

The Impact of Tariff Structure on Customer Retention, Usage, and Profitability of Access Services

Raghuram Iyengar

The Wharton School of the University of Pennsylvania, Philadelphia, Pennsylvania 19104,
riyengar@wharton.upenn.edu

Kamel Jedidi

Columbia Business School, Columbia University, New York, New York 10027, kj7@columbia.edu

Skander Essegaiier

College of Administrative Sciences and Economics, Koç University, 34450 Istanbul, Turkey,
sessegaiier@ku.edu.tr

Peter J. Danaher

Department of Marketing, Monash University, Caulfield East, Victoria 3145, Australia,
peter.danaher@monash.edu

Past research in marketing and psychology suggests that pricing structure may influence consumers' perception of value. In the context of two commonly used pricing schemes, pay-per-use and two-part tariff, we evaluate the impact of pricing structure on consumer preferences for access services. To this end, we develop a utility-based model of consumer retention and usage of a new service. A notable feature of the model is its ability to capture the pricing structure effect and measure its impact on consumer retention, usage, and pricing policy.

Using data from a pricing field experiment for a new telecommunication service, we find that consumers derive lower utility from consumption under a two-part tariff than pay-per-use pricing, resulting in lower retention of customers and lower usage of the service. Specifically, our demand analysis shows that a two-part tariff structure leads to an average decline of 10.5% in the annual retention rate and an average decrease of 38.7% in yearly usage relative to pay-per-use pricing after controlling for income effects. Despite the higher customer churn and lower usage, we find that the two-part tariff is still the profit-maximizing pricing structure. However, our results show that if firms ignore the pricing structure (or access fee) effect, then they would overcharge customers for the access fee and undercharge them for the per-minute price. Translated in terms of profitability, the failure to account for the access fee effect leads to a reduction of 11% in firm profit.

Key words: nonlinear pricing; tariff structure; discrete/continuous choice models

History: Received: March 7, 2010; accepted: March 23, 2011; Duncan Simester, guest editor-in-chief, and then Preyas Desai served as the editor-in-chief and Scott Neslin served as associate editor for this article.

Published online in *Articles in Advance* July 15, 2011.

1. Introduction

The pricing of subscription services has become a key issue for many firms, particularly in the technology sector (Danaher 2002, Essegaiier et al. 2002, Lambrecht and Skiera 2006, Narayanan et al. 2007, Sundararajan 2004). For example, when launching their new cell phone service in the United States, Virgin Mobile considered three pricing options: (i) matching the industry standard of a two-part tariff (TPT), where consumers are charged a monthly access fee and a marginal price per use with price levels set equal to competitors; (ii) the same as in (i) but with lower price levels than competitors; and (iii) introducing a pay-per-use (PPU) scheme with no monthly fee and no contract (McGovern 2003). Such a decision faced by Virgin Mobile—TPT versus PPU—is fundamental

to many subscription services as they grapple with a need to understand how consumers may react to alternative pricing structures. Other examples include the pricing of online music by iTunes (Mark and White 2007) and Amazon's Prime service that limits overnight shipping charges to just \$3.99 on all items to those who pay an annual fee.

Past work on nonlinear pricing suggests that a TPT pricing structure will allow firms to generate higher profits than a PPU structure (Oi 1971). However, this result relies on the assumption that, given the same level of marginal price, the demand curve is the same under both PPU and TPT pricing structures. In contrast, research in psychology and marketing suggests that myriad pricing-related factors may affect consumer demand above and beyond the traditional

marginal price perspective. For example, the price format can influence consumers' perception of value and their consumption (e.g., Prelec and Loewenstein 1998). Similarly, Heath and Fennema (1996) find that subjects assigned to a TPT condition—being a combination of an entry fee to a bar and a per-game charge—expect to play less pool games than those assigned to the PPU condition (no entry fee and only a per-game charge) even after controlling for differences in income (i.e., the effect of entry fee) across the two conditions. Other examples include the systematic effects of payment schedule—monthly or annual payment—on service usage and retention (Soman and Gourville 2001) and price-ending format on consumers' purchase decisions (Anderson and Simester 2003). Overall, this research stream suggests that “pricing can transform, as well as capture, the utility of an offer” (Bertini and Wathieu 2008, p. 236).

The purpose of this paper is to assess the impact of pricing structure on consumer demand for access services and firm profitability. We investigate this issue in the context of PPU and TPT, two common pricing structures for access services (e.g., Danaher 2002, Essegai et al. 2002). We develop a utility-based model of consumer usage and retention of a subscription service and allow the model parameters to vary with the type of pricing structure faced by the customer. We test the model on data from a field experiment for a cellular phone service where participants are randomly assigned to a TPT or PPU plan. Access and usage prices are *varied* at intervals over a 13-month period, with subjects able to relinquish the service at any time. A unique feature of the field experiment is the joint manipulation of pricing structure and price levels. We capitalize on this variation to assess the impact of pricing structure on consumer usage and retention, estimate demand parameters for each pricing structure, and derive optimal PPU and TPT tariffs.

Our results indicate that consumers derive lower utility from using the service when they are under a TPT plan compared with a PPU plan, resulting in both lower retention of customers and lower usage of the service. This negative impact of TPT pricing (which we label the “access fee effect”) is true even after controlling for income effects, heterogeneity across customers, and observable and unobservable time-varying factors. Our results also show an average decline of 10.5% in the annual retention rate and a 38.7% decrease in the yearly usage relative to PPU pricing after controlling for income effects (10.5% of this drop in usage is due to consumer churn and 28.2% is due to reduction in usage conditional on retention). Despite the higher customer churn and lower usage, we find that TPT is still the profit-maximizing pricing structure for the firm. However,

we find that the firm would overcharge customers for the access fee and undercharge them for the per-minute price if it ignores the access fee effect. In terms of profitability, the failure to account for this effect leads to an 11% reduction in firm profit.

Previous research on multipart pricing of services has focused on consumers' preferences for tariffs and usage behavior. Several studies have shown that consumers have a “flat-fee bias,” namely, a preference for unlimited usage in return for a fixed monthly fee (e.g., Kridel et al. 1993, Lambrecht and Skiera 2006). By contrast, Miravete (2003) reports evidence for consumers exhibiting a PPU bias. In addition, a number of empirical studies have used either observational data (e.g., Iyengar et al. 2007) or natural experiments (e.g., Ascarza et al. 2011, Narayanan et al. 2007) to analyze tariff choice and usage behavior. Recent work has also considered three-part tariffs, where the access fee permits a certain usage allowance beyond which a usage price is charged (e.g., Iyengar et al. 2007, 2008; Lambrecht et al. 2007). In this paper, we exploit the richness of the field experiment to examine how customers differentially respond to PPU versus TPT pricing once they are on one of these plans.¹ In addition, we evaluate the impact of pricing structure on customer retention and usage and firm profitability. To the best of our knowledge, none of this prior research has investigated these issues together.

The rest of this paper is organized as follows. In §2, we describe the field experiment. In §3, we discuss our proposed model. Section 4 contains the results of our estimation. In §5, we explain our results and describe how we rule out alternative explanations. In §6, we determine the impact of pricing structure on retention and consumption. Section 7 contains the results on pricing policy optimization. Section 8 concludes the paper.

2. Field Experiment

The data come from a field experiment of a product trial for a new subscription-based telecommunication service offered by a firm in an OECD (Organisation for Economic Co-operation and Development) country and have been previously used in Danaher (2002).² The firm has a monopoly over the provision of fixed line services but faces competition in the wireless services. Being the first of its kind in this country, the

¹ In this paper, we do not study consumer choice between a PPU or a TPT plan. Instead, we examine consumption behavior given a customer subscription to a PPU or TPT plan.

² Danaher (2002) focuses on deriving the optimal pricing of the new subscription service regardless of the pricing structure effect. In contrast, this paper focuses on assessing the *moderating* impact of the pricing structure on consumer preferences, retention, and usage, as well as its implications for optimal pricing.

new service integrates features that combine the benefits of a fixed phone line and a cellular phone. We use these data to understand how customers differentially respond to PPU versus TPT pricing and to assess the impact of pricing structure on customer retention, usage, and optimal pricing.

Trialists were residential, fixed line customers who agreed to try this new “combined” fixed–cellular line service. At the time of the service trial, the overwhelming majority of trialists did not have a cellular phone. The duration of the experiment consisted of 13 monthly periods (starting from October). Each trialist was randomly assigned to either a PPU condition or a TPT condition. Once assigned to one of the tariff conditions, a trialist remained in that condition for the duration of the experiment. Thus the data do not suffer from self-selection biases.

Whereas the tariff structure (the experimental condition) assigned to a trialist did not change throughout the experiment, the price levels within a tariff condition were modified by the firm. Before a price change, a customer was informed about the new price levels he or she would be facing for the coming period: a consumer in the PPU condition was informed about the next period per-minute price, and a consumer in the TPT condition was informed about both the per-minute price and the fixed fee for the following period. Note that this method of disseminating price-related information to customers precludes them from forming any expectations of when a price change may occur and by how much the price would change. The experimental price variability allows us to estimate a customer’s price sensitivity and determine whether it differs by tariff type.

Trialists were able to drop out of the service at the end of any of these 13 periods. Any time a customer decides to drop out, he or she would inform the company. As such, the company is able to observe, over the duration of the trial period, if and when a trialist from the pool decides to drop the new service. At the end of each period, trialists who still subscribe to the service face a keep/drop decision: stay with the new service or drop it.

Trialists were not aware of other types of pricing schemes being offered to other customers for the same service. Also, because the experiment was conducted during a prelaunch trial phase for the product, there was no specific advertising or promotional activity. Thus, across-customer comparisons are very unlikely. Table 1 contains a description of the field experiment and the price changes that took place over the 13-month period.³

³ We removed two outliers from the data because they had abnormally excessive usage during the experiment.

The average usage in the first month in Table 1 is statistically insignificant across conditions ($p > 0.1$). Thus, if the experiment was run for one month, one may conclude that there is no significant difference in usage between PPU and TPT. This explains why the sponsoring firm ran the experiment for 13 months. The longer duration was required because it takes time for consumers to adapt to the new service, and the firm was also interested in the impact of tariff structure on customer churn, which is unlikely to be observed in a short period of time.

Table 1 suggests that TPT pricing leads to higher churn than PPU. This result, however, cannot solely be attributed to the pricing structure effect because customers in different conditions faced different price levels. The model we propose will separate out these two effects.

Note that there are two different TPT conditions—TPT1, wherein there were increases in the fixed fee but the per-minute price remains the same (\$0.10 per minute); and TPT2, wherein there were increases in both the fixed fee and the per-minute price. As our interest lies in exploring the access fee effect, we consider a trialist to be in a TPT condition if he or she is either in a TPT1 or in a TPT2 condition. Pooling trialists in TPT1 and TPT2 conditions into a single TPT condition does not pose any significant problems because our proposed model controls for the effect of pricing differences through a budget constraint. Furthermore, our model accounts for consumer heterogeneity using a hierarchical Bayes specification. Thus any residual differences between consumers in TPT1 and those in TPT2 are captured by the individual-level parameters.⁴

Finally, it is important to mention that the field experiment data also included a second PPU condition (not shown in Table 1) where customers were charged \$0.10 per minute for the entire duration of the experiment. We do not include this condition in our analysis because there is no pricing variability in these data to allow the estimation of consumer-level price sensitivity. However, we can use this condition to assess the empirical validity of our model.

3. Model

In each period t ($t = 1, \dots, T$), customer i chooses whether to keep or drop the service. This decision depends on the consumer utility from anticipated usage. Conditional on renewal, the customer

⁴ In our empirical analysis, we found no significant differences between the parameter estimates of consumers in these two conditions. In addition, we performed two separate model estimations: PPU versus TPT1 and PPU versus TPT2. The empirical results show that our findings are statistically indistinguishable across these two comparisons.

Table 1 Field Experiment Description

Treatment group	PPU	TPT1	TPT2
Number of trialists	75	69	70
Number of dropouts	19	28	32
Percent dropout ^a	25.3	40.6	45.7
Avg. first-month usage	321.8 min.	371.6 min.	266.0 min.
Standard deviation	363.6	404.2	250.4

Prices changes over time	Access \$/month	Usage \$/min.	Access \$/month	Usage \$/min.	Access \$/month	Usage \$/min.
Oct (First month)	0	0.1	10	0.1	10	0.1
Nov–Jan	0	0.15	15	0.1	15	0.15
Feb–Apr	0	0.3	25	0.1	25	0.3
May–Oct	0	0.6	35	0.1	35	0.6

^aPercent dropout is the cumulative percentage of customers in the treatment group who dropped out of the service during the whole experiment.

then uses the service over the period. Our modeling approach reflects the joint nature of the choice and usage decisions. We build on the traditional unified framework of discrete/continuous consumer choices in which both choices flow from the same underlying utility maximization decision (Hanemann 1984). In §3.1, we present the consumer utility function. In §§3.2 and 3.3, we discuss how to use this utility function to model the keep/drop and usage decisions, respectively. We conclude the model development in §3.4 by describing how we capture the impact of pricing structure while controlling for the effects of observed and unobserved sources of consumer heterogeneity and time dependence.

3.1. Utility Function

Let $Y_{it} = 1(0)$ if customer i keeps (drops) the service in period t . Then, for a usage level q_{it} by customer i in period t and consumption z_{it} of an outside good (representing the customer expenditures on all other telecommunication services), we specify the following Stone–Geary (Geary 1950, Stone 1954) utility function to represent customer i 's utility for *keeping* the service:

$$U_{it}(Y_{it} = 1, q_{it}, z_{it}) = \gamma_{it} + \alpha_{it} \text{Log}(q_{it} + \tau^1) + \beta_{it} \text{Log}(z_{it} + \tau^0), \quad (1)$$

where γ_{it} is an intercept parameter, $\alpha_{it} > 0$ and $\beta_{it} > 0$ are satiation parameters, and τ^0 and τ^1 are translation parameters that allow the indifference curves to intersect the axes. Following Kim et al. (2007), we set $\tau^0 = 0$ to ensure that there is always an interior solution for the composite good and fix $\tau^1 = 1$ to identify α_{it} .⁵

⁵ The restriction $\tau^1 = 1$ ensures that the indifference curves will have finite nonzero slope at the axes, creating the possibility of a corner solution, whereby an active subscriber would optimally choose not to consume any minutes of usage (despite having access to the service). See Deaton and Muellbauer (1980, p. 65).

In Equation (1), customer i derives his or her utility from accessing and using the service and from consuming the composite good. If the customer decides to drop the service, his or her utility function is then given by

$$U_{it}(Y_{it} = 0, q_{it} = 0, z_{it}) = \beta_{it} \text{Log}(z_{it}). \quad (2)$$

Note that we set the intercept of Equation (2) to zero because only differences in intercepts are identified in discrete choice models. Consequently, γ_{it} in Equation (1) captures the additional utility or disutility a customer derives from having the option to access the service.

The Stone–Geary utility function we specify in Equation (1) has a number of desirable features. First, it allows diminishing marginal returns from consumption and therefore captures satiation. This is important in our context because more than one unit of usage time is consumed. Second, as we subsequently show, the specification allows the utility function to depend on income (a factor that Danaher 2002 found to be significant) and leads to a tractable demand function. For these reasons, the Stone–Geary utility function has been adopted in past research (e.g., Du and Kamakura 2008, Wales and Woodland 1983).⁶ Next, we describe how the specified utility function guides the keep/drop decision and the subsequent usage decision.

⁶ For comparison, we also tested the following utility function: $U_{it}(Y_{it} = 1, q_{it}, z_{it}) = \gamma_{it} + \alpha_{it} q_{it} + \beta_{it} \text{Log}(z_{it})$. This specification, which also preserves income and has a tractable demand function, results in a poorer model fit than that for the Stone–Geary utility function. This suggests that capturing diminishing return is important for cell phone usage. More importantly, our empirical results are unaffected by the choice of the utility function. In particular, consumers in the TPT condition obtain lower utility from the same level of consumption compared with those in the PPU condition ($\theta^c = -0.40$ with 95% C.I. $[-0.58, -0.23]$).

3.2. Keep/Drop Decision

Let I_i be customer i 's income, and let φ_i be the proportion of his or her income allocated to telecommunication category spending. Thus, $\varphi_i I_i$ represents his or her telecommunication budget. We assume that this budget is constant over the duration of the experiment.⁷ Let x_i be an indicator variable that takes a value of 1 (0) if customer i is assigned to a TPT (PPU) plan. Let F_{it} be the access fee, and let p_{it} be the per-minute usage rate for the plan at time t . Note that $F_{it} = 0$ when $x_i = 0$ (i.e., customer i has a PPU plan). Then assuming a unit price for the composite good z_{it} , customer i faces the following budget constraint:

$$p_{it}q_{it} + z_{it} + x_i F_{it} \leq \varphi_i I_i. \quad (3)$$

A utility-maximizing consumer will exhaust the budget, in which case $z_{it} = \varphi_i I_i - x_i F_{it} - p_{it}q_{it}$.

Note that we use the telecommunication budget $\varphi_i I_i$ as the relevant income. In such a case, z_{it} represents all other telecommunication goods. This specification is consistent with the finding in past work that consumers often set budgets for categories of expenses (e.g., entertainment) and track expenses against their budget (see e.g., Thaler 1985). It is also consistent with the Allenby et al. (2004) specification of a category-specific budgetary allotment. One implication of using a category-specific budget is that the consumer's expenditure on the service can no longer be considered negligible.⁸ Because it is not quasilinear, the Stone–Geary specification has the desirable property of allowing for such an income effect. In addition, a log specification captures the decreasing marginal utility of income: a higher price has a lower impact on the utility of high-income consumers than it has on the utility of low-income consumers (Sudhir 2001). Thus, higher (lower) income consumers are more (less) likely to stay with a more expensive plan.

The decision to continue with the service depends on the maximum level of utility that a consumer can derive from it. To determine this level of utility, a rational customer maximizes his or her utility

⁷ It is possible that the budget changes over time because of, for example, category expansion. We empirically tested for this possibility by specifying a time-varying budget (i.e., $\varphi_i I_i$). This specification, however, did not improve model fit nor did it affect the empirical results. In particular, consumers in the TPT condition obtain lower utility from the same level of consumption compared with those in the PPU condition ($\theta^c = -0.57$ with 95% C.I. $[-0.83, -0.22]$).

⁸ A recent report indicates that consumers on average spend approximately 3%–4% of their after-tax income on telecommunications (Cox and Alm 2008). Assuming that consumers spend 4% of their after-tax income on telecommunications, a descriptive analysis of our data set reveals that the average expenditure on the service among our trialists is of the order of 10% of their total telecommunication budget, which is far from negligible.

in Equation (1) subject to the budget constraint in Equation (3). Solving this utility maximization problem gives the following optimal quantity:

$$q_{it}^* = \frac{\alpha_{it}}{\alpha_{it} + \beta_{it}} \frac{\varphi_i I_i - x_i F_{it}}{p_{it}} - \frac{\beta_{it}}{\alpha_{it} + \beta_{it}}. \quad (4)$$

Substituting q_{it}^* and $z_{it} = \varphi_i I_i - x_i F_{it} - p_{it}q_{it}^*$ into Equation (1), we obtain the maximum utility that customer i obtains in period t if he or she continues with the service. This is given by

$$V_{it}(Y_{it} = 1, p_{it}, F_{it}, I_i) = \gamma_{it} + \alpha_{it} \log[q_{it}^* + 1] + \beta_{it} \log[\varphi_i I_i - x_i F_{it} - p_{it}q_{it}^*], \quad (5)$$

where q_{it}^* is given by Equation (4). Note that pricing (access fee F_{it} and per-minute usage price p_{it}) affects consumer utility directly through the budget constraint and indirectly through the optimal quantity q_{it}^* . Similarly, the indirect utility from dropping the service is given by

$$V_{it}(Y_{it} = 0, I_i) = \beta_{it} \log(\varphi_i I_i). \quad (6)$$

Thus the difference between the indirect utilities derived from keeping the service versus dropping it, which drives the customer keep/drop decision, is

$$w_{it}(p_{it}, F_{it}, I_i) = \gamma_{it} + \alpha_{it} \log[q_{it}^* + 1] + \beta_{it} \log \left[\frac{\varphi_i I_i - x_i F_{it} - p_{it}q_{it}^*}{\varphi_i I_i} \right]. \quad (7)$$

To capture random error in the customer keep/drop decision, let ν_{it} be a normally distributed error term. The random utility for the keep/drop decision for customer i in period t is

$$W_{it}(p_{it}, F_{it}, I_i) = w_{it}(p_{it}, F_{it}, I_i) + \nu_{it}. \quad (8)$$

We assume that ν_{it} has a mean of 0 and a variance of σ_{u1}^2 (σ_{u2}^2) if customer i is assigned to the PPU (TPT) condition. We specify tariff-specific variances to examine whether there is any differential impact of PPU and TPT on consumer uncertainty. Thus, the probability that customer i keeps the service, $P(Y_{it} = 1 | p_{it}, F_{it}, I_i) = P(W_{it} > 0)$, in the PPU condition is $\Phi(w_{it}/\sigma_{u1})$ and in the TPT condition is $\Phi(w_{it}/\sigma_{u2})$, where $\Phi(\cdot)$ denotes the cumulative normal distribution function.

As in any random utility model, the satiation parameters (α_{it}, β_{it}) and the variance parameters ($\sigma_{u1}^2, \sigma_{u2}^2$) are not separately identified. For identification, we can fix $\beta_{it} = 1$ or set either of the variance parameters equal to 1. However, only the ratio $\alpha_{it}/(\alpha_{it} + \beta_{it})$ is estimable (see Equation (4)). Thus, the parameters α_{it} and β_{it} cannot be separately identified. Therefore, to ensure that both the utility and the quantity components of our model are identified, we follow

Allenby et al. (2004) and arbitrarily set $\beta_{it} = 1$ (alternatively, we could set $\alpha_{it} = 1$).

Recall that $0 < \varphi_i < 1$ represents the proportion of income that customer i allocates to telecommunication spending. Our field experiment has data on customers' monthly income but lacks information on the proportion each customer allocates to telecommunication spending. We estimate this latter quantity directly from the data by specifying

$$\varphi_i = \frac{1}{(1 + e^{\pi + \chi_i})}, \quad (9)$$

where π is a parameter that reflects the average percentage expenditure of total income allocated to the telecommunication category, and χ_i is a person-specific random effect, normally distributed with a mean of 0 and a variance of σ_χ^2 .

3.3. Usage Decision

Because of the temporal difference between the time the consumer makes the keep/drop decision and the time he or she consumes the service (which happens during the period), the optimal quantity q_{it}^* in Equation (4) may deviate from the observed quantity q_{it}^{Actual} . Following common practice (e.g., Danaher 2002), we postulate a log-log relationship between these two quantities.⁹ To avoid taking logs of a zero quantity in a particular month, let $Q_{it} = q_{it}^{\text{Actual}} + 1$, and let $Q_{it}^* = q_{it}^* + 1$. Then we postulate the following log-log regression of Q_{it} on Q_{it}^* :

$$\begin{aligned} \text{Log}(Q_{it}) &= \text{Log}(Q_{it}^*) + \eta_{it} \\ &= \text{Log}\left(\frac{\alpha_{it}}{\alpha_{it} + 1}\right) + \text{Log}(\varphi_i I_i - x_i F_{it} + p_{it}) \\ &\quad - \text{Log}(p_{it}) + \eta_{it}, \end{aligned} \quad (10)$$

where η_{it} is a random error term.

To account for time dependence, we decompose the random component η_{it} as follows:

$$\eta_{it} = \varepsilon_t^q + \omega_{it}, \quad (11)$$

where ε_t^q is a period-specific shock, and ω_{it} is random white noise. We assume that ω_{it} is normally distributed with a zero mean and a variance of σ_{q1}^2 (σ_{q2}^2) when the customer is assigned to the PPU (TPT) condition. We specify tariff-specific variances to examine whether there is any differential impact of PPU and

TPT on consumer usage uncertainty. To account for time dependence in customer usage, we specify the following state-space formulation:

$$\varepsilon_t^q = \rho^q \varepsilon_{t-1}^q + \kappa_t^q, \quad \kappa_t^q \sim N(0, \sigma_{\kappa_t^q}^2), \quad (12)$$

where the initial condition parameter ε_0^q and the decay parameter ρ^q ($0 < \rho^q < 1$) are to be estimated from the data. This error specification represents a stationary, autoregressive process that accounts for time dependence in usage.

Overall, Equations (8) and (10) specify how customers make the keep/drop and usage decisions. We now discuss how we capture the impact of pricing structure in these models.

3.4. Impact of Tariff Structure

We use a varying-parameter approach to model the impact of tariff structure (PPU versus TPT) on customers' base-level utility for the service, γ_{it} , and their utility from consumption, α_{it} (see Equation (1)), while controlling for both observed and unobserved sources of variance resulting from customer heterogeneity and time dependence.¹⁰

3.4.1. Impact on Base-Level Utility. Let QTR_{jt} be a quarterly dummy for period t ($j = 1, \dots, 4$), where QTR_{1t} corresponds to the January–March quarter, QTR_{2t} to the April–June quarter, QTR_{3t} to July–September quarter, and QTR_{4t} to the October–December quarter.¹¹ Recall that $x_i = 1(0)$ if customer i is assigned to a TPT (PPU) plan. Then we reparametrize the base-level utility parameter γ_{it} as follows:

$$\gamma_{it} = \theta^u x_i + \sum_{j=1}^3 \mu_j^u \text{QTR}_{jt} + \sum_{j=1}^3 \delta_j x_i \times \text{QTR}_{jt} + \varsigma_i^u + \varepsilon_t^u, \quad (13)$$

where θ^u captures the tariff-structure effect, μ_j^u ($j = 1, \dots, 3$) are parameters to capture quarterly seasonality/time effects (with QTR_{4t} as the baseline quarter), and δ_j ($j = 1, \dots, 3$) captures how these effects

¹⁰ The tariff structure may also affect how consumers react to marginal prices. We estimated an ad hoc model where TPT consumers are allowed to react differentially to marginal price. Specifically, the budget constraint is specified as $(1 + \varsigma x_i)p_{it}q_{it} + z_{it} + x_i F_{it} \leq \varphi_i I_i$, where ζ is a parameter that captures the differential effect. We find that this additional parameter is not significant ($\zeta = 0.01$ with 95% C.I. $[-0.01, 0.03]$).

¹¹ We follow Danaher (2002), who uses quarterly dummies to capture seasonality. To check the robustness of our results, we estimated our model using two other parsimonious operationalizations of seasonality: (i) a quadratic specification with month and month squared, and (ii) a sine-cosine specification to capture any cyclical variations over time. Both cases resulted in worse model fit than that of our proposed model, which uses a quarterly dummy specification for seasonality. In addition, our key results are not affected by how seasonality is measured.

⁹ A descriptive analysis of the usage data shows that usage (in minutes per month) in each of the three groups is skewed (with the mean always exceeding the median), and some consumers have extremely high usage. Mosteller and Tukey (1977) recommend a log transformation. Analysis of the log-transformed usage data shows that, for each group, the mean and median are very close, and the histograms of the usage are normally distributed.

vary across tariff conditions. The individual-specific parameter s_i^u captures unobserved influences that affect customer i 's preference for the service. We assume that $s_i^u \sim N(0, \sigma_u^2)$. Finally, ε_t^u is a time-specific error term that captures the unobserved temporal influences on consumer preferences.

Note that because we observe only 13 months of data, the quarterly dummies confound seasonality and trend effects. In addition, as the marginal and access fee prices were varied over time, these quarterly dummies may seem to confound the experimental price changes as well. However, this is not the case because the price changes are not exactly synchronized over the quarters (see Table 1). For instance, in January (QTR1), we see that the marginal price is \$0.15 per minute, whereas in February and March (also in QTR1), the marginal price is \$0.30 per minute. More precisely, the correlations of QTR1 with the access fee (F) and with the per-minute price (p) are, respectively, -0.15 and -0.29 . The corresponding correlations for QTR2 are 0.11 and 0.16 ; those for QTR3 are 0.16 and 0.37 . Thus, given these low correlations, the quarterly dummies do not appear to confound the price changes over time.

Given our context of a new service, consumers can have uncertainty about the quality of the service provider, which they may learn about over time. Iyengar et al. (2007) show that quality uncertainty and learning can be captured using a state-space specification. Following their approach, we adopt a state-space specification:

$$\varepsilon_t^u = \rho^u \varepsilon_{t-1}^u + \kappa_t^u, \quad \kappa_t^u \sim N(0, \sigma_{\kappa^u}^2), \quad (14)$$

where the initial condition parameter ε_0^u and the decay parameter ρ^u ($0 < \rho^u < 1$) are to be estimated. This specification allows for past service valuation to affect current valuation.

3.4.2. Impact on Utility from Consumption. Recall that the satiation parameter α_{it} captures the effect of consumption on the utility of the service. As $\alpha_{it} > 0$, we reparametrize it as follows:

$$\alpha_{it} = \exp\left(\mu_0^c + \theta^c x_i + \lambda \log(q_{i,t-1}^{\text{Actual}}) + \sum_{j=1}^3 \mu_j^c \text{QTR}_{jt} + \zeta_i^c\right), \quad (15)$$

where μ_0^c is an intercept, and θ^c is a moderating parameter that captures the impact of tariff structure on α_{it} . A positive (negative) θ^c indicates that consumers in the TPT condition obtain higher (lower) utility with the same level of consumption compared with those in the PPU condition. This parameter allows us to assess the impact that the mere presence of an access fee in a TPT scheme, compared with a

PPU scheme, may have on consumer utility for the service.

In addition to the quarterly dummies, we also include in Equation (15) past consumption ($q_{i,t-1}^{\text{Actual}}$) as a covariate to capture the effect of habit formation (see Pollak and Wales 1992). Thus, a positive (negative) λ coefficient indicates that the utility from consumption α_{it} increases (decreases) with past usage. As the optimal quantity q_{it}^* (see Equation (4)) is a function of α_{it} , such a reparametrization captures how customers update their belief about how much to consume in a future month. Such information guides the service renewal decision (see Equation (8)). Finally, ζ_i^c is an individual-specific random variable to capture consumers' unobserved heterogeneity. We assume that $\zeta_i^c \sim N(0, \sigma_c^2)$.

We adopt a Bayesian framework for simulation-based inference. The details of the Bayesian estimation are available from the authors upon request. Before discussing the estimation results, we briefly discuss identification of the model parameters.

3.5. Identification

The identification of the model parameters is dependent on consumers' observed churn and consumption decisions as a result of exogenous price changes. For the 13-month duration of the experiment, we observe whether consumers defect or stay with the service. For those who stay, we observe their monthly consumption. We also observe the monthly access fees and per-minute prices that the firm charged across the TPT and PPU conditions.

Consumers' monthly churn allows us to identify the intercept (γ_{it}) and the variance terms ($\sigma_{u1}^2, \sigma_{u2}^2$) of the utility function. The latter parameters are identified because we fixed the outside good's satiation parameter to 1 (i.e., $\beta_{it} = 1$). Both the monthly churn and service usage data, along with the price changes, allow us to identify the effect on utility from consumption (α_{it}) and the proportion of income allocated to the telecommunication service (φ_{it}). In addition, because consumers are randomly assigned to the PPU and TPT conditions, we can identify the access fee effect as captured by the parameters θ^u and θ^c . That is, we can separately identify utility intercepts and consumption effects for each experimental condition.

The identification of consumer unobserved heterogeneity is possible because of the panel structure of the data (repeated observations), but it is aided by the changing price levels during the experiment. With every price change, consumers' utilities of staying with the service are changed, leading to different probabilities of churn and different levels of consumption. These variations in the data permit the identification of the distributions of heterogeneity.

The state-space error structure for both the keep/drop and usage decisions are identified because of the

time-series nature of the churn and usage data. For the same reason, the quarterly effects (μ_j s) and the habit formation parameter (λ) are also estimable. It is important to note that, because we observe each quarter only once, the quarterly effects confound any possible seasonality and trend effects.

Finally, note that our descriptive results in Table 1 indicate that TPT pricing leads to higher churn than does PPU pricing. This result, however, cannot solely be attributed to the impact of pricing structure because customers in different conditions faced different price levels. Moreover, it would be hard to disentangle these underlying effects by only considering summary statistics from data. As a consequence, a formal model is needed to assess the impact of structure on consumer behavior. The identification of the parameters θ^u and θ^c then relies on our choice of the Stone–Geary utility function. This function is ideal for our context: it allows for diminishing marginal returns from consumption, has a tractable demand function, and has good predictive validity (see §4.2). Furthermore, our main results are robust to the choice of utility function (see footnote 6).

4. Estimation Results

We estimate our model and four nested models that we now describe. First, to assess the effect of price structure, we estimate a nested model where PPU and TPT pricing do not differentially affect the consumer utility function (i.e., we set the parameters θ^u , θ^c , δ_1 , δ_2 , and δ_3 to 0). A comparison of the fit of this null model with our model fit can give an indication of the effect size of tariff structure on consumer decisions. We refer to this model as the “no price-structure effect model.” The second nested model, which we call “no time dependence model,” assumes no unobserved temporal dependence in consumer preferences. A comparison of the fit of this model with that of our model shows the benefit of accounting for the unobserved sources of time dependence, which we capture through the state-space specification. To assess the magnitude of habit formation, we estimate a third nested model where we set the parameter λ in Equation (15) to 0. We call this null model “no habit formation model.” Finally, to assess the magnitude of heterogeneity in our data, the fourth model constrains all model parameters to be fixed across customers. We refer to this model as the “no heterogeneity model.”

4.1. Model Comparison

We use Markov chain Monte Carlo (MCMC) methods to estimate the five models described above. For each model, we ran sampling chains for 100,000 iterations and assessed convergence by monitoring the time series of the draws. We report the results based

Table 2 Model Performance Comparison

Model	LML ^a	MAD-churn (%)	MAD-usage
Proposed	−2,902.4^b	0.65	0.11
No price structure effect	−2,926.3	0.90	0.14
No time dependence	−3,028.5	1.03	0.26
No habit formation	−3,115.4	0.97	0.15
No heterogeneity	−3,564.7	1.10	0.15

^aLML denotes log-marginal likelihood.

^bSelected model.

on 70,000 draws retained after discarding the initial 30,000 draws as burn-in iterations.

4.1.1. Goodness of Fit. To compare models, we use the Bayes factor (BF), which accounts for model fit and penalizes model complexity. Table 2 reports the log-marginal likelihoods (LMLs) for all the models. Kass and Raftery (1995) suggest that a value of $\log BF = (\text{LML}_{M1} - \text{LML}_{M2})$ greater than 5 provides strong evidence for the superiority of a model. Hence the LML results provide strong evidence for the superiority of our model relative to all other models. Thus, it is important to account for heterogeneity, time dependence, and habit formation effects in the model. More importantly, we note that the no price structure effect model is significantly worse in fit than our model. This shows that TPT and PPU pricing has a significant differential effect on a consumer’s decision to keep/drop the service and how much to consume. Later, we discuss the results from a demand analysis that illustrates the consequences of ignoring this effect on pricing the new service.

We use each model’s estimated parameters to assess its fit. We use two measures of fit. The first is the mean absolute deviation for percentage churn (MAD-churn), which captures how closely the model predicts actual monthly churn. (Note that we cannot use the mean absolute percentage error because it cannot be calculated for months with no churn.) The second is MAD-usage, which assesses how well the model predicts actual monthly usage (in logarithm scale). Based on these MAD criteria, Table 2 shows that our proposed model (in bold) has the best fit to the data.

Figure 1(a) contrasts actual and predicted churn over time under PPU (panel a) and TPT (panel b) pricing for our model. Note that we only show predictions for months 2 to 13 because the data from month 1 were not used for model estimation as a result of the lagged quantity specification. Similarly, Figure 1(b) compares actual and predicted monthly service usage in minutes per trialist (in logarithm scale) under PPU (panel a) and TPT (panel b) pricing. Both figures clearly show that our model predictions fit the churn and usage data very well regardless of the pricing condition.

Figure 1(a) Actual vs. Predicted Churn by Pricing Condition





Figure 1(b) Actual vs. Predicted Usage by Pricing Condition (in Log Scale)



4.2. Predictive Validity

Recall that there was a PPU condition in the experiment that did not experience any price change (we dropped this condition from the analysis because the data do not allow us to estimate individual-level price sensitivity parameters). However, we can make use of the data for this condition to assess the out-of-sample predictive validity of our model.

For each MCMC draw, we use the population-level distribution to generate parameters for each consumer in the holdout PPU condition. We use these parameters to predict consumer-level churn and usage behavior. Figure 2(a) contrasts actual and predicted churn, and Figure 2(b) compares actual and predicted monthly service usage in minutes per trialist (in logarithm scale) averaged across the MCMC draws.

Both figures clearly show that our out-of-sample predictions fit the churn and usage data well. In particular, the MAD for percentage churn is 1.90 and that for log usage is 0.23. In addition, we checked for any systematic bias in our churn and usage predictions.

Figure 2(a) Actual vs. Predicted Churn for the Holdout Condition




Figure 2(b) Actual vs. Predicted Usage for the Holdout Condition (in Log Scale)




Table 3 Parameter Estimates: Posterior Means and Posterior 95% Confidence Interval

Description	Param. label	Posterior mean ^a	95% posterior interval
Impact on base-level utility			
TPT indicator (x_i)	θ^u	0.73	(-1.31, 2.23)
QTR1	μ_1^u	0.93	(-1.21, 2.79)
QTR2	μ_2^u	0.18	(-1.98, 2.39)
QTR3	μ_3^u	-0.44	(-2.39, 1.34)
QTR1 * x_i	δ_1	-1.20	(-2.84, 0.38)
QTR2 * x_i	δ_2	-1.56	(-3.36, 0.07)
QTR3 * x_i	δ_3	-1.09	(-2.79, 0.56)
Heterogeneity variance	σ_u^2	4.50	(2.25, 8.90)
Impact on utility from consumption			
Intercept	μ_0^c	-3.14	(-3.33, -2.89)
TPT indicator (x_i)	θ^c	-0.57	(-0.67, -0.46)
Habit formation ($q_t - 1$)	λ	0.53	(0.50, 0.55)
QTR1	μ_1^c	-0.19	(-0.35, -0.08)
QTR2	μ_2^c	0.22	(0.09, 0.33)
QTR3	μ_3^c	0.40	(0.26, 0.52)
Unobserved heterogeneity	σ_c^2	0.39	(0.25, 0.55)
Telecom budget			
Mean proportion of income	π	3.19	(3.17, 3.21)
Heterogeneity variance	σ_π^2	0.79	(0.61, 1.01)
Utility variance			
PPU condition	σ_{u1}^2	1.03	(0.39, 2.60)
TPT condition	σ_{u2}^2	0.89	(0.36, 2.14)
Usage variance			
PPU condition	σ_{q1}^2	0.91	(0.82, 1.01)
TPT condition	σ_{q2}^2	0.75	(0.68, 0.81)
State space: Churn			
Initial value	ε_0^u	5.95	(2.08, 9.72)
Decay parameter	ρ^u	0.67	(0.34, 0.89)
Variance	$-\sigma_{\kappa^u}^2$	0.93	(0.33, 2.11)
State space: Usage			
Initial value	ε_0^q	-0.98	(-4.27, 2.27)
Decay parameter	ρ^q	0.43	(0.11, 0.83)
Variance	$\sigma_{\kappa^q}^2$	0.29	(0.14, 0.62)

^aPosterior mean for parameter. All “significant” coefficients are highlighted in boldface.

We find an average relative bias for quantity to be 0.01, and the average relative bias for churn is 0.18. These biases are very low in magnitude. Thus our model has good out-of-sample predictive validity, and there is no systematic bias in predictions.

4.3. Parameter Estimates

We now discuss the parameter estimates from our model. Table 3 summarizes the posterior distributions of the parameters by reporting their posterior means and posterior 95% confidence intervals. The results show that the parameter θ^u is positive but insignificant. In addition, none of the interaction effects of TPT with QTR1–QTR3 is significant. These results indicate that, on average, the pricing structure does not have a significant impact on the base-level utility of the service. The pricing structure, however, appears to affect consumer utility from consumption. Specifically, the moderating parameter θ^c is significant and negative. This indicates that consumers in the TPT condition obtain *lower* utility from the same level of

consumption compared with those in the PPU condition. Later, we quantify the size of this “access fee effect” by examining its impact on customer churn and consumption.¹²

The results also show a significant and positive habit formation effect. Thus, the utility from consumption increases with an increase in customer usage of the service over time, which results in lower likelihood of churn and higher expected usage. Table 3 shows that all usage quarterly dummies are significant. This suggests that consumption varies significantly over time. Using Equation (9), we find that customers allocate an average of 3.9% ($=1/(1 + \exp(3.19))$) of their income to telecommunication spending, consistent with the 3%–4% range reported in Cox and Alm (2008). We also find that the utility and usage variances do not vary significantly

¹²Technically, the access fee effect combines the effects of pricing structure on both the base-level utility and the utility from consumption.

across the two pricing conditions. This indicates that the pricing structure does not differentially affect consumer churn and usage uncertainties. Finally, as reflected by the large improvement in LML (see Table 2), we find significant unobserved heterogeneity in customers' base-level utility, utility from consumption, and budget allocated to telecommunication spending as measured by the parameters σ_u^2 , σ_c^2 , and σ_χ^2 , respectively.

Recall that we specify a state-space formulation to govern the time dependence of customer churn and usage over time. We find the initial value of the base utility, ε_0^q , to be positive and significant (see Table 3). This suggests that customers draw positive utility from having access to the service when they initially joined. In contrast, the initial value parameter for usage ε_0^q is insignificant. We also find significant decay parameters ρ^u and ρ^q . This suggests that both the utility and usage variances are stabilizing over time. Finally, the magnitude of the state-space variance parameters $\sigma_{\kappa^u}^2$ and $\sigma_{\kappa^q}^2$ suggests that the observed habit formation and quarterly effects do not fully capture time dependence.

In sum, we find consumers derive lower utility for TPT compared with PPU plans, resulting in both lower retention of customers and lower usage of the service. This negative impact on retention and usage is true even after controlling for income effects, heterogeneity across customers, and observable and unobservable time-varying factors.

5. Discussion

We now provide theoretical support for our results and discuss possible alternative explanations and how we rule out such alternatives.

5.1. Theoretical Support

Our finding that TPT customers derive lower utility from the same level of consumption compared with those in the PPU condition supports predictions from the price-partitioning literature (e.g., Bertini and Wathieu 2008, Chakravarthi et al. 2002, Cheema 2008). This literature suggests that consumers pay more attention to partitioned versus consolidated prices. In addition, payments are more painful when they are decoupled from the benefits of consumption (Prelec and Loewenstein 1998).¹³ Furthermore, Hamilton and Srivastava (2008) provide empirical evidence of an

inverse relationship between price sensitivity and consumption utility. Thus, as there are no tangible benefits associated with the access fee, consumers in the TPT condition are likely to pay more attention to price and derive less utility from consumption than those in the PPU condition, where the marginal price is directly associated with consumption.¹⁴

Our result is also consistent with prospect theory and mental accounting, which suggests that consumers tend to perceive multiple prices as more punishing than a single price of equal amount (Kahneman and Tversky 1979, Thaler and Johnson 1990). This theory provides a cogent argument in favor of PPU, resulting from a characterization of access services as a bundle of two components. When consuming access services, consumers purchase a bundle of two related but different product components, each with its specific benefit and value driver. The first one is the "access component," that is, the right to access the service if and when desired over the length of the subscription period (option value); and the second one is the "usage component," being the time spent actually using the service (usage value). Under a TPT pricing structure, the price of a service is *partitioned*, which facilitates a mental accounting assignment of the per-minute price to the usage component of the product bundle and the access fee to its access component. By contrast, under a PPU pricing structure, there is a single consolidated price for the product bundle, and multiple losses are integrated (the component prices), which according to mental accounting would contribute to a *rise* in the evaluation for the PPU service. Mental accounting would then postulate that a TPT has lower consumer utility because two price components are more punishing than one (Thaler 1985).

Finally, at first glance, our finding may seem inconsistent with the sunk cost fallacy (Arkes and Blumer 1985). It is important to note, however, that the sunk cost fallacy has been found in the context of flat-fee pricing and *not* in the context of TPT pricing. We use the findings in Heath and Fennema (1996) to clarify the distinction and to show that our results are consistent with theirs. In a between-subject experiment, Heath and Fennema asked subjects to state how many pool games they would expect to play given an entry fee to a bar and a per-game price for playing pool. The entry fee was varied at three levels (\$0, \$2, and

¹³ The flat-fee bias literature (e.g., Lambrecht and Skiera 2006) suggests that consumers prefer a flat-fee plan over a PPU plan even when it is the more expensive option. This is because, under a flat-fee plan, consumers know the size of the bill ahead of time, which facilitates any mental prepayment of the expense in their mind (Prelec and Loewenstein 1998). This theory does not apply in our context of TPT for two reasons. First, unlike a flat-fee tariff, consumers do not know the size of the final bill ahead of time.

Second, unlike a flat fee, a TPT has a per-unit charge that acts like a "meter" and maintains the tight coupling between payment and consumption.

¹⁴ Our model captures the joint effect of the impact of pricing structure on consumption utility and price sensitivity because the satiation parameter (α_{it}) and the budget coefficient (β_{it}) cannot be separately identified (see Equation (4)).

\$6), and the per-game price was varied at two levels (\$0 and \$0.50), leading to six conditions. Note that these conditions include the following: only entry fee and no-per game price (a flat-fee condition), no-entry fee and only a per-game price (a PPU condition), and both an entry fee and a per-game price (a TPT condition). In the flat-fee condition, they found that as the flat fee increased, subjects expected to play more games, which is consistent with the sunk cost fallacy. However, this was not the case in the TPT condition. As the entry fees increased, subjects expected to play fewer games. This result persisted even after controlling for the income effect, which they manipulated by “endowing” subjects in another experimental condition with the equivalent of the entry fee. Thus, the sunk cost fallacy does not appear to hold in a TPT context. Heath and Fennema also find that subjects in the TPT condition expected to play fewer games than those in a PPU condition even after adjusting for income effect. This result is clearly consistent with our access fee finding.

5.2. Alternative Explanations

Recall that all customers faced price increases during the experiment. Those in the TPT condition faced access fee and/or marginal price increases, whereas those in the PPU condition faced only marginal price increases. Thus the access fee effect could alternatively be explained by consumers’ differential reactions to price increases. For example, TPT customers may experience higher usage regret than PPU customers because they faced both access fee and marginal price increases. They may also experience a sticker shock or reference price effect. The access fee effect could also be explained by consumers’ forward-looking expectations that stem from changing prices. We discuss how we rule out these alternative explanations.

5.2.1. Differential Reactions to Price Increases.

In principle, the psychological effect of price increases should cancel out because every subject in the experiment experienced a certain price increase, regardless of whether the subject is assigned to a PPU or TPT condition. However, because the price increases are more pronounced in the TPT than the PPU condition, the access fee effect may be accentuated.

To investigate this issue, recall that the experiment has two two-part tariff conditions—TPT1, where the marginal usage price is fixed and the access fee is increased over the duration of the experiment; and TPT2, where both the marginal and access fees increased. If the magnitude of price increases were the underlying mechanism, then we should expect to find a stronger access fee effect in the TPT2 condition than in the TPT1 condition. We investigated this possibility in two ways. First, we tested for differences between

the individual-level parameters for consumers in TPT1 and those in TPT2, but we found no significant difference in these estimates across the two samples ($p > 0.05$). Second, we estimated our model separately for TPT1 versus PPU and TPT2 versus PPU. We find that our key result still holds and is not statistically different across the two conditions. Specifically, for TPT1 we find $\theta^c = -0.77$ (95% C.I. $[-1.03, -0.43]$), and for TPT2, $\theta^c = -0.40$ (95% C.I. $[-0.63, -0.17]$). Thus the magnitude of the price increases does not appear to accentuate the access fee effect.

Our external validation results (see §4.2) also provide additional evidence that consumers in the TPT and PPU conditions did not react differentially to price changes. Recall that the out-of-sample consumers did not experience any price change. Thus, if price increases are an issue, then we should expect to observe a systematic bias in our predictions. However, our results show that our predictions do not suffer from any systematic biases.

5.2.2. Reference Price Effects. To account for such possible effects, we extend our model by including price increases (sticker shocks) as additional covariates. Let Rp_{it} and RF_{it} be the marginal and access fee reference prices in month t , respectively. Then we measure the access fee sticker shock by $\Delta F_{it} = F_{it} - RF_{it}$ and the marginal price sticker shock by $\Delta p_{it} = p_{it} - Rp_{it}$. Following the reference price literature (e.g., Neslin and van Heerde 2009), we measure reference prices in three ways: the first month price, the last month price, and an average past price based on all previous months. Note that the use of last month price to measure reference price may be ambiguous as ΔF_{it} takes a zero value for both TPT and PPU conditions when the access fee does not change from month to month. However, ΔF is always positive when the reference price is defined otherwise.

We estimated our model with the two sticker shock measures as covariates. Specifically, the intercept of the utility function (γ_{it}) and the coefficient of consumption (α_{it}) are as follows:

$$\gamma_{it} = \theta^u x_i + \sum_{j=1}^3 \mu_j^u \text{QTR}_{jt} + \sum_{j=1}^3 \delta_j x_i \times \text{QTR}_{jt} + \omega_1^u \Delta F_{it} + \omega_2^u \Delta p_{it} + s_i^u + \varepsilon_i^u, \quad (16a)$$

$$\alpha_{it} = \exp \left(\mu_0^c + \theta^c x_i + \lambda \log(q_{i,t-1}^{\text{Actual}}) + \sum_{j=1}^3 \mu_j^c \text{QTR}_{jt} + \omega_1^c \Delta F_{it} + \omega_2^c \Delta p_{it} + \zeta_i^c \right). \quad (16b)$$

Regardless of how we measure reference prices, the results show that our finding of lower utility from consumption for participants in the TPT condition still

holds. For example, when the reference price is measured by the average past price, we find $\theta^e = -0.42$ with a 95% confidence interval $(-0.56, -0.30)$. We also find that the access fee sticker shock ΔF_{it} is not significant in both the intercept (Equation (16a)) and the coefficient of consumption (Equation (16b)). Thus the access fee effect is valid and does not confound the sticker-shock effect.

5.2.3. Forward-Looking Expectations. Our results show a significant and positive habit formation effect (see Table 3). A possible interpretation of this finding is that consumers use the price changes that they observe in the past to form expectations about future price changes. Because price increases in the TPT condition are higher in magnitude than those in the PPU condition, TPT consumers may reduce their consumption at a faster rate than PPU consumers to avoid an expensive “addiction” in the future. Thus, the differences in price expectations between the TPT and PPU conditions may explain the lower utility from consumption we find for TPT customers.

It is feasible to embed price expectations into our model, but there are a few design aspects about our field experiment that make this extension unnecessary. First, it is important to note that phone minutes cannot be stored for the new service in our experiment (i.e., no rollover minutes). Thus, there is no benefit in a consumer deliberately using the phone less this month to store minutes for later use (at a possibly lower rate). Second, participants in the experiment were not prewarned of the entire schedule of price changes. They were told in the prior month about pending tariff changes, but they were not told about changes beyond that month. This makes it difficult for participants to anticipate later tariff changes and to change their phone calling behavior in anticipation of tariff increases, because they were unaware of the frequency and amount of the future price changes. Finally, if reference prices are any indication of future price expectations, then our results show that such variables are insignificant, and the access fee effect is robust. For these reasons we do not specifically model consumer expectations.

To summarize, although it is beyond the scope of this paper to ascertain which underlying mechanism(s) is responsible for the access fee effect, we believe that our key result supports predictions from the price partitioning literature and is consistent with prospect theory and mental accounting. Furthermore, the results of our additional analyses provide sufficient evidence that the access fee effect is not an artifact of the experimental design. Our investigation is focused on empirically documenting the existence of an access fee effect and assessing the magnitude of its impact on service demand and profitability, which we discuss next.

6. Impact of Tariff Structure on Demand

We now examine how the demand for the new service is affected by tariff structure and by variations in the per-minute rate and access fee. For each tariff j , we use the MCMC draws of the parameters to predict each customer’s monthly keep/drop decision and his or her usage for the month if he or she stays with the service. It is important to note that a customer will exit the service permanently if the model predicts churn in a particular month, and his or her usage is then set to zero thereafter. Note also that the parameters we use to predict demand vary depending on whether tariff j is TPT or PPU.

Let R_{kj} be the percentage of retained customers by the end of the year computed from the k th MCMC draw of the parameters given tariff j ($k = 1, \dots, K$). Then we measure annual retention under tariff j (\bar{R}_j) by the average percentage of retained customers across the MCMC draws ($\bar{R}_j = \sum_k R_{kj}/K$). Similarly, let C_{kj} be the predicted average yearly usage of the service in the sample using the k th MCMC draw of the parameters given tariff j . Then we measure annual usage under tariff j (\bar{C}_j) by the average predicted yearly usage across the MCMC draws ($\bar{C}_j = \sum_k C_{kj}/K$). In this simulation we use $K = 300$ simulated draws of the parameters.


6.1. Impact on Annual Retention

Figure 3 shows how annual retention varies as a function of tariff structure (PPU versus TPT), per-minute rate, and access fee. Regardless of pricing structure, annual retention decreases with increasing per-minute rate and access fee. More importantly, TPT pricing results in uniformly lower annual retention than PPU pricing, even when the access fee is as low as \$10 per month.

The lower retention rate in the TPT condition stems from two effects. One, for the same level of marginal price, customers in the TPT condition have to pay an additional access fee F . This will decrease their optimal consumption level (see Equation (4)) and therefore their utility from retaining the service compared with the PPU customers. We call this the “income effect.” Two, in our empirical analysis, we found that TPT customers derive lower utility from consumption compared with those in the PPU condition (i.e., lower α_{ii}). We call this the “access fee effect.”

To separate out these effects, we performed an additional TPT simulation where we adjusted for the income effect by “endowing” customers with the same amount of the access fee. Thus any remaining difference in annual retention between PPU and TPT is due to the access fee effect. The results of this simulation are shown in Figure 3 under the retention curve labeled “TPT–income adjusted.” For instance,

Figure 3 Impact of Pricing Structure on Annual Retention



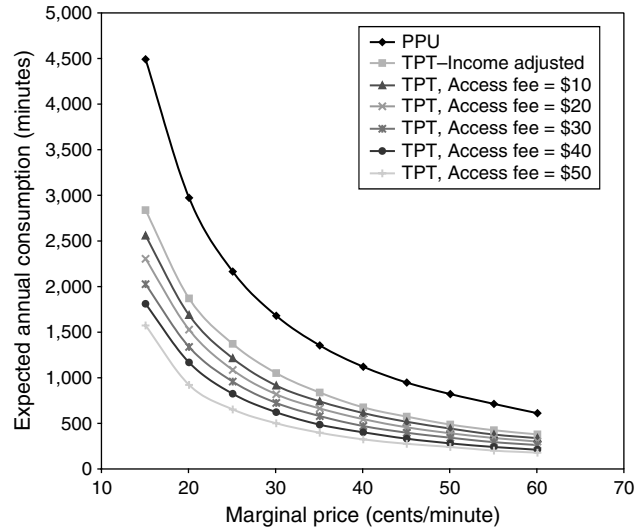
consider the retention values for a per-minute price of $p = \$0.20$. Under this price, a PPU tariff leads to a 68% retention rate, a TPT tariff with an access fee of $F = \$20$ yields a retention rate of 58%, and the income-adjusted TPT has a retention rate of 61%. Thus, the total 10% difference ($=68\% - 58\%$) can be split into 7% ($=68\% - 61\%$) because of the access fee effect and the remaining 3% because of the income effect. As expected, the contribution of the income effect increases with the level of access fee. For example, when $p = \$0.20$ and $F = \$50$, the retention rate is 42%, which is 26% lower than the retention rate of 68% in the PPU tariff of $\$0.20$. Of this 24% difference, 7% ($=68\% - 61\%$) is due to the access fee effect, and the remaining 19% ($=61\% - 42\%$) is due to the income effect. A similar analysis can be performed for every per-minute price in Figure 3. Overall, across all these price points, we find that TPT leads to a decline of 10.5% in the retention rate relative to PPU that is due purely to the access fee effect.

6.2. Impact on Annual Usage

Figure 4 shows how tariff-structure (PPU versus TPT), per-minute rate, and access fee affect annual usage per customer. As expected, regardless of pricing structure, annual usage declines with increasing per-minute rate and access fee. More surprisingly, however, TPT pricing yields uniformly lower annual usage than PPU pricing.

The difference in usage between PPU and TPT pricing can, as before, be decomposed into income and access fee effects using the “TPT-income adjusted” consumption curve in Figure 4. Again, consider a per-minute price of $p = \$0.20$. Under this price, the annual usage per customer in the PPU tariff is 2,963 minutes; a TPT with an access fee of $\$20$ yields an annual

Figure 4 Impact of Pricing Structure on Annual Usage



usage of 1,518 minutes; and the corresponding “TPT-income adjusted” is 1,860 minutes. Thus, of the total difference of 1,445 minutes ($=2963 - 1518$), the access fee effect accounts for 1,103 ($=2,963 - 1,860$) minutes, whereas the income effect accounts for the remaining 342 ($=1,860 - 1,518$) minutes. Thus, TPT pricing (with $F = \$20$) yields a 37% ($=(2,963 - 1,860)/2,963$) reduction in usage that is due purely to the access fee effect. As for retention, the contribution of the income effect increases with the level of access fee. Overall, across all price points, we find that TPT pricing leads to an average decrease of 38.7% in usage relative to PPU pricing that is due mainly to the access fee effect. Further decomposition shows that 10.5% of this drop in usage is due to consumer churn; the remaining 28.2% is due to reduction in usage conditional on retention.

In summary, we find that TPT pricing leads to lower retention and usage compared with PPU pricing over and above what is expected by the income effect. Overall, we find that the access fee effect accounts for 10.5% of the annual retention drop and 38.7% of the annual usage decline under the TPT scheme compared with PPU pricing.

7. Impact of Tariff Structure on Optimal Pricing Policy

Recall that our major goals from using the field experiment data are to (i) understand how customers differentially respond to PPU versus TPT pricing; (ii) decide on which pricing structure to use when launching the new service; and (iii) depending on the structure chosen, determine the optimal price level(s). The demand analysis in the previous section has studied the first issue. In this section, we address the latter two.

Table 4 Optimal Pricing Policy

	Proposed model		No price structure model	
	PPU	TPT	PPU	TPT
Per-minute price (\$)	0.50	0.35	0.40	0.25
Access fee	—	\$40	—	\$55
Annual usage	810 min.	513 min.	1,034 min.	1,065 min.
Annual retention (%)	62	45	61	46
Annual profit per customer (\$)	283	318	258	410

For the purpose of illustration, suppose the variable cost for offering the new service is \$0.15 per minute.¹⁵ We perform the following grid search to identify the optimal pricing plan for the new service. We evaluate all possible combinations of the pricing factors at discrete points: (1) access fee in increments of \$5, ranging from \$10 to \$60 per month; (2) per-minute rate in increments of 5 cents, ranging from 15 cents per minute to 60 cents per minute; and (3) pricing structure (TPT or PPU). We then identify the pricing plan with the highest annual profit for the service provider. For every pricing plan, we use our proposed model parameters to predict customer annual usage by following the same approach we discussed in the previous section. Table 4 describes the pricing plan with the highest yearly profit per customer, one obtained assuming PPU pricing and the other assuming TPT pricing. The profit calculations are conditional on customer acquisition and do not account for acquisition cost.¹⁶ For comparison, Table 4 also reports these same quantities obtained using the “no price structure” model, which ignores the access fee effect.

Despite the higher customer churn and lower usage, we find that TPT is still the profit-maximizing pricing structure for our proposed model. As expected, the optimal per-minute price of this TPT policy (\$0.35/min) is lower than that of the PPU policy (\$0.50/min). For PPU, our proposed model recommends a higher per-minute price than does the no price structure model. This is because of the higher PPU utility from consumption obtained from our model compared to that of the no price structure model. For TPT, our model identifies a lower access fee compared with the one suggested by the no price

structure model. This difference stems from the lower TPT utility from consumption in our proposed model, which is due to the access fee effect that the no price structure model fails to capture.

Overall, ignoring the access fee effect results in underestimating the profitability of PPU pricing and overestimating it for TPT pricing. This falsely accentuates the difference between TPT and PPU profits. Using the proposed model, the optimal TPT scheme provides 12% $(=(318 - 283)/283)$ higher annual profit per customer compared with the PPU tariff. In contrast, using the no price structure model, the optimal TPT scheme gives almost 59% $(=(410 - 258)/2,584)$ higher annual profit per customer compared with the PPU tariff.

To further explore the managerial implications from ignoring the access fee effect, we use the proposed model to assess the profitability of the optimal pricing plans identified by the no price structure model. This mimics a scenario in which a firm may erroneously set the optimal prices without considering the access fee effect, when, in reality, their customers behave as indicated by our proposed model. Thus, this analysis indicates the magnitude of profit reduction that will ensue from using a misspecified model. We find that the firm should make an annual profit of \$284 and \$271 per customer under TPT and PPU pricing, respectively. Given that TPT pricing is optimal for both models, the firm would therefore be forgoing a \$34 $(=318 - 284)$ profit per customer per year. Put differently, the failure to account for the access fee effect leads to about an 11% $(=(318 - 284)/318)$ reduction in the firm’s profit.

Our profit estimates are biased upward because the model parameters for profit maximization are estimated with error (see Mannor et al. 2007). We adapted the Mannor et al. cross-validation approach to assess the magnitude of this bias. We randomly divided our sample of customers into two subsamples, a calibration sample and a validation sample. We used the calibration sample estimates from our proposed model to derive optimal PPU and TPT pricing tariffs. We then used the optimal tariffs to calculate the annual profit per customer in both the calibration and validation samples. The difference between the annual profits calculated from the calibration ($\text{Profit}_{\text{cal}}$) and the validation sample ($\text{Profit}_{\text{val}}$) provides an estimate of the bias. We find that the optimal PPU and TPT tariffs based on the calibration sample are the same as those obtained from the full sample (see Table 4), albeit with lower annual profit per customer as expected. Specifically, $\text{Profit}_{\text{cal}}$ for PPU (TPT) is \$274.13 (\$316.24), and $\text{Profit}_{\text{val}}$ for the PPU (TPT) is \$262.90 (\$308.80). Thus, we find an upward bias of 3.9% $(=(274.13 - 262.90)/283)$ for PPU and 2.3% $(=(316.24 - 308.80)/318)$ for TPT.

¹⁵ We can calculate the variable cost per minute by dividing the operating costs (revenue – EBITDA) of the provider by the total number of mobile minutes. For instance, for AT&T, by September 2008, the total operating costs were \$10,239 million, and the total number of minutes consumed by subscribers was 54,995 million minutes. This leads to a cost per minute of around \$0.18 per minute. Other service providers such as Verizon have similar costs.

¹⁶ It is possible that it may cost more to induce a customer to sign up for a TPT plan compared with a PPU plan. This will further diminish the difference in profitability from the optimal TPT and PPU plans.

It may not be surprising that TPT is the optimal tariff structure because PPU is a special case of a two-part tariff when the access fee is zero. Therefore, the profit from PPU cannot exceed that from TPT (Oi 1971). This result, however, assumes that the demand curve is *invariant* to the pricing structure. That is, usage at a given per-unit price is the same regardless of whether the consumer has a PPU or a TPT plan. However, as we have shown in this paper, the price structure *can* influence consumers' consumption. In such a case, it is possible that a TPT scheme may not be optimal. See the electronic companion to this paper, available as part of the online version that can be found at <http://mktsci.pubs.informs.org/>.

8. Conclusion

Past research in psychology and marketing suggests that the pricing structure a firm chooses can alter consumers' value for a product or a service. In this paper, we focus on two commonly used pricing structures (two-part tariff and pay-per-use) and empirically quantify their impact on consumer preferences. We develop a utility-based model of consumer usage and retention of a subscription service that allows the model parameters to vary with the type of pricing structure faced by the customer (PPU or TPT). An important feature of our model is its ability to disentangle the access fee effect from the income effect and to measure its impact on consumer retention and usage, as well as on pricing policy. Moreover, our model controls for various sources of observed and unobserved consumer heterogeneity and time dependence.

Using data from a pricing field experiment for a new telecommunication service, we find that consumers have significantly lower utility from consumption for TPT compared with PPU pricing plans, resulting in both lower retention of customers and lower usage of the service. Our demand analysis shows that a TPT structure leads to an average decline of 11% in the annual retention rate and an average decrease of 38.7% in yearly usage relative to PPU pricing, even after controlling for income effects. Despite the higher customer churn and lower usage, we find that TPT is still the profit-maximizing pricing structure. In particular, our results show that a firm would overcharge customers for the access fee and undercharge them for the per-minute price if it ignores the access fee effect. In terms of profitability, the failure to account for such an effect leads to a reduction of 11% in firm profit.

In this paper we considered the impact of PPU and TPT pricing structures on the demand for a new access service offered by a monopolist. Future research should generalize this investigation to include other

pricing structures (e.g., three-part tariffs) and services that operate in a competitive environment. Our field experiment's design randomly assigned customers to PPU and TPT pricing conditions. As such, we cannot model consumers' choice between plans. Future research could examine the impact of pricing structure not only on retention and consumption but on consumer tariff choice as well. A related area of research concerns the use of CRM databases or natural experiments. Although such data are more realistic, they suffer from consumer self-selection issues and limited variability in prices over time compared with field experiment data. Finally, our model did not include consumer expectations about future price changes. It would be fruitful to extend the model in a dynamic, structural fashion to examine the differential impact of pricing structure on consumer demand.

9. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://mktsci.pubs.informs.org/>.

Acknowledgments

The authors benefited greatly from comments by Marco Bertini; Amar Cheema; the *Marketing Science* review team; and seminar attendees at the Olin Business School, Washington University at St. Louis; the Kenan-Flagler Business School, University of North Carolina; the Kellogg School of Management, Northwestern University; the Fall 2009 Wharton Brown Bag series, the 2009 Koç-Özyeğin-Sabancı Marketing Workshop, the 2010 Eighth Triennial Choice Symposium, and the 2009 and 2010 INFORMS Marketing Science Conferences.

References

- Allenby, G. M., T. S. Shively, S. Yang, M. J. Garratt. 2004. A choice model for packaged goods: Dealing with discrete quantities and quantity discounts. *Marketing Sci.* 23(1) 95–108.
- Anderson, E., D. Simester. 2003. Effect of \$9 price endings on retail sales: Evidence from field experiments. *Quant. Marketing Econom.* 1(1) 93–110.
- Arkes, H. R., C. Blumer. 1985. The psychology of sunk cost. *Organ. Behav. Human Decision Processes* 35(1) 124–140.
- Ascarza, E., A. Lambrecht, N. J. Vilcassim. 2011. "When talk is 'free': The effect of tariff structure on usage under two- and three-part tariffs. Working paper, Columbia University, New York.
- Bertini, M., L. Wathieu. 2008. Research note—Attention arousal through price partitioning. *Marketing Sci.* 27(2) 236–246.
- Chakravarthi, D., R. Krish, P. Paul, J. Srivastava. 2002. Partitioned presentation of multicomponent bundle prices: Evaluation, choice and underlying processing effects. *J. Consumer Psychol.* 12(3) 215–229.
- Cheema, A. 2008. Surcharges and seller reputation. *J. Consumer Res.* 35(1) 167–177.
- Cox, W. M., R. Alm. 2008. You are what you spend. *New York Times* (February 10), <http://www.nytimes.com/2008/02/10/opinion/10cox.html>.
- Danaher, P. J. 2002. Optimal pricing of new subscription services: Analysis of a market experiment. *Marketing Sci.* 21(2) 119–138.

- Deaton, A., J. Muellbauer. 1980. *Economics and Consumer Behavior*. Cambridge University Press, Cambridge, UK.
- Du, R. X., W. A. Kamakura. 2008. Where did all that money go? Understanding how consumers allocate their consumption budget. *J. Marketing* 72(6) 109–131.
- Essegaier, S., S. Gupta, J. Z. Zhang. 2002. Pricing access services. *Marketing Sci.* 21 139–159.
- Geary, R. C. 1950. A note on "A constant-utility index of the cost of living." *Rev. Econom. Stud.* 18(1) 65–66.
- Hamilton, R. W., J. Srivastava. 2008. When 2 + 2 is not the same as 1 + 3: Variations in price sensitivity across components of partitioned prices. *J. Marketing Res.* 45(4) 450–461.
- Hanemann, M. W. 1984. Discrete/continuous models of consumer demand. *Econometrica* 52(3) 541–561.
- Heath, C., M. G. Fennema. 1996. Mental depreciation and marginal decision making. *Organ. Behav. Human Decision Processes* 68(2) 95–108.
- Iyengar, R., A. Ansari, S. Gupta. 2007. A model of consumer learning for service quality and usage. *J. Marketing Res.* 44(4) 529–544.
- Iyengar, R., K. Jedidi, R. Kohli. 2008. A conjoint approach to multi-part pricing. *J. Marketing Res.* 45(2) 195–210.
- Kahneman, D., A. Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47(2) 263–292.
- Kass, R. E., A. E. Raftery. 1995. Bayes factors. *J. Amer. Statist. Assoc.* 90(430) 773–795.
- Kim, J., G. M. Allenby, P. E. Rossi. 2007. Product attributes and models of multiple discreteness. *J. Econom.* 138(1) 208–230.
- Kridel, D. J., D. E. Lehman, D. L. Weisman. 1993. Option value, telecommunications demand and policy. *Inform. Econom. Policy* 5(2) 125–144.
- Lambrecht, A., B. Skiera. 2006. Paying too much and being happy about it: Existence, causes, and consequences of tariff-choice biases. *J. Marketing Res.* 43(2) 212–223.
- Lambrecht, A., K. Seim, B. Skiera. 2007. Does uncertainty matter? Consumer behavior under three-part tariffs. *Marketing Sci.* 26(5) 698–710.
- Mannor, S., D. Simester, P. Sun, J. N. Tsitsiklis. 2007. Bias and variance approximation in value function estimates. *Management Sci.* 53(2) 308–322.
- Mark, K., R. E. White. 2007. Apple's iPod system: iPod, iTunes and FairPlay. Case 9B06M080, Richard Ivey School of Business, University of Western Ontario, London, ON, Canada.
- McGovern, G. 2003. *Virgin Mobile USA: Pricing for the Very First Time*. Harvard Business School Publishing, Boston.
- Mosteller, F., J. W. Tukey. 1977. *Data Analysis and Regression: A Second Course in Statistics*. Addison-Wesley, Reading, MA.
- Miravete, E. J. 2003. Choosing the wrong calling plan? Ignorance and learning. *Amer. Econom. Rev.* 93(1) 297–310.
- Narayanan, S., P. K. Chintagunta, E. J. Miravete. 2007. The role of self-selection, usage uncertainty and learning in the demand for local telephone service. *Quant. Marketing Econom.* 5(1) 1–34.
- Neslin, S. A., H. J. van Heerde. 2009. Promotion dynamics. *Foundations Trends Marketing* 3(4) 177–268.
- Oi, W. Y. 1971. A Disneyland dilemma: Two-part tariffs for a Mickey Mouse monopoly. *Quart. J. Econom.* 85(1) 77–96.
- Pollak, R. A., T. J. Wales. 1992. *Demand System Specification and Estimation*. Oxford University Press, New York.
- Prelec, D., G. F. Loewenstein. 1998. The red and the black: Mental accounting of savings and debt. *Marketing Sci.* 17(1) 4–28.
- Soman, D., J. T. Gourville. 2001. Transaction decoupling: How price bundling affects the decision to consume. *J. Marketing Res.* 38(1) 30–44.
- Stone, R. 1954. Linear expenditure systems and demand analysis: An application to the pattern of British demand. *Econom. J.* 64(255) 511–527.
- Sudhir, K. 2001. Competitive pricing behavior in the auto market: A structural analysis. *Marketing Sci.* 20(1) 42–60.
- Sundararajan, A. 2004. Nonlinear pricing of information goods. *Management Sci.* 50(12) 1660–1673.
- Thaler, R. H. 1985. Mental accounting and consumer choice. *Marketing Sci.* 4(3) 199–214.
- Thaler, R. H., E. J. Johnson. 1990. Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Sci.* 36(6) 643–660.
- Wales, T. J., A. D. Woodland. 1983. Estimation of consumer demand systems with binding non-negativity constraints. *J. Econom.* 21(3) 263–285.