

EVA ASCARZA, PETER EBBES, ODED NETZER, and MATTHEW DANIELSON*

Customer relationship management (CRM) campaigns have traditionally focused on maximizing the profitability of the targeted customers. The authors demonstrate that in business settings characterized by network externalities, a CRM campaign that is aimed at changing the behavior of specific customers propagates through the social network, thereby also affecting the behavior of nontargeted customers. Using a randomized field experiment involving nearly 6,000 customers of a mobile telecommunication provider, they find that the social connections of targeted customers increase their consumption and become less likely to churn, due to a campaign that was neither targeted at them nor offered them any direct incentives. The authors estimate a social multiplier of 1.28. That is, the effect of the campaign on first-degree connections of targeted customers is 28% of the effect of the campaign on the targeted customers. By further leveraging the randomized experimental design, the authors show that, consistent with a network externality account, the increase in activity among the nontargeted but connected customers is driven by the increase in communication between the targeted customers and their connections, making the local network of the nontargeted customers more valuable. These findings suggest that in targeting CRM marketing campaigns, firms should consider not only the profitability of the targeted customer but also the potential spillover of the campaign to nontargeted but connected customers.

Keywords: social effects, network effects, field experiment, mobile

Online Supplement: <http://dx.doi.org/10.1509/jmr.15.0442>

Beyond the Target Customer: Social Effects of Customer Relationship Management Campaigns

*Eva Ascarza is Associate Professor of Marketing, Columbia Business School, Columbia University (e-mail: ascarza@gsb.columbia.edu). Peter Ebbes is Associate Professor of Marketing, HEC Paris (e-mail: ebbes@hec.fr). Oded Netzer is Associate Professor of Marketing, Columbia Business School, Columbia University (e-mail: onetzer@gsb.columbia.edu). Matthew Danielson is Principal Research Software Engineer, Ampler Inc. (e-mail: mdanielson@amplero.com). The authors thank Olly Downs and his team at Ampler Inc. for their efforts in making this research possible. The manuscript benefited from comments by Kamel Jedidi; the audience of HKUST, Carlson, Imperial College, Free University Amsterdam, Bocconi, Harvard Business School, Chicago Booth, Georgetown University, Drexel University and the brown bag seminar series and DNA mixer at the Columbia Business School; and the participants of *Marketing Science* 2015, the Workshop in Consumer Analytics at San Pedro de Atacama (Chile) conference, the Yale Customer Insights conference, the 2016 Choice Symposium on Customer Retention, and the European Marketing Academy conference. Peter Ebbes acknowledges research support from Investissements d'Avenir (ANR-11-IDEX-0003/LabexEcodec/ANR-11-LABX-0047). Coeditor: Randolph Bucklin; Associate Editor: Duncan Simester.

At the heart of customer relationship management (CRM) is the concept of customer centricity. Customer centricity emphasizes the idea that firms should recognize that customers are different and target only those customers for whom the marketing effort will pay off (Blattberg, Kim, and Neslin 2008; Fader 2012; Rust and Verhoef 2005). However, increasingly, customers have a variety of means to connect and interact with one another. The number of business settings in which customers are directly connected to other customers through the firm's product or service is rapidly increasing. Examples include communication settings (such as traditional telecom providers, but also more recent services such as WeChat, WhatsApp, or Snapchat), cloud storage and file sharing services (e.g., Dropbox, Google Drive), "sharing economy" marketplaces (e.g., Uber, Airbnb), payment services (e.g., PayPal, Venmo), and online games. In these settings, network effects and network externalities (e.g., Katz and Shapiro 1985)

are often present.¹ Consequently, in such business settings, marketing campaigns that are targeted to specific customers with the goal of changing the behavior of those customers may also indirectly affect the behavior of other, nontargeted, customers. While the social interaction literature has shown that social effects exist in a variety of marketing contexts, the CRM literature and practice has largely ignored such social effects in designing and evaluating CRM campaigns.

The objective of this research is to investigate whether targeted CRM campaigns that are aimed at changing the behavior of specific customers also affect the behavior of the targeted customers' connections, who are not targeted themselves. Our focus on CRM campaigns excludes referral campaigns (Biyalogorsky, Gerstner, and Libai 2001; Chae et al. 2016; Schmitt, Skiera, and Van den Bulte 2011), which are directly aimed at creating a social effect by giving incentives to existing customers or prospective customers as well as their connections. Instead, our focus is on typical CRM campaigns that (1) have no explicit social component to them; (2) have been individually targeted to specific customers, possibly on the basis of those customers' past behavior; (3) are not transferable; and (4) cannot be shared. Examples include targeted retention campaigns and cross-selling and upselling tactics, which are generally characterized by offering incentives (e.g., discounts, free consumption units, premium services) to particular customers.

On the one hand, one could expect a positive effect of such typical CRM campaigns on nontargeted customers because these customers may derive more value from the product or service due to positive network externalities. There are several reasons why we may expect positive network externalities. First, a successful CRM campaign may grow the user base because new users start using the service. At the same time, a CRM campaign may incentivize existing users to stay with the firm. In both cases, the focal customer experiences a larger and potentially more active (ego) network (e.g., Aral and Walker 2011; Nitzan and Libai 2011). Second, network externalities are often associated with higher level of product usage. For example, Aral and Walker (2011) show that in a social networking website, network externalities lead to higher levels of adoption of a new feature as well as an increase in usage of that feature among connected friends. Similarly, Trusov, Bodapati, and Bucklin (2010) show a social effect on users' activity in a social network site. Manchanda, Packard, and Pattabhiramaiah (2015) find that stronger ties in online customer communities lead to higher levels of expenditure within the community. Taken together, these results suggest that CRM campaigns that increase usage among the targeted customers may also increase the usage level of the focal customers' connections.

On the other hand, one could also expect a negative (social) effect of the campaign because the benefit/incentive (e.g., discount, free consumption) is offered only to the targeted customers and is not available to the nontargeted customers. In turn, if the targeted customer talked about the benefits of the campaign with his or her connections (i.e., initiated word of mouth about the campaign), then the nontargeted customers might become dissatisfied with the service due to perceived "peer-induced" unfairness (Li and Jain 2016; Nguyen and

Simkin 2013). Such decrease in satisfaction among nontargeted customers could result in a reduction of consumption and an increase in churn among those customers.

The CRM literature and practice has traditionally ignored potential social effects of marketing campaigns, thus implicitly assuming that such effects either do not exist or are too small to be of managerial relevance. An exception can be found in Lemon and Seiders (2006), who call for firms to consider not only the core, or targeted, customers but also what they call the "augmented customers." They postulate that firms' marketing actions affect not only the targeted customers but also the augmented customers. In this research, we investigate this issue and explore which types of customer behavior (e.g., usage, churn) of nontargeted customers that are connected to targeted customers can be (socially) affected by a CRM campaign. We estimate the direction and magnitude of these social effects and discuss the possible mechanisms through which such an effect propagates through the network. Furthermore, we quantify the economic value of the social spillover effect. Our calculations suggest that the dollar value of the campaign spillover may be substantial and should not be ignored by CRM practice.

To investigate the potential social effect of CRM campaigns, we run a field experiment in the context of a telecommunication service provider. We randomize a targeted CRM campaign among current customers such that the focal customers are offered free money to top up, or refill, their prepaid telephone plan. We then analyze the activity (cell phone usage and churn) of both the egos (i.e., targeted customers) and their alters (i.e., customers connected to the targeted customers who themselves were not targeted; hereafter, we use the terms "focal" and "ego" interchangeably, as well as "connections" and "alters").

One important benefit of implementing the experiment in a telecommunication setting is that we can ensure that the campaign incentive is made available to the targeted (focal) customers only and *not* to their connections; unlike coupons or referral-type promotions, this top-up credit is nontransferable.

We empirically demonstrate that the effect of the targeted CRM campaign propagates beyond the targeted customers in terms of both usage and churn. In particular, we find that the (nontargeted) connections of egos in the treatment group have significantly higher consumption levels than the (nontargeted) connections of egos in the control group. We show that the campaign caused a 35% increase in usage among the targeted customers. On top of that, the campaign caused a 10% increase in usage among their connections, *who were not targeted themselves*. Furthermore, we find that the campaign reduced churn among the connections of targeted customers. On the basis of our findings, we estimate that the incremental profit of the "social spillover effect" is, on average, \$.85 per connection across the 12 weeks following the campaign.

One question that naturally arises is, if connections of targeted customers did not receive any direct benefit from the campaign, why do they consume more and why do they churn less than connections of nontargeted customers? We investigate this question by leveraging the randomization of our research design using an instrumental variable (IV) regression approach. We show, consistent with a network externality account, that an increase in communication between the focal (targeted) customers and their connections causes an increase in the consumption of the connections and reduces their churn.

¹In addition to the aforementioned examples, many more traditional business settings, such as retailers, banks, and gyms, may also exhibit network externality.

Furthermore, the strength of social relationship between the ego and his or her alters moderates the magnitude of the social effect. In particular, the campaign spillover effect is larger for egos and alters with stronger ties to each other. These findings are all in support of positive network externalities of the CRM campaigns.

Our research complements the work on CRM and database marketing (e.g., Boulding et al. 2005; Neslin et al. 2006; Reinartz, Krafft, and Hoyer 2004) by quantifying the effects of CRM campaigns beyond the target customer. Our work is also related to the literature on social influence, which has mainly focused on the contagious effect of new product introduction and customer acquisition (e.g., Haenlein and Libai 2013; Iyengar, Van den Bulte, and Valente 2011; Schmitt, Skiera, and Van den Bulte 2011). In that respect, our research is most closely related to the work of Nitzan and Libai (2011), who demonstrate that churn behavior may be contagious. We differ from their work in three important ways. First, our focus is not on the contagion of churn per se but rather the propagation of a *change* in customer behavior (including both churn and usage) in response to a targeted marketing campaign. Second, we leverage a randomized field experiment to estimate the causal effect of the campaign on the nontargeted connected customers. Finally, using our modeling approach, we quantify the magnitude of the campaign spillover as well as the monetary value of this social effect.

Our findings have clear implications for marketers. In targeting CRM marketing campaigns, firms need to consider not only the profitability generated by the targeted customers, but also the potential spillover of the campaign to nontargeted, but connected, customers. As we discuss in the final section, we believe that this implication is not limited to telecommunication but generalizable to many other contexts in which network externalities are present.

This article continues as follows. In “Research Setting,” we describe and discuss the research framework and the field experiment. In “Results,” we quantify the impact of the targeted promotions on the targeted customers, as well as on their (nontargeted) alters. Then, we examine the mechanism underlying the social spillover effect from the CRM campaign. “Managerial Relevance of the Social Effect” focuses on the managerial implications of this research, quantifying the consumption spillover and estimating the monetary value of the social effect. Finally, we conclude in “General Discussion” with a discussion of the theoretical and practical implications of our work.

RESEARCH SETTING

Identifying Social Effects

To investigate the social effect of targeted marketing campaigns, one has to look beyond the targeted customers and measure the changes in activity among the customers connected to them. With the appropriate individual-level data and sufficient variation in marketing actions, firms can easily measure the effectiveness of their promotions on the targeted customer (e.g., Gupta 1988; Neslin, Henderson, and Quelch 1985). However, measuring the effect of the promotion on the nontargeted, but connected, customers (i.e., measuring the social effect) is more complicated because identification of social influence from observational data is challenging (e.g., Manski 1993; Nair, Manchanda, and Bhatia 2010; Nitzan and

Libai 2011; Shalizi and Thomas 2011). Consider a firm running a marketing campaign and suppose that the targeted customers increase their activity after the campaign. Suppose that the firm also observes an increase in activity among the customers connected to the targeted customers. Can one directly attribute the increase in activity of the connections to a social effect of the marketing campaign? The answer is no; such observed similarity in behavior between a targeted customer and his or her nontargeted connections after the campaign could also be explained by homophily (i.e., unobserved similarities in customers’ preferences, such as common preferences for a particular service provider), correlated unobservables (i.e., a common shock affecting the behavior of the connected customers, such as improved quality of the service/product in a certain area or other unobserved marketing actions of the firm or its competitors), or simultaneity/reflection (i.e., alters’ behavior affecting focal’s behavior). Thus, using observational data alone, it is difficult to conclude that the observed changes in the nontargeted connected customers’ behavior were *caused* by the marketing campaign.

To address the challenges in making causal claims with respect to social effects, we conducted a randomized field experiment in collaboration with a telecommunication provider in which a set of randomly selected focal customers received a marketing promotion that incentivized them to change their own behavior (treatment group), while other focal customers did not receive the promotion (control group). This intervention induces exogenous variation in the behavior of the focal customers, which we leverage to identify the causal effects of interest. In particular, we investigate the social effect of the marketing campaign by comparing the behavior of the customers connected to the focal customers in the treatment group with the behavior of customers connected to the focal customers in the control group. Importantly, none of these connections received the treatment themselves. Thus, we employ a “peer encouragement design,” which is characterized by randomly encouraging certain behaviors in a set of nodes—that is, the egos—to analyze the effects of the encouragement (treatment) on the nodes’ peers—that is, the alters (e.g., Aral 2016; Bapna and Umyarov 2015; Hinz et al. 2011). The randomized nature of our research design addresses the potential issues of homophily, correlated unobservables, and simultaneity/reflection.

Randomized Field Experiment

The field experiment was conducted in Australia, where the penetration of cell phones is higher than 130%. During the period of the experiment, there were three main providers in this market; the company we collaborate with was the second largest in terms of market share, with a customer base of approximately 10 million people. This is a “calling party pay” market; in other words, the person who initiates the call/text incurs all the costs and the receiver is not charged.

Focal customers (egos). Customers selected to participate in this marketing campaign belong to a 28-day non-rollover prepaid plan with unlimited in-network voice, domestic short message service texts, and access to major social network platforms (e.g., Facebook, LinkedIn). That is, every time a customer adds (a minimum amount of) credit to his or her account, the customer has unlimited in-network activity for a period of 28 days. During that time, the balance/credit can be

used to call out of network or internationally and to download or upload data. If a customer reaches zero balance or does not recharge the account within 28 days, then his or her balance is set to 0 and the account is suspended. Suspended accounts can receive calls and texts but cannot initiate any type of communication. Once a customer is suspended, he or she can receive calls/texts for a period of six months. At any time during that period, the customer can become active by adding credit to the account. However, if the customer does not add credit to the account within six months, then the provider cancels the service and the account is terminated. We note that we do not observe this passive churn among our focal customers because our data set is shorter than six months. This contrasts with churn among postpaid customers at the same company who cancel the service by calling the provider (active churn).

The customers selected for this experiment (i.e., focal customers in the treatment and control groups) were customers who were active at the time of the intervention (i.e., they had credit and could initiate any kind of communication) but would be suspended in the following week unless they added credit to their account. The campaign's goal was to prevent the targeted customers from going into suspension by encouraging them to top up their account. Given that the focus of our study is the social effect of marketing promotions, we only consider focal customers who had at least one connection on the same network (further details about the definition of a connection are provided in the "Connected customers (alters)" section).

Intervention. We select 1,041 customers with the characteristics described earlier and randomly split the sample into a treatment (63%) and a control (37%) group. Customers in the treatment group received a text message offering free credit if the customer replied "Yes" to the text message. All customers received the same credit incentive. Once the promotional text was sent, the promotion would be valid for seven days, and upon acceptance, the bonus credit would expire after seven days. It should be noted that the intervention was only based on the targeted customers' calling plan and behavior prior to the campaign and not in any way based on the behavior of the targeted customers' connections. In that respect, this was a typical CRM campaign aimed at increasing engagement among current customers.

Connected customers (alters). One of the advantages of using a telecommunication network for our study is that we can perfectly observe the communications between customers, which eliminates the need for constructing (sometimes noisy) proxies for connections between individuals (Hill, Provost, and Volinsky 2006; Nitzan and Libai 2011). At the moment of the intervention, we identify all customers who have communicated with each focal customer (either treatment or control) at least twice, by either call or text, in the month prior to the experiment. The edges in our ego networks are unweighted and undirected.² Defining the ego networks before the campaign ensures exogeneity of the network with respect

to the treatment. We consider only the connections who belong to the same telecommunication provider because they are the only customers for whom we can observe behavior both before and after the treatment. Note that we only track the behavior of the first-degree connections of the focal customers. In theory, the social effect could also reach second-degree or higher-order connections. However, looking only at first-degree social effects allows us to simplify the analysis, avoid network sampling issues (Ebbes, Huang, and Rangaswamy 2015), and be conservative in our estimated social effect of the marketing campaign.

As with any experiment conducted in a network, we face the challenge of contamination or interference (e.g., Aral 2016; Eckles, Karrer, and Ugander 2014; Fienberg 2012) if the control group is exposed to the treatment through their connections. In our experiment, contamination could occur if (1) the alters were treated directly or (2) the alters were connected to customers, other than the ego, who were treated or who were connected to treated customers. In both these cases, the stable unit treatment value assumption (Rubin 1980) would be violated as the egos in the control group could be exposed to the treatment indirectly through the alters. This issue would be particularly problematic in the case of first-degree contamination, which happens when the alter is treated, or second-degree contamination, which happens when the alter's connections are treated. To minimize this concern, we cross-reference all egos and alters in the sample and find that only .60% of alters (i.e., 32 out of 5,308) received a marketing promotion during the time of the study. Moreover, we looked at the alters' first-degree connections and find that 1.60% of alters (85 out of 5,308) had (at least) one connection who received the marketing promotion. While this is a small subset of our entire network, we removed the complete ego networks to which these alters belong. Thus, our final sample includes 961 ego customers and 4,700 alters, which we use for all subsequent analyses.³ Importantly, we can assure that this reduced sample is free of first- and second-degree contamination. Admittedly, there is still a possibility of higher-degree contamination in the control group if, for example, a third-degree connection of a focal customer in the control group (or, equivalently, a second-degree connection of an alter in the control group) was treated. However, given the small incidence of treated alters, the size of the total network (~10 million customers), the scope of the experiment (~1,000 focal customers), and the fact that the contamination would need to be transmitted through at least three nodes, we believe that the likelihood of such higher-degree contamination and its potential biasing effect on estimating the treatment effect is very low.

Behavioral Data

We collect two types of activity data: (1) individual-level activity data of the egos and (2) individual-level activity data of the alters. We track the behavior of egos and alters for 16 weeks, including 4 weeks before and 12 weeks after the intervention.

²We set a threshold of at least two communications to avoid "random" call recipients (e.g., a taxi, a restaurant) from being part of the ego network. While the weight and directionality of the edges are not used for the network formation, we leverage this information in the section "Managerial Relevance of the Social Effect" when we examine the impact of the strength of ties.

³Note that instead of removing only the "contaminated" alter, we use a more conservative approach by removing the entire ego network containing such a node. We also reran our analyses on the larger sample without removing the entire ego networks but removing only the contaminated alters. Our main conclusions did not change.

Ego behavior. For each ego customer, we observe weekly activity in the form of texts, calls, and minutes. Each of these metrics is split by inbound and outbound activity (e.g., calls the customer makes and calls the customer receives). We also observe whether an ego customer is suspended in a particular week (i.e., whether the customer can initiate calls/text). Regarding churn behavior, customers can churn at any time by calling the provider to close their account, permanently deactivating the SIM card associated with the account. However, this behavior is very rare among the ego customers because of the type of (prepaid) plan these customers have. Therefore, we ignore churn behavior among ego customers.⁴ Note that churn behavior is relevant for the alters because many of these customers are on postpaid plans. In postpaid plans, churn is more prevalent and serves as an important key performance indicator for the provider. Thus, for ego customers, we focus on measures of activity, including suspension, minutes, calls, and texts. Table 1 shows summaries of ego behavior (weekly averages) during the four weeks before the intervention.

From Table 1 it follows that the usage variables have skewed distributions, as is evident from considerable differences between the mean and median of these variables. We address this issue in our difference-in-differences (hereafter “diff-in-diffs”) regression approach by using the log-transformed variables.

Alter behavior. We observe weekly usage, suspension, and churn for each of the alters before and after the experiment. Table 2 summarizes the alters’ activity during the four weeks prior to the intervention. For suspension behavior, we report, for each ego, the proportion of alters of the ego who were suspended at the moment of the intervention. As we do for the ego behavior variables, we log-transform the alters’ usage variables.

Comparing Tables 1 and 2, we see that, on average, alters are more active than egos. This difference is not surprising because the egos selected for the experiment were customers who were at risk of being suspended at the time of the experiment, which suggests that these customers were less active during the weeks prior to the intervention.

In addition to usage and suspension, we also observe the size of each ego network, that is, the number of alters each ego had at the moment of the intervention, as well as the number of connections each of the alters had during the four weeks prior to the intervention. On average, egos had 4.89 alters, and alters had 6.57 connections.⁵ Note that these metrics are “static” (i.e., they do not change over time) and exogenous to the treatment because they were computed according to behavior prior to the intervention.

⁴We do not treat suspension as churn because many of the suspended customers reactivate after some weeks of suspension. In our sample, 69.5% of the egos were suspended at some point during the 12 weeks following the experiment, of which 34.7% reactivated. Thus, suspension in this context relates more to usage than to churn.

⁵Due to limitations in the company’s database, the ego and alter degree (i.e., number of connections at the moment of the intervention) are not computed in the same way. A connection of an ego is defined as a customer who communicated with the ego at least twice during the four weeks prior to the experiment, whereas to calculate the number of connections of an alter, we count for each week the number of customers who communicated with the alter at least once in that week. We then average the number of connections of each alter across the four weeks prior to the campaign. Thus, a connection of an alter is defined at the weekly level (and we compute it for each of the four weeks prior to the experiment), whereas a connection of an ego is defined at the monthly level (and we compute it once, at the moment of the intervention).

Table 1
DESCRIPTIVE STATISTICS OF EGO BEHAVIOR BEFORE THE EXPERIMENT

	<i>M</i>	<i>SD</i>	<i>Percentiles</i>			<i>N</i>
			<i>25th</i>	<i>50th</i>	<i>75th</i>	
Minutes inbound	9.6	40.3	.2	1.9	7.7	961
Minutes outbound	35.1	62.2	4.8	19.5	43.9	961
Calls inbound	3.4	6.0	.3	1.3	4.0	961
Calls outbound	22.5	28.2	5.3	14.3	28.3	961
Texts inbound	35.9	95.3	1.0	6.3	26.5	961
Texts outbound	71.3	154.1	3.8	16.3	62.8	961

Notes: Usage metrics are weekly averages (during the four weeks before the intervention) that are then averaged across customers.

RANDOMIZATION

Before we can rely on the experiment to estimate the social effect of the campaign, we need to verify that the customers assigned to the treatment group (i.e., those to be targeted by the campaign) are similar in terms of their usage prior to the campaign to those assigned to the control group (i.e., those not targeted by the campaign). We do not expect any difference between these two groups of customers because of the random assignment between treatment and control. Table 3 shows the sample means of the treatment and control groups as well as statistical tests for the difference in means between the two groups for the different customer activities. We find that for all types of behaviors, the average activity in the control and treatment groups are not statistically different (*p*-values shown in the rightmost column in Table 3). We note that the results in Table 3 are for the log-transformed usage variables. We obtain similar result when we replicate the analyses for the untransformed (before log) variables (see Web Appendix A1). Thus, we conclude that the randomization between the control and experimental groups was well executed.

RESULTS

We now turn to investigate the effect of the marketing intervention on customer behavior. While our main goal is to measure the effect of the treatment on the alters, we first analyze the effect of the marketing campaign on the targeted customers. It is important to establish the effect of the marketing campaign on the targeted customers because it would be unrealistic to expect any social spillover without an effect of the campaign on the focal customers.

Effect of the Marketing Campaign on Targeted Customers (Egos)

We evaluate the effect of the marketing campaign on the focal customers by analyzing two managerially relevant behaviors, namely, suspension and usage (i.e., minutes, calls, and texts). Recall that the purpose of the campaign was to keep these customers active (prevent suspension) and increase their usage levels. We first present several model-free analyses before statistically estimating the treatment effect through diff-in-diffs regression models.

Model-free analyses for ego usage and suspension. We start by looking at suspension among ego customers (Table 4). As expected, the campaign was successful at preventing suspension: while 47.5% of the customers in the control group

Table 2
DESCRIPTIVE STATISTICS OF ALTER BEHAVIOR BEFORE
THE TREATMENT

Usage	M	SD	Percentiles			N
			25th	50th	75th	
Minutes inbound	58.8	102.7	10.9	30.2	68.0	4,700
Minutes outbound	69.3	132.4	10.7	33.0	77.7	4,700
Calls inbound	25.5	33.2	7.0	15.8	31.5	4,700
Calls outbound	46.5	66.2	11.3	27.0	55.8	4,700
Texts inbound	169.7	255.7	28.0	74.5	198.2	4,700
Texts outbound	127.3	222.0	10.8	43.5	141.9	4,700
Suspension						
Percentage of alters suspended	12.1	21.1	.0	.0	17.9	961

Notes: Usage metrics are weekly averages (during the four weeks before the intervention) that are then averaged across customers. Suspension is computed at the moment of the intervention, then averaged across customers.

were suspended in the week following the intervention, only 35.4% of the customers in the treatment group were. The difference between the two groups is statistically significant ($p < .01$). We also compare the number of customers who are suspended at the end of the observation period (week 12). We observe that treated customers are less likely to be suspended than those in the control group (48.4% vs. 57%; $p < .01$) even 12 weeks after the intervention, implying that the lack of suspension caused by the campaign persisted even after the customers used all the free minutes. Next, we investigate ego usage (or consumption), considering only “outbound” consumption (i.e., calls that the ego initiates) because it better reflects decisions made by the focal customer.⁶ Table 4 shows the difference between pre- and postintervention activity for the treatment and control egos.

We compute for each ego the difference between the ego’s posttreatment weekly consumption and his or her average weekly consumption in the four weeks before the intervention (with both the original and the log variables). As shown in Table 4, customers in both the treatment and control groups decreased their usage in the 12 weeks following the campaign. This trend is consistent with the targeting selection of the marketing campaign that was aimed at customers who were about to be suspended and with the overall downward trend in usage that the focal firm experienced. More importantly, for all usage variables (minutes, calls, and texts), the decrease in activity is smaller, on average, for customers in the treatment condition than for those in the control condition. We find that the treatment effect is positive, and statistically different from zero, for number of minutes and number of calls, but not for number of texts. Therefore, the campaign had an overall positive effect on the usage of the targeted egos during the 12 weeks following the campaign.

Because the treatment offered customers free money for making calls or sending texts, the question is whether

Table 3
RANDOMIZATION CHECK IN ALL OBSERVED VARIABLES IN THE
FOUR WEEKS BEFORE THE EXPERIMENT

	Control		Treatment		Difference		
	M	SE	M	SE	Difference	SE	p-Value
<i>Ego Usage (log)</i>							
Inbound texts	1.30	.07	1.32	.05	-.03	.08	.76
Outbound texts	2.73	.08	2.75	.06	-.03	.10	.79
Inbound minutes	1.02	.05	1.01	.04	.01	.06	.91
Outbound minutes	2.54	.07	2.59	.05	-.04	.08	.58
Inbound calls	2.14	.09	2.17	.07	-.04	.11	.73
Outbound calls	2.81	.10	2.91	.07	-.10	.12	.40
<i>Alter Usage (log)^a</i>							
Inbound texts	2.43	.09	2.54	.07	-.11	.12	.37
Outbound texts	2.53	.10	2.59	.07	-.06	.12	.65
Inbound minutes	1.99	.08	2.10	.06	-.11	.09	.25
Outbound minutes	2.39	.09	2.44	.07	-.05	.11	.65
Inbound calls	3.18	.11	3.25	.08	-.08	.14	.59
Outbound calls	2.67	.11	2.75	.08	-.07	.14	.61
<i>Alter Suspension</i>							
Percentage of alters suspended	13.16	1.19	11.46	.82	1.70	1.41	.23
<i>Other Covariates</i>							
Degree (number of alters)	5.37	.46	4.61	.20	-.76	.44	.08
Number of connections (of the alters)	5.39	.18	5.12	.14	-.27	.23	.24

^aTo check the randomization at the randomized unit level, we test the differences in alter usage and degree with a between-effect regression (i.e., averaging alter usage at the ego level across alters and weeks and regressing the treatment dummy on those averages). We also estimate these differences at the alter level, including a random effect for egos. The random effect regressions provided similar results.

the observed positive treatment effect on usage across the 12 weeks post intervention is solely driven by the monetary incentive given as the treatment. To investigate this, we analyze the treatment effect at the weekly level for each week following the intervention. If the effect is fully driven by the free money, we should observe a decrease in suspension and an increase in usage only during the first two or three weeks following the intervention. Recall that the customer had one week to accept the offer and another week to use the free money. Thus, any effect observed three or more weeks after campaign cannot be attributed to the free money offered in the campaign. Figure 1 shows the average differences in weekly consumption (individual differences) for the entire posttreatment period for the treatment (solid line) and control groups (dashed line). It can be seen that the treatment group exhibits higher usage levels immediately after the intervention (consistent with customers using the free credit), but this effect persists until the end of the observation window (12 weeks after the treatment), indicating that the intervention increased ego consumption beyond the economic incentive of the promotion. Figure 1 also shows the weekly suspension rate in the treatment and control groups, confirming that the campaign also had a lasting effect in preventing suspension well beyond the first week. In summary, Figure 1 illustrates that the campaign had a positive effect in both the short and the long run for both usage and suspension.

⁶While customers could choose not to answer certain calls, the “inbound” calls/texts are mainly determined by the customer who initiates the activity, not by the receiver. We leverage the information obtained from incoming communications when characterizing types of relationships between egos and alters in follow-up analyses later in the article.

Table 4
AVERAGE EGO USAGE AND SUSPENSION
POSTINTERVENTION

	Control	Treatment	Difference	p-Value
Suspended status in week 1	47.6%	35.4%	-12.2%	.00
Suspended status in week 12	57.8%	48.4%	-9.4%	.00
Difference in minutes	-13.17	-6.62	6.54	.04
Difference in calls	-8.89	-5.50	3.40	.02
Difference in texts	-30.46	-28.79	1.68	.83
Difference in log(minutes)	-1.23	-.96	.27	.00
Difference in log(calls)	-1.11	-.89	.21	.01
Difference in log(texts)	-1.34	-1.16	.18	.05

Notes: Usage behavior includes all outgoing communications initiated by the ego during the 12 weeks following the intervention.

We investigate these effects more formally in the next subsection.

Diff-in-diffs regression model results for ego usage. To statistically test for the observed differences in usage between the treatment and control groups, we estimate linear regression models that include the individual difference between pre- and postcampaign usage as the dependent variable, and treatment and week dummies as the independent variables. We use the log-transformed variables to account for the skewed distributions of the activity variables (see Table 1).⁷ We split the data into two observation windows to investigate the “short-term” (weeks 1–6) and “long-term” (weeks 7–12) effects of the campaign. More specifically, for each metric of ego usage, we estimate the following diff-in-diffs models for the short and long term, respectively:

$$(1) \Delta y_{it}^{ego} = \alpha_0 + \alpha_1 T_i + \sum_{\tau=2}^6 \alpha_{\tau} D_{\tau t} + \varepsilon_{it} \text{ for } t = 1, \dots, 6 \text{ and}$$

$$(2) \Delta y_{it}^{ego} = \beta_0 + \beta_1 T_i + \sum_{\tau=8}^{12} \beta_{\tau-6} D_{\tau t} + \varepsilon_{it} \text{ for } t = 7, \dots, 12,$$

where Δy_{it}^{ego} is the difference between individual i 's usage in period t and his or her (individual) average usage during the four weeks prior to the treatment. The term T_i is an indicator variable that takes the value 1 if the customer was part of the treatment group and 0 if the customer was part of the control group; $D_{\tau t}$ is a dummy variable for week t that equals 1 when $\tau = t$ and 0 otherwise. The error term ε_{it} has mean 0 and variances $\sigma_{\varepsilon,s}^2$ and $\sigma_{\varepsilon,l}^2$ for the short term and the long term, respectively. We use panel-corrected standard errors to account for potential serial correlation in the error terms in Equations 1 and 2 (e.g., Hoechle 2007). We find positive and statistically significant coefficients of the treatment dummies, which confirms the “model-free” evidence seen in Table 4 and Figure 1 that the treated egos used the telecommunication service more than the nontreated egos both in the short term (Table 5) and in the long term (Table 6).

The analysis in this subsection demonstrates that treated customers overall consume more than nontreated customers.

⁷Web Appendix A2 replicates Figure 1 for the log-transformed variables that were used in the diff-in-diffs regression models. A full description and motivation of the diff-in-diffs regression approach is given in Web Appendix A3.

Furthermore, the difference in consumption caused by the campaign extends beyond the “free money” given to the targeted customers, as the effect lasts for up to 12 weeks after the campaign, whereas the credit incentive was available only for 2 weeks. Next, we investigate the social effects of the marketing intervention. That is, we investigate whether the campaign affected the alters, that is, customers who were not targeted themselves but were connected to those who were targeted by the CRM campaign.

Effect of the Marketing Campaign on Nontargeted Customers (Alters)

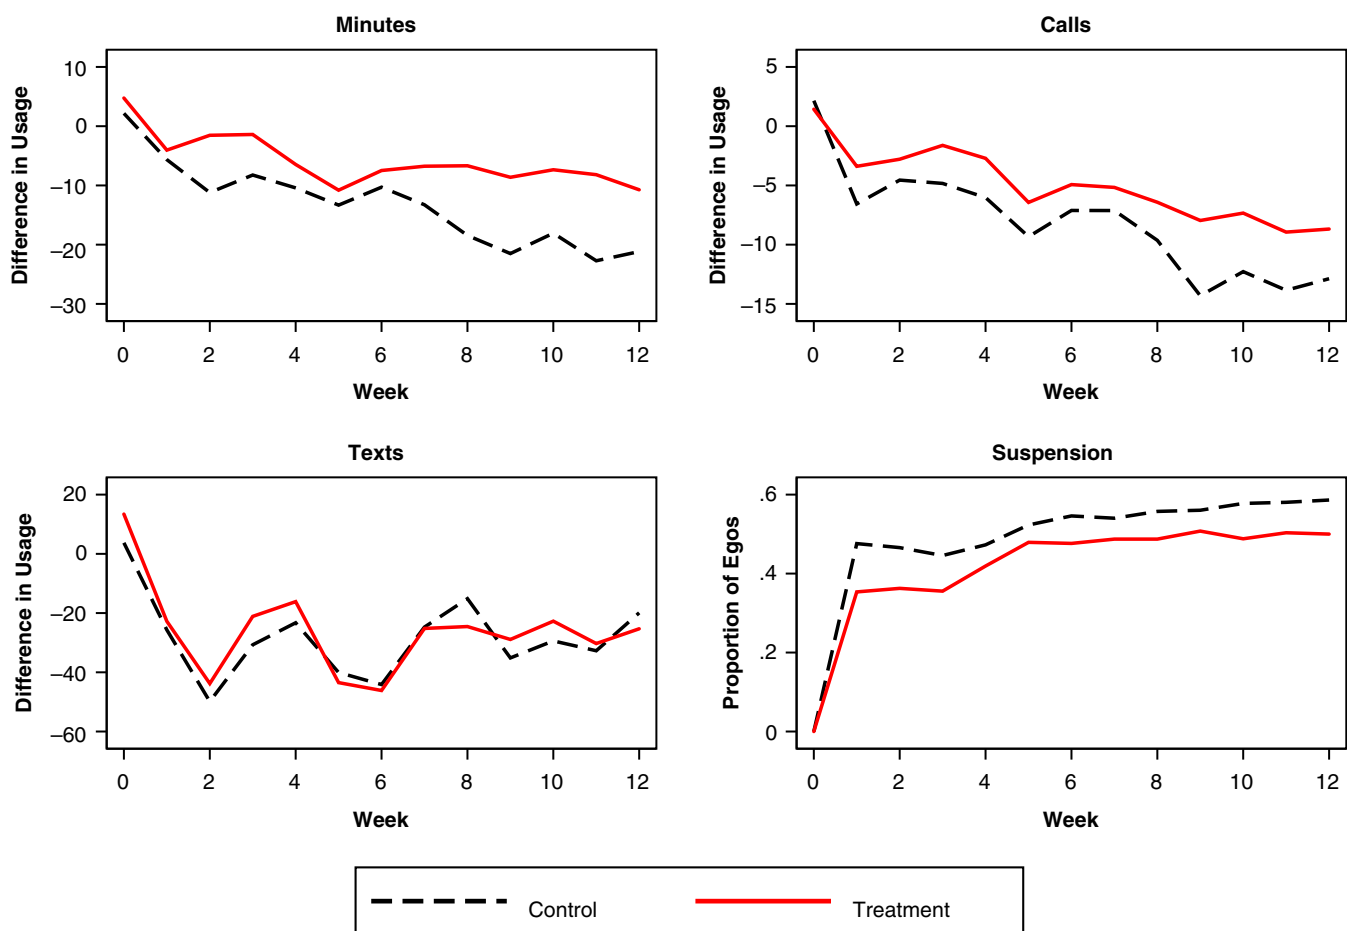
The main goal of our study is to quantify the impact of the targeted marketing intervention on nontargeted connected customers, in terms of both whether they are more likely to stay with the firm (i.e., not to churn) and whether they change their level of activity (usage). While most CRM marketing campaigns are designed and evaluated with respect to their effect on the target customers only—thus implicitly assuming that social spillover effects do not exist—we postulate that in contexts in which network externalities are present, a targeted marketing campaign will likely propagate through the network, therefore also indirectly affecting (connected) customers who were not originally targeted.

Consistent with the social interaction literature (e.g., Aral and Walker 2011; Nitzan and Libai 2011; Trusov et al. 2010), and given that the campaign positively influences ego usage and negatively influences ego suspension, one may expect a positive spillover of the marketing campaign from the egos to their alters. However, unlike contagious effects in the adoption of new products, traditional targeted CRM campaigns (that are not referral-type campaigns) are not likely to generate word of mouth (WOM) about the campaign itself. Even if a campaign does create WOM, one could expect to find the opposite effect. That is, if the ego discusses the campaign with his or her alters, then the alters would find out that they did not receive the benefits of the campaign, which could lead to a negative effect of the campaign on the nontargeted customers due to perceived unfairness (Li and Jain 2016; Nguyen and Simkin 2013).

We propose that in the presence of network externalities, a successfully targeted marketing campaign is likely to influence the behavior of the nontargeted connected customers in a positive way, but not necessarily through WOM. First, customers derive value from having more connections belonging to their network because, among other reasons, calling/texting customers in-network is cheaper (Nitzan and Libai 2011). Hence, if a campaign is successful at retaining targeted customers, then retention might also be higher among customers connected to the targets. Second, customers derive higher value when other customers who are connected to them use the service more (Trusov et al. 2010). As a consequence, a campaign that increases consumption among targeted customers might also increase consumption among their connections, especially when the connections perceive a more active network (e.g., getting more calls from targeted customers).

In summary, we would expect that the targeted marketing campaign causes the alters to (1) churn less and (2) exhibit higher levels of activity. Regarding the latter, we acknowledge that in the context of telecommunication, observed increase in activity among alters could be due to pure reciprocity in

Figure 1
POSTTREATMENT EGO USAGE AND SUSPENSION BY TREATMENT CONDITION



Notes: "Usage" refers to differences between weekly consumption before and after the intervention; "suspension" refers to average number of customers in suspended status in a given week.

calling/texting behavior. That is, an alter's usage might increase just because the alter was returning the calls/texts received from his or her ego, not because the alter derived higher value from the network. To ensure that a potential positive difference in usage is not purely due to reciprocity in calls, we also create a more conservative metric for alter usage that ignores the calls that the alter makes to the ego (indicated by "excluding ego" in Table 7). Using this metric as an outcome variable also helps us separate the possible confound (Manski 1993) between the variable used to define the network (in our case, calls or texts between egos and alters prior to the intervention) and the outcome variable (alter usage *excluding* the calls made to the ego).

As we did for the egos, we first present "model-free" evidence of the treatment effect on alter usage, suspension, and churn (at both the aggregate and disaggregate levels), followed by diff-in-diffs regression models to estimate the treatment effect.⁸ Note that in contrast to most studies that investigate

social effects, we run a randomized field experiment in which the intervention is exogenously manipulated. In addition, we make sure that each alter is connected to only one ego customer. As such, we can estimate the causal effect of the treatment on alter behavior by simply comparing the activity and churn of the alters whose egos are in the treatment group with the activity and churn of the alters whose egos are in the control group.

Model-free analyses for alter usage, churn, and suspension. Figure 2 shows the average outbound activity of the alters (including and excluding communications with the ego) during the posttreatment weeks, as well as their churn and suspension rates.⁹ Alters are grouped according to whether or not their ego received the treatment. For example, the solid lines in the figures represent the weekly average (across alters) of the individual

⁸For the remainder of the article, we use "minutes" as a variable representing alter usage. The results of analyzing "calls" and "texts" are similar (see Web Appendix A4).

⁹The model-free time series plots for the rest of the activities, as well as the log-transformed variables, are presented in Web Appendix A2. We note that, like the time series plots for the egos, the usage trends are negative. That is, during the 12 weeks postintervention, the average usage declines compared with the four weeks before the intervention. Discussions with the data provider confirmed that the mobile operator was experiencing an overall decline in usage across the customer base during this period.

Table 5
SHORT-TERM EFFECT OF TREATMENT ON EGO USAGE
(WEEKS 1–6 AFTER THE TREATMENT)

	Outbound Usage		
	Minutes	Calls	Texts
Treatment	.235*** (.045)	.188*** (.039)	.150*** (.047)
Constant	-.823*** (.061)	-.675*** (.052)	-.795*** (.062)
Week dummies	Yes	Yes	Yes
Observations	5,702	5,702	5,702

*** $p < .01$.

Notes: Data reflect short-term effects of treatment on ego usage, with robust standard errors in parentheses. The number of observations is 6 (weeks) \times 961 (egos), excluding egos who canceled their contract in a particular week.

differences between pre- and postcampaign consumption for alters whose egos were treated, whereas the dotted line are the averages corresponding to those alters whose egos are in the control condition.

Figure 2 (top row) suggests that alters whose egos were treated tend to make calls for more minutes than alters whose egos belong to the control group, and this positive difference in consumption is persistent over time. Moreover, this difference seems not to be driven by reciprocity in calls (the alter calling back the ego) because this pattern is robust to excluding calls from alter to ego. Regarding suspension and churn, recall that in this context, suspension (a status observed among some prepaid users) implies that the user cannot initiate any type of communication but can receive calls/texts for a period of up to six months. At any time during that period, a customer can move from a suspension to an active state by adding credit to his or her account. Churn, on the other hand, means that the customer has completely terminated the relationship with the firm and the phone has been disconnected. As can be seen from Figure 2 (bottom row), both suspension and churn are lower for the alters of treated (ego) customers than for the alters of control customers. This difference seems more pronounced during weeks 3–6 for alter suspension, whereas the difference in alter churn appears to be larger during weeks 8–12.

As we did for the egos, we compare the (aggregate) behaviors across conditions by subtracting averaged consumption of each alter in the sample ($N = 4,700$) during the four weeks before the treatment from the alter's observed weekly consumption after the intervention. We look at behavior in both the short term (weeks 1–6) and the long term (weeks 7–12). Table 7 shows the average differences across the two experimental conditions. With respect to suspension (percentage of alters who are suspended in week 6 or 12) and churn (percentage of alters who canceled their service by week 6 or 12), consistent with the pattern observed in Figure 2 (bottom row), suspension and churn are lower for the alters who are connected to treated customers. We find that the effect of the treatment on alter suspension and churn is stronger in the later weeks (weeks 7–12) than the earlier weeks (weeks 1–6) following the experiment. This pattern is to be expected, given that the treatment must affect ego behavior first in order to affect alter behavior later. We find statistically significant long-term effects of the campaign on the alters of treated egos for both suspension and churn. Regarding usage, and similar to

Table 6
LONG-TERM EFFECT OF TREATMENT ON EGO USAGE (WEEKS
7–12 AFTER THE TREATMENT)

	Outbound Usage		
	Minutes	Calls	Texts
Treatment	.335*** (.048)	.257*** (.041)	.230*** (.050)
Constant	-1.306*** (.064)	-1.173*** (.056)	-1.294*** (.067)
Week dummies	Yes	Yes	Yes
Observations	5,625	5,625	5,625

*** $p < .01$.

Notes: Data reflect long-term effects of treatment on ego usage, with robust standard errors in parentheses. The number of observations is 6 (weeks) \times 961 (egos), excluding egos who canceled their contract in a particular week.

the results for the egos, we find a positive treatment effect; average usage is significantly higher for the alters of treated egos than for the alters of control egos, even when we exclude the minutes when the alter talks to the ego. Similar to the effect of the campaign on alters' suspension and churn, we find a stronger effect of the campaign on alters' usage in the long term than in the short term. Note that the alters did not receive any financial incentive from the campaign; thus, the observed increase in activity of the alters cannot be explained by the financial incentive offered in the campaign.

In summary, the model-free analyses suggest that there is a potential spillover effect of the CRM campaign on the alters, particularly for churn in the long run, but also for usage. We next formally test for these effects using diff-in-diffs regression models.

Diff-in-diffs regression model results for alter usage and probit results for suspension and churn. We test whether the observed differences from our model-free analyses between the treatment and control groups in alter consumption are statistically significant by estimating a diff-in-diffs regression model, similar to the one used for the ego analysis. As before, we use the log-transformed variables for the usage variables. For these activities, we estimate the following regression models:

$$(3) \Delta y_{ijt}^{\text{alter}} = \gamma_0 + \gamma_1 T_i + \sum_{\tau=2}^6 \gamma_{\tau} D_{\tau} + \xi_{ijt} \text{ for } t = 1, \dots, 6 \text{ and}$$

$$(4) \Delta y_{ijt}^{\text{alter}} = \delta_0 + \delta_1 T_i + \sum_{\tau=8}^{12} \delta_{\tau-6} D_{\tau} + \xi_{ijt} \text{ for } t = 7, \dots, 12,$$

where $\Delta y_{ijt}^{\text{alter}}$ is the difference between the pre- and post-intervention values of the log of the usage of alter j , who is connected to ego i , and T_i is an indicator variable that takes the value 1 if ego i (connected to alter j) is part of the treatment group and 0 otherwise. The weekly dummies are defined as in Equations 1 and 2, and ξ_{ijt} is an error term with mean 0 and variance $\sigma_{\xi_{i,t}}^2$ or $\sigma_{\xi_{j,t}}^2$, for the short- and long-term effects, respectively. We use panel-corrected standard errors to account for potential serial correlation in the model error terms in Equations 3 and 4 (e.g., Hoechle 2007). We refer to Web Appendix A3 for a more detailed explanation of the models used. The estimation results are presented in the first two columns of Tables 8 and 9. These results confirm the results

Table 7
AVERAGE ALTER USAGE, SUSPENSION, AND CHURN
POSTINTERVENTION

	Control	Treatment	Difference	p-Value
<i>Short Term (Weeks 1–6)</i>				
Suspended status in week 6	25.7%	23.5%	–2.2%	.09
Churned by week 6	1.7%	1.4%	–.3%	.48
Difference in minutes	–9.65	–5.94	3.71	.11
Difference in minutes (excluding ego)	–10.00	–5.78	4.22	.06
Difference in log minutes	–.68	–.60	.08	.04
Difference in log minutes (excluding ego)	–.68	–.60	.08	.04
<i>Long Term (Weeks 7–12)</i>				
Suspended status in week 12	30.7%	27.6%	–3.1%	.02
Churned by week 12	3.7%	2.4%	–1.3%	.01
Difference in minutes	–20.21	–12.60	7.60	.02
Difference in minutes (excluding ego)	–19.35	–12.06	7.29	.02
Difference in log minutes	–1.01	–.90	.10	.03
Difference in log minutes (excluding ego)	–.99	–.89	.11	.03

depicted in Figure 2. Alters of treated egos have significantly higher consumption than alters of control egos, in both the short and the long run. These differences are statistically significant even when the calls from the alter to his or her ego are excluded. The magnitude of the treatment effect on alters is stronger in the long run (Table 9) than in the short run (Table 8), suggesting that some time is needed after the promotion for the effect of the campaign to propagate from egos to alters. For the sake of brevity, we show the regression results only for minutes called. We obtain similar results when estimating numbers of calls and texts; these results are given in Web Appendix A4.

Alter suspension and churn. We formally test for the differences in suspension and churn by estimating a binary probit model for each of the behaviors in both the short and the long term.¹⁰ More specifically, we estimate the following models:

$$(5) \quad \text{Prob}(y_{ijt}^{\text{alter}}) = P\left(\pi_0 + \pi_1 T_i + \sum_{\tau=2}^6 \pi_{\tau} D_{\tau i} + v_{ijt} > 0\right)$$

for $t = 1, \dots, 6$ and

$$(6) \quad \text{Prob}(y_{ijt}^{\text{alter}}) = P\left(\theta_0 + \theta_1 T_i + \sum_{\tau=8}^{12} \theta_{\tau-6} D_{\tau i} + v_{ijt} > 0\right).$$

for $t = 7, \dots, 12$.

Consistent with the notation in Equations 3 and 4, y_{ijt}^{alter} is a binary variable indicating whether alter j of ego i is suspended/churned in week t , and T_i indicates whether ego i is assigned

to the treatment group. The weekly dummies are defined as in Equations 1–4, and v_{ijt} are normally distributed with mean 0 and variance $\sigma_{v,s}^2$ or $\sigma_{v,l}^2$, for the short and long term, respectively. The results for the probit regressions (last two columns of Tables 8 and 9) confirm that alters of treated egos exhibit statistically significantly lower churn in the long term. This finding is consistent with previous work that has shown that a decrease in usage often precedes customer churn (e.g., Ascarza and Hardie 2013; Neslin et al. 2006). Regarding suspension behavior, although there is a negative effect in both the short and the long run, this effect is not statistically significantly different from zero. Thus, combining the results from Figure 2 and Tables 8 and 9, we have empirically demonstrated that the CRM marketing campaign had a positive impact on nontargeted connected customers of the treatment group: these customers have higher usage and lower churn than the nontargeted connected customers of the control group.

Investigating the Social Effect of Targeted Promotions

We conceptualize the pattern of our findings so far in Figure 3, Panel A. We have shown that the marketing intervention affects not only the targeted customers (represented by arrow A in Figure 3, Panel A) but also those who are connected to them (represented by arrow B in Figure 3, Panel A). In particular, we find that alters of treated egos have higher consumption and lower long-term churn than alters whose egos were not treated. The latter finding is particularly interesting because, unlike the egos, alters were not directly targeted by the company, and they did not receive any direct benefit from the campaign. Thus, although we empirically find that treatment has a statistically significant effect on the alters' behavior, we cannot claim that the treatment itself has a *direct* effect on the alters, as depicted by arrow B in Figure 3, Panel A. Instead, the effect of the campaign must have propagated to the alters through the behavior of their egos (as depicted in Figure 3, Panel B).¹¹

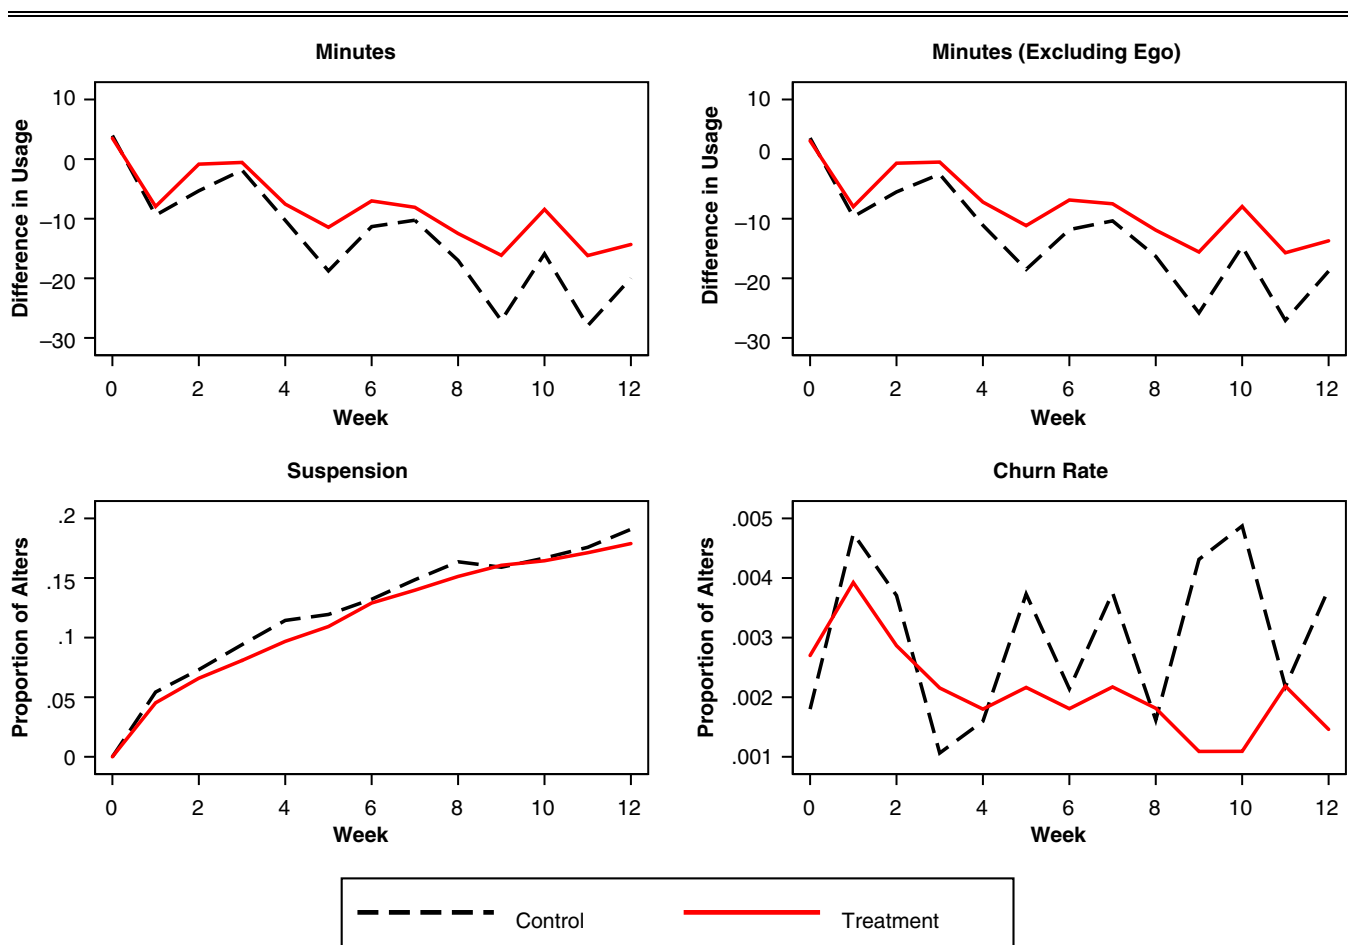
As discussed in the section "Effect of the Marketing Campaign on Targeted Customers (Egos)," we postulate that the propagation of the campaign from egos to alters is due to increased consumption of the egos and, specifically, increased consumption between the ego and his or her alters. This type of indirect effect is consistent with the presence of network externalities among the alters, who, after the campaign, face a more active local telecommunication network from which they derive higher value, and thus they increase their consumption. In this section, we empirically test the causal link between ego usage and alter usage and churn. Moreover, we provide further empirical support for network externality effects by exploring the moderating effect of the strength of the tie between the ego and the alter on the campaign's effect on the alter.

The effect of increased activity of egos on alters' usage and churn. We investigate whether increased communication between the egos and their alters in the weeks immediately

¹⁰Note that, unlike usage, the models for suspension and churn are not estimated in differences from the precampaign period. Therefore, it is possible that not controlling for individual-specific effects in the estimation would lead to inefficient (but still consistent) parameter estimates. We estimate a panel-data model that clusters the data at the alter level to appropriately estimate the standard errors of the estimated regression effects.

¹¹One might be tempted to run a standard mediation analysis of the effect of treatment on alters' behavior through ego behavior. However, it should be noted that such mediation is unnecessary because a direct effect of treatment on alters is theoretically implausible because alters were never directly exposed to the treatment. Moreover, as we will discuss later, one cannot run such a mediation analysis because the mediator (ego behavior) is endogenous.

Figure 2
POSTTREATMENT ALTER USAGE, SUSPENSION, AND CHURN BY TREATMENT CONDITION



Notes: “Usage” refers to differences between weekly consumption before and after the intervention; “suspension” refers to average number of customers in suspended status in a given week; “churn” refers to average number of customers who cancel their service in a given week.

following the campaign (i.e., short term) causes an increase in usage and a decrease in churn among alters in subsequent weeks (i.e., long term). That is, we are interested in estimating the dashed arrow in Figure 3, Panel B. A simple regression model that regresses the alter usage on ego usage will likely suffer from an endogeneity bias due to the presence of omitted variables that could affect the usage of both egos and alters. For

instance, a drop in network coverage quality in a certain area could lead to both egos and alters (living nearby) decreasing their consumption and, in some cases, subsequently churning. Similar arguments could be made if one considered the effect of competitors running promotional campaigns or changes in demand around the holiday season. While we can control for some unobserved shocks that are common to all users in the

Table 8
SHORT-TERM EFFECT OF TREATMENT ON ALTER USAGE AND CHURN (WEEKS 1–6 AFTER THE TREATMENT)

	<i>Outbound Minutes</i>		<i>Did Suspend</i>	<i>Did Churn</i>
	<i>Total</i>	<i>Total (Excluding Ego)</i>		
Treatment	.076*** (.019)	.077*** (.019)	-.053 (.043)	-.045 (.076)
Constant	-.598*** (.026)	-.600*** (.026)	-1.624*** (.039)	-2.605*** (.088)
Week dummies	Yes	Yes	Yes	Yes
Observations	27,987	27,987	27,987	27,987

*** $p < .01$.

Notes: Data reflect short-term effects of treatment on alter usage, with robust standard errors in parentheses. A linear (diff-in-diffs) regression is used for usage, and a probit regression for suspension and churn. The number of observations is 6 (weeks) \times 4,700 (alters), excluding alters who canceled their contract in a particular week.

Table 9
LONG-TERM EFFECT OF TREATMENT ON ALTER USAGE AND CHURN (WEEKS 7–12 AFTER THE TREATMENT)

	<i>Outbound Minutes</i>		<i>Did Suspend</i>	<i>Did Churn</i>
	<i>Total</i>	<i>Total (Excluding Ego)</i>		
Treatment	.098*** (.022)	.100*** (.022)	-.025 (.040)	-.236*** (.080)
Constant	-.847*** (.030)	-.835*** (.030)	-1.051*** (.032)	-2.646*** (.096)
Week dummies	Yes	Yes	Yes	Yes
Observations	27,598	27,598	27,598	27,598

*** $p < .01$.

Notes: Long-term effects of treatment on alter usage, with robust standard errors in parentheses. An ordinary least squares (diff-in-diffs) regression is used for usage, and a probit regression for suspension and churn. The number of observations is 6 (weeks) \times 4,700 (alters), excluding alters who canceled their contract in a particular week.

network, such as holiday season effects, it is practically impossible to control for all unobserved common shocks that are particular to every pair of an ego and an alter. As a consequence, simply regressing changes in alter usage (dependent variable) on changes in ego usage (independent variable), even when controlling for other observed factors, can lead to biased estimates of the regression parameters. It should be noted that our analysis thus far of the (causal) effect of the marketing promotion on the alters' usage and churn does not suffer from endogeneity because the treatment variable is exogenous by design and is therefore uncorrelated with any unobservable. The endogeneity problem only emerges when one tries to establish a causal link between ego usage and alter usage or churn.

To address this challenge, we employ an IV approach and use the experimental treatment dummy variable as an IV for the (endogenous) ego usage variable. There are two main reasons why the treatment dummy variable is a good candidate for an IV in this analysis. First, treatment is randomized and thus by construction is uncorrelated with any omitted variables in the regression. Second, as the analysis in "Effect of the Marketing Campaign on Targeted Customers (Egos)" shows, the treatment significantly influenced ego usage. We use the control function approach (Germann, Ebbes, and Grewal 2015; Petrin and Train 2010) to estimate the model. We choose weeks 1–6 (short term) to measure egos' behavior and weeks 7–12 (long term) to measure alters' behavior. Further details about the IV model, variable specifications, estimations, and robustness checks, are provided in Web Appendix A5. We would like to highlight that the using the treatment dummy as the IV helps us split the overall effect of treatment on alter behavior from "Effect of the Marketing Campaign on Nontargeted Customers (Alters)" into the effect of treatment on the communication between the ego and the alters and the effect of the communication between the ego and the alters on alter behavior. If treatment were the only regressor, we could have arithmetically calculated the effect of the communication between the ego and the alters on alter usage for the linear case, using the results from the previous two sections. However, because we also use weekly dummies as control variables, we must use an IV regression.

Table 10 shows the results of the IV regression analyses for the different types of activities (usage and churn) and different specifications of the communication between the ego and the alters. First, as indicated by the last row of Table 10 (first-stage t-statistic), we find that, as expected, the instrument has a strong and significant positive effect on the endogenous

variables (ego-to-alter usage). More importantly, the results for the three specifications of ego activity are convergent; an increase in short-term ego usage (minutes called, number of calls, and number of texts) postintervention leads to an increase in alter usage and a reduction in alter churn.

Thus, these results corroborate that the marketing campaign has a spillover effect that propagates to nontargeted users *through the increased usage* of the targeted customers. As discussed earlier, we postulate that the increase in ego usage (due to the marketing campaign) induces a more active local network around the egos, generating a positive network externality for the alters. Next, we provide additional support for this account by investigating the role of tie strength in moderating the treatment effect.

The moderating effect of tie strength. If, indeed, the indirect effect of the targeted promotion on the nontargeted customers (i.e., alters) propagates through the egos, then the treatment effect should be stronger for dyads of egos and alters that have stronger ties (Manchanda et al. 2015). We investigate this conjecture by quantifying the moderating role of the strength of the relationship between egos and alters on the effect of treatment on alters' usage (arrow B in Figure 3, Panel A).¹² We operationalize tie strength ($Strength_{ij}$) as the average number of minutes (in logs) that alter j called ego i during the four weeks prior to the intervention. We measure tie strength prior to the campaign to ensure independence between the measure of tie strength and treatment. Extending Equation 4, we estimate the following model:


$$(10) \quad \Delta y_{ijt}^{alter} = \varphi_0 + \varphi_1 T_i + \varphi_2 Strength_{ij} + \varphi_3 T_i \times Strength_{ij} + \sum_{\tau=8}^{12} \varphi_{\tau-4} D_{\tau t} + \varepsilon_{ijt},$$

for $t = 7, \dots, 12$. This equation adds two terms to Equation 4, the main effect of $Strength_{ij}$ and the interaction between $Strength_{ij}$ and the treatment variable. As dependent variable we take the (differenced) weekly number of outgoing minutes the alter talked to any connection other than the ego (Δy_{ijt}^{alter}).

As can be seen in Table 11, we find a significant and positive interaction effect between tie strength and the campaign

¹²An alternative approach would be to test the moderating role of tie strength on the effect of ego usage on alter usage (arrow C in Figure 3, Panel B). One challenge with such an approach is that both independent variables (ego usage and its interaction with tie strength) are endogenous. While in theory this could be estimated with a single instrument, this analysis is likely to be inefficient and less robust (Wooldridge 2007).

Figure 3
SCHEMATIC DIAGRAM OF THE EFFECT OF TREATMENT ON EGO
AND ALTER BEHAVIOR



treatment ($\phi_3 > 0$), indicating that the social effect presented in “Effect of the Marketing Campaign on Nontargeted Customers (Alters)” is even larger for ego–alter pairs that had a stronger connection prior to the campaign. Because the dependent variable excludes communication from each alter to his/her ego, reciprocity in calls cannot account for this effect. In a separate analysis, we have also operationalized tie strength as the number of minutes the ego called the alter before the intervention. We find a similar positive interaction effect between treatment and tie strength (detailed results of this analysis are in Web Appendix A6).

An alternative variable that is expected to moderate the social effect is the number of connections each alter has. All else being equal, one would expect the ego to play a more prominent social role for alters who have fewer connections. Accordingly, we posit a negative sign for the interaction between treatment and number of alter’s connections. With reference to Equation 10, we substitute $Strength_{ij}$ with the number of connections alter j has, operationalized as the average number of customers of the focal provider that alter j communicated with in the four weeks prior to the campaign. The results of this regression are given in Table 12 and are consistent with the results reported in Table 11. Specifically, across the two activity types, we find that the lower the role the ego plays in the alter’s network (i.e., when the alter has more nodes in his or her network), the weaker the treatment effect. In summary, using two measures of social importance (tie strength and alter’s number of connections), we find that the stronger the social tie between egos and alters, the stronger the propagation of the campaign from the targeted customer to the (nontargeted) connections.

MANAGERIAL RELEVANCE OF THE SOCIAL EFFECT

Thus far, we have shown that ego usage (“Effect of the Marketing Campaign on Targeted Customers [Egos]”) and alter usage (“Effect of the Marketing Campaign on

Nontargeted Customers [Alters]”) increase *because* of the treatment. We have also demonstrated the presence of a statistically significant propagation of the targeted campaign to nontargeted customers (“Investigating the Social Effect of Targeted Promotions”). In this section, we quantify the managerial relevance of the social impact of the CRM campaign. That is, how big is the spillover effect from egos to alters?¹³ We calculate two metrics to quantify the social effect of the campaign. The first metric is the magnitude of the consumption spillover, which we compute as the percentage increase in usage among the customers who were directly targeted (egos), relative to the percentage increase in usage among their alters. The second metric relates to the monetary value of the social effect, which we compute as the incremental value the firm receives from the nontargeted customers (i.e., alters) due to having targeted their connections (i.e., egos).

Quantifying the Consumption Spillover

Using the estimated models for the effect of treatment on ego usage (Tables 5 and 6) and the effect of treatment on alter usage (Tables 8 and 9), we compute the percentage increase in usage for both egos and alters due to the treatment for the 12 weeks following the campaign. Using the average consumption for each customer for the 4 weeks prior to the campaign as baseline, we convert the parameter estimates to percentage increase by transforming the diff-in-diffs regression specification into usage levels. Following these calculations, we find that the campaign caused a 34.8% increase in number of minutes called by egos in the 12 weeks following the campaign. The corresponding increase in alter usage is 9.7%. That is, the spillover effect of the campaign on alter usage is approximately 28% ($.097/.348$). It is important to recall that while the egos received an economic incentive to increase their usage, the alters did not.

Using the same approach, we also compute the size of the spillover effect for different levels of tie strength. To do so, we use the parameter estimates of the model that incorporates the interaction between treatment and tie strength (Table 11) and calculate the percentage increase in usage for alters whose tie strength is one standard deviation above and below the population mean. For alters who have stronger relationships with their egos, the increase in usage is 14.3% (which translates to a spillover of $41\% = .143/.348$), whereas for those with weaker ties, the increase is 5.3%, which corresponds to a spillover of 16%.

Measuring the Financial Value of the Spillover Effect

Most commonly, CRM marketing campaigns are evaluated according to the lift in profitability of the targeted customers relative to the incurred costs of the campaign. In this article, we have demonstrated that a marketing campaign can also affect the usage and churn—and, thus, the profitability—of the customers connected to the targeted customers, suggesting that there is an additional value obtained from the alters that should be considered when firms evaluate the return on investment of their targeted campaigns. Here, we quantify that incremental value by comparing postcampaign profit obtained from the

¹³Note that we can make that comparison only for usage, and not for churn, because in our context egos belong to a prepaid plan in which churn is hardly ever observed.

Table 10
EFFECT OF SHORT-TERM EGO-TO-ALTER USAGE ON LONG-TERM ALTER USAGE AND CHURN (INSTRUMENTAL VARIABLE REGRESSIONS)

Ego to Alter (Regressor)	Alter Usage (Excluding Ego) as Dependent Variable					
	Minutes		Calls		Texts	
	M	Churn	M	Churn	M	Churn
Minutes	3.204*** (.737)	-7.525*** (2.749)				
Calls			1.765*** (.623)	-7.494*** (2.632)		
Texts					.891* (.520)	-5.392*** (1.878)
Intercept	-.0315 (.171)	-4.533*** (.654)	-.458*** (.105)	-3.999*** (.448)	-.520** (.242)	-5.274*** (.888)
Week dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	27,598	27,598	27,598	27,598	27,598	27,598
First-stage t-statistic	3.957	3.957	5.901	5.901	4.083	4.083

* $p < .10$.

** $p < .05$.

*** $p < .01$.

Notes: Data reflect effect of short-term ego-to-alter usage on long-term alter usage and churn using the control function approach, with standard errors in parentheses. The regressor ego-to-alter usage is operationalized as the average of (differenced) ego usage during weeks 1–6. The dependent variable of alter usage is operationalized as the average of (differenced) alter usage during weeks 7–12. Bootstrapping is used to estimate the standard errors.

alters of treated customers with that of the alters of customers in the control group.

To do so, and given that we did not obtain information about the profitability of individual customers, we need to rely on certain assumptions about average measures of profitability for the customers in our sample. For details on all the assumptions made regarding consumption levels, discount factor, operating margins, and calculations, we refer the reader to Web Appendix A7. Using these assumptions, we estimate that an alter whose ego was treated generates \$.85 more profit than an alter whose ego was not treated. In other words, beyond the effect of the marketing campaign on the targeted customers, this campaign also increases the profits of the nontargeted (but connected) customers by \$.85 per alter. Given that egos have, on average, five alters each, the campaign generates an extra \$4.25 in profits per targeted customer from social spillover.

We acknowledge that these “back-of-the-envelope” calculations of additional spillover profitability are approximative in that they are based on average levels of revenue and are dependent on several assumptions. Nevertheless, we believe that this analysis highlights that the social effect of CRM campaigns can have a substantial positive financial impact when network externalities are present.

GENERAL DISCUSSION

In this article, we quantify the social effect of CRM marketing campaigns. We show that a CRM campaign that is aimed at changing the behavior of some customers can propagate through the social network of the targeted customers and thus also affect the behavior of nontargeted, but connected, customers. In the context of telecommunication, we find that the social connections of targeted customers were more likely to increase their consumption and less likely to churn due to a campaign that was not targeted at them and did not offer them any incentives to change their behavior. In particular, we estimate a social multiplier of 1.28. That is, the spillover effect of the campaign to nontargeted customers is 28% of the effect of the campaign on the targeted customers. Financially, this propagation translates to an additional profit of \$0.85 per nontargeted customer who is connected to a targeted customer.

Using a randomized field experiment, we estimate the causal effect of a CRM campaign on both the targeted and the nontargeted customers. We further leverage the experimental design using an IV regression to estimate the causal effect of the activity of the egos on the activity of their alters. We show that the effect of the campaign propagates from egos to alters through an increase in the activity from the targeted customer to his or her alters. Furthermore, we observe a stronger social effect for dyads with stronger ties. Although we do not observe the content of a conversation between an ego and an alter, it is unlikely that WOM is the main driver of the propagation of the CRM campaign. In fact, if the ego were to discuss the campaign with the alters, we would expect a negative effect of the campaign on alters because the targeted campaign would not be available to them.

We put forward a network externality explanation, which is consistent with our finding that customers increase the usage of the service and are less likely to churn when their (local) network becomes more active. Network externality research has shown that a larger and more active network often leads to

Table 11
LONG-TERM EFFECT OF TREATMENT ON ALTER USAGE (WEEKS 7–12 AFTER THE TREATMENT) MODERATED BY TIE STRENGTH

	Outbound Minutes	
	Total	Total (Excluding Ego)
Treatment	.096*** (.023)	.098*** (.023)
Tie strength	-.216*** (.020)	-.179*** (.019)
Tie strength × Treatment	.052** (.025)	.059** (.024)
Constant	-.844*** (.030)	-.833*** (.030)
Week dummies	Yes	Yes
Observations	27,598	27,598

** $p < .05$.

*** $p < .01$.

Notes: Data reflect long-term effects on alter usage, with robust standard errors in parentheses. Tie strength is operationalized as the number of minutes the alter called the ego before the intervention.

Table 12
LONG-TERM EFFECT OF TREATMENT ON ALTER USAGE
(WEEKS 7–12 AFTER THE TREATMENT) MODERATED BY THE
NUMBER OF ALTER CONNECTIONS

	Outbound Minutes	
	Total	Total (Excluding Ego)
Treatment	.0491** (.023)	.050** (.022)
Number of connections	-.224*** (.018)	-.225*** (.018)
Number of connections × Treatment	-.084*** (.023)	-.087*** (.023)
Constant	-.819*** (.030)	-.807*** (.030)
Week dummies	Yes	Yes
Observations	27,598	27,598

** $p < .05$.

*** $p < .01$.

Notes: Data reflect long-term effects on alter usage, with standard errors in parentheses.

higher value to the network members (e.g., Aral and Walker 2011; Nitzan and Libai 2011). Thus, the decrease in churn and suspension among alters can be easily attributed to such network externality effects. What is less obvious is why—conditioned on not churning, and excluding communication with the ego—alters increase their usage after their ego has been treated. That is, why would an alter call his or her other connections more because his or her (treated) egos called them more? While, given the nature of the data, we cannot uniquely pinpoint the underlying mechanism of this finding, we postulate that the increased activity in the network (because of the increased usage from the ego to the alter) allows the alters to perceive a higher value of their network, which subsequently leads to higher levels of usage (Aral and Walker 2011; Manchanda et al. 2015). For example, given the increasing number of alternative methods of communication available to customers (e.g., WhatsApp, WeChat, Skype, Google Hangouts, multiple SIM cards), an increased perceived value of one of the communication networks can motivate the alter to use that network more often as the primary mode for communication. We leave the investigation of the specific mechanisms underlying how network externality affects usage and churn for future research.

Our research has clear implications for marketing managers. In business contexts in which customers are connected, targeted campaigns might have higher return on investment than what is currently believed. Moreover, our findings suggest that firms should leverage social effects in deciding which customers to target. On the one hand, the CRM practice has focused primarily on targeting customers according to expected lift in profitability of the targeted customer. On the other hand, the social contagion literature has, for the most part, ignored profitability and primarily focused on targeting “hubs” with strong social influence. Our results suggest that firms should consider a combination of these two effects and target customers with the highest lift in *social profitability* due to the campaign. Beyond the profitability of the campaigns, a firm operating close to its capacity limits should consider the social impact of its targeted actions. For example, in contexts with capacity constraints (e.g., wireless providers in developing countries) or in cases in which utilization capacity directly links to customer satisfaction (e.g., gyms), companies should

anticipate increased activity not only from the targeted customers but also from those connected to them.

Our research contributes to the broader CRM literature (e.g., Berger et al. 2002; Kumar, Lemon, and Parasuraman 2006; Rust and Verhoef 2005) that has focused on measuring the impact of marketing actions on (targeted) individual customers. In this research, we quantify the effects of marketing actions *beyond the target customer* and show how, in the presence of network externalities, the impact of marketing activities on firm profitability might be higher than otherwise estimated. Our work likewise complements extant work on social influence in new product introduction and customer acquisition (e.g., Iyengar, Van den Bulte, and Valente 2011; Schmitt, Skiera, and Van den Bulte 2011). Consistent with the findings of Nitzan and Libai (2011), our research confirms that social influence is not limited to new behaviors (e.g., adoption of new products) but is also present in marketing campaigns aimed at changing the behavior of existing customers. More broadly, our work complements the research on the spillover effects of marketing actions. Previous research has shown that marketing campaigns can spill over to brands that are “connected” to the focal brand (e.g., Chae et al. 2016; Erdem and Sun 2002; Rutz and Bucklin 2011). In this research, we extend the notion of a marketing action spillover from one customer to another.

We chose to investigate the propagation of social CRM campaigns in the context of a telecommunication firm. There are several reasons for this choice. First, the telecommunication context allows us to directly observe the customer’s network. Second, the telecommunication industry is of major interest to CRM academics and practitioners (Rivera and Van der Meulen 2014). Finally, the context of telecommunication is characterized by strong network externality effects. Indeed, we expect a weaker spillover effect in applications that are characterized by low network externality, such as, for example, consumer packaged goods applications. That being said, we believe that our findings have implications for industries other than telecommunication, such as file-sharing services, peer-to-peer marketplaces, payment services, and online games. Elaborating on the generalizability of our results, there are three main conditions needed for a business setting to observe and leverage our results: (1) a reasonable amount of network externalities, (2) the ability to individually target marketing actions, and (3) observation of the customers’ social network. While the first condition is a necessary condition for the effect to occur, the remaining two conditions are needed for the firm to measure the social effect and leverage our findings. We encourage firms across different sectors to better develop capabilities that will allow them to measure social interactions and individually target their marketing actions.

The data we had access to impose some limitations on our research. First, we investigate a conservative propagation of the campaign only to first-degree connections. Future research could investigate whether campaigns propagate beyond the first degree. However, in looking beyond first-degree effects, potential contamination and interference in network experiments becomes more challenging to handle (Aral 2016). Second, the campaign we observed was a successful one in terms of affecting the targeted customers. It is likely that less successful campaigns will have limited propagations. In

some cases, marketing campaigns may even have a negative effect on the targeted customers (e.g., Ascarza, Iyengar, and Schleicher 2016). Do campaigns with negative direct effects generate negative spillover effects? Given the documented network effect of churn (Nitzan and Libai 2011) and the WOM effect of negative information (Moldovan and Goldenberg 2004), one may expect negative propagation for such unsuccessful campaigns. We encourage future researchers to investigate these questions as well as measure the degree of the propagation of CRM campaigns in different business settings.

In summary, we provide empirical evidence that CRM campaigns can have a spillover effect beyond the target customer. This finding has implications for the targeting and evaluation of such campaigns. We hope that this research will serve as a stepping-stone in changing the view in the CRM community from considering profit primarily in terms of customer value to considering it in terms of customer *social* value.

REFERENCES

- Aral, Sinan (2016), "Networked Experiments" in *The Oxford Handbook of the Economics of Networks*, Yann Bramoullé, Andrea Galeotti, and Brian Rogers, eds. Oxford, UK: Oxford University Press.
- Aral, Sinan, and Dylan Walker (2011), "Creating Social Contagion Through Viral Product Design: A Randomized Trial of Peer Influence in Networks," *Management Science*, 57 (9), 1623–39.
- Ascarza, Eva, and Bruce Hardie (2013), "A Joint Model of Usage and Churn in Contractual Settings," *Marketing Science*, 32 (4), 570–90.
- Ascarza, Eva, Raghuram Iyengar, and Martin Schleicher (2016), "The Perils of Proactive Churn Prevention Using Plan Recommendations: Evidence from a Field Experiment," *Journal of Marketing Research*, 53 (February), 46–60.
- Bapna, Ravi, and Akhmed Umyarov (2015), "Do Your Online Friends Make You Pay? A Randomized Field Experiment in an Online Music Social Network," *Management Science*, 61 (8), 1902–20.
- Berger, Paul D., Ruth N. Bolton, Douglas Bowman, Elten Briggs, Vasanth Kumar, Arun Parasuraman, et al. (2002), "Marketing Actions and the Value of Customer Assets: A Framework for Customer Asset Management," *Journal of Service Research*, 5 (1), 39–54.
- Biyalogorsky, Eyal, Eitan Gerstner, and Barak Libai (2001), "Customer Referral Management: Optimal Reward Programs," *Marketing Science*, 20 (1), 82–95.
- Blattberg, Robert C., Byung-Do Kim, and Scott A. Neslin (2008), *Database Marketing: Analyzing and Managing Customers*. New York: Springer.
- Boulding, William, Richard Staelin, Michael Ehret, and Wesley J. Johnston (2005), "A Customer Relationship Management Roadmap: What Is Known, Potential Pitfalls, and Where to Go," *Journal of Marketing*, 69 (October), 155–66.
- Chae, Inyoung, Andrew T. Stephen, Yakov Bart, and Dai Yao (2016), "Spillover Effects in Seeded Word-of-Mouth Marketing Campaigns," *Marketing Science*, 36 (1), 89–104.
- Ebbes, Peter, Zan Huang, and Arvind Rangaswamy (2015), "Sampling Designs for Recovering Local and Global Characteristics of Social Networks," *International Journal of Research in Marketing*, 33 (3), 578–99.
- Eckles, Dean, Brian Karrer, and Johan Ugander (2014), "Design and Analyses of Experiments in Networks: Reducing Bias from Interference," working paper, Data Science Team, Facebook.
- Erdem, Tülin, and Baohong Sun (2002), "An Empirical Investigation of the Spillover Effects of Advertising and Sales Promotions in Umbrella Branding," *Journal of Marketing Research*, 39 (November), 408–20.
- Fader, Peter (2012), *Customer Centricity*, 2nd ed. Philadelphia: Wharton Digital Press.
- Fienberg, Stephen E. (2012), "A Brief History of Statistical Models for Network Analysis and Open Challenges," *Journal of Computational and Graphical Statistics*, 21 (4), 825–39.
- Germann, Frank, Peter Ebbes, and Rajdeep Grewal (2015), "The Chief Marketing Officer Matters!" *Journal of Marketing*, 79 (May), 1–22.
- Gupta, Sunil (1988), "Impact of Sales Promotions on When, What, and How Much to Buy," *Journal of Marketing Research*, 25 (November), 342–55.
- Haenlein, Michael, and Barak Libai (2013), "Targeting Revenue Leaders for a New Product," *Journal of Marketing*, 77 (May), 65–80.
- Hill, Shawndra, Foster Provost, and Chris Volinsky (2006), "Network-Based Marketing: Identifying Likely Adopters via Consumer Networks," *Statistical Science*, 21 (2), 256–76.
- Hinz, Oliver, Bernd Skiera, Christian Barrot, and Jan U. Becker (2011), "Seeding Strategies for Viral Marketing: An Empirical Comparison," *Journal of Marketing*, 75 (November), 55–71.
- Hoechle, Daniel (2007), "Robust Standard Errors for Panel Regressions with Cross-Sectional Dependence," *Stata Journal*, 7 (3), 281–312.
- Iyengar, Raghuram, Christophe Van den Bulte, and Thomas W. Valente (2011), "Opinion Leadership and Social Contagion in New Product Diffusion," *Marketing Science*, 30 (2), 195–212.
- Katz, Michael L., and Carl Shapiro (1985), "Network Externalities, Competition, and Compatibility," *American Economic Review*, 75 (3), 424–40.
- Kumar, Vineet, Katherine N. Lemon, and Arun Parasuraman (2006), "Managing Customers for Value: An Overview and Research Agenda," *Journal of Service Research*, 9 (2), 7–94.
- Lemon, Katherine N. and Kathleen Seiders (2006), "Making Marketing Accountable: A Broader View," in *Does Marketing Need Reform?* Jagdish N. Sheth and Rajendra S. Sisodia, eds. Armonk, NY: M.E. Sharpe, 201–08.
- Li, Krista J., and Sanjay Jain (2016), "Behavior-Based Pricing: An Analysis of the Impact of Peer-Induced Fairness," *Management Science*, 62 (9), 2705–21.
- Manchanda, Puneet, Grant Packard, and Adithya Pattabhiramaiah (2015), "Social Dollars: The Economic Impact of Customer Participation in a Firm-Sponsored Online Customer Community," *Marketing Science*, 34 (3), 367–87.
- Manski, Charles F. (1993), "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies*, 60 (3), 531–42.
- Moldovan, Sarit, and Jacob Goldenberg (2004), "Cellular Automata Modeling of Resistance to Innovations: Effects and Solutions," *Technological Forecasting and Social Change*, 71 (5), 425–42.
- Nair, Harikesh S., Puneet Manchanda, and Tulikaa Bhatia (2010), "Asymmetric Social Interactions in Physician Prescription Behavior: The Role of Opinion Leaders," *Journal of Marketing Research*, 47 (October), 883–95.
- Neslin, Scott A., Sunil Gupta, Wagner Kamakura, Junxiang Lu, and Charlotte H. Mason (2006), "Defection Detection: Measuring and Understanding the Predictive Accuracy of Customer Churn Models," *Journal of Marketing Research*, 43 (May), 204–11.
- Neslin, Scott A., Caroline Henderson, and John Quelch (1985), "Consumer Promotions and the Acceleration of Product Purchases," *Marketing Science*, 4 (2), 147–65.
- Nguyen, Bang, and Lyndon Simkin (2013), "The Dark Side of CRM: Advantaged and Disadvantaged Customers," *Journal of Consumer Marketing*, 30 (1), 17–30.

- Nitzan, Irit, and Barak Libai (2011), "Social Effects on Customer Retention," *Journal of Marketing*, 75 (November), 24–38.
- Petrin, Amil, and Kenneth Train (2010), "A Control Function Approach to Endogeneity in Consumer Choice Models," *Journal of Marketing Research*, 47 (February), 3–13.
- Reinartz, Werner, Manfred Krafft, and Wayne D. Hoyer (2004), "The Customer Relationship Management Process: Its Measurement and Impact on Performance," *Journal of Marketing Research*, 41 (August), 293–305.
- Rivera, Janessa, and Robert van der Meulen (2014), "Gartner Says CRM Will Be at the Heart of Digital Initiatives for Years to Come," press release (February 12), <http://www.gartner.com/newsroom/id/2665215>.
- Rubin, Donald B. (1980), "Discussion of Basu's 'Randomization Analysis of Experimental Data: The Fisher Randomization Test,'" *Journal of the American Statistical Association*, 75 (371), 591–93.
- Rust, Roland T., and Peter C. Verhoef (2005), "Optimizing the Marketing Interventions Mix in Intermediate-Term CRM," *Marketing Science*, 24 (3), 477–89.
- Rutz, Oliver J., and Randolph E. Bucklin (2011), "From Generic to Branded: A Model of Spillover in Paid Search Advertising," *Journal of Marketing Research*, 48 (February), 87–102.
- Schmitt, Phillip, Bernd Skiera, and Christophe Van den Bulte (2011), "Referral Programs and Customer Value," *Journal of Marketing*, 75 (January), 46–59.
- Shalizi, Cosma R., and Andrew C. Thomas (2011), "Homophily and Contagion Are Generically Confounded in Observational Social Network Studies," *Sociological Methods & Research*, 40 (2), 211–39.
- Trusov, Michael, Anand Bodapati, and Randolph E. Bucklin (2010), "Determining Influential Users in Internet Social Networks," *Journal of Marketing Research*, 47 (August), 643–58.
- Wooldridge, Jeffrey (2007), "Difference-in-Differences Estimation," Lecture 10 of *What's New in Econometrics?* Guido Imbens and Jeffrey Wooldridge, lecturers. Course presented at National Bureau of Economic Research Summer Institute 2007, <http://www.nber.org/minicourse3.html>.