EVA ASCARZA, RAGHURAM IYENGAR, and MARTIN SCHLEICHER*

Facing the issue of increasing customer churn, many service firms have begun recommending pricing plans to their customers. One reason behind this type of retention campaign is that customers who subscribe to a plan suitable for them should be less likely to churn because they derive greater benefits from the service. In this article, the authors examine the effectiveness of such retention campaigns using a large-scale field experiment in which some customers are offered plan recommendations and some are not. They find that being proactive and encouraging customers to switch to cost-minimizing plans can, surprisingly, increase rather than decrease customer churn: whereas only 6% of customers in the control condition churned during the three months following the intervention, 10% did so in the treatment group. The authors propose two explanations for how the campaign increased churn, namely, (1) by lowering customers' inertia to switch plans and (2) by increasing the salience of past-usage patterns among potential churners. The data provide support for both explanations. By leveraging the richness of their field experiment, the authors assess the impact of targeted encouragement campaigns on customer behavior and firm revenues and derive recommendations for service firms.

Keywords: churn/retention, field experiment, pricing, tariff/plan choice, targeting

Online Supplement: http://dx.doi.org/10.1509/jmr.13.0483

The Perils of Proactive Churn Prevention Using Plan Recommendations: Evidence from a Field Experiment

Managing customer attrition, or churn, is a key challenge in customer relationship management (e.g., Blattberg, Kim, and Neslin 2008; Reinartz, Krafft, and Hoyer 2004; Reinartz and Kumar 2000). For firms, a high level of attrition, coupled with an increasing cost of new customer acquisition, can have severe long-term financial consequences. In industries such as telecom, cable television, credit cards, online services, insurance, and health clubs, the importance of customer retention cannot be overstated. For example, the annual churn rate for wireless telephone providers is approximately 15%–30% worldwide, which has been estimated to cost organizations up to $10 billion annually.[1] As a result, companies are increasingly managing customer retention proactively by identifying valuable customers who are likely to churn and taking appropriate action to retain them (e.g., by providing incentives to stay).

*Eva Ascarza is Assistant Professor of Marketing, Columbia Business School, Columbia University (e-mail: ascarza@columbia.edu). Raghuram Iyengar is Associate Professor of Marketing, the Wharton School, University of Pennsylvania (e-mail: riyengar@wharton.upenn.edu). Martin Schleicher is Assistant Professor of Marketing, IAE Business School, Universidad Austral (e-mail: mschleicher@iae.edu.ar). The authors thank several managers at the participating company for their efforts in making this research possible. The authors benefited from comments by Jean-Pierre Dubé; Pete Fader; Bruce Hardie; Kamel Jedidi; Don Lehmann; Oded Netzer; Nick Reinholtz; Christophe Van den Bulte; the audiences at Cheung Kong Graduate School of Business, Stanford Graduate School of Business, INSEAD, Emory, HEC, Wharton, Chicago Booth, and University of Colorado, Boulder; participants in the brown bag seminar series at the Columbia Business School; and the participants of the Marketing in Israel 2013, Theory + Practice in Marketing 2014, and Marketing Science 2014 conferences. Joel Huber served as guest editor and Duncan Simester served as associate editor for this article.

[1]For U.S. estimates, see http://www.statista.com/statistics/283511/average-monthly-churn-rate-top-wireless-carriers-us/. See http://www.jacada.com/images/WhitePapers/pdfs/45.100.0206-Customer-Retention-Strategies.pdf.

ISSN: 0022-2437 (print)
1547-7193 (electronic)

Journal of Marketing Research
Vol. LIII (February 2016), 46–60
DOI: 10.1509/jmr.13.0483

Pricing offers an important means for firms to manage the value that their customers receive. Most service firms pursue a menu-based pricing strategy and give customers the flexibility to choose a plan that matches their preferences. For example, a customer who wishes to visit a health club occasionally may choose to pay per visit, whereas another who expects to visit the health club frequently may choose a yearly plan. The success of a menu-based pricing strategy rests on the assumption that consumers will select an appropriate plan for themselves. However, prior research has cast doubt on this assumption. For example, in a systematic analysis of data on cell phone usage and expenditures between 2001 and 2003, Bar-Gill and Stone (2009) estimate that 42.5 million consumers overspent each year, with at least 20% of their annual cell phone spending as overspending. This phenomenon is not unique to telecom services. Customers choosing cable/satellite packages, health care plans, credit cards, insurance packages, and gym memberships exhibit similar behavior (e.g., DellaVigna and Malmendier 2006).[2] Such overspending is positively associated with customer churn (e.g., Ater and Landsman 2013; Iyengar, Ansari, and Gupta 2007; Lambrecht and Skiera 2006).

Given these insights, service firms have tried to improve retention by helping their customers choose suitable plans. For example, health insurance providers offer plan cost estimators that allow customers to compare, for each member's usage level, their out-of-pocket cost estimates for different plans. Some telecommunications companies are helping customers manage their usage by sending text messages as they reach their prepaid monthly quota of minutes.[3] Previous work has suggested that such modifications may be futile because people have intrinsic limits to navigating nonlinear pricing schedules (e.g., Gopalakrishnan, Iyengar, and Meyer 2015). More recently, service firms have begun contacting existing customers directly and offering plans more appropriate for their levels of usage.[4] For example, Verizon in the United States offers recommendations of plans based on consumers' estimated monthly usage.[5] Other providers, such as Vodafone and AT&T, offer similar services.

The purpose of this article is to assess the effectiveness of proactive recommendation programs for managing customer retention. We conducted a large-scale field experiment involving 65,000 customers of a wireless telecommunication service. Participants were randomly assigned to one of two groups: (1) a group that received an encouragement to switch to a service plan that was predicted to save them money on the basis of their past behavior (the treatment group), or (2) a group that did not receive such an encouragement (the control group). The field experiment was conducted over a six-month period, with the intervention (i.e., the campaign of encouraging customers to switch plans) applied to the participants in the treatment group at the end of the third month.

Our results indicate that being proactive and encouraging customers to switch to better plans can, surprisingly, increase rather than decrease customer churn. More specifically, we compare churn rates across conditions and find that, whereas only 6% of customers in the control group churned during the three months following the intervention, 10% did so in the treatment group.

We propose two explanations for the increase in churn. The first explanation relates to the change in customer inertia due to the firm's intervention. As prior research has documented, there is ample evidence of customer inertia in access-based services, with customers failing to switch plans even in the absence of any commitment contract (e.g., Ascarza, Lambrecht, and Vilcassim 2012; Goettler and Clay 2011; Iyengar, Ansari, and Gupta 2007). This customer inertia is likely to be reduced when customers are encouraged to switch to alternative plans. We propose that such lowered inertia may also prompt customers to explore competitive offerings, resulting in an increase in customer churn (Wieringa and Verhoef 2007). The second explanation is that the encouragement to switch plans also increases the salience of past-usage patterns (e.g., overspending on the current plan). We propose that such emphasized salience of past usage increases the (latent) propensity to churn, especially among customers who were already at risk of leaving the company (Blattberg, Kim, and Neslin 2008). Our data provide support for both accounts.

Drawing on these findings, we discuss conditions under which firms should run these types of encouragement campaigns. By leveraging the richness of our field experiment, we assess the potential impact on customer retention if the company were to run more targeted campaigns. More precisely, after segmenting customers on the basis of characteristics easily observed by the firm, we quantify the consequences of proactively contacting different groups of customers and identify the segments who should and should not be targeted for such campaigns.

Our research complements extant work on customer relationship management that has investigated drivers of customer retention (e.g., Bolton 1998; Lemon, White, and Winer 2002; Neslin et al. 2006; Nitzan and Libai 2011). We take one step further and investigate customers' reaction to the firm's retention efforts. This study also complements the empirical literature on price discrimination (e.g., Ascarza, Lambrecht, and Vilcassim 2012; Danaher 2002; Iyengar et al. 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' individual pricing recommendations on customer behavior.

We proceed as follows. First, we describe the research setting. Next, we quantify the impact of encouraging customers to switch plans on the customers' demand for access services and examine two possible explanations for the observed phenomenon of higher churn in the treatment group. We assess the robustness of our findings and discuss alternative explanations for the behavior we observe in the

[2]Companies that specialize in billing management have asserted that most people in the United States are overpaying by $1,000 or more for each of these expenses.

[3]See http://news.verizonwireless.com/news/2012/10/pr2012-10-02.html; http://www.t-mobile.com/Company/CompanyInfo.aspx?tp=Abt_Tab_ConsumerInfo.

[4]To the best of our knowledge, these practices have been limited to upgrades (i.e., companies encourage customers to switch to higher-priced plans), which is the type of recommendation we investigate in this work. We discuss the case of downgrades in the "Robustness of the Findings and Discussion of Additional Explanations" section.

[5]See http://www.verizonwireless.com/b2c/splash/shareEverythingCalculator.jsp?intcmp=vzw-vnt-se-shareeverything.

data. Then, we focus on the managerial implications of this research and quantify the impact of running more targeted campaigns on customer behavior and firm revenues. We conclude with a discussion of the implications of our research.

RESEARCH SETTING

To assess how customers respond to an encouragement to switch to cost-minimizing plans, we would ideally have a research setting that satisfies several conditions. First, we need a reasonable fraction of customers who, on the basis of their past behavior, would potentially benefit (i.e., save money) by switching plans. Second, the encouragement (vs. no encouragement) must be randomized across customers. Third, we must have data on consumers' demand for service both before and after the campaign and on their churn behavior after the campaign. Finally, we must observe or control for any marketing efforts deployed.

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

Field Experiment

Pricing plan. We selected customers enrolled in a specific type of plan for the study. This plan involves the purchase of a fixed amount of credit every month, which entitles the customer to an initial balance equal to the amount paid plus a percentage bonus (e.g., if a customer signed up for the $30 plan, his or her 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 in which the price per unit of consumption is zero until the customer uses up the initial credit, after which time a per-minute charge is applied. At the end of each month, customers are automatically enrolled in the same plan for the following month (i.e., the default is to continue with the same service contract). Customers have no contractual obligations with the provider (i.e., they can churn) and are allowed to change the base plan (upgrade or downgrade) at the end of every month. At the time of the study, the company offered six different plans, with the fixed amount of monthly credit ranging from $19 to $63.

Competitive environment. During the period of our study, there were three firms in the market. The firm we collaborated with is the second largest in terms of market share. All three firms offered similar sets of pricing plans (both prepaid and postpaid) with very little price fluctuation over time. In particular, there was no specific type of plan (e.g., all-inclusive prepaid plan) that was offered by only one or two of the phone providers.

Customers. Customers were eligible for inclusion in the experiment if they (1) had a certain level of monthly revenue (where we define "revenue" as "fees collected by the firm from a customer"; in this case, more than $47 per month for each of three months prior to the campaign), such that the plans offered in the campaign would be beneficial to them; (2) had been with the company for more than seven months (to ensure some level of self-awareness about individual usage); and (3) had not received any targeted promotional activities in the preceding three months. The last condition ensures that there were no marketing activities targeted to eligible customers for a few months before the campaign that might have changed their sensitivity to the reallocation campaign. There were also no other activities targeted toward these customers during the campaign itself. Of the customers who satisfied all three criteria, 64,147 were randomly selected to be included in the experiment. We randomly assigned 10,058 customers (15.7%) to the control group and the remaining 54,089 to the treatment group.

Intervention. Customers in the treatment condition were contacted (via phone) by the telecommunications provider and encouraged to upgrade to one of the two highest-priced fixed monthly credit plans: plans with a fixed monthly credit of $47 and $63, respectively. For ease of discussion, we call the plan with a fixed monthly credit of $47 ($63) the lower (higher) featured plan. For all customers, including those in the control condition, one or both of the featured plans would have been better than their current plan (see the summary statistics in the next section).[6] To incentivize customers to upgrade to the suggested plans, the company offered an additional credit of $15 for each of the following three months if they agreed to upgrade to one of the featured plans. Customers could accept the offer during that call (thus upgrading to the selected featured plan automatically) or by contacting the provider within one month. Those who did not accept the promotion by the end of that month could still switch to any plan (including the featured ones) in later months, but without receiving any monetary incentive. Customers in the control group were not contacted and did not receive any encouragement or incentive to switch plans. However, they were free to switch to any of the plans offered by the firm, including the featured plans, at any time. Figure 1 shows the study design as well as the sequence of decisions and incentives the customers faced after the intervention.

Data

For each customer, we observe (1) plan choice and (2) monthly revenue for three months before and after the campaign, as well as (3) churn behavior for the three months after the campaign. Plan choice is categorized by the amount of monthly fixed credit (i.e., prepaid credit included with the plan) associated with each plan. Monthly revenue can then be computed as fixed credit plus any additional credit purchased (i.e., credit added during the month). The total revenue per customer comprises fees paid for voice, text messages, and data use, which we do not measure separately. Note that once a customer churns (leaves the provider), we do not observe his or her behavior in following months.

[6]By "better," we mean that, on the basis of the customer's past usage, the featured plans would have been cost minimizing in all three months prior to the campaign.

Figure 1
DESIGN OF THE STUDY

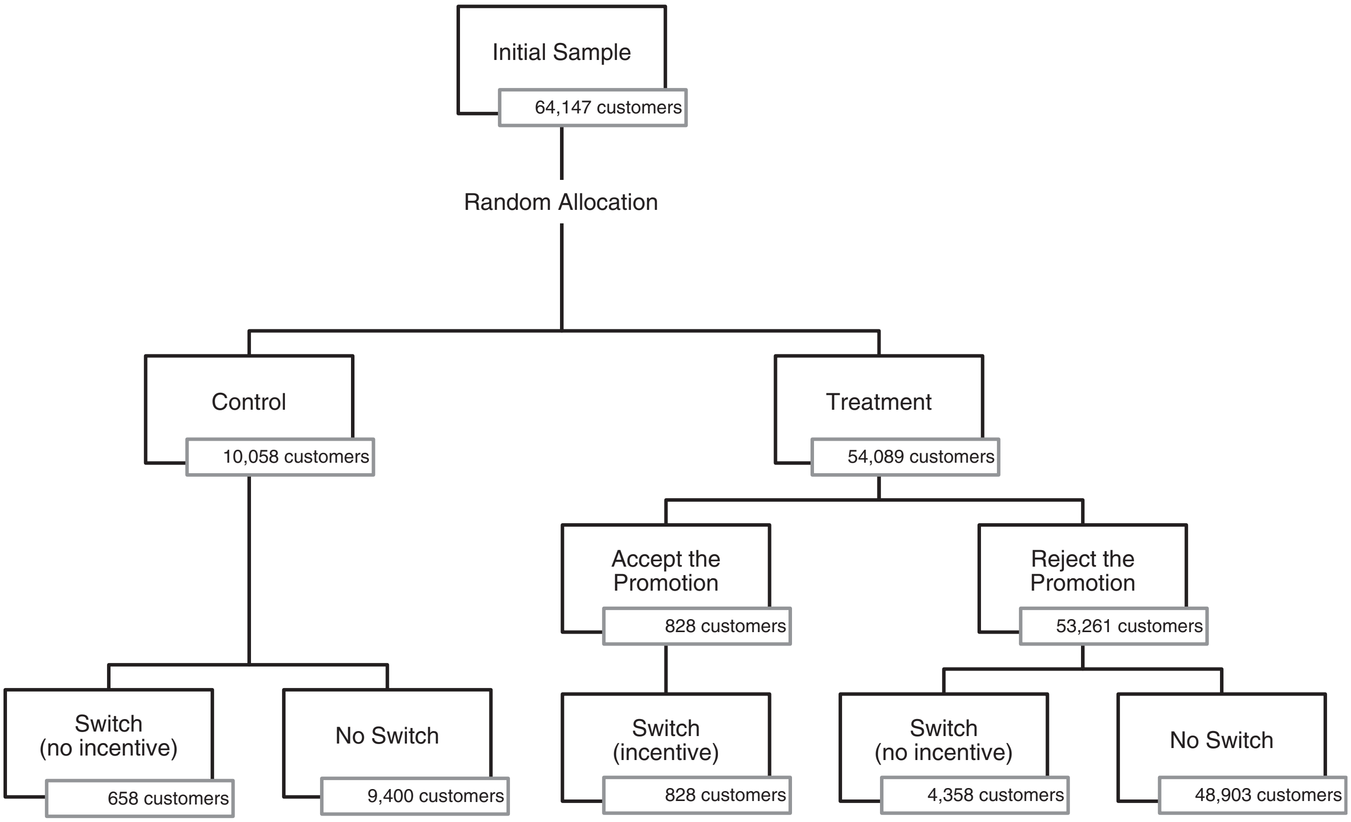

Precampaign descriptives. Table 1 shows the descriptive statistics for the two groups prior to the campaign. The distribution of monthly (fixed) fee is almost identical across groups. Regarding monthly revenue, customers also exhibit similar distributions across the two conditions: the average revenue (during the three months prior to the campaign) is \$112 for customers in the control group and \$114 for those in the treatment group. The percentiles for the average monthly revenue are almost identical in both conditions.

Table 1
PRECAMPAIGN DESCRIPTIVE STATISTICS FOR PARTICIPANTS IN CONTROL AND TREATMENT GROUPS

	Plan (\$)		*Average Revenue (\$)*	
	Control	*Treatment*	*Control*	*Treatment*
Mean	37.89	37.38	112.04	113.95
10th percentile	35.95	35.95	66.54	64.90
25th percentile	35.95	35.95	78.82	77.34
50th percentile	39.95	39.95	98.61	98.20
75th percentile	39.95	39.95	128.88	131.27
90th percentile	39.95	39.95	170.97	178.39

Notes: “Plan” denotes the fixed credit offered by each plan. “Average revenue” is the average monthly revenue collected from an individual customer during the three months before the campaign.

Postcampaign descriptives. Table 2 shows the extent of plan switching after the campaign.[7] In general, the campaign encouraged switching behavior: 6.54% of customers in the control group switched plans during the first three months after the campaign, as compared with 9.59% in the treatment group ($p < .001$). Regarding the plan that customers switched to, 35.5% of the switchers in the treatment condition chose one of the two featured plans, whereas only 22.5% did so in the control condition. Recall that the latter group was not exposed to any encouragement to switch plans but did so on their own initiative. Finally, the percentage of switchers who then switched again (“U-turners”) 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 between the two ($p > .1$).

IMPACT OF THE CAMPAIGN ON CUSTOMER BEHAVIOR

In this section, we analyze the impact of the treatment on customer demand for access services. We focus on two managerially relevant behaviors: namely, churn (i.e., whether customers switch to a different service provider) and usage (i.e., revenue). First, we quantify the aggregate impact of the

[7]Note that “switching” means subscribing to a different plan offered by the focal provider. It does not include switching to a plan offered by a competitor, which we denote by “churn.”

Table 2
POSTCAMPAIGN PLAN SWITCHING

	Control	*Treatment*
Number of customers	10,058	54,089
Number of switchers (% of customers)	658 (6.54%)	5,186 (9.59%)
. . .to a higher featured plan (% of switchers)	71 (10.79%)	1,041 (20.07%)
. . .to a lower featured plan (% of switchers)	77 (11.70%)	797 (15.37%)
. . .to a nonfeatured plan (% of switchers)	510 (77.51%)	3,348 (64.56%)
U-turners (% of switchers)	14 (2.13%)	135 (2.60%)

Notes: Featured plans are the two plans that were recommended in the experiment. These were the highest-fee plans offered by the company, at \$63 and \$47 per month, respectively. “U-turners” denotes customers who switched to one of the featured plans but later switched again to a different plan.

campaign on customer behavior. Then, we investigate whether the treatment effect is heterogeneous across customers. Finally, we synthesize the findings and propose two explanations for how the treatment affected behavior.

Aggregate Effect of the Treatment

As we conduct a randomized experiment, we measure the effect of the encouragement to switch plans by directly comparing customer behavior across the two experimental conditions. Given that all customers were active at the time of the campaign, we measure postcampaign churn behavior by determining the percentage of customers who left the service at any time during the three months after the campaign. Regarding revenue, we examine individual differences in revenue before and after the campaign—that is, we estimate the difference in differences. Specifically, for each customer, we calculate the difference between his or her average consumption during the three months before and the three months after the campaign. We use two metrics for revenue: net and conditional revenue. When calculating conditional revenue, we consider only periods in which the customer is active (i.e., when a customer churns, the revenue for subsequent months is not accounted for), whereas for net revenue, we consider all periods (i.e., when a customer churns, the revenue for subsequent months is set to zero).[8] We compare customers on these two behaviors across the two experimental conditions (see Table 3).

The encouragement clearly increased customer churn; whereas only 6.4% of customers in the control group left the company during the first three months after the campaign, 10.0% did so in the treatment group ($p < .001$). Average conditional revenue is not different across the two groups. Customers in both conditions show a negative trend in revenue after the campaign (the average individual difference is –4.4), but the magnitude of this trend is identical in the two groups ($p = .997$).[9] When we examine net revenue, it appears that customers in the treatment condition consume significantly less than those in the control condition; however, if we take into account the other two metrics (i.e., conditional revenue and churn rate), we find that this difference is driven mostly by the effect of churn. In summary, while the encouragement seemed not to affect individual usage, it had a devastating effect on customer churn.

Before we analyze customer heterogeneity in the response to the campaign, it would be worthwhile to discuss a general word of caution to apply when analyzing these types of experimental campaigns. Facing the data at hand, an analyst might be tempted to analyze the campaign by comparing only the customers who accept the promotion with those in the control group (naively thinking that customers who reject the offer should not be included in the analysis). If the analyst did so in this case, he or she would find that those who accepted the encouragement had lower churn rate (5.9% vs. 6.4%) and higher revenue change (.7 vs. –4.5) than those who were not exposed to the promotion (see Table 4). As a consequence, the analyst might infer, incorrectly, that the campaign was successful. But herein lies a word of caution: treatment-group customers selected themselves into the groups that did or did not accept the encouragement to change plans, and, due to this self-selection, they cannot be directly compared with customers in the control group.[10] This issue of self-selection relates to the broad literature on treatment compliance and the difference between the effect of the “intent to treat” (the randomized encouragement in our study) and the “treatment on treated” (the behavior of people who accepted the campaign; see, e.g., Barnard et al. 1998; Hirano et al. 2000). Moreover, a promotion could have a negative effect on behavior if, for example, customers perceive that the promotion has no value (Simonson, Carmon, and O’Curry 1994). In other words, simply offering a promotion might alter behavior even for customers who do not accept such an offer. As a consequence, to assess the net effect of a campaign, researchers must also track and analyze the behavior of customers who reject the offer. Given these two issues, in the remainder of this section, we compare the churn rate across the treatment and control groups (as presented in Table 3) to explore in further detail the causal impact of the encouragement on customer churn.

Customer Heterogeneity in Campaign Response

It is likely that customers differ in their responses to the campaign. Understanding such customer heterogeneity in

[8]The main difference between net and conditional revenue is that the latter better isolates the effect on customer usage independently from churn behavior (which we study separately). However, conditional revenue has missing values for customers who churn in the first or second month after the campaign; thus, it might be influenced by this selection effect.

[9]The decline in usage could reflect the selection criteria in the experiment—customers who were overspending in the past three months were chosen to participate, and their behavior may show regression to the mean.

[10]For example, in the case we study, it is very possible that customers who were more inclined to leave the company were less likely to accept the recommendation because they would not benefit from the change as much.

Table 3
AGGREGATE EFFECT OF THE ENCOURAGEMENT ON CHURN AND REVENUE

	Control	*Treatment*	*Difference*	*SE*	p-*Value*
Number of customers	10,058	54,089			
Percentage of customers churning	6.4%	10.0%	−3.6%	.3%	<.001
Change in net revenue	−$8.3	−$11.9	−$3.6	$.6	<.001
Change in conditional revenue	−$4.4	−$4.4	$.0	$.5	.997

Notes: Percentage of customers churning is measured for the three months after the experiment. The change in revenue is computed as the (individual) difference in monthly average consumption before and after the experiment.

campaign response is important from a theoretical perspective because it may help in uncovering potential underlying mechanisms for how the campaign influenced customer behavior. It is important managerially as well because it can help in finding better targets for future campaigns.

To determine the customer characteristics that would yield the most insights into how customers respond differently to the campaign, we leverage previous work on customer relationship management in services (e.g., Ascarza and Hardie 2013; Iyengar, Ansari, and Gupta 2007; Lambrecht and Skiera 2006; Lemmens and Croux 2006). Broadly, this stream of work has explored which customer characteristics affect their demand for services. For example, Lemmens and Croux (2006) and Ascarza and Hardie (2013) find that customers who have a downward sloping trend for usage are more likely to churn. Lemmens and Croux (2006) also find that monthly variability in usage is a good predictor for churn. Lambrecht and Skiera (2006), Iyengar, Ansari, and Gupta (2007), and Ater and Landsman (2013) find that consumers whose usage is above their plan allowance are more likely to churn. This finding is consistent with the retention strategy suggested by Joo, Jun, and Kim (2002). There, the authors argue that having subscribers on the plan most beneficial to them can increase their loyalty and satisfaction levels, a view also shared by Wong (2010, 2011).

On the basis of these findings, we choose to measure the following customer characteristics: (1) overage (i.e., consumption above the amount credited by the customer's plan), measured as average revenue minus the monthly fee for the plan; (2) individual usage variability, measured as the coefficient of variation in a customer's monthly revenue; and (3) trend, measured as the percentage change in revenue from the first to third month of the study. (Note that all metrics are obtained from customers' behavior prior to the campaign.) Table 5 contains the summary statistics for these metrics.

Table 4
AGGREGATE EFFECT OF THE ENCOURAGEMENT ON CHURN AND REVENUE, BY CAMPAIGN ACCEPTANCE

		Treatment	
	Control	*Accepted*	*Rejected*
Number of customers	10,058	828	53,261
Percentage of customers churning	6.4%	5.9%	10.1%
Change in net revenue	−$8.3	−$3.7	−$12.1
Change in conditional revenue	−$4.4	−$.7	−$4.5

Notes: Percentage of customers churning is measured for the three months after the experiment. The change in revenue is computed as the (individual) difference in monthly average consumption before and after the experiment.

We group customers on the basis of each of the characteristics and compare the effect of the treatment across segments. Table 6 shows that the effect of the treatment on postcampaign behavior is heterogeneous across customers. Comparing the effect of the campaign on churn behavior (i.e., percentage of customers in the treatment condition who churned minus percentage of customers in the control condition who churned), we find that the treatment increased churn to a lesser extent, and sometimes even decreased it, for customers with lower (vs. higher) overage, lower variability, and positive trend. For example, while the average difference in churn rate between customers in the treatment and control groups is 3.6 percentage points, those customers with a positive trend show a difference of only 2.4 percentage points. The effect of the treatment on customer revenue was also heterogeneous. The campaign increased revenue for customers with lower overage, higher variability, and positive trend. Nevertheless, the magnitude of the treatment effect on revenue is rather small, ranging from −$2.4 to $2.5 in monthly revenue across all groups. As a consequence, in the remainder of the article we focus on the effect of the treatment on customer churn.

These results suggest that the treatment differentially affected customers depending on their precampaign behavior. To assess the statistical significance of the reported effects, we estimate a probit model using churn as the dependent variable. In particular, we model the probability of a customer leaving the company as follows:

$$(1)\qquad \text{Prob}\left(\text{Churn}_i|\tilde{\beta}, T_i, X_i^c\right) = \text{Prob}\left(\beta_0 + \beta^T T_i + \beta^C X_i^c + \beta^{TC} T_i X_i^c + \varepsilon_i > 0\right),$$

where T_i is a dummy variable that takes a value of 1 if the customer received the treatment and 0 otherwise. The variable X_i^c contains all the usage-based characteristics (i.e., overage, variability, and trend) as well as a plan dummy to further control for idiosyncratic differences across customers who subscribed to different plans.[11] The vector $\tilde{\beta}$ contains the estimated parameters, including the constant

[11]The plan dummy indicates whether the customer's precampaign rate plan was the one whose fee is closest to those of the featured plans (i.e., fee = $39). We also estimated models including dummy variables for all four plans (all nonfeatured plans), but there were no differences in behavior across the three lowest-fee plans. Thus, we keep only a single plan dummy for model parsimony.

Table 5
SUMMARY STATISTICS OF PRECAMPAIGN CUSTOMER USAGE

Variable	*Description*	*Mean*	*SD*	*Min*	*Max*
Overage ($)	Average consumption above the amount allowed by the customer's plan	76.19	58.90	7.26	1,577.88
Variability	Individual coefficient of variation of usage	.21	.13	.00	1.54
Trend (%)	Individual percentage change in usage	−.03	.19	−1.39	1.09

Notes: All variables are computed on the basis of usage observations before the experiment was run.

(β_0), the main effects (β^T,β^C), and the interaction effects (β^{TC}). Finally, ε_i denotes the error term, which is assumed to be normally distributed with mean 0 and unit variance. We mean-center all continuous variables (i.e., overage, variability, and trend) so that the main effect of the treatment (β^T) in the model with full set of interactions represents the one corresponding to the "average" customer.

Table 7 shows the results for the churn model. In the first column, we replicate the significant positive effect of the treatment on the propensity to churn (see Table 3); we show its robustness by adding multiple control variables (second column). The results in the second column show that customers with higher levels of overage, higher variability, and negative trend, as well as those in plans with higher fixed fees, are more likely to churn. The relationships between past-usage variables and churn are consistent with previous empirical research (e.g., Ascarza and Hardie 2013; Iyengar, Ansari, and Gupta 2007; Lambrecht and Skiera 2006; Lemmens and Croux 2006). In the third column of Table 7, we show results of a test for heterogeneity in the treatment effect by adding interactions between the treatment dummy and all precampaign characteristics. (Note that in this model, the main effects β^C capture the relationship between the covariates and churn for customers in the control condition.)

The results in Table 7's third column show that there is substantial heterogeneity in the treatment effect. The interaction between variability and treatment is significant ($\beta_2^{TC} = .412$, $p < .05$), implying that the treatment was more influential—that is, in this case, given that the dependent variable is churn, more harmful—for customers with high levels of variability. Similarly, the interaction with trend is negative, though only significant at the 10% level ($\beta^{TC} = -.0839$, $p < .1$), suggesting that the treatment was more beneficial for the firm when applied to customers with higher trend. There was also heterogeneity in the campaign response depending on the plan the customers subscribed to; all else being equal, the treatment was more beneficial when applied to customers in the $39 plan (the plan whose allowance was closest to those of the featured plans). In conclusion, the effect of the treatment on customer churn is significantly positive ($\beta^T = .291$, $p < .001$) and is moderated by usage-related characteristics prior to the campaign.

Discussion of the Results and Further Analyses

To summarize, we have shown that the campaign encouraging customers to switch plans increased churn significantly: whereas only 6% of customers in the control condition churned during the three months after the campaign, 10% in the treatment group did so. Moreover, we find substantial heterogeneity in the response to the campaign: customers with high levels of variability in past usage and steeply negative trend are those for whom the treatment was most harmful to the firm. In this subsection, we provide two possible explanations for the increase in churn: lower customer inertia and increased salience of usage-related variables (e.g., overage, variability).

Table 6
POSTCAMPAIGN CHURN AND REVENUE BEHAVIOR, BY CUSTOMER CHARACTERISTICS

		Churn (%)			*Revenue Change ($)*		
	N	*Control*	*Treatment*	*Difference*	*Control*	*Treatment*	*Difference*
All customers	64,147	6.4	10.0	3.6	−4.4	−4.4	.0
By Overage							
≤$60	31,570	6.4	9.8	3.4	1.2	3.7	2.5
>$60	32,577	6.4	10.3	3.9	−9.9	−12.3	−2.4
By Variability							
≤.2	35,018	6.1	8.8	2.7	−1.1	−1.9	−.8
>.2	29,129	6.7	11.5	4.8	−8.3	−7.5	.8
By Trend							
Negative	35,776	6.3	10.9	4.6	−5.8	−6.6	−.9
Positive	28,371	6.6	9.0	2.4	−2.6	−1.7	1.0

Notes: Overage is defined as the (individual) average amount of consumption above amount allowed by the customer's plan. Variability is defined as the (individual) coefficient of variation of revenue. Trend is measured as the proportional change in consumption during the three months prior to the experiment. The cutoff points for overage and variability were selected as the closest integer value to the median split. Note that all measures for segmenting customers are based on precampaign behavior.

Table 7
HETEROGENEOUS EFFECT OF THE ENCOURAGEMENT ON CHURN

	Churn (Main Effect)	*Churn (Controls)*	*Churn (Heterogeneity)*
Treatment (β_T)	.243*** (.021)	.247*** (.021)	.291*** (.034)
Overage/1,000 (β_1^C)		.402*** (.112)	.282 (.357)
Variability(β_2^C)		.657*** (.052)	.290* (.152)
Trend (% increase) (β_3^C)		−.290*** (.036)	−.062 (.104)
$39 plan dummy ($\beta_4^C$)		.059*** (.014)	.132*** (.041)
Treatment × Overage/1,000(β_1^{TC})			.137 (.376)
Treatment × Variability(β_2^{TC})			.412** (.162)
Treatment × Trend(β_3^{TC})			−.084* (.043)
Treatment × $39 plan dummy($\beta_4^{TC}$)			−.257** (.111)
Constant (α)	−1.522*** (.019)	−1.565*** (.021)	−1.606*** (.032)
Observations	64,147	64,141	64,141
Pseudo R-squared	.4%	1.2%	1.3%

*$p < .1$.
**$p < .05$.
***$p < .01$.
Notes: Standard errors appear in parentheses. Results are from a probit model with churn behavior as dependent variable. For easier interpretation, overage, variability, and trend have been mean-centered. Variable overage is rescaled to avoid standard errors of zero. The variable "plan dummy" takes the value of 1 if the customer subscribes to the $39 plan and 0 otherwise.

Lowering customer inertia. There is extensive evidence that inertia prevents customers from switching plans and providers. For example, Ascarza, Lambrecht, and Vilcassim (2012) and Iyengar, Ansari, and Gupta (2007) show that customers fail to switch service plans even in the absence of any commitment contract. The lack of switching is frequently attributed to customer inertia (e.g., Handel 2013; Miravete and Palacios-Huerta 2014); for example, customers might not consider the possibility of switching in every single period. Research has also shown that customer inertia partly explains why customers stay with the same brand or provider. Rust, Lemon, and Zeithaml (2004) separate the effect of inertia and other switching drivers and show that the former has a significant impact on brand choice. Wieringa and Verhoef (2007) find that, for a large segment of customers (71% of the customers analyzed in their study), the decision not to switch providers does not depend on the attractiveness of other alternatives, the customer's own usage rate, or the type of contract. The authors label this segment the "inertial nonswitchers."

Thus, consistent with previous research, we would expect that some fraction of the lack of switching observed in our data could be attributed to customer inertia. Recall, however, that customers in the treatment condition were encouraged to switch plans. Such a treatment would likely lower their inertia by, for example, making customers realize that it was easy to switch. We postulate that once customers realize that they can easily switch plans within the company, they might also explore competitors' offerings, thus increasing their risk of churning. Note that we can only speculate on this idea because we did not observe whether people who left the focal company opened an account with any of the competitors.[12] We can, however, test for the underlying mechanism—that is, lower switching inertia due to the treatment—on the basis of the amount of switching across the focal provider's plans.

If we compared switching behavior between treatment and control conditions, we would find that, during the months following the campaign, customers in the treatment group switched plans significantly more than those in the control group (see Table 2). However, given that customers in the treatment group were offered $15 if they switched to a featured plan within a month, finding more switching in the treatment group does not necessarily imply lower inertia among customers in that group because some of the switching may have been due to the monetary incentive. To overcome this confound, we compare the level of post-campaign switching between customers in the treatment group who *rejected* the offer and those in the control group. The rationale for this comparison is as follows: If the campaign reduced customers' inertia to switch, then customers who rejected the campaign (and thus did not receive the $15 incentive) should still have a higher propensity to switch plans than people in the control group. This is indeed the case: while 6.54% of customers in the control condition switched plans during the three months after the campaign, 8.18% of customers who rejected the promotion switched plans during the same time period ($p < .001$).

Note that, as discussed in the previous section, this comparison likely suffers from self-selection bias given that we compare the control group only with customers who reject the campaign. In other words, if the decision of accepting/rejecting the campaign is related to the propensity/desire to switch plans, then the comparison should suffer from selection bias. We control for such selection using propensity score analysis. Briefly, we model the probability of accepting the promotion among all customers in the treatment group. We then use the model estimates to predict the individual propensity of accepting the promotion for all customers, including those in the control group. Next, we estimate a probit model with switching as a dependent variable and the treatment and the propensity score as covariates (for details, see the Appendix). The results show that customers in the

[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.

treatment group who rejected the campaign were significantly more likely to switch plans in the next three months than customers in the control group. In other words, even after we controlled for possible self-selection, the customers in the treatment group who did not accept the promotion were more likely to switch plans than those in the control group. We therefore conclude that the encouragement campaign lowered customers' inertia.

Enhanced sensitivity to past usage. Another explanation for the observed phenomenon is that the encouragement to switch plans enhances the sensitivity to past service consumption, which, for some customers, increases the propensity to churn. This explanation is suggested by the results presented in Table 7, in which we show that usage-related characteristics such as previous variability in monthly consumption and usage trend moderate the effect of the treatment on churn behavior. That is, highly volatile customers and those with downward trend exhibit a stronger effect of the encouragement (i.e., higher increase in the likelihood of churning). We further corroborate this intuition of enhanced sensitivity to usage characteristics by analyzing revenue as the dependent variable. If the treatment enhanced sensitivity to past consumption, then customers' future consumption should also be affected, and we should find similar moderating effects of past-usage variables when analyzing posttreatment usage (or revenue). We find that this is the case. In particular, we show that the treatment has strongest effect on future revenue for customers with lower levels of usage, higher variability, and higher trend as well as those on high-fixed-fee plans. (The Web Appendix provides details about the analysis and the full set of results.) Therefore, when we consider all the usage-related variables together, the results indicate that the encouragement enhanced customers' sensitivity to past consumption.

Why did highlighting past usage increase customer churn? Blattberg, Kim, and Neslin (2008, p. 626) posit that contacting customers who are at risk of churning might stimulate a "latent need" to churn. This idea is elaborated in a case study by Berson et al. (2000, pp. 282–95) in which the authors illustrate that proactively contacting customers can increase churn for some customers by drawing their attention to the amount that they are paying for their level of usage. In our context, we provide empirical evidence that treatment has a greater impact for customers who are predicted to be more likely to churn—that is, those with higher overage, higher variability, and lower trend (see the second and third columns in Table 7).[13] In other words, enhancing sensitivity to past usage among customers who are at high risk of churning increases their propensity to do so. In summary, we provide some evidence for why the treatment increased churn: it did so by altering the inertia that prevented customers from switching (and churning) and by enhancing customers' sensitivity to past usage, especially for those customers who were at higher risk of churning.

[13]Note that the interaction between overage and treatment is not significant in Table 7. Further investigation shows that this null effect is due to the assumption of linearity in our model specification. In the Web Appendix, we explore different ways to operationalize overage while allowing for nonlinearity in the interaction effect. We find that there is a significant interaction effect between overage and treatment.

ROBUSTNESS OF THE FINDINGS AND DISCUSSION OF ADDITIONAL EXPLANATIONS

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

Allowing for Temporal Dynamics and Unobserved Heterogeneity in Customer Churn

Our analysis thus far has measured churn as a single observation per customer. This approach might appear simplistic a priori because it does not control for time dynamics (e.g., seasonality) and does not allow for unobserved customer heterogeneity. Given the exogenous randomization in the field experiment, 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 reestimate the churn model (Table 7) as a discrete-time hazard model, using multiple observations per customer. We control for time effects (via monthly time dummies) and for unobserved customer heterogeneity (via random effect). The results of these tests are available in the Web Appendix. We replicate our main results: The encouragement to switch plans increased customer churn significantly, especially for customers with high variability and decreasing trend.

Modeling Churn and Plan Switching Simultaneously

The analyses in the previous two sections implicitly assume that churn behavior is not related to other actions that a customer might undertake, for example, switching plans. However, it is possible that these two decisions (i.e., whether to churn and whether to switch plans within the company) are related to each other. Furthermore, given the nature of our treatment—encouragement to switch to a featured plan—it is likely that consumers simultaneously considered switching plans and/or churning when they were presented with the encouragement offer. We check the robustness of our findings regarding the effect of the treatment on customer churn by jointly modeling churn and postcampaign switch.

We use a multinomial probit model with the same specification as the full model presented in Table 7, with the sole difference that customers are allowed to churn, switch plans, or do nothing. The results of this analysis corroborate the robustness of our main findings. (The details of this analysis as well as the full set of results appear in the Web Appendix.)

Selection on Observables

When analyzing changes in customer inertia, we employed a propensity score model to address self-selection issues arising from observed differences between customers who rejected the promotion and customers who were not exposed to the promotion. It is possible, however, that the observed customer characteristics might not completely capture all forms of self-selection. For example, it is plausible that customers who expected to switch to a specific plan because, say, they were trying to reduce their consumption were also more likely to reject the promotion because they expected their usage to decrease in the future. We determine 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 A1) and the model for switching plans (Equation A2). The rationale is as follows: If there are common unobservable factors affecting both decisions, then the residuals of our models—which include only observable characteristics—should be correlated. We find that the magnitude of correlation is very small, –.0121 (p = .005).

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

Churn Triggered by Any Firm Intervention

We have proposed two theoretical explanations for why the reallocation campaign might result in higher churn rates, namely, lowering customer inertia and enhancing sensitivity to past usage. Further analysis of the pre- and postcampaign behavior empirically supports these two mechanisms. One could argue, however, that another possible cause of the increase in churn is the mere intervention of the firm. That is, simply contacting customers might have prompted people to leave. This may be the case if customers do not welcome any firm-initiated contact.

To address this concern, we obtain information about another retention campaign run by the focal firm during approximately the same time of our study. The two campaigns are very different in nature. The goal of this other, so-called commitment campaign was to secure long-term commitment contracts among postpaid customers. All customers who were selected for the commitment campaign were enrolled in a different plan than those in the reallocation campaign. That is, there is no overlap between customers in the two campaigns. For the commitment campaign, the firm selected 150,038 postpaid customers, of whom 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 during the three months after the commitment campaign and compare customers who rejected the campaign promotion with those in the control condition. (We exclude customers who accepted the campaign because accepting it implied a commitment to stay with the same provider for 18 months.) Churn rate is slightly lower among customers in the control group (.41% vs. .54%). However, the two figures are not significantly different.[14] We therefore conclude that it seems unlikely that mere intervention from the firm drove the increase in churn we observe in the reallocation campaign.

Plan Upgrades and Downgrades

In the analysis of the reallocation plan described previously, 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 (to a plan with a higher fixed monthly fee) or a downgrade (to a plan with a lower fixed monthly fee). We determine the fraction of upgrades and downgrades in each treatment condition. For the control condition, from the set of 10,058 customers, 567 customers upgraded (5.64%), and only 91 downgraded (.90%). Similarly, from the set of 54,089 customers in the treatment condition, 4,713 upgraded (8.71%), and only 473 downgraded (.87%). There are two points to note. First, given that the campaign targeted customers who were overconsuming, it is not surprising that there were few downgrades. Thus, our label of a plan switch primarily captures upgrades. Second, there is no difference in the frequency of downgrade across the control and treatment conditions (p = .765), indicating that the campaign did not make people choose lower fixed monthly credit plans.

PROFITABILITY OF TARGETED CAMPAIGNS

In this section, we leverage the richness of the field experiment to quantify the profitability of targeted campaigns. For illustration, we select two easily observed usage variables as a basis for segmentation, namely, overage and usage variability. From the full sample of customers (N = 64,147), we select the customers who had both high observed variability (top 20%) and high overage (top 20%) before the intervention, resulting in a subsample of 2,565 customers. We proceed similarly to create additional subsamples of customers with high variability/low overage, low variability/high overage, and low variability/low overage. (Given the findings from previous sections, the effect of the experiment in the last subset of customers should be the most beneficial for the firm.).

Using the data from our field experiment, we then compare, for each subset of customers, actual postcampaign behavior between control and treatment conditions. Note that because we have a field experiment and customers are selected on the basis of their precampaign behavior, our results do not suffer from endogenous self-selection. Furthermore, because the control/treatment split was fully random, the proportion of treatment-group customers in each subsample very much resembles the distribution in the full population. Regarding postcampaign behavior, we measure customer churn and revenue as presented previously in our analysis. Briefly, churn is computed as the proportion of customers who churned in the three months after the campaign, and revenue is computed as the difference between each customer's average monthly consumption during the three months before and the three months after the campaign. Consistent with the main analysis, we use conditional revenue, that is, revenue conditional on a customer being active. Finally, we compute the total postcampaign consumption (which we call "net value"), which corresponds to the total revenue collected from each customer during the three months after the campaign. Table 8 shows the potential

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

Table 8

POSTCAMPAIGN BEHAVIOR FOR TARGETED CAMPAIGNS

	Control	Treatment	Difference	SE
A: High Variability/Low Overage				
Churn (%)	7.25	15.01	7.76	1.90
Revenue ($)	−.59	4.25	4.84	1.56
Net value ($)	198.22	192.33	−5.89	5.73
N = 2,564, of which 15% are in control group				
B: High Variability/High Overage				
Churn (%)	6.78	14.26	7.48	1.83
Revenue ($)	−33.27	−43.63	−10.37	4.90
Net value ($)	488.46	446.60	−41.86	16.53
N = 2,565, of which 16% are in control group				
C: Low Variability/Low Overage				
Churn (%)	8.79	8.13	−.66	1.55
Revenue ($)	5.50	7.32	1.82	1.23
Net value ($)	188.01	194.77	6.75	4.56
N = 2,565, of which 14% are in control group				
D: Low Variability/High Overage				
Churn (%)	7.48	9.26	1.78	1.63
Revenue ($)	−12.76	−17.22	−4.46	3.55
Net value ($)	521.78	510.69	−11.10	16.34
N = 2,565, of which 14% are in control group				

Notes: Churn percentage is based on the number of customers who left the company during the three months after the campaign. Revenue difference is computed as the difference between the average monthly revenue per person before and after the treatment. Net value is the total revenue per person during the three months after the experiment.

effects on these three metrics if the company were to run campaigns targeted at different segments of customers.

This analysis confirms that a campaign targeted toward customers with low levels of variability leads to much lower churn. For example, Table 8, Panel B, shows that targeting customers who are at high levels for both variables results in 14.3% churn in the treatment group, an increase of 7.5 percentage points over the control group. In contrast, targeting consumers with low variability and high overage leads to 9.3% churn in the treatment group, only 1.8 percentage points higher than the churn observed in the control group for that customer segment (Table 8, Panel D). When we average the churn rates across conditions when the variability is low or high, we find that churn is approximately 8.6% in the low variability condition and 13.4% in the high variability condition.

As we expected, the campaign targeting customers who are low on both variables leads to the lowest level of churn in the treatment group (8.1%), which is also lower than the rate in the corresponding control group (8.8%). In addition to decreasing churn, such targeting is also beneficial to the firm in that it increases revenue (Table 8, Panel C). Finally, such a campaign leads to an increase of approximately $6.8 in net value (i.e., monthly revenue per customer). In summary, the targeted campaign in Table 8, Panel C, is significantly better than the mass campaign on all metrics of resulting customer behavior (see Table 3).

GENERAL DISCUSSION AND CONCLUSIONS

Given the impact of customer churn on long-term profitability, firms frequently employ campaigns aimed at increasing retention among their current customers. One way to retain customers is to ensure that they derive the appropriate level of benefits from their spending. For service companies, pricing offers an important means for managing the value that customers receive from the service; by offering a menu-based pricing strategy, firms give customers the flexibility to choose the plan best suited to their preferences/needs. To help manage their churn, firms are increasingly beginning to recommend more appropriate plans to their customers. We examine the effectiveness of such recommendation programs for managing customer churn by conducting a large-scale field experiment in which some customers receive an encouragement to switch to cost-minimizing plans, while others do not. We find that encouraging customers to switch to plans more beneficial to them can, surprisingly, be harmful for a firm in that it may significantly increase—rather than decrease—customer churn. We propose two explanations for this phenomenon, namely, lowered customer inertia and enhanced sensitivity to past consumption. Our analysis provides empirical evidence for both these mechanisms.

The managerial implications of our research are noteworthy. First, we offer guidance to service providers (e.g., telecommunications, utilities, financial services) concerned about the mismatch of customers to service plans and its consequences in future churn. We show how selecting the right customers to target would have a much higher impact on profitability than encouraging every "suboptimal" customer to switch plans. In particular, we demonstrate how firms could target customers on the basis of easily observed characteristics, such as level of overage and individual usage variability, to improve the outcome of reallocation campaigns.

It is important to note that in this research, we evaluated an "upgrade" campaign. That is, the campaign encouraged consumers who were currently overconsuming to upgrade to a more appropriate plan for them. In contrast, there are many contexts in which consumers may purchase plans that have a higher coverage than what they optimally need (see, e.g., 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 both of our proposed mechanisms would still apply. Nevertheless, we acknowledge that there could also be differences in behavior due to the different motivations that affect upgrade versus downgrade decisions.

Second, we call for caution when analyzing experimental data. As marketers increasingly use "A/B" testing to assess the impact of their marketing actions, we encounter cases in which the customers in the treatment condition who did not undertake a certain action (e.g., those who rejected the recommendation in the current context) are naively ignored. We have illustrated how erroneous this approach might be.

Our research complements the existing literature on customer churn/retention (e.g., Bolton 1998; Lemon, White, and Winer 2002; Neslin et al. 2006; Nitzan and Libai 2011) by investigating not only the drivers of customer churn but also how customers respond to the firm's retention efforts. In particular, we offer a word of caution to marketers concerned about customer churn: Retention campaigns might have the unintended consequence of increasing churn. In the case investigated here, we find that customers with decreasing trends and higher levels of variability in past consumption are those among whom the campaign most increased churn.

This study also complements the literature on price discrimination (e.g., Ascarza, Lambrecht, and Vilcassim 2012; Danaher 2002; Iyengar et al. 2011; Narayanan, Chintagunta, and Miravete 2007) that has evaluated the impact on firm profitability of offering differing pricing contracts. Our results show that firm profitability is also affected (not necessarily improved) when firms proactively encourage customers to switch to plans more beneficial to them. Our study also adds to the broader literature on product recommendations by offering a new setting in which customers may exhibit reactance (Bodapati 2008; Fitzsimons and Lehmann 2004; Häubl and Murray 2006). Our results suggest that it may be best for retailers to refrain from giving recommendations after consumers make a purchase because it may induce consumers to overly focus on their past behavior. A caveat is in order: 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 setting in which the status quo may be to not 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 (for an excellent summary, see Johnson et al. 2012). One of the factors determining the success of any choice architecture is the degree of freedom that consumers perceive that they have when making their decisions (Sunstein and Thaler 2003). In line with this view, a libertarian (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 health care and financial decisions (Kling et al. 2011; Thaler and Benartzi 2004). 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, making recommendations can have severe negative consequences. That nudges have a heterogeneous impact is consistent with previous work. For example, Costa and Kahn (2013) find that although informing households about relative energy use led to an overall average 2% decrease in usage, liberals reduced their energy consumption while conservatives increased theirs.

Our analysis can be extended in several ways. First, our six-month window enables us to clearly pin down the pre- and postexperiment behavior. However, a longer time frame would allow the researcher to disentangle other dynamics in customer churn and revenue and quantify the impact of this type of promotion on customer lifetime value. Moreover, we do not observe whether the customers who churn from the focal provider switch to a competing provider. To offer a full support of the notion that these customers switch to a competing provider, we would need individual-level multifirm data, which, unfortunately, is difficult to obtain. Furthermore, we acknowledge that other factors (e.g., lack of trust, reactance) might also contribute to the increase in churn we observe. Attitudinal measures (e.g., satisfaction with the service, company trust) would be a useful complement to our data to provide more support for the proposed mechanisms. We hope that future studies will address these issues.

APPENDIX: CONTROLLING FOR SELF-SELECTION IN THE SWITCHING MODEL

We use propensity score analysis (Rosenbaum and Rubin 1983) to control for possible sources of self-selection. Propensity score has been widely used in statistics (e.g., Rosenbaum and Rubin 1983, 1984, 1985), biostatistics/medicine (e.g., Austin 2008; Brookhart et al. 2006; D'Agostino 1998; Rubin 1997), 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 nonbalanced 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 customer relationship management programs (Mithas, Krishnan, and Fornell 2005), among others.

We proceed in three stages: First, we select all customers in the treatment condition and model the probability (or propensity) of their accepting the recommendation to switch plans given other observed variables. Second, drawing on the

Table A1

PARAMETER ESTIMATES OF THE PROPENSITY MODEL FOR CUSTOMERS WHO ACCEPTED THE PROMOTION

	Probability of Accepting
Overage	.002*** (.001)
Overage^2	-3.34×10^{-6}*** (1.22×10^{-6})
Variability	−.631*** (.225)
Trend	.227*** (.077)
Overage × Variability	.002* (.001)
$39 plan dummy	.186*** (.028)
Constant	−2.313*** (.0487)
Log-likelihood	−4,228.762
Observations	54,083

*$p < .1$.
***$p < .01$.
Notes: Standard errors appear in parentheses. Results are from a probit model with "accepting the recommendation" as dependent variable. The variables overage, variability, and trend have been mean-centered. Only customers from the treatment group are included in this analysis.

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 estimate differences in postcampaign switching behavior between customers in the control group and those who rejected the promotion, while controlling for the estimated propensity to accept/reject the campaign.

Stage 1: Propensity Model

For a customer i, let Accept_i be an indicator variable that takes a value of 1 if the customer accepts the promotion and 0 otherwise. Then, the probability (i.e., propensity) of the customer accepting the promotion is modeled as follows:

$$\text{(A1)} \qquad \text{Prob}\,(\text{Accept}_i|\delta, \gamma, Z_i) = \text{Prob}\,(\delta + \gamma Z_i + \epsilon_i > 0),$$

where δ is a constant, Z_i contains observed customer-specific characteristics, γ 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 a customer accepting the promotion using a probit model. Regarding Z_i, we consider all usage- and plan-related characteristics discussed in the analysis section. We estimate several model specifications including transformations and interactions for the observed variables and select the propensity score model on the basis of fit.[15]

Table A1 reports the estimates for the best-fitting propensity model. The estimates from the propensity model have face validity. We find that higher overage makes it more likely that a consumer will accept the encouragement but with diminishing returns (i.e., the quadratic term is negative). Customers who are on higher-priced plans and who show a positive trend in their consumption are more likely to accept the encouragement. Conversely, the higher the variability in their past usage, the more likely they are to reject it.

Stage 2: Computing Propensity Scores

We consider all customers in our sample and use the estimates of our model to compute the propensity score (i.e., the probability that each customer would have accepted the promotion given his or her observed characteristics). Note that this step now includes the customers in the control group, who were not considered in our estimation of the propensity model because the promotion was never offered to them. Nevertheless, we can compute their likelihood of accepting the promotion if it had been offered, given the propensity model and their observed characteristics.

[15]As such, the estimates for the individual propensity scores rely on the functional form of the selected model. We report results for all model specifications in the Web Appendix.

Stage 3: Effect of Encouragement (Excluding Those Customers Who Accepted the Promotion)

Finally, to compare the switching propensity between those who reject the promotion and those in the control group, we use regression adjustment with the propensity score (e.g., Austin and Mamdani 2006). In particular, we run a probit model with switching as a dependent variable and with treatment and the individual propensity score as independent variables. In particular, we model the probability of switching as follows:

$$\text{(A2)} \qquad \text{Prob}\left(\text{Switch}_i|\tilde{\varphi}, T_{i,}, \text{Score}_i\right) = \text{Prob}\left(\varphi^0 + \varphi^T T^i + \varphi^S \text{Score}_i + \nu_i > 0\right),$$

where φ^0 is a constant, T^i takes the value of 1 if the customer is in the treatment group and 0 otherwise, Score_i is the individual's propensity to accept the offer (i.e., $\text{Score}_i = \delta + \gamma Z_i$ from Equation A1), and the parameters φ^T and φ^S capture the sensitivity to the treatment and the score, respectively. Finally, ν_i denotes the error term, which is normally distributed with mean 0 and unit variance. Table A2 shows the results for such regression.

Results confirm that customers who reject the encouragement are more likely to switch plans than those customers who did not receive any encouragement (first column in Table A2). This effect is not statistically different when we control for selection on observables (second column in Table A2).

Table A2

EFFECT OF THE ENCOURAGEMENT ON PLAN SWITCHING

	Probability of Switching	*Probability of Switching (Controlling for Selection)*
Treatment	.118*** (.021)	.104*** (.021)
Propensity score		−37.030*** (1.486)
Constant	−1.511*** (.019)	−.958*** (.029)
Observations	63,319	63,313
Pseudo R-squared	.09%	2.00%

***$p < .01$.
Notes: Standard errors appear in parentheses. Results are from a probit model with "probability of switching" as dependent variable. We computed "propensity score" on the basis of the parameter estimates from Table A1. We excluded customers who accepted the promotion from this analysis.

REFERENCES

Ascarza, Eva and Bruce Hardie (2013), "A Joint Model of Usage and Churn in Contractual Settings," *Marketing Science*, 32 (4), 570–90.

———, 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 (December), 882–99.

Ater, Itai and Vardit Landsman (2013), "Do Customers Learn from Experience? Evidence from Retail Banking," *Management Science*, 59 (9), 2019–35.

Austin, Peter C. (2008), "A Critical Appraisal of Propensity-Score Matching in the Medical Literature Between 1996 and 2003," *Statistics in Medicine*, 27 (12), 2037–49.

——— 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), 2084–106.

Bar-Gill, Oren and Rebecca Stone (2009), "Mobile Misperceptions," *Harvard Journal of Law & 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 & Research*, 27 (2), 285–317.

Berson, Alex, Stephen Smith, and Kurt Thearling (2000), *Building Data Mining Applications for CRM*. New York: McGraw-Hill.

Blattberg, Robert C., Byung-Do Kim, and Scott A. Neslin (2008), *Database Marketing: Analyzing and Managing Customers*. New York: Springer.

Bodapati, Anand V. (2008), "Recommendation Systems with Purchase Data," *Journal of Marketing Research*, 45 (February), 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.

Bronnenberg, Bart J., Jean-Pierre Dubé, and Carl F. Mela (2010), "Do Digital Video Recorders Influence Sales?" *Journal of Marketing Research*, 47 (December), 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), 1149–56.

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 (3), 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), 2265–81.

Danaher, Peter J. (2002), "Optimal Pricing of New Subscription Services: Analysis of a Market Experiment," *Marketing Science*, 21 (2), 119–38.

Dehejia, Rajeev H. and Sadek Wahba (2002), "Propensity Score-Matching Methods for Non-Experimental Causal Studies," *Review of Economics and Statistics*, 84 (1), 151–61.

DellaVigna, Stefano and Ulrike Malmendier (2006), "Paying Not to Go to the Gym," *American Economic Review*, 96 (3), 694–719.

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 (July), 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 (August), 633–52.

Gopalakrishnan, Arun, Raghuram Iyengar, and Robert J. Meyer (2015), "Consumer Dynamic Usage Allocation and Learning Under Multi-Part Tariffs," *Marketing Science*, 34(1), 116–33.

Handel, Benjamin R. (2013), "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 103 (7), 2643–82.

Häubl, Gerald and Kyle B. Murray (2006), "Double Agents: Assessing the Role of Electronic Product Recommendation Systems," *MIT Sloan Management Review*, 47 (3), 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), 1161–89.

———, ———, 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 (April), 131–43.

Iyengar, Raghuram, Asim Ansari, and Sunil Gupta (2007), "A Model of Consumer Learning for Service Quality and Usage," *Journal of Marketing Research*, 44 (November), 529–44.

———, 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), 820–36.

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," *PLOS One*, 8 (12), 1–6.

———, Suzanne B. Shu, Benedict G.C. Dellaert, Craig Fox, Daniel G. Goldstein, Gerald Häubl, et al. (2012), "Beyond Nudges: Tools of a Choice Architecture," *Marketing Letters*, 23 (2), 487–504.

Joo, Young-Hyuck, Jong-Kun Jun, and Byung-Do Kim (2002), "Encouraging Customers to Pay Less for Mobile Telecommunication Services," *Journal of Database Marketing*, 9 (4), 350–60.

Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel (2011), "Comparison Friction: Experimental Evidence from Medicare Drug Plans," *Quarterly Journal of Economics*, 127 (1), 199–235.

Lambrecht, Anja and Bernd Skiera (2006), "Paying Too Much and Being Happy About It: Existence, Causes, and Consequences of Tariff-Choice Biases," *Journal of Marketing Research*, 43 (May), 212–23.

Lemmens, Aurélie and Christophe Croux (2006), "Bagging and Boosting Classification Trees to Predict Churn," *Journal of Marketing Research*, 43 (May), 276–86.

Lemon, Katherine N., Tiffany Barnett White, and Russell S. Winer (2002), "Dynamic Customer Relationship Management: Incorporating Future Considerations into the Service Retention Decision," *Journal of Marketing*, 66 (January), 1–14.

Miravete, Eugenio J. and Ignacio Palacios-Huerta (2014), "Consumer Inertia, Choice Dependence, and Learning from Experience in a Repeated Decision Problem," *Review of Economics and Statistics*, 96 (3), 524–37.

Mithas, Sunil, Mayuram S. Krishnan, and Claes Fornell (2005), "Why Do Customer Relationship Management Applications Affect Customer Satisfaction?" *Journal of Marketing*, 69 (October), 201–09.

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 (May), 204–11.
Nitzan, Irit and Barak Libai (2011), "Social Effects on Customer Retention," *Journal of Marketing*, 75 (November), 24–38.
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.
——— and V. Kumar (2000), "On the Profitability of Long-life Customers in a Noncontractual Setting: Am Empirical Investigation and Implications for Marketing," *Journal of Marketing*, 64 (October), 17–35.
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.
——— and ——— (1984), "Reducing Bias in Observational Studies Using Subclassification on the Propensity Score," *Journal of the American Statistical Association*, 79 (387), 516–24.
——— and ——— (1985), "Constructing a Control Group Using Multivariate Matched Sampling Methods That Incorporate the Propensity Score," *American Statistician*, 39 (1), 33–38.
Rubin, Donald B. (1997), "Estimating Causal Effects from Large Data Sets Using Propensity Scores," *Annals of Internal Medicine*, 127 (8), 757–63.
Rust, Roland T., Katherine N. Lemon, and Valerie A. Zeithaml (2004), "Return on Marketing: Using Customer Equity to Focus Marketing Strategy," *Journal of Marketing*, 68 (January), 109–27.
Simonson, Itamar, Ziv Carmon, and Suzanne O'Curry (1994), "Experimental Evidence on the Negative Effect of Product Features and Sales Promotions on Brand Choice," *Marketing Science*, 13 (1), 23–40.
Sunstein, Cass R. and Richard H. Thaler (2003), "Libertarian Paternalism Is Not an Oxymoron," *University of Chicago Law Review*, 70 (4), 1159–202.
Thaler, Richard H. and Shlomo Benartzi (2004), "Save More Tomorrow: Using Behavioral Economics to Increase Employee Savings," *Journal of Political Economy*, 112 (1), S164–87.
——— and Cass R. Sunstein (2008), *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New Haven, CT: Yale University Press.
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), 174–86.
Wong, Ken Kwong-Kay (2010), "Should Wireless Carriers Protect Residential Voice Subscribers from High Overage and Underage Charges? Insights from the Canadian Telecommunications Market," *Telecommunications Policy*, 34 (8), 461–81.
——— (2011), "Getting What You Paid For: Fighting Wireless Customer Churn with Rate Plan Optimization," *Database Marketing & Customer Strategy Management*, 18 (2), 73–82.