How firms can go wrong by offering the right service contract:

Evidence from a field experiment

Eva Ascarza
Columbia Business School
3022 Broadway, Uris Hall
New York, NY 10027-6902
ascarza@columbia.edu

Raghuram Iyengar
The Wharton School of the University of Pennsylvania
3730 Walnut Street
Philadelphia PA, 19104-6340
riyengar@wharton.upenn.edu

Martin Schleicher
Columbia Business School
3022 Broadway, Uris Hall
New York, NY 10027-6902
mschleicher13@gsb.columbia.edu

October 12, 2013

Acknowledgements. We thank several managers at the participating company for their efforts in making this research possible. The authors benefited from comments by Pete Fader, Don Lehmann, Oded Netzer, Christophe Van den Bulte, and the audience members at the Summer 2013 brown bag seminar series at the Columbia Business School.

How firms can go wrong by offering the right service contract:

Evidence from a field experiment

Abstract

Past evidence reveals that customers make ex-post mistakes when choosing service plans. Fearing a negative impact from customers' overspending mistakes on long-term profits, some firms are becoming proactive and now recommend optimal tariffs to their existing customers. In this paper, we use a randomized field experiment to examine the profitability of encouraging existing customers to switch to better plans. We find that encouraging customers to switch to cost-minimizing plans can actually harm the firm. The primary source for this negative effect is the change in behavior among customers who decide to *reject* the firm's recommendation. For this set of customers, churn notably increases, resulting in substantial losses. We propose two mechanisms for such increase in churn, namely lower inertia and customer regret. Our data provide empirical evidence for both drivers in the context we study. We also explore the impact of hypothetical targeted campaigns. The results suggest that selecting the right customers to target has a higher impact on profitability than allocating customers into optimal tariffs.

1. Introduction

Firms typically offer their customers a choice among several pricing plans to better match their valuations. The success of this price discrimination strategy rests on the assumption that consumers will select the "right" plan – a plan that maximizes their welfare. However, plenty of evidence from past research casts doubt on this assumption revealing customers' (ex-post) mistakes when choosing service plans. For example, in a systematic analysis of data on cell-phone usage and expenditures between 2001 and 2003, Bar-Gill and Stone (2009) estimated that each year 42 million consumers make overspending mistakes, costing at least 20% of their annual spend. These errors translate into an aggregate annual welfare loss of almost $12 billion across the U.S. population. Relatedly, the Federal Communications Commission (FCC) received about 760 complaints in the first six months of 2010 with customers being unhappy with their monthly bill. Bill shock accounted for about as many complaints as early termination fees or mystery charges on bills.[1] This phenomenon is not unique to telecom services; customers choosing cable/satellite packages, health care plans, credit cards, insurance packages, or gym memberships also exhibit similar behavior.[2]

Fearing a negative impact from customers' overspending mistakes on long-term profits, some firms are beginning to act proactively. For example, health insurance providers now offer plan cost estimators which allow customers to compare, for each member's usage level, the out-of-pocket cost estimate for different plans. Some telecommunications companies are helping customers to better manage their usage by sending text messages when they reach their free

[1] http://transition.fcc.gov/stage/Bill-Shock-White-Paper.pdf

[2] Companies specialized in billing management claim that most people in the U.S. are overpaying $1,000 or more in each of these expenses (http://www.truaxis.com/statementrewards-for-consumers/)

monthly quota of minutes.[3] While this practice is implemented to ease bill shock, research showing that people have intrinsic limits to navigate non-linear pricing schedules implies such modifications may be futile (e.g., Gopalakrishnan, Iyengar and Meyer 2013). More recently, service firms have begun encouraging customers to switch to more appropriate plans. For instance, Verizon in the U.S. offers recommendations of service plans to customers based on their estimated monthly usage.[4] Other providers like Vodafone and AT&T offer similar services.

How effective are such programs? There is some doubt based on past research. On the one hand, encouraging customers to switch plans might have a positive impact on customer satisfaction. If a customer has overspent in the past, she may become dissatisfied with the firm's contract offerings, hence her probability of leaving the service (i.e., churn) might increase (Bolton 1998; Bolton and Lemon 1999). As a consequence, allocating customers to optimal tariffs would, potentially, decrease churn, which can have a substantial impact on customer and firm value (Ascarza and Hardie 2012; Gupta, Lehmann and Stuart 2004).

On the other hand, encouraging customers to switch plans may lower customer loyalty. First, encouraging customers to switch to alternative plans may lower their inertia in tariff choice, as documented in past research (e.g., Ascarza, Lambrecht and Vilcassim 2012; Goettler and Clay 2012; Iyengar, Ansari and Gupta 2007). This might in turn induce customers to seek alternative providers (Wieringa and Verhoef 2007). Second, the encouragement itself might enhance consumers' awareness of their sub-optimal behavior, inducing consumers to experience regret as they compare their current spending with potential savings under a better plan — "what

[3] http://news.verizonwireless.com/news/2012/10/pr2012-10-02.html; http://www.t-mobile.com/Company/CompanyInfo.aspx?tp=Abt_Tab_ConsumerInfo
[4] http://www.verizonwireless.com/b2c/splash/shareEverythingCalculator.jsp?intcmp=vzw-vnt-se-shareeverything

could have been" (Bell 1982). This experienced regret may lead to brand switching (Inman et al. 1997; Tsiros and Mittal 2000; Zeelenberg 1999), hence increasing customer churn.

The purpose of this paper is to assess how encouraging existing customers to switch to better plans affects their behavior and hence the firm's profitability. We conducted a large scale field experiment involving 65,000 customers of a cellular phone service. Participants were randomly assigned to whether they received an encouragement to switch to a service contract that is predicted to save them money based on their past behavior (the test group) or not (the control group). The field experiment was conducted over a 6-month period with the recommendation (i.e., "reallocation" campaign) being given to the participants in the test group at the end of the third month. A notable feature of the field experiment is that while the presence and absence of recommendation was randomized, participants were free to accept or reject the recommendation. Such "broken randomized experiments", where participants may show noncompliance with the randomly assigned treatment, are common in the social sciences (see Barnard et al. 1998 for a detailed discussion). We explicitly address consumer self-selection using propensity score matching (e.g., Angrist, Imbens and Rubin 1996) to determine the true impact of the encouragement to switch on customer behavior.

Our results indicate that being proactive and encouraging customers to switch to better tariffs can harm firm profitability. This negative effect holds even after controlling for observed differences across customers and for self-selection due to noncompliance. The primary driver for the negative effect is the change in behavior among customers in the test group who decide to *reject* the recommendation. For these customers, we observe that churn notably increases, resulting in big losses for the firm. On the positive side, the smaller fraction of customers who accept the firm's encouragement do exhibit a marginal increase in revenue. These parallel and opposite effects indicate that selecting the right customers to target with the marketing campaign (in this case, not targeting the wrong customers) is crucial.

We propose two explanations for the increase in churn. First, the encouragement may lower the inertia that prevents customers from switching plans—customer inertia is important in access-based services as customers fail to switch plans even in the absence of any commitment contract (e.g., Goettler and Clay 2012). Once customers realize that it is easy to switch plans within the company, they may explore competitive offerings, which may have been overlooked otherwise. Second, the campaign may emphasize the sub-optimality of consumers' past decisions, making them aware of how much they could have saved. As a consequence, the encouragement to switch might induce customers to feel regret and they may leave to avoid such a feeling (e.g., Inman, Dyer and Jia1997). We analyze plan switching and churn behavior among the customers in our data and provide empirical evidence for both these drivers. Based on these findings, we discuss the conditions under which firms should run this type of encouragement and also examine which customers should and should not be targeted. By leveraging the richness of our field experiment, we assess the impact on customer behavior and firm revenue if the company were to run targeted encouragement campaigns.

Our research complements empirical literature on price discrimination (e.g., Ascarza et al., 2012; Danaher 2002; Iyengar, Jedidi, Essegaier and Danaher 2011; Narayanan, Chintagunta and Miravete 2007) that has evaluated the impact of offering differing pricing contracts on firm profitability, but has not investigated the impact of firms' targeted pricing recommendations on customer behavior. It complements the extant work on customer relationship management that has investigated drivers of customer retention (e.g., Bolton 1998, Lemon, Barnett-White and Winer 2002; Neslin et al. 2006; Nitzan and Libai 2011). Our research is an addition to the broader literature on product recommendations (e.g., Fitzsimons and Lehmann 2004) investigating drivers and consequences of customers' reactance. More generally, our study relates to the work on choice architecture (e.g., Thaler and Sunstein 2008; Johnson et al. 2012) by examining the impact of a direct encouragement aimed to alter consumers' choices.

We proceed as follows. In Section 2, we describe the research setting. In Section 3, we describe the analysis and quantify the impact of encouraging customers to switch plans on their demand for access services. In Section 4 we examine two possible drivers for the observed phenomenon of higher churn among consumers that reject the recommendation. Section 5 assesses the robustness of our findings and discusses alternative explanations for the behavior we observe in the data. Section 6 focuses on the managerial implications of this research and shows the impact of running more targeted campaigns on customer behavior and firm revenues. Section 7 concludes.

2. Research setting

To assess how customers respond to an encouragement to switch to better tariffs, the research setting should ideally satisfy several conditions. First, for an encouragement to work, a reasonable fraction of customers must be on a suboptimal plan. Second, the encouragement must be randomized across customers who are suboptimal in their usage. Third, one must have data on consumers' demand for service both before and after the campaign and their churn behavior after the campaign. Finally, any marketing efforts deployed must be observed or otherwise controlled for.

We secured the cooperation of a wireless communications firm in South America to meet these stringent conditions. Similar to other companies in its industry, the company was keen on recommending plans to customers based on their usage. Managers were hopeful that customers would find this beneficial, resulting on lower churn. They realized, however, that there was uncertainty regarding how consumers' behavior would change based on any such encouragement.

2.1 Field Experiment

Pricing Plan: Customers enrolled in a specific type of plan were selected for the study. The plan involves the purchase of a fixed amount of credit every month, which entitles the customer with an initial balance of the amount purchased plus an additional percentage bonus (e.g., if a customer signed up for the $30 plan, their initial credit each month would be $40). At any point during the month, if the balance reaches zero then the customer can purchase additional credit, with no bonus. This type of plan is similar to a three-part tariff plan as the price per unit of consumption is zero until the customer uses her initial credit, then a per-minute price is applied. At the end of each month, customers are automatically enrolled in the same plan (i.e., the default is to continue with the same service contract). However, customers have no contractual obligations with the provider (i.e., they can leave at any time) and are also allowed to change the base plan (upgrade or downgrade) at the end of every month. At the time of the study the company offered 6 different plans, with the fixed amount of monthly credit ranging from $19 to $63.

Customers: Customers were eligible to be included in the experiment upon satisfying three criteria. These criteria were having (a) a certain level of revenue (more than $47 per month, during all three months prior to the campaign) such that the plans offered in the campaign would be beneficial and the encouragement may work, (b) been with the company for more than seven months to ensure some level of self-awareness about individual usage, and (c) not received any targeted promotional activities in the last three months. The latter condition ensures that there were no marketing activities towards eligible customers for a few months before the campaign. There were no other activities targeted towards these customers during the campaign as well. Of the customers that satisfied all three criteria, 64,147 were randomly selected to be included in the experiment. 10,058 customers (15.7%) were randomly assigned to the control group, and the remaining 54,089 were assigned to the treatment condition.

Intervention: Customers in the treatment condition were contacted by the operator and encouraged to upgrade to one of the two highest fixed monthly credit plans — plans with a fixed monthly credit of \$47 and \$63. For ease of discussion, we will term the plan with a fixed monthly credit of \$47 (\$63) as the lower (higher) featured plan. For all customers, including those in the control condition, at least one or both of the suggested plans were better than their current plan (please see the summary statistics in the next section).[5]

All customers were contacted by phone, and were asked to accept or reject the promotion during that month.[6] Those who accepted the promotion were automatically upgraded to the featured plan of their choice whereas those who rejected it stayed on their current plan. They could, though, switch to both non-featured and featured plans in later months. Similarly, while those in the control group were not contacted by the operator and did not receive any encouragement to switch plans, they were free to switch to any of the plans offered by the firm,

[5] By better we mean that the featured plan would have been cost-minimizing for the consumer in *all three* months prior to the campaign.
[6] To incentivize customers to upgrade to the one of the two suggested plans, the company offered an additional credit of \$15 for the following three months. No long-term contract commitments were required to accept the promotion. Thus, the reported churn rate after this specific campaign is a conservative estimate of the outcome from other encouragement campaigns that do not include such an incentive.

including the featured plans. Of the 54,089 customers who received the encouragement, 828 (1.53%) accepted it.

2.2 Data

For each customer, we observe monthly revenue for three months before and after the campaign. Note that if a customer leaves the company after the campaign, their revenue is zero for each month after churn. Monthly revenue from each customer can be split into (a) fixed monthly credit (i.e., automatic top-up at the beginning of the month) and (b) additional credit that the customer may purchase (i.e., credit added during the month).[7]

Pre-Campaign Descriptives: Table 1 shows the descriptive statistics for each of the two groups prior to the campaign. Both groups are very similar in all observed characteristics. The distribution of monthly (fixed) fee is almost identical across groups. The monthly revenue averaged across customers is \$112 (\$114) in the control (treatment) group. Customers also exhibit very similar level of within-individual revenue variation across the two conditions: the coefficient of variation of revenue (using the three months prior to the campaign) is 0.21 for both groups. Table 1 corroborates that customer were indeed randomly assigned to the two groups.

Insert Table 1 here

Post-Campaign Descriptives: Table 2 shows the extent of plan switching after the campaign. In general, the campaign was successful in encouraging switching behavior: 6.5% of customers in the control group switched tariffs during the first three months after the campaign, as compared to 9.6% in the treatment group ($p < 0.001$). Regarding the plan that customers switch to, 35.5% of the switchers in the treatment condition chose one of the two featured plans, whereas only

[7] This revenue comes from voice, SMS and data, which we do not observe separately.

22.5% did so in the control condition. Recall that the latter were not exposed to any encouragement to switch plans but some customers did on their own initiative. Finally, the percentage of customers who switched again is low in both conditions (2.1% and 2.6% of the switchers in the treatment and control groups, respectively) and is not significantly different ($p > 0.1$).

Insert Table 2 here

While the campaign had a positive impact on plan switching, note that only around 10% of customers in the treatment group did so. In other words, 90% of customers refused to switch tariffs, even knowing that their current plan was not cost-minimizing. Also, most of the switchers (64.56%) opt for a plan that was not recommended, suggesting that the mere act of contacting people, rather than the specific featured plans, could drive switching behavior.

3. Profitability of the intervention

In this section, we analyze the impact of the encouragement on customer usage and churn. First, we measure the impact of the campaign on customers' revenues and churn propensity. Next, we address self-selection issues and investigate the effect of the encouragement depending on whether customers accepted or rejected the recommendation.

3.1 Impact of the campaign on churn and revenues

As we conduct a randomized experiment, we can compare customer behavior before and after the campaign across the two experimental conditions. For the analysis regarding churn, given that all customers were active at the time of the campaign, we compare post-campaign churn across conditions. Churn is computed as the proportion of customers who churned, at any time, during the three months after the campaign. To analyze changes in revenue, we compute, for each customer, the difference between her average consumption during the three months before and after the campaign. We do so in two ways. We calculate the difference in conditional revenue, which considers only customers who stayed with the provider for all three months after the campaign. We also compute the difference in net revenue, which uses the data from all customers, including those who churned. Note that in those cases, revenue is set to zero for all months after churn. Table 3 summarizes the results.

Insert Table 3 here

The results suggest that there is a decrease in customer revenue which is independent of the campaign — those in the control group decreased their revenue by, on average, \$3.90 per month.[8] The campaign seems to have little impact on customer revenue since those in the treatment condition decreased their consumption by \$3.87, on average. In contrast, the campaign had a significant effect on customer churn. While churn in the control group was around 6%, the proportion of churners in the treatment group was 10%. This drop in the number of customers has notable consequences for net revenue—while the control group decreased net revenue by \$9.73, the treatment group did so by \$13.64.

To statistically assess the significance of the treatment, we estimate two different regressions with a treatment dummy as a covariate. We use a logistic regression to model churn behavior and a linear regression to model the difference in the average revenue from each customer before and after the campaign. The latter is again determined in two ways — conditional revenue and net revenue. Table 4 contains the results.

[8] Even accounting for the observed decrease in revenue, the large majority of customers would have saved money (ex-post) should they have switched to one of the featured plans. For example, looking at the control group, only 2.53% of customers had revenue below the \$47 offered in the lower featured plan. On the contrary, 86% of them would have saved money in *all three* periods after the campaign.

The results suggest that the campaign clearly increased churn among customers in the treatment group. In addition, if one considers the revenue from consumers who stayed with the company for all three months after the campaign, there is no significant change in their revenue. Finally, when looking at net revenue, there seems to be a negative impact of treatment driven by the revenue loss due to the increase in churn.

Insert Table 4 here

On the surface, these results are disappointing from the firm's perspective as one of the benefits it was hoping for was customers to stay longer with the firm after being exposed to the intervention. Such reasoning on the part of the firm, however, assumes that the campaign will appeal to all customers. This may not be the case especially since the majority of customers rejected the offer. We split the above results based on whether customers accepted or rejected the campaign and explore the effects on revenue and churn rate for both groups. Table 5 presents the results.

Insert Table 5 here

There are notable differences in post-campaign behavior based on whether or not a customer accepted the offer to switch to a more optimal plan. In particular, *revenue* increases mostly for customers who *accept* the offer (a $0.33 increase in monthly revenue vs. a $3.90 decrease) whereas c*hurn* increases for customers who *rejected* the offer (10.11% vs. 6.40%). Regarding defection among the customers who accepted the promotion hence allocated themselves in optimal plans, we observe that churn is slightly lower than that of the control condition (5.92% vs. 6.40%).

Overall, these results suggest that the campaign worked only for customers who accepted the recommendation, while it failed for those who rejected it. However, one must be cautious when making such claim since even though the field experiment randomly allocates customers into control and treatment conditions, there is no control over who decides to accept the offer; ultimately it is the customer's choice. As a consequence, we cannot assume a priori that individuals belonging to each of these groups (either accept or reject) are equivalent to those in the control condition. To further investigate this issue, we look at the differences in usage behavior (before the campaign) between the control and the self-selected conditions arising from customers who accept or reject the offering. Table 6 shows the descriptive statistics for these three groups of customers.

Insert Table 6 here

We observe that customers who accepted the recommendation have higher revenues ($13 more on average) and a slightly lower variability. In other words, the allocation of customers in the accepted or rejected groups was not random. As a consequence, in order to compare the post-campaign behavior between these two groups and the control group we first need to control for self-selection within customers in the treatment condition. Such self-selection from respondents within a randomized experiment results in what is generally called "broken randomized experiments" (Barnard et al. 1998; Hirano et al. 2000), where there is not guarantee that individuals in the treatment condition will comply with the treatment. In the next section we describe how we control for this issue in the current context.

3.2 Controlling for self-selection

We use propensity score matching (Rosenbaum and Rubin 1983) to control for possible sources of self-selection. Propensity score is been widely used in statistics (e.g., Rosenbaum and Rubin 1983, 1984, 1985), biostatistics/medicine (e.g., Rubin 1997; D'Agostino 1998; Brookhart et al. 2006; Austin 2008) and economics (e.g., Dehejia and Wahba 2002; Hirano, Imbens and Ridder

2003). More recently, it has been also applied in the marketing literature to compare non-balanced samples in the context of DVR adoption (Bronnenberg, Dubé and Mela 2010), market entry (Huang et al. 2012), participation in referral programs (Garnefeld et al. 2013), online and offline behavior (Gensler, Leeflang and Skiera 2012), and the implementation of CRM programs (Mithas, Krishan, and Fornell 2005), among others.

Unlike most of the studies in marketing, which rely on non-experimental data, we utilize data from a field experiment and so observe customers who did not have the option to accept the promotion (i.e., those in the control group). Therefore, our implementation of the propensity score resembles that of the field studies with noncompliance, or "broken randomized experiments," (e.g., Hirano et al. 2000) with the exception that, in our case, some customers never had an opportunity to comply. We account for these characteristics in our selection model.

We proceed in three stages: First, we select all customers in the treatment condition and model the probability (or propensity) of accepting the recommendation given other observed variables. Second, based on the obtained model parameters and the observables, we calculate the probability of accepting the promotion for all customers in our sample, including those in the control condition. Finally, we employ a matching technique to balance the samples of the two pairs of customer groups (control vs. accepted, and control vs. rejected) to compare post-campaign behavior between those who accept or reject the promotion and the control group.

Stage 1- Propensity model: For a customer *i*, let $accept_i$ be an indicator variable that takes a value of 1 if she accepts the promotion and 0 otherwise. Then, the probability (i.e., the propensity) of accepting the promotion is modeled as follows:

$$\text{Prob}(\text{accept}_i) = \text{Prob}(\alpha + X_i * \beta + \varepsilon_i > 0), \quad (1)$$

where α is a constant, X_i contains observed customer-specific characteristics, β is the sensitivity to these characteristics, and ε_i is normally distributed with mean 0 and unit variance. That is, we model the probability of accepting the promotion using a probit model.

Regarding X_i, we select those customer characteristics that are likely to be related with the decision of accepting/rejecting the promotion. Consistent with economic theory, customers' expected savings should be positively related with switching to a more optimal plan. As a proxy for individual's expected savings, we use the difference between the customer's current consumption and the fee of the lowest of the offered alternatives. In addition, current plan is also likely to affect the decision to upgrade plans. If a customer is used to pay a higher plan fee, she might be more willing to switch to a higher fee plan, given that her monthly fixed expenditure is already high. Finally, we also expect usage variability to affect this decision. Customer with stable patterns of usage will forecast their usage more easily, thus will be more certain of the potential savings of the new plan. Moreover, a customer with very high variance in usage might anticipate that there will be low usage months in which the new plan might not be cost-minimizing.

Using the available data, we compute, for each customer, *Distance* (measured as the difference between the average revenue prior to the campaign to the fee of the lower featured plan), *Plan fee* (i.e., the fixed amount in their current plan), and *variability* (measured as the individual coefficient of variation in revenue in the three months prior to the campaign).We estimate several model specifications including interactions for the observed variables and select the propensity score model on the basis of fit.[9] Table 7 reports the estimates for the best fitting propensity model.

Insert Table 7 here

The estimates from the propensity model have face validity. We find that higher expected savings make it more likely that a consumer will accept the encouragement but with diminishing returns. Customers that are on higher priced plans are more likely to accept the encouragement and, the higher the variability in their past usage, the more likely they are to reject it.

Stage 2- Computing propensity scores: We now consider all customers in our sample and use the estimates of our model to compute the propensity score (i.e., the probability to accept the promotion given the observed variables) for each customer. Note that this step now includes the customers in the control group who were not considered to estimate the propensity model because the promotion was never offered to them. Nevertheless, we can compute their likelihood of doing so given their observed characteristics.

Stage 3- Effect of encouragement in each of the groups: Finally, in order to compare the outcomes of interest between those who accept or reject the promotion and the control group, we employ a matching technique to balance the samples. We proceed in two ways for the two different dependent variables: (1) To measure the effect in revenue, we construct a matching estimator that matches each group of customers (either from the accepted or rejected condition) with those in the control group with similar propensities. In particular, we use a (non-parametric) kernel matching such that each treated observation is compared with a weighted average of the

[9] The results for all specifications are reported in the Appendix

outcomes of the non-treated (i.e., control) customers.[10] (2) To assess the effect of the encouragement on customer churn, we use regression adjustment with the propensity score (e.g., Austin and Mamdani 2006). In particular, we run a logistic regression with churn as a dependent variable and the encouragement treatment and the propensity score as independent variables. Table 8 shows the results for such analyses.

Insert Table 8 here

Results show that, after controlling for self-selection, customers who accepted the recommendation increased their revenue significantly more than those in the control group (albeit the magnitude is small, \$4 per customer per month). Such an increase in customers' revenue is consistent with past work which suggests that, after switching to pricing plans with more available allowance, consumers increase their usage more than what can be economically rationalized based on their past behavior (Ascarza et al. 2012; Gopalakrishnan et al. 2013). Regarding churn, our analysis shows that after controlling for self-selection due to noncompliance, the campaign had a very negative effect on churn. Whereas the customers who accepted the promotion seem to churn slightly less than those in the control group (the difference is not statistically significant), those who rejected the company's recommendation reacted very negatively, increasing their churn propensity significantly. Given that in our setting the average monthly churn rate is approximately 2% (we observe that 6.40% of customers in the control condition churn within three months), a 1.647 increase in churn across three months corresponds to, approximately 0.6 percentage points increase in monthly churn rate.[11]

---

[10] This matching method leverages all available data hence improves efficiency. To test the robustness of our results, we replicate the analysis using nearest neighbor matching which, unlike the kernel method, minimizes bias as a trade-off of efficiency. The results obtained using both methods are statistically equivalent.

[11] In Section 6.1 we assess the long-term revenue impact of the encouragement campaign.

4. Possible drivers of churn behavior

We propose two underlying drivers for the increase in churn behavior in more detail. The first driver is related to lowering the inertia that prevents customers to switch plans while the second driver is related to regret from salience of suboptimal behavior. Our data suggest that both drivers are at work.

4.1 Lowering customer inertia

Past research provides evidence that customers fail to switch service plans even in the absence of any commitment contract (e.g., Ascarza et al. 2012; Goettler and Clay 2012; Iyengar et al. 2007). The lack of switching is frequently attributed to customer inertia (e.g., Dubé, Hitsch and Rossi 2010). In the present context, it is likely that the reallocation campaign decreases the inertia to switch plans by, for example, making customers realize how easy it is to switch. Once customers realize that it is easy to switch plans within the company, they may also explore competitive offerings. Note that we can only speculate on the latter as we do not observe whether people who leave the company open an account with any of the competitors in the market.[12] We can, however, test for the underlying mechanism based on the amount of switching we observe across plans that the company offers.

Given that customers could change tariffs at any time, we compare the level of switching between those who rejected the offer and the control group during the three months following the campaign.[13] The rationale for this comparison is as follows: if the campaign reduced the inertia to switch, customers who rejected the campaign should still have a higher propensity to switch

[12] Over the duration of the field experiment, none of the competitors changed prices or ran specific promotional activities aimed to capture customers from other carriers.
[13] We do not include customers who accepted the promotion, for whom switching rate = 100%.

plans than people in the control group, who were not given any encouragement. This is indeed so (see Table 9). While 6.54% of customers in the control condition switched plans during the three months after the campaign, 8.18% of customers who rejected the promotion switched plans during the same time period.

Insert Table 9 here

Note that such a comparison could potentially suffer from self-selection bias given that we only compare the control group with customers who reject the campaign. Similar to the analysis presented when measuring the effect on revenue and churn, we control for self-selection using the propensity score for accepting the promotion (see Equation 1) and run a logistic regression with switching as a dependent variable and the treatment and the propensity score as covariates. The results (Table 10) show that customers in the treatment group who rejected the campaign were significantly more likely to switch plans in the next three months as compared to customers in the control group. This is the case even after controlling for self-selection from noncompliance. We therefore conclude that the campaign reduced customers' switching inertia.

Insert Table 10 here

4.2 Regret from salience of sub-optimal behavior

The unpleasant experience of finding out that a forgone alternative would have resulted in a better outcome is known as regret (Bell 1982; Tsiros 1998). In our context, the campaign emphasizes the sub-optimality of consumers' past decisions and makes them aware of how much they could have saved if they had been on a different, better, plan. Therefore, it is likely that the encouragement induced customers to feel regret. Past research in consumer behavior has shown that experiencing regret has a negative impact on customer satisfaction (Inman et al. 1997) and a direct effect on brand switching (Tsiros and Mittal 2000).

To find evidence of such claim in our data, we analyze factors that are related with the decision to churn. Our data do not include self-reported measures of customer regret, unfortunately, but have information that can be used as a proxy. We can use the variable *distance* for this purpose. Recall that *distance* measures the difference between the average revenue prior to the campaign and the fee of the lower featured plan. We posit that the relationship between *distance* and churn will be different, depending on whether customers have either accepted or rejected the recommendation. For customers who accept the recommendation, *distance* should not be a relevant predictor for churn. This is because these customers have already acted upon and neutralized their feeling of regret by accepting the encouragement. In contrast, for customers who reject the offered plans and therefore continue to experience regret, *distance* should be a significant driver of churn. Finally, since it is the encouragement what makes past suboptimal behavior salient, then for customers in the control group (who did not receive any encouragement and hence regret was not made salient), regret might not be a significant driver of churn.

To corroborate the above arguments, we analyze the relationship between *distance* at the moment of the campaign and subsequent churn in each of the three groups of customers. We proceed as follows. For a customer *i*, let $churn_i$ be an indicator variable that takes a value of 1 if she leaves the service at any time within three months after the campaign and 0 otherwise. We model the probability of churn as follows:

$$\text{Prob}(\text{churn}_i) = \text{Prob}(\gamma + Z_i * \delta + \eta_i > 0), \qquad (2)$$

where γ is a constant, Z_i contains observed customer-specific characteristics and δ is the sensitivity to these characteristics, and η_i is normally distributed with mean 0 and unit variance.[14]

Regarding Z_i , we use the same set of covariates as in the model for the propensity to accept the promotion (see Table 7). As before, we estimate various model specifications by adding quadratic and interaction terms. We select the simplest model (without interaction variables) as it provides consistently the best fit across the three groups (control, accept, and reject). Table 11 contains the estimates from the churn model for all three sets of customers.

Insert Table 11 here

Looking at the results across the three groups we find a pattern that is consistent with the claim that the campaign induced regret. Specifically, we find that *distance* is significant *only* for customers who rejected the recommendation (middle column). Consistent with regret theory, only for these customers for whom the possible savings were made salient by the firm's recommendation and who did not act upon it (i.e., rejected the encouragement), the potential benefit from other plans significantly correlates with churn behavior.

5. Discussion of assumptions and robustness of the results

We discuss the validity of some of our methodological choices, test the robustness of our findings and explore additional explanations for the phenomenon we observe in the data.

5.1 Model selection on unobservables

We employ a propensity score model to address self-selection issues arising from customers in the treatment group either accepting or rejecting the encouragement (e.g., Rosenbaum and Rubin

[14] It is plausible that the stochastic components of Equations (1) and (2) are correlated as there may be unobserved common factors driving the decision to accept the campaign and to churn. We explore this issue in Section 5.

1983; 1984; 1985). Such models address the selection based only on observables. It is possible that the observed customer characteristics might not completely capture all forms of self-selection. For example, it is plausible that customers who expect to churn in the future are less likely to accept the promotion since their expected savings will not be as high. We test whether such unobservable factors are a concern in our context by testing the correlation between the residuals from the model for the propensity to accept the promotion (Equation 1) and the model for churn (Equation 2). The rationale is as follows. If there were common unobservable factors affecting both decisions, then the residuals of our models—which include only observable characteristics—should be correlated. We find that this correlation is 0.0019 ($p = 0.6602$) for all customers who received the campaign. If we separate customers based who accept and reject the campaign, then the correlations are –0.023 ($p = 0.5094$) and 0.0038 ($p = 0.3837$) respectively.

As a consequence, while our propensity score model may not fully capture all factors related with the decision of accepting the promotion, this analysis provides empirical support that the missing variables are uncorrelated with subsequent churn behavior. We therefore feel confident that our results do not suffer from selection bias.

5.2 Allowing for temporal dynamics and unobserved heterogeneity in customer churn

Our analysis thus far has measured churn as a single observation per customer which takes a value of 1 if the customer churns during the 3 months after the campaign was run, and 0 otherwise. This approach might appear simplistic a priori since it does not control for time dynamics (e.g., trend, seasonality) and does not allow for customer heterogeneity. Given the exogenous randomization in the field experiment, and the fact that the propensity score already controls for customer self-selection, we do not believe that these two factors will alter the effects we observe. Nevertheless, we check the robustness of our results by leveraging the longitudinal aspect of our data. We estimate all regressions (Tables 8 and 11) with a panel data with multiple observations per customers in which we control for time effects (via time trend) and for unobserved customer heterogeneity (via random effects). In all cases, we replicate the results presented in previous sections. The results are available from the authors.

5.3 Churn triggered by *any* firm intervention

We have proposed two theoretical explanations as of why the reallocation campaign results in higher churn rates, namely lowering inertia and customer regret. Further analysis of the pre- and post-campaign behavior empirically supports these two mechanisms. However, one could argue that another possibility for the increase in churn is the mere intervention from the firm. That is, simply contacting customers through the campaign might have prompted people to leave. Thus, it is not specifically the reallocation promotion that made people aware of their suboptimal behavior but potentially any intervention may show an impact similar to what we observe in this research.

To address this concern, we obtain information about another campaign run by the focal firm during, roughly, the same time of our study. The two campaigns are very different in nature. The goal of this new — so called "commitment" — campaign was to secure long-term commitment contracts among post-paid customers. All customers selected for the commitment campaign were enrolled in a different plan than the one analyzed earlier. That is, there is no overlap between customers in the two different campaigns. For this campaign, the firm selected 150,038 post-paid customers, of which, 6.1% were randomly allocated to the control group. Those in the treatment condition received a text from the company offering a new handset at a discounted price if the customer agreed to enroll in a long-term (18 month) commitment contract. Among the customers who received the promotion, 1.2% accepted the offer.

We focus on churn behavior after the campaign and compare customers who rejected the campaign with those in the control condition. We find no significant difference between these two groups. Churn rate is slightly lower among customers in the control group (0.41% versus 0.54%). However these two figures are not significantly different.[15] We, therefore, conclude that it seems unlikely that the mere intervention from the firm is driving the increase in churn we observe in the reallocation campaign.

5.4 Plan upgrades and downgrades

In the analysis described above, we label a plan switch as the choice of any plan that is different from the current plan. A more granular approach could categorize a switch as either an upgrade (a higher fixed monthly credit) or a downgrade (a lower fixed monthly credit). Note that any customer that accepts the recommendation is by definition upgrading their plan. For the other two groups (i.e., control group or rejectors), we determine the fraction of upgrades and downgrades. For the control condition, out of 10,058 customers, 567 customers upgraded (5.6%) and only 91 downgraded (0.9%). Similarly, from the set of 53,261 customers who rejected the campaign, 3,885 upgraded (7.3%) and only 473 downgraded (0.9%). There are two points to note. First, given that the campaign had targeted customers who were over-consuming, it is not surprising that there is little plan downgrade. Thus, our label of a plan switch primarily captures upgrades. Second, there is no difference in the level of downgrade across the control and rejected

[15] We assess the statistical significance of these comparisons in two ways: (1) applying propensity score in the same fashion as in section 3, and (2) computing confidence intervals around the group proportions using bootstrap. Both approaches provide a convergent set of results and they are available from the authors.

conditions indicating that the campaign did not make people choose lower fixed monthly credit plans.

6. Managerial implications

In this section we measure the impact of the campaign on firm long-term profitability by combining the effects on customer churn and revenue. Next, we elaborate on the proposed mechanisms affecting churn and discuss how companies could design more profitable campaigns through targeting.

6.1 Long-term impact of the encouragement campaign on firm profitability

We have provided strong evidence that encouraging customers to switch to more optimal tariffs has two significant effects: On the one hand, customers who accept the encouragement increase their revenue, hence become more profitable to the firm. On the other hand, customers who reject the recommendation show much higher levels of churn than those in the control condition. Past work on customer lifetime value (e.g., Gupta et al. 2004) has noted that changes in customer retention have a much larger influence on their lifetime value than do changes in margins. To assess the extent to which this is the case in our setting, we combine the obtained effects on revenue and churn and compute the change in customer lifetime value (CLV) for customers in the control and treatment conditions.

Given the context of our analysis, we assume a monthly revenue of \$112 (see Table 1), a 2% monthly churn rate (consistent with the data on churn in the control group — see Table 3) and a

15% operating margin on customer revenues.[16] These metrics result in an estimate of lifetime value of $549 for a customer in the control group. Now, let us consider two scenarios for the treatment group. In the first scenario — i.e., accept scenario— in which a customer has higher revenue ($116 = $112 + $4) but no increase in churn rate, and the second scenario — i.e., reject scenario — in which customer's revenue in unchanged and the monthly churn rate is 2.6%.[17] Combining these metrics we obtain that, for a customer who accepts the recommendation, her (post-campaign) lifetime value to the firm is $568, whereas for a customer who reject it, her lifetime value is only $455. Putting these numbers in a bigger scale, let us assume that there are 100,000 customers who qualify for an encouragement campaign similar to the one analyzed in this research. The overall value of the customer base is $54.9M if the company does nothing, and is $56.8M ($45.5M) in the higher revenue (higher churn) scenarios.[18] For our context then, looking at the two decisions altogether, the effect of the marginal increase in revenue on overall firm profitability is much weaker than that of the significant increase in churn. As a consequence, an encouragement campaign like the one analyzed here should be run with caution since the risk from its rejection may overshadow the benefit from its acceptance.

6.2 Design more profitable campaigns through targeting

Our results, thus far, show the following. First, the mass campaign that the company ran was not successful — it increased the amount of customer churn (mostly from customers who rejected the campaign) while the impact on customer revenue was marginal. Second, from our analysis of

---

[16] See http://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/margin.html for margins across various industries. We use the operating margin corresponding to telecom services. In addition, assuming a constant monthly churn rate, a 6.4% churn over a period three months corresponds to around a 2% monthly churn.

[17] Assuming a base churn rate of 2%, an increase of 1.65 in the proportion of customers who churn over a three-month period (Table 9) corresponds to a monthly churn rate of 2.6%.

[18] Note that, mainly due to data limitations, we abstract from social effects in customer retention (Nitzan and Libai 2011). If those effects were present in this context and we incorporated those in the customer lifetime value framework, then the difference between the two scenarios would be even more pronounced.

the acceptance of the campaign and churn behavior, two customer variables are important: (a) *variability* (measured as the individual coefficient of variation in revenue prior to the campaign) and (b) *distance* (measured as the difference between the average revenue prior to the campaign to the fee of the lower featured plan). Note that our results suggest that customers with low levels of variability are more likely to accept the campaign and, if they reject it, are less likely to churn later (see Tables 7 and 11). The impact of distance on customer behavior is more nuanced: customers with high level of distance are more likely to accept the promotion but are also more likely to churn later if they reject it. In what follows, we determine the net impact of the encouragement if the company had targeted customers based on these two observed characteristics. This analysis can help in assessing the impact from a targeted campaign if the firm implements one in the future.

We leverage the richness of the field experiment to determine the impact of targeted campaigns based on variability and distance measured at the moment of the campaign. From the full sample of customers (N = 64,147), we select the customers with highest observed variability (top 20^{th} percentile) and highest distance (top 20^{th} percentile) before the intervention, resulting in a sub-sample of 2,565 customers. We then compare their actual post-campaign behavior between control and treatment conditions using the data from our field experiment. Given that the control/treatment split was random, the proportion of treated customers in this subsample very much resembles that from the full population. We proceed similarly to create a subsample of customers with high variability-low distance, low variability-low distance and so forth.

Regarding the measures of post-campaign behavior, we determine customer churn and revenue similar to what we presented earlier in Section 3.1. Briefly, churn is computed as the proportion of customers who churned in the three months after the campaign. To analyze changes in revenue, we compute, for each customer, the difference between their average consumption during the three months before and after the campaign. For calculating the difference in conditional revenue, customers who stayed with the provider for all three months after the campaign are included. For determining the difference in net revenue, we use the data from all customers, including those who churned. In those cases, revenue is set to zero for all months after churn. Tables 12a-12d show the differences in customer behavior between control and treatment conditions, resulting from the four targeted campaigns.

This analysis confirms that a targeted reallocation campaign in which customers with low levels of variability are selected results in much lower churn. For instance, Table 12b shows that targeting customers who are at high levels on both variables results in 14.17% churn in the treatment group. In contrast, targeting consumers with low variability and high distance leads to 9.53% churn (Table 12d). Upon averaging the churn rates across conditions when the variability is low or high, we find that churn is around 8.8% in the low and 14.3% in the high variability conditions.

Insert Table 12 here

While a high level of distance increases the propensity for consumers to accept the encouragement, these results suggest that the negative impact on churn from consumers who reject the promotion is substantially enough to prefer targeting customers with a low level of distance. Upon averaging the churn rates across conditions when the distance is low or high, we find that churn is around 11.2% in the low and 11.9% in the high distance conditions.

A campaign targeting customers that are low on both variables leads to the lowest level of churn in the treatment group (8.05%), which is also lower than the rate in the corresponding control group (8.76%). Such targeting is beneficial for both conditional revenue and net revenues from the treatment group as compared to the control group (Table 12c). Finally, this targeted campaign is significantly better than the mass campaign on all three metrics of resulting customer behavior (see Table 3).

In sum, our insights into underlying driving mechanisms of regret and inertia are useful to develop potential targeted campaigns.

7. General discussion and conclusions

Firms seeking to maximize profits often rely on price discrimination to better match prices to customers' valuations. Whether this price discrimination strategy is useful depends on whether consumers will select the plan that maximizes their welfare. However, much past evidence reveals that customers do make ex-post mistakes when choosing service plans. In this paper we examine the profitability of encouraging customers to switch to better plans. We use data from a field experiment in which some customers received an encouragement to switch to more optimal plans, while others did not. We find that encouraging customers to switch to better tariffs might, in turn, be very harmful for the firm. The main source for this negative effect is the increase in churn behavior from customers who decide to reject the recommendation. On the contrary, customers who accept the firm's recommendation increase their usage, hence providing higher revenue to the firm. The results are robust to consumers' self-selection due to noncompliance, temporal dynamics and unobserved heterogeneity. We propose two theoretically founded mechanisms to explain how the firm's encouragement may increase customer churn, namely lowering inertia and customer regret. Our data provides evidence that both mechanisms are at work.

We offer clear guidance to service providers (e.g., telecommunications, utilities, financial services) concerned about the mismatch of customers to service plans. In particular, we show how selecting the right customers to target (in this case, avoiding the wrong customers) would have a much higher impact in profitability than the mere act of matching customers to optimal tariffs. Furthermore, firms should target customers based on easily observed characteristics like distance (or proximity between the individual's usage and the offered plan) and individual usage variability in order to improve the outcome of reallocation campaigns.

It is important to note that in this research we evaluated an "upgrade" campaign. That is, it encouraged consumers who were currently over-consuming to upgrade to a better plan. In contrast, there are many contexts in which consumers may purchase plans that have a higher coverage than what they optimally need (see for example Johnson et al. 2013). In such cases, it may be best for consumers to downgrade their service contracts. Would campaigns that encourage customers to downgrade have an impact on churn similar to what we have documented? We speculate that the impact would be comparable primarily because our two proposed mechanisms of lowered inertia to switch and regret would still apply, with the latter resulting from paying a higher monthly fixed fee than what customers should. Nevertheless, we acknowledge that there could be also differences in behavior due to the different motivations that affect upgrade vs. downgrade decisions.

Our research complements the extant empirical literature on price discrimination (e.g., Ascarza et al., 2012; Danaher 2002; Iyengar et al. 2011; Narayanan et al. 2007) that has evaluated the impact of offering differing pricing contracts on firm profitability. Our results show that firm profitability is also affected (not necessarily improved) when firms proactively encourage customers to switch to better tariffs. It also adds to the broader literature on product recommendations by offering a new setting in which customers may exhibit reactance (Bodapati 2008; Häubl and Murray 2006; Fitzsimons and Lehmann 2004). Our results suggest that it may be best for retailers to refrain from giving any recommendations after consumers make a purchase as it may induce regret among those who compare the recommendation with their past purchase. A caveat is in order as our research was undertaken in a recurrent (subscription-based) setting in which customers already subscribed to the service with a contractual plan as opposed to a one-shot transactional one where the status quo may be not to own the product.

More broadly, we add to the literature on choice architecture and its impact on consumer behavior (e.g., Thaler and Sunstein 2008). Choice architecture broadly refers to the notion that the way that choices are structured and described can have a substantial impact on what is chosen (see Johnson et al. 2012 for a summary). One of the factors determining the success of any choice architecture is the degree of freedom that consumers perceive to have when making their decisions (Sunstein and Thaler 2003). In line with this idea, a libertarian paternalistic (or soft paternalistic) policy suggests that choice architects can influence people's behavior to make their lives better but people should also be allowed to accept any recommendations on their own volition. The impact of nudges has been evaluated in contexts such as healthcare and financial decisions (Thaler and Benartzi 2004; Kling et al. 2012). The firm's proactive action in our context can be construed as soft paternalism: customers do not necessarily have to choose any of the options that the firm believes is good for them. Our results suggest that even when people are making choices in relatively libertarian contexts, there can be severe negative consequences from recommendations mainly from people who reject them. That nudges have a differential impact depending on the characteristics of decision makers is consistent with past work (e.g., Costa and Kahn 2013). Costa and Kahn (2013) find that while informing households about relative energy use led to an average 2% decrease in usage, liberals reduced their consumption while Republicans increased theirs. It will be interesting to test the robustness of our finding in other contexts especially since the mechanisms are quite general.

Our analysis can be extended in several ways. First, our 6-month window allows us to clearly pin down the pre- and post-experiment behavior. However, a longer time frame would allow the researcher to disentangle other dynamics in customer churn and revenue. Moreover, we do not observe whether the customers who churn from the focal provider switch to the competition. To offer a full support of that claim we would need individual-level multi-firm data which, unfortunately, is difficult to obtain. Furthermore, while we provide some support for both mechanisms, our research does not speak as to which factor—either inertia or regret—is more important on driving churn. Self-awareness (e.g., actual knowledge of one's usage) and attitudinal measures (e.g., satisfaction with the service) would be perfect complements to our data to clearly disentangle regret from inertia. We trust that future research will address these issues.

References

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin (1996), "Identification of causal effects using instrumental variables," *Journal of the American Statistical Association,* 91(434), 444-455.

Ascarza, Eva, and Bruce Hardie (2013), "A joint model of usage and churn in contractual settings," *Marketing Science,* 32(4), 570-590.

Ascarza, Eva, Anja Lambrecht, and Naufel J. Vilcassim (2012), "When talk is "free": The effect of tariff structure on usage under two and three-part tariffs," *Journal of Marketing Research,* 49(6), 882-899.

Austin, Peter C. (2008), "A critical appraisal of propensity-score matching in the medical literature between 1996 and 2003," *Statistics in Medicine 27*(12), 2037-2049.

Austin, Peter C., and Muhammad M. Mamdani (2006), "A comparison of propensity score methods: A case-study estimating the effectiveness of post-AMI statin use," *Statistics in Medicine, 25*(12), 2084-2106.

Bar-Gill, Oren, and Rebecca Stone (2009), "Mobile misperceptions," *Harvard Journal of Law and Technology,* 23(1), 49-118.

Barnard, John, Jiangtao Du, Jennifer L. Hill, and Donald B. Rubin (1998), "A broader template for analyzing broken randomized experiments," *Sociological Methods and Research, 27*(2), 285-317.

Bell, David E. (1982), "Regret in decision making under uncertainty," *Operations Research, 30*(5), 961-981.

Bodapati, Anand V. (2008), "Recommendation systems with purchase data," *Journal of Marketing Research,* 45(1), 77-93.

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

Bolton, Ruth N., and Katherine N. Lemon (1999), "A dynamic model of customers' usage of services: usage as an antecedent and consequence of satisfaction," *Journal of Marketing Research,* 36(2) 171-186.

Bronnenberg, Bart J., Jean-Pierre Dubé, and Carl F. Mela (2010), "Do digital video recorders influence sales?," *Journal of Marketing Research,* 47(6), 998-1010.

Brookhart, M. Alan, Sebastian Schneeweiss, Kenneth J. Rothman, Robert J. Glynn, Jerry Avorn, and Til Stürmer (2006), "Variable selection for propensity score models," *American Journal of Epidemiology*, 163(12), 1149-1156.

Costa, Dora L., and Matthew E. Kahn (2013), "Energy conservation nudges and environmental ideology: Evidence from a randomized residential electricity field experiment," *Journal of the European Economic Association,* 11, 680-702.

D'Agostino, Ralph B. (1998), "Tutorial in biostatistics: propensity score methods for bias reduction in the comparison of a treatment to a non-randomized control group," *Statistics in Medicine*, 17(19), 2265-2281.

Danaher, Peter J. (2002), "Optimal pricing of new subscription services: Analysis of a market experiment," *Marketing Science*, 21(2), 119-138.

Dehejia, Rajeev H., and Sadek Wahba (2002), "Propensity score-matching methods for non-experimental causal studies," *Review of Economics and Statistics*, 84(1), 151-161.

Dubé, Jean-Pierre, Günter J. Hitsch, and Peter E. Rossi (2010), "State dependence and alternative explanations for consumer inertia," *The RAND Journal of Economics*, 41(3), 417-445.

Fitzsimons, Gavan J., and Donald R. Lehmann (2004), "Reactance to recommendations: When unsolicited advice yields contrary responses," *Marketing Science*, 23(1), 82-94.

Garnefeld, Ina, Andreas Eggert, Sabrina V. Helm, and Stephen S. Tax (2013), "Growing existing customers' revenue streams through customer referral programs," *Journal of Marketing*, 77(4), 17-32.

Gensler, Sonja, Peter Leeflang, and Bernd Skiera (2012), "Impact of online channel use on customer revenues and costs to serve: Considering product portfolios and self-selection," *International Journal of Research in Marketing*, 29(2), 192-201.

Goettler, Ronald, and Karen Clay (2011), "Tariff choice with consumer learning and switching costs," *Journal of Marketing Research*, 48(4), 633-652.

Gopalakrishnan, Arun, Raghuram Iyengar, and Robert M. Meyer (2013), "Consumer dynamic usage allocation and learning under multi-part tariffs: theory and empirical evidence," Working paper. University of Pennsylvania.

Gupta, Sunil, Donald Lehmann, and Jennifer Ames Stuart (2004) "Valuing customers," *Journal of Marketing Research*, 41(1), 7-18.

Häubl, Gerald, and Kyle B. Murray (2006), "Double agents: assessing the role of electronic product recommendation systems," *Sloan Management Review,* 47, 8-12.

Hirano, Keisuke, Guido W. Imbens, and Geert Ridder (2003), "Efficient estimation of average treatment effects using the estimated propensity score" *Econometrica*, 71(4), 1161-1189.

Hirano, Keisuke, Guido W. Imbens, Donald B. Rubin, and Xiao-Hua Zhou. (2000), "Assessing the effect of an influenza vaccine in an encouragement design," *Biostatistics*, 1(1), 69-88.

Huang, Qingyi, Vincent R. Nijs, Karsten Hansen, and Eric T. Anderson. (2012), "Wal-Mart's impact on supplier profits," *Journal of Marketing Research*, 49(2), 131-143.

Inman, J. Jeffrey, James S. Dyer, and Jianmin Jia (1997), "A generalized utility model of disappointment and regret effects on post-choice valuation," *Marketing Science*, 16(2), 97-111.

Iyengar, Raghuram, Asim Ansari, and Sunil Gupta (2007), "A model of consumer learning for service quality and usage," *Journal of Marketing Research*, 44(4), 529-544.

Iyengar, Raghuram, Kamel Jedidi, Skander Essegaier, and Peter J. Danaher (2011), "The impact of tariff structure on customer retention, usage, and profitability of access services," *Marketing Science*, 30(5), 820-836.

Johnson, Eric J., Suzanne B. Shu, Benedict GC Dellaert, Craig Fox, Daniel G. Goldstein, Gerald Häubl, Richard P. Larrick, J. W. Payne, E. Peters, D. Schkade, B. Wansink and E. Weber (2012), "Beyond nudges: tools of a choice architecture," *Marketing Letters,* 23, 487-504.

Johnson, Eric J., Ran Hassin, Tom Baker, Allison T. Bajger, and Galen Treuer (2013), "Can consumers make affordable care affordable? The value of choice architecture," Working paper, University of Pennsylvania, Institute for Law and Economics Research, Paper No. 13-28.

Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel. (2011), "Comparison friction: Experimental evidence from Medicare drug plans," *The Quarterly Journal of Economics,* 127, 199-235.

Lemon, Katherine N., Tiffany Barnett White, and Russell S. Winer (2002), "Dynamic customer relationship management: incorporating future considerations into the service retention decision," The *Journal of Marketing*, 66(1), 1-14.

Mithas, Sunil, Mayuram S. Krishnan, and Claes Fornell (2005), "Why do customer relationship management applications affect customer satisfaction?," *Journal of Marketing*, 69(4), 201-209.

Narayanan, Sridhar., Pradeep K. Chintagunta, and Eugenio J. Miravete (2007). "The role of self selection, usage uncertainty and learning in the demand for local telephone service," *Quantitative Marketing and Economics*, 5(1), 1-34.

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(2), 204-211.

Nitzan, Irit, and Barak Libai (2011), "Social effects on customer retention," *Journal of Marketing*, 75(6), 24-38.

Rosenbaum, Paul R., and Donald B. Rubin (1983), "The central role of the propensity score in observational studies for causal effects," *Biometrika*, 70(1), 41-55.

Rosenbaum, Paul R., and Donald B. Rubin (1984), "Reducing bias in observational studies using subclassification on the propensity score," *Journal of the American Statistical Association*, 79(387), 516-524.

Rosenbaum, Paul R., and Donald B. Rubin (1985), "Constructing a control group using multivariate matched sampling methods that incorporate the propensity score," *The American Statistician*, 39(1), 33-38.

Rubin, Donald B. (1990), "Comment: Neyman (1923) and causal inference in experiments and observational studies," *Statistical Science*, 5(4), 472-480.

Rubin, Donald B. (1997), "Estimating causal effects from large data sets using propensity scores" *Annals of Internal Medicine*, 127, 757-763.

Sunstein, Cass R., and Richard H. Thaler (2003), "Libertarian paternalism is not an oxymoron" *The University of Chicago Law Review,* 70, 1159-1202.

Thaler, Richard H., and Shlomo Benartzi (2004), "Save more tomorrow: using behavioral economics to increase employee savings," *Journal of Political Economy, 112*(1), S164-S187.

Thaler, Richard H., and Cass R. Sunstein (2008), *Nudge: Improving decisions about health, wealth and happiness*. Yale University Press.

Tsiros, Michael (1998), "Effect of regret on post-choice valuation: The case of more than two alternatives" *Organizational Behavior and Human Decision Processes*, 76(1), 48-69.

Tsiros, Michael, and Vikas Mittal (2000), "Regret: A model of its antecedents and consequences in consumer decision making," *Journal of Consumer Research*, 26(4), 401-417.

Wieringa, Jaap E., and Peter C. Verhoef (2007), "Understanding customer switching behavior in a liberalizing service market an exploratory study," *Journal of Service Research*, 10(2), 174-186.

Zeelenberg, Marcel (1999), "Anticipated regret, expected feedback and behavioral decision making," *Journal of Behavioral Decision Making*, 12(2), 93-106.

TABLES

Table 1: Pre-campaign Descriptive Statistics

	Control group	Treatment group
Plan (in $)		
Mean	*37.89*	*37.38*
p10	35.95	35.95
p25	35.95	35.95
p50	39.95	39.95
p75	39.95	39.95
p90	39.95	39.95
Avg. Revenue		
Mean	*112.04*	*113.95*
p10	66.54	64.90
p25	78.82	77.34
p50	98.62	98.22
p75	128.88	131.29
p90	170.97	178.40
Coef. of variation Revenue		
Mean	*0.21*	*0.21*
p10	0.07	0.07
p25	0.12	0.12
p50	0.19	0.19
p75	0.28	0.27
p90	0.38	0.38

Table 2: Post-campaign Plan Switching

	Control		Treatment	
Number of customers	10,058		54,089	
Number of switchers (% of customers)	658	6.54%	5,186	9.59%
to higher featured plan (% of switchers)	71	0.71%	1,041	1.35%
to lower featured plan	77	0.77%	797	1.07%
to a non-featured plan	510	5.07%	3,348	7.17%

Table 3: Post-campaign Behavior for Churn and Revenue

	Control	Treatment
Conditional Revenue difference (in $)	–3.90	–3.87
Net Revenue difference (in $)	–9.73	–13.64
Churn (in %)	6.40	10.04

Conditional Revenue difference is based on only customers who stayed with the company for all three months after the campaign. Net Revenue difference is based on all customers (when a customer churns, her revenue is set to zero for all subsequent months). The churn percentage is based on the number of customers leaving the company during the three months after the campaign.

Table 4: Effect of the Campaign on Churn and Revenue

	Churn Behavior	Conditional Revenue difference ($)	Net Revenue difference($)
Treatment	**0.490*** **	0.0295	**–3.635*** **
	(0.043)	(0.471)	(0.573)
Constant	**–2.682*** **	**–3.903*** **	**–8.308*** **
	(0.0407)	(0.431)	(0.526)
Observations	64,147	58,068	64,147

Standard errors in parentheses. *** p <0.01

The churn column contains the results from a logistic regression that models churn behavior after the campaign. Both revenue columns contain results from a linear regression in which the difference in average revenue from each customer before and after the campaign is modeled. In the conditional revenue column, data from customers who stay with the provider for all three months after the campaign is included. In the net revenue column, data from all customers is included with revenue for a customer being zero for all months after churn.

Table 5: Post-campaign Behavior, by Campaign Acceptance

	Control	Did not accept	Accept
Conditional Revenue difference (in $)	–3.90	–3.94	0.33
Churn*(in %)	6.40	10.11	5.92

* during the three months after the campaign

Table 6: Descriptive Statistics at the Time of the Campaign, by Acceptance Group

	Control	Rejected	Accepted
Plan (in $)			
Mean	*37.89*	*37.37*	*38.01*
p10	35.95	35.95	35.95
p25	35.95	35.95	35.95
p50	39.95	39.95	39.95
p75	39.95	39.95	39.95
p90	39.95	39.95	39.95
Avg. Revenue			
Mean	*112.04*	*113.76*	*126.33*
p10	66.54	64.84	69.43
p25	78.82	77.22	84.44
p50	98.62	98.01	110.30
p75	128.88	130.97	147.20
p90	170.97	178.11	196.72
Coef. of variation Revenue			
Mean	*0.21*	*0.21*	*0.20*
p10	0.07	0.07	0.07
p25	0.12	0.12	0.11
p50	0.19	0.19	0.18
p75	0.28	0.28	0.26
p90	0.38	0.38	0.36

Table 7: Parameters Estimates of the Propensity Model

	Propensity model
Distance	0.002***
	(0.000)
Distance^2	–3.56e-06***
	(0.000)
Variability	–0.606***
	(0.168)
Plan fee	0.025***
	(0.005)
Distance x Variability	0.003*
	(0.001)
Constant	–3.114***
	(0.179)
Observations	54,089

Standard errors in parentheses. *** p <0.01, * p <0.1

Table 8: Effect of the Encouragement on Revenue and Churn (by Acceptance Group)

	Difference between…	
	Accepted and Control	Rejected and Control
Monthly **revenue** difference before and after campaign (in $)	**4.006*** **	–0.087
	(1.675)	(0.349)
Matching method	Kernel matching	Kernel matching
% of customers who **churn** in the first 3 months after the campaign	–0.105	**1.647*** **
	(0.153)	(0.071)
Matching method	Propensity score as covariate	Propensity score as covariate

Standard errors in parentheses. *** p <0.01

Table 9: Post-campaign Plan Switching: Comparing the Control Group with the Customers who Rejected the Encouragement

	Control		Reject	
Number of customers	10,058		53,261	
Number of switchers	658	6.54%	4,358	8.18%

Table 10: Effect of the Campaign on Switching Behavior: Comparing the Control Group with the Customers who Rejected the Encouragement

	Rejected vs. Control
Switch (% of customers)	**0.167*****
St. Error	(0.044)
Matching method	Propensity score as covariate

Table 11: Churn Model (by Customer Type)

	Control	Reject	Accept
Distance	0.0002	**0.0003*****	0.0012
	(0.0004)	(0.0001)	(0.0009)
Price Plan	**0.0251*****	**0.0167*****	0.0210
	(0.0071)	(0.0023)	(0.0254)
Variability	**0.305****	**0.840*****	0.410
	(0.1480)	(0.0529)	(0.577)
Constant	–2.562***	–2.110***	–2.551***
	(0.272)	(0.089)	(0.974)
Observations	10,058	53,261	8280

Standard errors are in parentheses. *** p<0.01, ** p<0.05

Table 12: Post-campaign Behavior for Churn and Revenue for targeted campaigns

(a) High Variability / Low Distance

	Control	Treatment	*Diff.*
Cond Revenue	–0.44	3.91	*4.35*
Net Revenue	–5.12	–4.72	*0.40*
Churn (in %)	7.65	14.42	*6.77*
N	2563		
% control	15%		

(b) High Variability / High Distance

	Control	Treatment	*Diff.*
Cond Revenue	–32.65	–43.29	*–10.64*
Net Revenue	–43.29	–64.58	*–21.29*
Churn (in %)	7.41	14.17	*6.76*
N	2565		
% control	16%		

(c) Low Variability / Low Distance

	Control	Treatment	*Diff.*
Cond Revenue	5.29	7.14	*1.85*
Net Revenue	1.00	2.70	*1.70*
Churn (in %)	8.76	8.05	*–0.71*
N	2565		
% control	14%		

(d) Low Variability / High Distance

	Control	Treatment	*Diff.*
Cond Revenue	–12.65	–17.46	*–4.81*
Net Revenue	–20.99	–31.66	*–10.67*
Churn (in %)	7.76	9.53	*1.77*
N	2565		
% control	14%		

Conditional Revenue difference is based on only customers who stayed with the company for all three months after the campaign. Net Revenue difference is based on all customers (when a customer churns, her revenue is set to zero for all subsequent months) The churn percentage is based on the number of customers leaving the company during the three months after the campaign.

APPENDIX

In this appendix we present the results of all model specifications considered in the propensity score model. Using the inputs discussed in Section 3.2, we estimate several model specifications including interactions for the observed variables. Table A-1 reports the estimates for these specifications.

Table A-1: Parameters Estimates of the Propensity Model

	Model				
	(1)	(2)	(3)	(4)	(5)
Distance	0.001***	0.002***	0.002***	–0.001	–0.002
	(0.000)	(0.000)	(0.000)	(0.003)	(0.003)
Distance^2		–3.41e06***	–3.56e-06***	–3.54e06***	–3.67e-06***
		(1.24e-06)	(1.21e-06)	(1.26e-06)	(1.24e-06)
Variability	–0.350***	–0.373***	–0.606***	–0.374***	–0.601***
	(0.111)	(0.112)	(0.168)	(0.112)	(0.168)
Plan fee	0.025***	0.025***	0.025***	0.017**	0.018**
	(0.005)	(0.005)	(0.005)	(0.007)	(0.007)
Distance x Variability			0.003*		0.003*
			(0.001)		(0.001)
Distance x Plan fee				0.000	9.63e-05
				(7.79e-05)	(7.76e-05)
Constant	–3.120***	–3.157***	–3.114***	–2.888***	–2.861***
	(0.178)	(0.178)	(0.179)	(0.268)	(0.268)
Observations	54,089	54,089	54,089	54,089	54,089

Standard errors in parentheses. *** p <0.01, ** p <0.05, * p <0.1

We select the most appropriate model on the basis of model fit (Table A-2). Based on the AIC and the BIC, the best fitting models are Models 2 and 3, respectively. The sole difference between these two models is whether the interaction between Distance and Variability is included. To compare these two model specifications we perform a LR test and conclude that Model 3 is the best fitting model (p = 0.039).

Table A-2: Fit Measures of the Propensity Models

	Model				
	(1)	(2)	(3)	(4)	(5)
Log Likelihood	–4246.7	–4240.2	–4238.5	–4239.3	–4237.7
Df	4	5	6	6	7
AIC	8501.4	8490.4	**8489.0**	8490.6	8489.4
BIC	8537.0	**8534.8**	8542.3	8544.0	8551.6
LR test (w.r.t. "next" simpler model)		13.047	3.413	1.781	1.588[1]
p-value		**0.000**	**0.039**	0.123	0.143
					3.22[2]
					0.044

(1) Compared to model 3. (2) Compared to model 4.