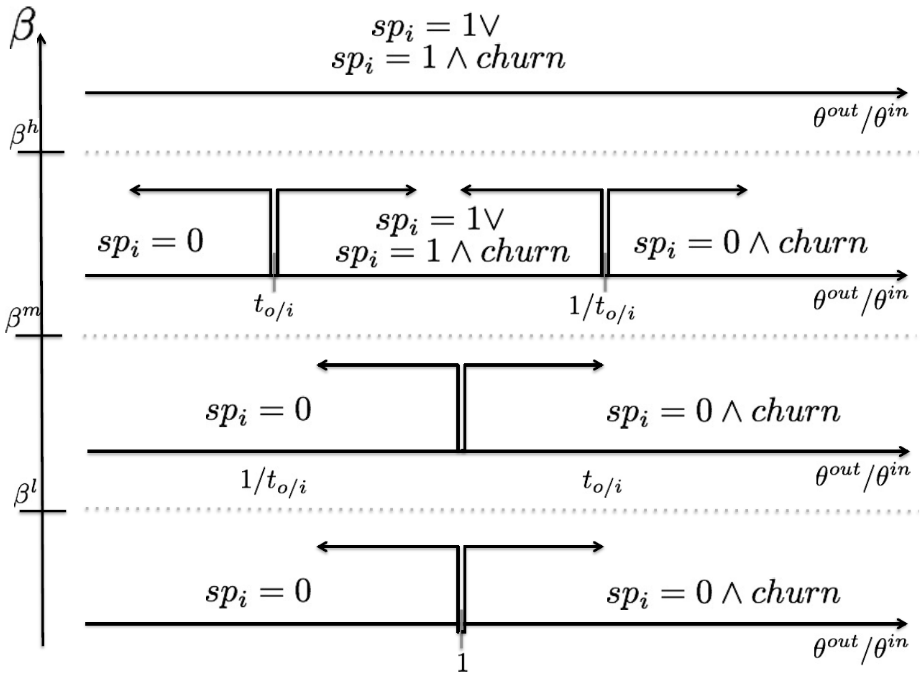

Figure 9 - Optimal tariff plan type ($sp = 0/1$) and churn decision as functions of the "peace of mind" effect (β) and of the ratio of aggregated utility inside (θ^{in}) and outside (θ^{out}) the carrier.

	$\theta^{out}/\theta^{in} < 1$	$\theta^{out}/\theta^{in} \geq 1$	Total
Total	7206	977	8183
	(44748)	(6140)	(50888)
$sp = 0$	3653	238	3891
	(19894)	(1358)	(21252)
$sp = 1$	4769	810	5579
	(24854)	(4782)	(29636)
$Churn$	0.021	0.024	0.022
	(0.144)	(0.152)	(0.145)
$Churn$ & $sp = 0$	0.013	0.012	0.012
	(0.111)	(0.108)	(0.111)
$Churn$ & $sp = 1$	0.028	0.027	0.028
	(0.165)	(0.163)	(0.165)

Number of distinct users (number of observations in parenthesis) for total, $sp = 0$ and $sp = 1$

Average rates (standard deviation in parenthesis) for $churn$, $churn$ & $sp = 0$ and $churn$ & $sp = 1$ (in terms of observations)

Table 1 - Number of users and churn rates as a function of number of friends inside and outside the carrier ($sp = 0$ indicates using a tariff plan that charges a lower price to call inside the carrier, $sp = 1$ indicates using a tariff plan that charges the same price to call inside and outside the carrier).

	(1)	(2)	(3)	(4)
		% costs with call inside the carrier		
% friends churn	-0.0559	-0.218***	-0.0319	-0.156***
	(-1.34)	(-3.63)	(-0.79)	(-2.82)
% friends join	0.0787***	0.140***		
	(2.66)	(3.25)		
Constant	0.925***	0.833***	0.930***	0.839***
	(204.23)	(17.20)	(220.69)	(17.00)
Tariff Plan Dummies	Yes	Yes	Yes	Yes
Month Dummies	Yes	Yes	Yes	Yes
Observations	24788	20732	24788	20732

t statistics in parentheses

* p < 0.1, ** p < 0.05, *** p < 0.01

Table 2 - Shift in calls inside and outside the carrier when friends churn (and the ego does not).

	(1)	(2)
	Churn	Churn
Churn fncfnf	0.000161	
	(0.89)	
Friends join	0.000559	0.00250
	(0.21)	(0.98)
Churn fncfnf 1path		-0.0000331
		(-0.17)
Tariff Plan Dummies	Yes	Yes
Month Dummies	Yes	Yes
Observations	47005	47005

t statistics in parentheses

* p < 0.1, ** p < 0.05, *** p < 0.01

Table 3 - Churn decisions of friends of friends do not affect churn decisions of egos when the subscribers that connect them do not churn.

	(O Probit)	(IV O Probit)	(IV O Probit)
	state	state	state
Main Result			
θ^{out}/θ^{in}	0.218***	0.312***	0.304***
	(7.81)	(6.69)	(6.88)
First Stage			
Churn ffnf		0.483***	
		(11.19)	
Churn ffnf 1path			0.526***

			(12.23)
Constant		-0.443***	-0.459***
		(-5.96)	(-6.63)
cut-off 1	-0.162***	-0.0672***	-0.0704***
	(-8.52)	(-2.73)	(-2.95)
cut-off 2	2.105***	2.198***	2.195***
	(100.07)	(87.59)	(89.23)
Marginal Effect			
$sp_i = 0$	-0.085***	-0.121***	-0.118***
	(-7.82)	(-6.70)	(-6.88)
$sp_i = 1$	0.074***	0.106***	0.103***
	(7.77)	(6.70)	(6.88)
$churn$	0.011***	0.015***	0.015***
	(7.92)	(6.53)	(6.74)
Estimated Probability			
$sp_i = 0$	0.411	0.411	0.411
$sp_i = 1$	0.568	0.568	0.568
$churn$	0.0206	0.0206	0.0206
Month Dummies	Yes	Yes	Yes
Observations	50888	50888	50888

t statistics in parentheses

* p < 0.1, ** p < 0.05, *** p < 0.01

Table 4 - Ordered Probit and Ordered IV Probit results.

	(IV O Probit)	(IV O Probit)
	state	state
Main Result		
θ^{out}/θ^{in}	0.063***	0.119***
	(0.03)	(0.02)
First Stage		
$Churn\ ffnf$	0.080***	
	(0.01)	
$Churn\ ffnf\ 1path$		0.121***
		(0.02)
Constant	0.382***	0.362***
	(0.01)	(0.01)
cut-off 1	-0.536***	-0.514***
	(0.03)	(0.026)
cut-off 2	1.872***	1.895***
	(0.03)	(0.02)
Marginal Effect		
$sp_i = 0$	-0.025***	-0.047***
	(0.013)	(0.009)
$sp_i = 1$	0.022***	0.042***
	(0.011)	(0.008)
$churn$	0.002***	0.004***
	(0.0011)	(0.0008)
Estimated Probability		
$sp_i = 0$	0.423	0.423
$sp_i = 1$	0.564	0.564
$churn$	0.0134	0.0134

	(IV O Probit)	(IV O Probit)
	state	state
Main Result		
$\theta^{out}/\theta^{in}_{strong}$	0.467***	0.453***
	(0.110)	(0.156)
$\theta^{out}/\theta^{in}_{weak}$	0.314***	0.359***
	(0.06)	(0.09)
First Stages		
$Churn\ ffnf\ strong$	0.019***	
	(0.001)	
$Churn\ ffnf\ weak$	0.022***	
	(0.001)	
$Churn\ ffnf\ strong\ 1path$		0.015***
		(0.002)
$Churn\ ffnf\ weak\ 1path$		0.017***
		(0.002)
Constant	-0.192**	-0.219***
	(0.10)	(0.10)
cut-off 1	-3.114***	-3.124***
	(0.247)	(0.247)
cut-off 2	0.647***	0.637***
	(0.244)	(0.0244)
Marginal Effect		
Strong friends:		
$sp_i = 0$	-0.180***	-0.174***
	(0.042)	(0.060)
$sp_i = 1$	0.179***	0.0174***
	(0.042)	(0.060)

	(IV O Probit)	(IV O Probit)
	state	state
Main Result		
$\theta^{out}/\theta^{in}_{weighted}$	0.901**	1.574***
	(0.44)	(0.403)
First Stage		
$Churn\ ffnf\ weighted$	0.446***	
	(0.09)	
$Churn\ ffnf\ 1path\ weighted$		0.394***
		(0.09)
Constant	0.022***	0.022***
	(0.002)	(0.003)
cut-off 1	-0.450***	-0.618***
	(0.03)	(0.032)
cut-off 2	1.782***	1.787***
	(0.02)	(0.02)
Marginal Effect		
$sp_i = 0$	-0.352***	-0.616***
	(0.171)	(0.158)
$sp_i = 1$	0.323***	0.561***
	(0.156)	(0.141)
churn	0.030***	0.055***
	(0.015)	(0.017)
Estimated Probability		
$sp_i = 0$	0.421	0.422
$sp_i = 1$	0.566	0.564

Month Dummies	Yes	Yes
Call Volume	Yes	Yes
Observations	50817	50817

$churn$	0.0004***	0.0004***
	(0.00012)	(0.00015)
Weak friends:		
$sp_i = 0$	-0.121***	-0.139***
	(0.024)	(0.035)
$sp_i = 1$	0.121***	0.0138***
	(0.024)	(0.035)
$churn$	0.0004***	0.0004***
	(0.0003)	(0.0001)
Estimated Probability		
$sp_i = 0$	0.397	0.397
$sp_i = 1$	0.603	0.603
churn	0.0002	0.0002
Month Dummies	Yes	Yes
Call Volume	Yes	Yes
Observations	47999	47999

$churn$	0.0128	0.0136
Month Dummies	Yes	Yes
Call Volume	Yes	Yes
Observations	50888	50888

Table 5 – Robustness checks for the definition of friends.

	Tariff Plan		Churn	
	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)
	state	state	state	state
Main Result				
θ^{out}/θ^{in}	0.317***	0.302**	0.101***	0.080***
	(0.05)	(0.05)	(0.04)	(0.04)
First Stage				
$Churn\ ffnf$	0.476***		0.486***	
	(0.05)		(0.05)	
$Churn\ ffnf\ 1path$		0.518***		0.523***
		(0.05)		(0.05)
Constant	-0.430***	-0.446***	-0.434***	-0.444***
	(0.08)	(0.07)	(0.10)	(0.10)
cut-off 1	-0.405***	-0.411***		
	(0.03)	(0.03)		
cut-off 2			1.673***	1.665***
			(0.04)	(0.04)
Marginal Effect				
$sp_i = 0$	-0.124***	-0.119***		
	(0.020)	(0.018)		
$sp_i = 1$	0.124***	0.119***	-0.008***	-0.006**
	(0.020)	(0.018)	(0.004)	(0.003)
$churn$			0.008***	0.006**
			(0.004)	(0.003)
Estimated Probability				
$sp_i = 0$	0.433	0.433		
$sp_i = 1$	0.567	0.566	0.964	0.964
$churn$			0.036	0.036
Month Dummies	Yes	Yes	Yes	Yes
Call Volume	Yes	Yes	Yes	Yes
Observations	49793	49793	49793	49793

Baseline: symmetric tariff plan	(IV MNL)	(IV MNL)
	state	state
Tariff plan lower price inside		
θ^{out}/θ^{in}	-0.686***	-0.637***
	(0.05)	(0.04)
1st Stage Residuals	0.534***	0.481
	(0.05)	(0.05)
Constant	-0.674***	-0.692***
	(0.03)	(0.03)
Churn		
θ^{out}/θ^{in}	0.265***	0.194**
	(0.10)	(0.10)
1st Stage Residuals	-0.878***	-0.821***
	(0.09)	(0.07)
Constant	-3.219***	-3.172***
	(0.11)	(0.10)
Month Dummies	Yes	Yes
Call Volume	Yes	Yes
Observations	50862	50862

Table 6 – Robustness checks for the ordering of outcomes.

	(O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)
	state	state	state	state	state	state	state	state	state	state	state
Remembering Effect	100%	90%	80%	70%	60%	50%	40%	30%	20%	10%	0%
θ^{out}/θ^{in}	0.263***	0.247***	0.233***	0.222***	0.212***	0.204***	0.197***	0.191***	0.186***	0.181***	0.177***

	(0.04)	(0.04)	(0.04)	(0.04)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
First Stage											
Churn ffnf	0.455***	0.550***	0.633***	0.702***	0.756***	0.797***	0.829***	0.852***	0.869***	0.882***	0.892***
	(0.04)	(0.04)	(0.04)	(0.04)	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)
Constant	-0.413***	-0.443***	-0.445***	-0.430***	-0.404***	-0.375***	-0.345***	-0.316***	-0.289***	-0.264***	-0.241***
	(0.07)	(0.06)	(0.05)	(0.04)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.01)	(0.01)
cut-off 1	-0.447***	-0.456***	-0.464***	-0.470***	-0.475***	-0.479***	-0.483***	-0.486***	-0.488***	-0.490***	-0.492***
	(0.03)	(0.03)	(0.05)	(0.03)	(0.03)	(0.03)	(0.02)	(0.03)	(0.03)	(0.03)	(0.03)
cut-off 2	1.958***	1.951***	1.945***	1.940***	1.936***	1.932***	1.929***	1.927***	1.924***	1.922***	1.921***
	(0.03)	(0.06)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Month Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Call Volume	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	50862	50862	50862	50862	50862	50862	50862	50862	50862	50862	50862

t statistics in parentheses

* p < 0.1, ** p < 0.05, *** p < 0.01

Table 7 - Effect of friends' churn on the state of the ego with remembrance discount factor (using “ffnf” as IV).

	(O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)	(IV O Probit)
	state	state	state	state	state	state	state	state	state	state	state
Remembering Effect	100%	90%	80%	70%	60%	50%	40%	30%	20%	10%	0%
θ^{out}/θ^{in}	0.243***	0.229***	0.218***	0.207***	0.199***	0.192***	0.186***	0.181***	0.176***	0.172***	0.169***
	(0.04)	(0.04)	(0.04)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
First Stage											
Churn ffnf	0.501***	0.595***	0.674***	0.738***	0.786***	0.823***	0.850***	0.870***	0.885***	0.896***	0.904***
	(0.04)	(0.04)	(0.04)	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Constant	-0.439***	-0.455***	-0.446***	-0.421***	-0.390***	-0.375***	-0.326***	-0.297***	-0.270***	-0.246***	-0.223***
	(0.07)	(0.06)	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.01)	(0.01)	(0.01)
cut-off 1	-0.456***	-0.464***	-0.471***	-0.476***	-0.481***	-0.484***	-0.487***	-0.490***	-0.492***	-0.494***	-0.494***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
cut-off 2	1.951***	1.944***	1.939***	1.934***	1.931***	1.928***	1.925***	1.923***	1.921***	1.919***	1.917***
	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Month Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Call Volume	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	50862	50862	50862	50862	50862	50862	50862	50862	50862	50862	50862

t statistics in parentheses

* p < 0.1, ** p < 0.05, *** p < 0.01

Table 8 - Effect of friends' churn on the state of the ego with remembrance discount factor (using “ffnf 1 path” as IV).