The Joint Sales Impact of Frequency Reward and Customer Tier Components of Loyalty Programs

Praveen K. Kopalle*
Scott A. Neslin
Baohong Sun
Yacheng Sun
Vanitha Swaminathan

May 28, 2009

*Praveen Kopalle is Associate Professor of Business Administration and Faculty Associate, Center for Digital Strategies at the Tuck School of Business at Dartmouth, Dartmouth College, Hanover, NH 03755; Ph: 603-646-3612, Fax: 603-646-1308, Email: Praveen.kopalle@dartmouth.edu.

Scott Neslin is Albert Wesley Frey Professor of business administration at the Tuck School of Business at Dartmouth, Dartmouth College, Hanover, NH 03755; Ph: 603-646-2841, Fax: 603-646-1308, Email: Scott.Neslin@dartmouth.edu.

Baohong Sun is Associate Professor of Marketing at the Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213; Tel: 412-268-6903; Fax: 412-268-7357; Email: bsun@andrew.cmu.edu.

Yacheng Sun is Assistant Professor of Marketing, Leeds School of Business, University of Colorado, Boulder. Bldg #4, UCB 419, Boulder, CO 80309; Tel: 303-492-6211; Email:Yacheng.Sun@Colorado.edu.

Vanitha Swaminathan is Associate Professor of Business Administration at the Katz school of business, University of Pittsburgh, 344 Mervis Hall, Pittsburgh, PA 15260; Ph: 412 648-1579; Email: vanitha@katz.pitt.edu.

The Joint Sales Impact of Frequency Reward and Customer Tier Components of Loyalty Programs

Abstract

In this paper, we estimate the joint impact of frequency of frequency reward and customer tier components of a loyalty program on customer behavior and the corresponding sales. We provide an integrated analysis and measurement of a loyalty program incorporating customers' purchase and cash-in decisions, points pressure and rewarded behavior effects, heterogeneity, and forward-looking behavior. We do so via developing and estimate a dynamic structural model of customer behavior. We focus on the analysis of a hotel loyalty program, but in order to establish generalizability, we demonstrate these effects in the context of an airline's loyalty program as well. Using the hotel data, we find that a two-segment solution fits the data best both in estimation and validation samples, revealing a "service-oriented" segment that highly values the customer tier component, and a "price-oriented" segment that is more price sensitive and highly values the frequency reward aspects of the loyalty program. We empirically illustrate points pressure and rewarded behavior effects. Our research models the strength of response to frequency reward and customer tier program. In addition, we conduct policy simulations that vary aspects of the reward structure and measure the differential impact for both the price oriented and service-oriented segments and develop frequency reward and customer tier program requirements to maximize firm revenues.

Keywords: Loyalty program, customer tier programs, frequency reward, database marketing, segmentation

1. Introduction

In a business environment marked by an increase in firm knowledge about its customers, managers and researchers are as aware as ever of the importance of brand loyalty (Gupta, Lehmann, and Stuart 2004). Accordingly, "loyalty programs", designed to maintain and enhance loyalty, have become "go-to" marketing programs for many companies (Deighton 2000; Lewis 2004; Liu 2007; Zeithaml, Rust, and Lemon 2001).

Advocates view loyalty programs as a means to soften price competition (e.g., Kim, Shi, and Srinivasan 2001; Klemperer 1987) or build a customer database (Butscher 1998; Reynolds 1995), and a dominant strategy when fixed costs are low (Kumar and Rao 2003). Since a minority of customers often contribute most to a firm's profits, it is logical to lavish attention on them (Peppers and Rogers 1997). Critics, however, cite high program costs (Dowling and Uncles 1997), question whether they really increase loyalty (Keiningham, Vavra, and Aksoy 2006; Hartmann and Viard 2006; Sharp and Sharp 1997; Shugan 2005), and see them as a potential prisoner's dilemma (Kopalle and Neslin 2003).

One factor muddling this debate is that loyalty programs consist of two distinct components: frequency reward and customer tier (Blattberg, Kim, and Neslin 2008). Differentiating the relative effectiveness of these two components is important in order to evaluate the success of any customer loyalty program. Frequency reward programs are of the form, "Buy X times, get something free." These are the original trading stamp and frequency reward programs. Customer tier programs are of the form, "Once you qualify for our Diamond tier, we will provide you with special services." Both components rely on accumulated customer sales to determine which customers get which rewards. However, they differ in the nature of the reward as well as the means by which customers attain them. A frequency program reward is meaningful but *ephemeral* – a free stay at a hotel, a free flight, etc. In contrast, customer tier programs offer a bundle of services such as dedicated phone lines, faster check-in service, etc., for an *extended time period*. Second, while frequency reward programs typically require customers to be proactive in trading-in their points to redeem, customer tier programs dispense their reward automatically. Once customers qualify for a certain tier, they are notified and treated according to their tier status.

The contribution of this paper is threefold. First, we provide an integrated analysis of the impact of two critical components (frequency reward and customer tier) of a loyalty program on customer behavior and corresponding sales. We develop a comprehensive dynamic model that incorporates the following: (a) *Forward-Looking Customers* – Both frequency rewards (e.g., free hotel stay) and customer tier benefit (e.g., elite status) aspects of loyalty programs encourage customers to consider the future ramifications of their current choices, because these choices bring them closer to receiving a reward. (b) *Obtaining the Reward* - For the frequency reward program, this requires an endogenous decision to "cash-in." In contrast, for customer tier programs, the reward is delivered automatically. (c) *Customer Heterogeneity* - some customers may feel a free stay is worth any effort, while others may not find it worthwhile. Some may not value additional services that are provided to "Diamond" customers, while others crave the royal treatment. (d) *Rewarded Behavior* - The reward indeed is short-term, but customer affect created by the reward can translate into an increase in loyalty. (e) *Points Pressure:* Our structural model will determine empirically whether customers are more likely to purchase as they get closer to the reward.

Second, we apply our model to two categories, hotels and airlines, and generate robust findings with respect to the differential customer response to frequency reward and customer tier components via structurally modeling the two loyalty program components. These substantive findings add to our knowledge base of customer response to loyalty programs. Third, we conduct policy simulations and determine the requirements for the frequency reward and customer tier components that maximize firm revenue. Taking advantage of the points pressure effect, our optimization suggests a little more stringent requirement for frequency rewards, relative to the extant policy. On the other hand, the requirement for the first tier is made a bit more lax as an inducement to the customer base. We believe these provide actionable guidelines for managers to help in designing an optimal loyalty program.

Our approach involves modeling the forward-looking behavior of customers using a theory-based, customer utility function (see Sun et al. 2003). This is more consistent with the customer decision making process, particularly in a high involvement category such as hotels. This in turn, may provide more accurate results and thus more robust and valid guidance for policy decisions (van Heerde, Dekimpe, and Putsis 2005). There is substantial evidence in previous research that consumers are indeed forward-looking (e.g., Erdem and Keane 1996, Erdem, Imai, and Keane 2003, Gönül and Srinivasan 1996, Sun, Neslin and Srinivasan 2003). In the loyalty program context as well, Lewis (2004) found evidence of forward-looking behavior. Prior loyalty programs research (i.e., Kivetz, Urmininsky, and Zheng 2006; Lewis 2004; Nunes and Drèze 2006) has also demonstrated the points pressure phenomenon, i.e., customers increasing their purchases in order to get a reward. Further, research by various authors (e.g., Bolton, Kannan, and Bramlett 2000; Lal and Bell 2003; Roehm, Pullins, and Roehm 2002; Taylor and Neslin 2005) provides evidence of rewarded behavior effects, i.e., positive carry-over of a reward in the last period to the next period. Further, researchers have examined heterogeneity among consumers in terms of rewards redemption within a latent class segmentation framework (i.e., Lewis 2004).

While the above research provides a solid foundation for understanding the impact of loyalty programs, most of it has focused on the frequency reward format. Also, while there has been a lot of research on customer acquisition and retention (Blattberg and Deighton 1996; Blattberg et al. 2008; Bolton et al. 2000; Gupta et al. 2004; Reinartz, Thomas, and Kumar 2005; Venkatesan and Kumar 2004), very little has been learned about customer response to customer tier programs, particularly in a dynamic setting (see Blattberg et al. 2008). In addition, no research to date has explicitly modeled the customer cash-in decision in a frequency reward program. An integrated analysis of a loyalty program containing both frequency reward and customer tier components is needed in order to disentangle their relative effects. Further, this analysis should capture possible cannibalization effects between free- and paid-stays. Our work integrates key phenomena and is therefore distinctive in its scope. Our results will enable us to answer questions such as: What are the total and separate revenue impacts of frequency reward and customer tier components? Are there segments, and if so, how do they differ? Do the benefits of the frequency reward outweigh the costs of cashing in? What are the magnitudes of the points pressure and rewarded behavior effects? What are the policy implications with respect to the joint design of loyalty programs?

We apply our model first to the loyalty program instituted by a major hotel chain, and then replicate it with an airline program. The analyses yield several interesting results: (1) The forward-looking model outperforms the myopic model and better explains customer decisions better both in calibration and holdout data; (2) Both analyses yield rich customer segments; customers vary more in their response to the frequency component than to the customer tier component; (3) Points pressure is significant with respect to both components; (4) There is a rewarded behavior effect with respect to the frequency reward, although not in all segments. We examine the effectiveness of both frequency reward and customer tier programs and develop a revenue-maximizing design which balances both points pressure and rewarded behavior effects.

2. Model

2.1. Background

The model is generally applicable to all loyalty programs comprised of both frequency reward and customer tier components. For the sake of convenience, we describe the model with our first application to a hotel's loyalty program. The hotel categorized each of their properties as "economy", "regular," or "luxury", and customers can cash-in points for an upgrade to better rooms at a luxury property or a free night's stay at one of the three properties. Customers are required to contact the hotel to cash in their points for the reward. Points are accrued via paid stays, they do not expire, and no points are earned on free stays. The customer tier component places customers in Platinum and Diamond tiers depending on how many paid stays they had during the past calendar year. Tier benefits include better service, priority check-in, lounge access, service recovery assistance, etc. Upon reaching a tier threshold during a calendar year, customers are informed of their status and provided with the service level commensurate with their tier for the rest of that calendar year as well as the following year.

2.2. Customer Decisions

We consider $i=1,\dots,I$ customers who participate in the loyalty program[1]. At time $t=1,\dots,T$ (i.e., biweekly), each customer makes one of $k = 0, 1, \dots, 7$ decisions. We define D_{ikt} to equal 1 if customer i makes decision k at time t; 0 otherwise. The decisions are

$$(1)\qquad k = \begin{cases} 0 \Rightarrow \text{Not stay with this hotel} \\ 1 \Rightarrow \text{Paid stay in an economy property} \\ 2 \Rightarrow \text{Paid stay in a regular property} \\ 3 \Rightarrow \text{Paid stay in a luxury property} \\ 4 \Rightarrow \text{Paid stay in a luxury property and cash in points for a room upgrade} \\ 5 \Rightarrow \text{Free reward stay in an economy property via cash in points} \\ 6 \Rightarrow \text{Free reward stay in a regular property via cash in points} \\ 7 \Rightarrow \text{Free reward stay in a luxury property via cash in points} \end{cases}$$

Not staying in the hotel[2] or staying at one of their properties but not cashing in points (k=0, 1, 2, 3) are always options for the customer. Customers must accrue enough points to choose either an upgrade or a free stay (k=4, 5, 6, 7); there are no other forms of point redemption. Points "inventory" accrues as follows:

(2) $INV_{it+1} = INV_{it} + PointsAccrued_{it} - PointsCashedIn_{it}$, where

INV_{it+1} = Customer i's cumulative point inventory at the beginning of period $t+1$

$PointsAccrued_{it}$ = Points earned by customer i in period t [3]

$PointsCashedIn_{it}$ = Points cashed in for either an upgrade or a free stay by customer i in t

---

[1] We do not model the question of whether to "participate". This is because firms make it free, very easy and encourage customers to sign up for the reward programs. However, some loyalty programs may require enrollment fees and effort. Modeling this aspect is an avenue for future research.

[2] Note $k=0$ includes "do not stay anywhere" and "stay at another hotel." Similar to Gönül and Srinivasan (1996) and Lewis (2004), we treat these situations as "not stay at this hotel" because it is not possible to distinguish them in typical loyalty program datasets where data are available for only the focal firm. Competitive action is therefore not explicitly included in the model, but is made implicit in that customer need to stay at a hotel and the attractiveness of competitive hotels versus the focal hotel arise according to a random process governed by the error terms in Equation (5). These error terms capture a customer's need to stay at the focal chain and makes some alternatives more attractive in some weeks than in others. The alternative-specific constants thus reflect an average preference in an average competitor environment and an average need for product category.

[3] The hotel we study allowed customers to purchase a very limited number of points. However, we find no evidence of this in our data, perhaps due to the generally high level of points inventories carried by customers.

Denote R_1 as the inventory requirement for an upgrade, and R_2, R_3, and R_4 for a free stay at economy, regular, and luxury properties ($R_4 > R_3 > R_2 > R_1$). If customer i has enough points for an upgrade, then the choice set, K_{it}, would be 0,1,..,4. If she is eligible for a free stay, she can choose from K_{it} = 0,1,..,7 depending on her point inventory. In essence, we have an economic model of consumer choice under cutoffs (see Swait (2001) for a detailed discussion).

Equations 1-2 depict the customer's decisions regarding staying and cash-ins. Membership in a customer tier, s, is determined as follows:

(3) $$E_{ist} = \begin{cases} 1 \text{ } if \text{ } T_s \leq INVP_{it} < T_{s+1} \\ 0 \text{ } otherwise \end{cases},$$ where,

E_{ist} = Indicator variable signifying that customer i is in customer tier s at period t

T_s = Required number of paid stays for tier s

$INVP_{it}$ = Number of paid stays by customer i at period t during a calendar year

At the beginning of each calendar year, the inventory of paid stays is reset to zero and accrues during the year according to the following equation:

(4) $$INVP_{it+1} = INVP_{it} + PStay_{it},$$ where,

$PStay_{it}$ = Paid stay added to customer i's corresponding inventory in period t

2.3. Customer Utility

Following is the utility function (U_{ikt}) for customer i, decision k, at time t:

(5) $$U_{ikt} = \begin{cases} 0 + \varepsilon_{i0t} \text{ } for \text{ } k = 0 \\ \alpha_{ik} + \beta_{i1} P_{kt} + \delta_{i1} \sum_{l=4}^{7} D_{ilt-1} + \delta_{i2} \sum_{l=1}^{7} D_{ilt-1} + \sum_{s=1}^{2} \lambda_{is} E_{ist} + \varepsilon_{ikt} \text{ } for \text{ } k = 1,..,7 \end{cases}$$

We let $\overline{U}_{ikt}$ to denote the deterministic part of the utility and anchor the corresponding customer utility at zero if the customer decides not to stay in the hotel (k = 0). Utility for staying at the hotel (k = 1,2,..7) depends on the alternative-specific preference associated with the decision, the price of the decision, whether the customer has cashed in and/or stayed with the hotel in the previous period, or is a member of a customer tier. We have separate alternative-specific constants (α's) for k = 1, 2, 3, 4. For the free stay (k = 5, 6, 7), we have one alternative-specific constant, α_5.[4] We expect α_{i4} and α_{i5} to be positive because upgrades and free stays should generate positive utility. The second term considers the effect of price (P_{kt} = 0, for k=5,6,7) and we expect β_{i1}<0. The third term represents the rewarded behavior effect, i.e., the carryover from last period to current period from receiving a reward. A positive δ_{i1} suggests the customer is "delighted" with the reward (Rust and Oliver 2000) and now has more positive affect toward the hotel.[5] In fact, this is one way that a frequency reward program can truly enhance loyalty.

The fourth term in equation (5) represents past choice carry-over. A positive δ_{i2} could occur if on average, when staying at this hotel, customer i has a positive experience that carries over to the next period. Or it could simply mean the customer is inertial.[6] While both the rewarded behavior and past choice terms are first-order, they can exert a multi-period effect. This is because they change utility in period t, which alters the chance the customer will stay with the hotel in period t, which through the δ's again alters utility in period t + 1, etc.

The fifth term in equation (5) represents the contribution to utility from membership in a customer tier. If the customer is in tier s, the contribution to utility will be λ_{is}. We expect $\lambda_{i2} > \lambda_{i1} > 0$ because the Diamond tier bring more services than the Platinum tier, but even the Platinum

[4] Note that it is not necessary to model alternative-specific constants free stays since the incremental utility of a free stay at property j for customer i is $\alpha_{i5} - (\alpha_{ij}+\beta_{i1}P_{jt})$ and this increases with j from 5 to 7.
[5] We assume the same δ_{i1} for the rewarded behavior effect, irrespective of which reward was received last period. We do this for two reasons. First is to simplify the model. Second, this reflects the assumption that within the range of rewards we study (upgrade or free stay), there is an equal carryover effect. Essentially, we assume "it's the thought that counts" when it comes to generating positive carryover affect due to the frequency reward.
[6] Theoretically, δ_{i2} could also be negative. This could indicate variety seeking or a budget constraint. We thank an anonymous reviewer for the budget constraint interpretation of a negative δ_{i2}.

tier delivers additional services compared to usual. The last term in equation (5), ε_{ikt} represents unobserved time-specific determinants of customer *i*'s utility for making decision *k*. We assume these can be modeled as independent, identically distributed extreme value random variables. This treatment is consistent with prior research (Chib, Seetharaman, Strijnev 2004; Chintagunta 2002; Erdem, Imai, Keane 2003; Sun 2005).

2.4. Forward-looking Customers

Customers must accumulate points or paid stays to qualify for a frequency reward or attain a higher customer tier. This naturally requires customers to be forward-looking. For example, customers who value a high tier will consider that each stay adds to the likelihood they will achieve this goal. The same goes for customers who value the cash-in reward. Depending on the price, a forward-looking customer may sometimes face an inter-temporal tradeoff by sacrificing the current utility of cash-in for a free stay (which does not accrue points or paid stays) and choosing a paid stay, thus aiming toward increasing tier status. This decision process can be represented by a dynamic model, whereby the customer makes current period decisions that maximize her long-term utility. We express the customer's decision problem as:

$$(6) \qquad \underset{D_{ikt}, k \in K_{it}}{Max} E\left\{ \sum_{t=1}^{\infty} \sum_{k \in K_{it}} [U_{ikt} D_{ikt}] \delta^{t-1} \right\}$$

where δ is the per period discount factor. The operator $E[\cdot]$ (we suppress the subscripts, *i,t,* for expositional simplicity) is the conditional expectation given the customer's information at time *t*. The value function (Bellman equation), expressing the maximum expected utility the customer can expect, given he or she is in period *t*, is given by:

$$(7)\ V_{it}(INV_{it}, INVP_{it}, D_{ikt-1}, \varepsilon_{ikt}) = \underset{D_{ikt}, k \in K_{it}}{Max} \{U_{ikt} + \delta E[V_{i,t+1}(INV_{i,t+1}, INVP_{it+1}, D_{ikt}, \varepsilon_{ikt+1})]\} = \underset{D_{ikt}, k \in K_{it}}{Max} V_{ikt}$$

The choice-specific value function, $V_{ikt} = U_{ikt} + \delta E[V_{i,t+1}(INV_{i,t+1}, INVP_{it+1}, D_{ikt}, \varepsilon_{ikt+1})]$, is the maximum long-term utility the customer can expect if he or she chooses decision *k*. The state variables are: inventory of total points (INV_{it}) given by equation (2), inventory of paid stays ($INVP_{it}$) given by equation (4), and past decisions, D_{ikt-1}.

2.5. Customer Expectation of Future Demand and Prices

In our model setup, the decision to stay in the focal hotel chain is an endogenous decision that is derived from the dynamic programming problem. However, when evaluating expected future utility, customers need to incorporate their expectations regarding (i) future prices at different types of properties, (ii) length of stay with the hotel, and (3) future tier status; these, in turn would determine their respective expected future points. We assume that a representative consumer i has rational expectations for future prices at each property type h, denoted p_{ih}, and the length of stay, denoted N_{ik}: more specifically, $p_{ih} \sim LTN(\bar{p}_{ih}, \sigma^2_{p_{ih}}), h = 1,...3.$ and $N_i \sim LTN(\bar{N}_i, \sigma^2_{N_i})$. We determine the respective means and standard deviations of these distributions from customer i's specific transaction history with the company. These distributions are then used in computing future value function, where we draw from corresponding normal distributions with the estimated means and variances (See Appendix for further details).

2.6. Heterogeneity and Estimation

We adopt latent class segmentation to incorporate heterogeneity (Kamakura and Russell 1989). Let $\omega_c = \{\alpha_{jc}, \beta_{1c}, \delta_{1c}, \delta_{2c}, \lambda_{sc}, j = 1,..,5; s = 1,2\}$ be the vector of coefficients to be estimated for each latent class *c*, and $\omega = \{\omega_1 ... \omega_C\}$. Then, the likelihood function of the sequence of choices of all customers is:

(8) $$L(\omega)=\sum_{c=1}^{C}\prod_{i=1}^{I}\prod_{t=1}^{T}\sum_{k=1}^{K}[D_{ikt}\bullet \Pr(D_{ikt}=1\,|\,INV_{it},INVP_{it},D_{ikt-1},\omega_c,i\in c)\bullet \Pr(i\in c)]$$

where c indexes latent classes and $Pr(i \in c)$ is the probability customer i is in latent class c. We estimate $Pr(i \in c)$, representing the percentage of customers in each segment. Given the extreme value distribution of the error term, the probability of observing customer i in latent class c making decision k at time t is:

(9) $$\Pr(D_{ikt}=1\,|\,INV_{it},INVP_{it},D_{ikt-1},\omega_c,i\in c)=\frac{e^{V_{ikt}}}{\sum_{j\in K_{it}} e^{V_{ijt}}}$$

Given the complexity of the dynamic programming problem, we adopt interpolation (Keane and Wolpin 1994) to estimate the model.

3. Data

Our data for the hotel application are from the loyalty program of a major hotel chain, covering 3907 customers over two years (1/1/02-12/31/03). Included are the dates of stay by customer, type of property stayed in (economy, regular, or luxury), price paid, number of points earned etc. Each trip includes a code indicating whether a reward was redeemed on that trip, and if so how many points were used. The database also provides the beginning inventory of points for each customer, and his or her tier status. The chain's policy does not allow points to be accumulated from other sources and our data includes all point transactions. An analysis of our data revealed that all customer stays involved hotel point accumulation, i.e., there were no occurrences where a customer paid for a stay at the hotel and no points were awarded to her account. Furthermore, point redemption from a customer's account at any time was always associated with that customer staying with the chain in that period. There were no cases where significant hotel points, unrelated to a paid stay, were added to a customer's account.

The hotel's frequency reward program involves 4 categories of frequency reward: free stay at an economy (5000 points), regular (8000 points), or luxury (12000 points) property, and a free upgrade at a luxury property (3000 points), where five points are earned per dollar spent, while Platinum and Diamond tier members respectively earn an additional 15% and 30% per dollar spent. For the customer tier, the requirements for number of paid stays for Platinum and Diamond tiers are 5 and 12 respectively. Customers assigned to higher tiers in a calendar year *YR* (based on paid stays in year *YR-1*) may be reassigned to a lower tier in year *YR+1*, if the number of *paid* stays with the hotel in year *YR* do not meet the requirements for the higher tiers. During a year, the average customer paid $151.00 per night, accumulated approximately 2448 points per stay, stayed 3.2 times, with 12.1% of stays at an economy property and 43.9% and 44% stays at regular and luxury properties. About 22.4% and 2.53% qualify for the Platinum and Diamond tiers respectively.

For estimating the model, we chose a random sample of 200 customers. We organized the data into biweekly periods because each period was then long enough to be a decision period yet short enough that multiple stays did not occur in a given period. Thus, there were 51 decision periods for each customer during the two-year time frame. About 98% of customers cashed in at least once during the two-year period; 72% cashed in exactly once. About 75% of customers had an initial points "inventory" of at least 10,000 points.

4. Results

4.1. Fit and Holdout Prediction

We estimated four models: (1) myopic model with two latent segments; (2) dynamic model with one segment; (3) dynamic model with two segments; and (4) dynamic model with three segments. The discount factor, δ, in Equation (7) is set to 0 in the myopic model while it is set at a weekly equivalent of .995 in the dynamic model (Lewis 2004; Sun 2005). Fit statistics for the four models are provided in Table 1. The two-segment dynamic model provides the best fit according to AIC and BIC.

[insert Table 1 about here]

The superiority of the two-segment dynamic model is supported by its performance in a holdout sample. We applied the model to 200 customers not used in the estimation. To assign holdout customers to segments, we use the estimated proportions of the two segments as our priors for segment membership (72.7% for the first segment and 27.3% for the second segment). We then update our belief for each customer by incorporating the likelihood of his/her decision sequence. We assign customers with posteriors greater than 72.7% to the second segment and others to the first segment. The corresponding fit statistics are shown in Table 1. The two-segment dynamic model again performs best.

4.2. The Two-Segment Solution

Table 1 also reports the mean parameter estimates and their standard errors for the two-segment model. We find: Segment 1 is more price sensitive than Segment 2. Interestingly, Segment 1 derives significant positive utility from the free-stay component of the frequency reward program and shows no significant preference for an upgrade while Segment 2 prefers the upgrade and no significant preference for a free stay. Segment 2's preference for a luxury property is positive and significant while for Segment 1, it is not significantly different from 0. The negative intercepts for an economy or a regular property for both segments suggests that the members don't have a positive affinity for these properties, all else equal. Members of Segment 2 have significant positive coefficients for both levels of the customer tier program, while the coefficient for only the Diamond tier is significant for Segment 1; however, this segment does have lower, but positive valuation for the Platinum tier.

Given the above observations, we label the first segment "Price-Oriented" in that they are more price sensitive and are more attracted to free stays. We label the second segment "Service-Oriented", since they derive positive value for both customer tiers, favor upgrades, and prefer a luxury hotel. The "price-oriented" segment is in the majority, 72.7%, of the sample, while the "service-oriented" segment is 27.3%. Further, there is support for the rewarded behavior effect for the price-oriented segment; the carryover coefficient for cash-in (δ_{i2}) is positive. There is not a significant carry-over for staying in both segments; it appears staying at the hotel in period t does not tend to increase the likelihood of staying with the hotel in period $t + 1$.

We provide further information on the "price-oriented" versus "service-oriented" segments in a holdout sample as another test (Table 2). The table shows, as expected, that the "service-oriented focused" segment stays more frequently, pays a higher price, stays at a luxury hotel, and often qualify for Diamond and Platinum tiers. Interestingly, the price-oriented segment is more likely to stay over the weekend, indicating they are more like a leisure segment.

[Insert Table 2 About Here]

4.3. Points Pressure and Rewarded Behavior Effects

Points Pressure. Points pressure refers to the build-up in purchase frequency as customers get closer to a reward. We find points pressure effects for both the frequency reward and customer tier components of the loyalty program. This is demonstrated in Figures 1a-1c and 2a-2c, where we use simulations of the model as well as the raw data to calculate the probability of a paid stay (i.e., choosing k=1, 2, 3, or 4) as the customer gets closer to earning the reward.

[Insert Figures 1a-1c Here]

[Insert Figures 2a-2c Here]

In Figure 1, we assess the points pressure effect with respect to the customer tier program. To isolate the tier program impact, we run the simulation assuming there is no frequency reward component. Figures 1a and 1b show that both segments increase their likelihood of paid stays as they get closer to earning membership in an average tier. The paid stay probability is the average probability across customers and across the two customer tier levels.

These results are driven by the positive coefficient ($\lambda > 0$) we calculate for customer tier membership. While tier classification is based on *previous* purchase levels, an alternative explanation for the strength of the tier effect is that tier membership is a proxy for high usage. However, Figure 1c shows an "out-of-model" demonstration of the points pressure effect. To create this figure, we simply use actual data to plot the percent of paid stays on the y-axis and the closeness to the next tier level on the x-axis. The points pressure effect is apparent. This shows that usage level is not constant, suggesting that tier effects are real and not a reflection of a usage pattern. Figures 2a-2b depict the points pressure effect for the frequency reward component (assuming no customer tier program). Per the parameter estimates, we may conclude that the points pressure for the "price-oriented" segment are driven by the free cash-in stays, while upgrades drive it for the "service-oriented" segment. Again, Figure 2c, based on actual data, is an out-of-model demonstration of points pressure due to the frequency reward program.

Rewarded Behavior. Rewarded behavior refers to positive carryover after a frequency reward is cashed in. Our goal here is to examine customer behavior prior to and after cashing in. Since different customers cash in at different points in time, we rescale periods such that period 0 refers to the period of cashing in and a period of -5 (+5) represents five periods prior to (after) cashing in. Figures 3a-3b are based on simulations assuming no customer tier program.

[Insert Figures 3a-3b Here]

Figure 3a illustrates the case of the “price-oriented” segment. The probability of staying, averaged across various reward levels, increases during the cash-in period (0) because the reward has positive utility. Then in the period after the reward is cashed in, we see that likelihood of staying is still elevated. This is the positive carryover from a reward. Consistent with points pressure effect, the stay probability is higher in several periods prior to the cash-in period.

Figure 3b shows the case of the “customer-tier focused” segment. There is an increase in staying probability during the cash-in period because the immediate utility of the reward is positive. However, once the customer has stayed for free, there is no rewarded behavior effect in the next period. This is because δ_l is not significant (see Table 1).

5. Policy Implications

5.1. Overall Impact of Loyalty Programs on Customer Behavior and Firm Revenues

We simulate the model to investigate the overall impact of the frequency reward and customer tier components of the loyalty program. We use the parameters in Table 1 to simulate four scenarios and report the simulated stays and the revenue: (1) the firm employs neither the frequency reward nor the customer tier component; (2) the firm employs both frequency reward and customer tier components; (3) the firm employs the frequency reward but not the customer tier component; (4) the firm employs the customer tier but not the frequency reward component. For each scenario, we simulate customer behavior for a year and record the steady state number of stays and revenue per customer. The number of stays without cashing in is calculated as the average, over customers and time periods, of the probability of choosing k = 1, 2, or 3, i.e., stay without cashing in. Similarly, we calculate the expected number of upgrades (*k*=4) and free stays (*k*=5, 6, or 7). The results are presented in Table 3 and yield several conclusions:

[Insert Table 3 Here]

First, without any program, "price-oriented" customers are less likely to stay with this hotel than the "service-oriented focused" customers (1.21 annual stays per customer versus 2.56). This is due to the lower alternative-specific constants for the price-oriented segment. Second, both segments respond to the customer tier program. This can be seen by comparing the base case scenario to the customer tier only scenario. For the price oriented segment, the simulated number of total stays per customer annually increases from 1.21 to 1.25. For the service-oriented segment, the number of stays increases from 2.56 to 2.59. This is because both segments have positive coefficients for the customer tier program and a positive points pressure effect as customers endeavor to build up points toward a higher tier. Third, both segments respond to the frequency reward program but for different reasons, while the price-oriented segment accumulates cash-in free stays, the service-oriented segment is motivated by cash-in upgrades. This can be seen clearly in the frequency reward only scenario. For the price- and service-oriented segments, the number of upgrade stays per year are 0.03 and .51 respectively, while the respective number of free says are 0.95 and 0.85. It is interesting that for the service-oriented segment, even though the utility for free stays is lower than that of upgrades, we still see a noticeable number of free stays. This is because for an upgrade to happen, the utility for an upgrade has to overcome the negative price effect of a purchase, while no such price effect exists for a cash-in free stay.

Fourth, while the frequency reward program generates free cash-ins, the total number of *paid* stays –stays with no cash-in plus stays with upgrades – increase. Focusing on the price-oriented segment, under the frequency reward only scenario, the annual number of stays per customer with no cash-in (from 1.21 to 1.24) and the number of cash-in upgrades (to 0.03) increase by a total of 0.06. The corresponding number of paid stays going up for the service-oriented segment is 0.56.

Fifth, very interestingly, there appears to be some synergy between the two program components for both segments. In the first segment, the number of paid stays increase by 0.06 (1.27-1.21) when the frequency program is offered alone, and by 0.04 (1.25-1.21) when the customer tier program is offered alone. If there were no cannibalization, with both programs the total number of paid stays would be 1.21+0.06+0.04 = 1.31, lower than the total number of paid stays when both programs are offered (1.34 = 1.31+0.03). The corresponding numbers for the second segment are 3.15 and 3.16. Finally, in terms of the number of paid stays, the firm is better off including both components of the program. This is further confirmed via the simulated revenue column where we take into consideration the distribution of length of stay.

5.2. Frequency Reward and Customer Tier Program Requirements to Maximize Revenues

Paid stays can be maximized by setting program requirements correctly. For example, if frequency reward requirements are lax, there will be minimal points pressure, but many cash-ins and hence ample rewarded behavior. If the requirements are tough, there will be more points pressure, but not many cash-ins so less rewarded behavior. The optimal design will strike the right balance. The customer tier program presents similar challenges. If the requirements are too lax, there is not much points pressure, but once the customer reaches a tier, paid stays increase because the customer receives continuously better service. If the requirements are too tough, we get more points pressure, but it takes customers longer to reach a tier, and many may not make it.

[Insert Table 4 About Here]

We conduct simulations to determine the frequency reward and customer tier requirements that would improve revenue per customer. We varied the free upgrade requirement from 1,000 points to 7,000, free stay requirement at economy, regular, and luxury properties from 3,000, 5,000, and 7,000 points respectively to 15,000 points, all in steps of 1,000 and with the constraint that the requirement for a higher class property is higher than for a lower class property. Furthermore, we also varied the paid stay requirement for Platinum and Diamond tiers from 1 and 4 respectively to 20 in steps of 1, requiring that the threshold for Diamond tier be higher than for the Platinum tier. In each cell, we simulated the behavior of 200 consumers. Table 4 presents the results for a subset of the design cells.

We found that while lower requirements encouraged more staying, it reduced revenue, primarily due to free cash-in stays. For example, when the point requirements for an upgrade and the three types of properties were 1000, 3000, 5000, and 7000 respectively, and the Platinum and Diamond tier requirements were 1 and 4, the annual revenue per customer was only $305.07; on the other hand, the revenue increases to $560.00 for the current requirements of 3,000, 5,000, 8,000, and 12,000 for the reward program and 5 and 12 paid stays for the customer tier. The revenue-maximization setting (with an average annual revenue of $734.14) is 6,000, 8,000, 10,000, and 13,000 for the reward program and 3 and 14 annual paid stays for the customer tier program. This is a bit more stringent for the reward program but takes advantage of the points pressure effect, somewhat less stringent for the Platinum tier (due to the lukewarm response from the price-oriented segment), and slightly more stringent for the Diamond tier (due to significant positive utility for segments). Of course our analysis might not maximize profits. The costs of providing extended services may be convex in the number of customers in a given tier. In other words, the cost to a hotel of providing personal service calls and faster check-ins may rise exponentially as a function of the number of customers promised this service.

6. Replication

We also applied our model to the loyalty program instituted by a major airline. With the airline, there are two classes of frequency rewards: a free upgrade or a free flight. Mileage is accrued via paid flights as well as through tie-in programs such as credit cards, and no mileage is earned on free flights. The customer tier component placed customers in Silver, Gold, and Platinum tiers depending on how many paid miles they had flown during the past calendar year.

Our data are from a major U.S. airline, covering 1,707 frequent travelers over two years (05/98-04/99). For estimating the model, we chose a random sample of 200 customers and organized the data into biweekly periods as before. About 90% of cash-ins in our data involved either an upgrade, requiring between 15,000 and 35,000 miles with a mode of 25,000 miles, or a free flight, requiring between 40,000 and 60,000 miles with a mode of 50,000 miles. Thus we set R_1 and R_2 to their most frequent values, 25,000 and 50,000 miles respectively. For the customer tiers, the mileage requirements were 25,000, 50,000, and 75,000.

Despite the unavailability of price information in this data, our analysis confirmed several results from the hotel application: (1) the outperformance of the forward-looking model in both calibration and holdout data, (2) existence of two distinct segments – in this case a "customer-tier focused" segment that values the customer tier component and finds the frequency program reward not worth the effort, and a "loyalty program enthusiasts" segment that values both components; The "customer-tier focused" segment is the vast majority in our sample. We conjecture this is due to the customers in our sample all being frequent fliers – averaging 0.71 flights per two-week period, or about 18 flights per year with the airline. More flights, even for free, are not very attractive to such people. Plus, for business travelers, the hassle costs of cashing in and finding a non-blacked-out flight are very high. (3) Both segments value the customer tier program and demonstrate points pressure in their efforts to qualify for a higher tier; (4) For the reward program, points pressure exists for the "loyalty program enthusiasts" segment, but not for the "customer-tier focused" segment. This is because the "customer-tier focused" segment does not value the reward. (5) There is a rewarded behavior effect with respect to the frequency reward for the smaller segment. Finally, there is no synergy between the two program components: the number upgrade cash-ins increases under the frequency reward program, but this is off-set by fewer, no cash-in paid flights.

7. Summary And Discussion

We evaluate the joint sales impact of the frequency reward and customer tier components of loyalty programs. We provide an integrated analysis and measurement of a loyalty program incorporating customers' purchase and cash-in decisions, points pressure and rewarded behavior effects, heterogeneity, and forward-looking behavior. The model captures forward-looking behavior, in that customers make current decisions of whether to stay or cash-in after accounting for the impact on their points inventory, which brings them closer to receiving a reward or maintaining their tier status. We model the customer's cash-in decision for free stays or upgrades, taking into account the immediate utility of these rewards as well as the effect of price and the long-term implications due to rewarded behavior and future inventory levels.

The results are quite interesting. We find as expected that customers are quite heterogeneous with respect to their response to these programs. We find a "price-oriented" segment that values the frequency reward program and the highest customer tier, and a "service-oriented" segment that finds the free stays not significant, yet highly values the upgrades and the customer tier program. The “price-oriented” segment is the majority in our sample. Given the leisure nature of this segment, free stays are quite attractive to such people. On the other hand, for business travelers (“service-oriented” segment), the upgrades, luxury type hotels, and the privileges associated with the customer tier are appealing.

We find that customers value the benefits of higher customer tiers more than lower tiers. This is more of a validity test than a finding, since one would expect customers to prefer the more services provided by higher tiers. We also find significant points pressure effects, whereby customers increase purchase likelihood as they get closer to a reward, and rewarded behavior effects, whereby cashing in the ephemeral frequency reward in the form of a free stay or an upgrade has a positive impact on purchase likelihood in the subsequent period.

Both loyalty program components increase the number of paid stays. The customer tier results in higher number of paid stays, a consequence of points pressure, encouraging customers to rise to a higher tier, and delivering positive utility once the customer gets to that tier. The frequency reward impact is more complex as it widens the customer’s choice set to include upgrades and free stays.

From a managerial perspective, our results have several implications: Loyalty programs consist of two distinct components – customer tier and frequency reward. The customer tier component comes closer to actually increasing loyalty (Shugan 2005), because it continually provides utility once the customer has attained a high tier. However the frequency reward component has at least a temporary impact on loyalty through the rewarded behavior effect. There also is a points pressure effect as customers increase their purchase rate to attain a higher tier or earn a reward. At the outset, the frequency reward program may appear unsuccessful because the number of free cash-in stays increase. However, the total number of paid stays also increases. To the extent a hotel has excess capacity, a loyalty program would be profitable since it expands revenue. The customer-tier component of the loyalty program is also successful in increasing the overall paid stays. Consistent with our joint revenue maximization results, it is interesting to observe that firms have recently been moving toward tougher requirements for frequency rewards (e.g., through the "inflation" in award points, *New York Times*, December 24, 2006; *Wall Street Journal*, December 5, 2006) while emphasizing the benefits of a customer tier programs (*New York Times*, December 26, 2006).

For researchers, the implications are: (1) It is feasible and advisable to simultaneously model both the frequency reward and customer tier components of loyalty programs. The effects of these programs are quite different so researchers need to make the distinction between the two when examining loyalty programs. (2) We find further evidence for points pressure and rewarded behavior in both the hotel and airline applications. Researchers need to recognize that both these phenomena can increase "loyalty" and incorporate them in future analyses of loyalty programs. (3) We demonstrate that a model with forward-looking consumers and endogenous cash-in decisions better approximate consumer reactions to loyalty programs.

There are several avenues for extending this research. First, our model may be applied to other types of loyalty programs to generalize the substantive results. Second, an ideal model should include competition, thereby expanding the customer's choice set. This is a limitation of our study. We could not incorporate competitive effects because as in most database marketing situations, we only had data on customer behavior with our client company. Third, future work should delve into the factors that decrease perceived costs and increase perceived value of the frequency reward component. For example, our result in the airline data of a negative net utility for cashing in, coupled with a rewarded behavior effect, requires further investigation into the psychological processes underlying this interesting finding. See Kivetz and Simonson (2002, 2003), Kwong and Soman (2006), and Roehm, Pullins, and Roehm (2002) for insights on the value side of the equation. Fourth, we need to understand better when firms should offer free product versus other types of rewards (e.g., Kim, Shi, and Srinivasan 2001) or combine currencies for cash-in (e.g., Drèze and Nunes 2004). In conclusion, while we believe our research adds both to the empirical knowledge base of the impact of loyalty programs, and to the methodological “tool kit” for analyzing these programs, there is still more work to be done to understand fully this ever-popular marketing activity.

REFERENCES

Bolton, R. N., P.K. Kannan, M. D. Bramlett. 2000. Implications of Loyalty Program Membership and Service Experiences for Customer Retention and Value. *Journal of the Academy of Marketing Science* 28 (Winter) 95-108.

Blattberg, R. C., B. Kim, S. A. Neslin. 2008. *Database Marketing: Analyzing and Managing Customers*, New York, Springer.

Blattberg, R. C., J. A. Deighton. 1996. Manage Marketing by the Customer Equity Test. *Harvard Business Review.* 74 (July-August) 136-44.

Butscher, S. A. 1998. Limited Loyalty Programs Create Strategic Databases. *Marketing News.* 31 (October 27) 13.

Chib, S., P. B. Seetharaman, A. Strijnev. 2004. Model of Brand Choice with a No-Purchase Option Calibrated to Scanner-Panel Data. *Journal of Marketing Research*. 41 (May) 184-196.

Chintagunta, P. 2002. Investigating Category Pricing Behavior in a Retail Chain. *Journal of Marketing Research.* 39 (May) 141-154.

Dowling, G. R., M. Uncles. 1997. Do Customer Loyalty Programs Really Work? *Sloan Management Review.* 38 (Summer) 71-82.

Deighton, J. A. 2000. Frequency Programs in Service Industries. *Handbook of Services Marketing and Management*, Dawn Iacobucci and Teresa A. Swartz (eds.), Thousand Oaks, Sage Publications, 401-408.

Drèze, X., J. C. Nunes. 2004. Using Combined-Currency Prices to Lower Consumers' Perceived Cost. *Journal of Marketing Research.* 41 (February) 59-72.

Erdem, T., S. Imai, M. P. Keane. 2003. .Brand and Quantity Choice Dynamics under Price Uncertainty. *Quantitative Marketing and Economics.* 1 (March). 5-64.

———, M. P. Keane. 1996. .Decision-making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets. *Marketing Science*. 15 (1) 1-20.

Gönül, F., K. Srinivasan. 1996. .Estimating the Impact of Consumer Expectations of Coupons on Purchase Behavior: A Dynamic Structural Model. *Marketing Science.* 15 (3) 262-279.

Gupta, S., D. R. Lehmann, J. A. Stuart. 2004. Valuing Customers. *Journal of Marketing Research.* 41 (February) 7-18.

Hartmann, W. R., V. B. Viard. 2006. Do Frequency Reward Programs Create Switching Costs? working paper, Graduate School of Business, Stanford University.

Kamakura, W. A., G. J. Russell. 1989. .A Probabilistic Choice Model for Market Segmentation and Elasticity Structure. *Journal of Marketing Research.* 26 (November) 379-390.

Keane, M. P., K. I. Wolpin. 1994. .The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence. *Review of Economics and Statistics.* 76 (4) 648-672.

Keiningham, T. L., T. G. Vavra, L. Aksoy. 2006. Managing through Rose-Colored Glasses. *Sloan Management Review.* 48 (Fall) 15-18.

Kim, B-D, M. Shi, K. Srinivasan. 2001. Reward Programs and Tacit Price Collusion. *Marketing Science.* 18 (3) 59-76.

Kivetz, R., I. Simonson. 2002. Earning the Right to Indulge: Effort as a Determinant of Customer Preferences Toward Frequency Program Rewards. *Journal of Marketing Research* 39 (2) 155-170.

———, I. Simonson. 2003. .The Idiosyncratic Fit Heuristic: Effort Advantage as a Determinant of Consumer Response to Loyalty Programs. *Journal of Marketing Research.* 40 (May) 454-467.

———, O. Urminsky, Y. Zheng. 2006. The Goal-Gradient Hypothesis Resurrected: Purchase Acceleration, Illusionary Goal Progress, and Customer Retention. *Journal of Marketing Research.* 43 (February) 39-58.

Klemperer, P. D. 1987. Markets with Consumer Switching Costs. *Quarterly Journal of Economics.* 102 (May) 375-394.

Kopalle, P., S. A. Neslin 2003. The Economic Viability of Frequency Reward Programs in a Strategic Competitive Environment. *Review of Marketing Science.* 1 (1) (accessed November 21, 2006) [available at http://www.bepress.com/romsjournal/vol1/iss1/].

Kumar, N., R. Rao 2003. Effect of Loyalty Card Programs on Supermarket Prices. working paper, School of Management, University of Texas at Dallas.

Kwong, J. Y., D. Soman. 2006. .The Role of Fluency in the Decision to Spend Loyalty Program Points. working paper, University of Toronto.

Lal, R., D. E. Bell. 2003. The Impact of Frequent Shopper Programs in Grocery Retailing. *Quantitative Marketing and Economics.* 1 (2). 179-202.

Lewis, M. 2004. .The Influence of Loyalty Programs and Short-Term Promotions on Customer Retention. *Journal of Marketing Research.* 41 (August). 281-292.

Liu, Y. 2007. The Long-Term Impact of Loyalty Programs on Consumer Purchase Behavior and Loyalty. *Journal of Marketing.* 71 (4) 19-35.

New York Times. 2006. I'm Elite and You're Not. December 26.

New York Times. 2006. Hotel Reward Programs: That 'free' Hotel Stay Gets Less Free. December 24.

Nunes, J. C., X. Drèze 2006. The Endowed Progress Effect: How Artificial Advancement Increases Effort. *Journal of Consumer Research.* 32 (March) 504-512.

Peppers, D., M. Rogers. 1997. *Enterprise One to One*. New York: Doubleday.

Reinartz, W, J. S. Thomas, V. Kumar. 2005. Balancing Acquisition and Retention Resources to Maximize Customer Profitability. *Journal of Marketing.* 69 (January) 63-79.

Reynolds, J. 1995. Database Marketing and Customer Loyalty: Examining the Evidence. *European Retail Digest*. Oxford Institute of Retail Management. Issue 7 (July) 31-38.

Roehm, M. L., E. B. Pullins, H. A. Roehm Jr. 2002. Designing Loyalty-Building Programs for Packaged Goods Brands. *Journal of Marketing Research.* 39 (May) 202-213.

Rust, R. T., R. L. Oliver. 2000. Should We Delight the Customer? *Journal of the Academy of Marketing Science.* 28 (Winter). 86-94.

Sharp, B., A. Sharp. 1997. .Loyalty Programs and Their Impact on Repeat-Purchase Loyalty Patterns. *International Journal of Research in Marketing.* 14 (5) 473-486.

Shugan, S. M. 2005. Brand Loyalty Programs: Are They Shams? *Marketing Science.* 24 (2) 185-193.

Sun, B. 2005. Promotion Effect on Endogenous Consumption. *Marketing Science.* 24 (3) 430-443.

Sun, B., S. A. Neslin, K. Srinivasan. 2003. Measuring the Impact of Promotions on Brand Switching under Rational Consumer Behavior. *Journal of Marketing Research.* 40 (November) 389-405.

Swait, J. 2001. A Non-Compensatory Choice Model Incorporating Attribute Cutoffs. *Transportation Research.* 35 (Part B) 903-928.

Taylor, G. A., S. A. Neslin. 2005. The Current and Future Impact of a Retail Frequency Reward Program. *Journal of Retailing.* 81 (4) 293-205.

van Heerde, H. J., M. G. Dekimpe, W. P. Putsis Jr. 2005. Marketing Models and the Lucas Critique. *Journal of Marketing Research.* 42 (February) 15-21.

Venkatesan, R., V. Kumar. 2004. Customer Lifetime Value Framework for Customer Selection and Resource Allocation Strategy. *Journal of Marketing.* 68 (October) 106-125.

Wall Street Journal. 2006. Making Use of Frequent-Flier Miles Gets Harder. December 5, D5.

Zeithaml, V. A., R. T. Rust, K. N. Lemon 2001. The Customer Pyramid: Creating and Serving Profitable Customers. *California Management Review* 43 (Summer) 118-142.

Table 1. Model Fit Statistics and Parameter Estimates

	Estimation Sample		**Holdout Sample**	
Number of Observations	10,200		10,200	
Myopic Model with 2 segments: LL	-6617.87		-7283.32	
AIC	13277.74		14608.64	
BIC	13429.57		14760.47	
Dynamic Model with 1 segments: LL	-7518.04		-9693.90	
AIC	15098.08		19409.80	
BIC	15322.21		19489.33	
Dynamic Model with 2 segments: LL	**-5687.17**		**-6388.21**	
AIC	**11416.34**		**12818.42**	
BIC	**11568.17**		**12970.25**	
Dynamic Model with 3 segments: LL	-5683.92		-6380.42	
AIC	11429.84		12824.84	
BIC	11653.97		13056.20	
Parameter Estimates (Dynamic Model with 2 segments)				
	Segment 1 (72.7%) "Price-Oriented"		**Segment 2 (27.3%) "Service-Oriented"**	
	Parameter Estimates	Standard Errors	Parameter Estimates	Standard Errors
Intercept, Economy hotel (base utility, α_1)	-4.039**	0.269	-3.421**	0.276
Intercept, Regular hotel (base utility, α_2)	-3.346**	0.242	-1.416**	0.166
Intercept, Luxury hotel (base utility, α_3)	-10.756	6.479	0.093*	0.0379
Price (β_1)	-0.102**	0.007	-0.029**	0.002
Utility of upgrade to a suite (α_4)	0.045	0.682	0.226**	0.064
Utility of cash-in for free stay (α_5)	0.032**	0.006	0.007	0.057
Utility of Platinum tier (λ_1)	0.074	0.064	.029**	0.006
Utility of Diamond tier (λ_2)	0.283**	0.012	0.254**	0.062
Inertia of cash-in (rewarded behavior, δ_1)	0.012*	0.006	0.000	0.006
Inertia of staying (state dependence, δ_2)	0.000	0.064	0.002	0.007

** $p < .01$, * $p < .05$

Table 2

Characteristics of Segments in a hold-out sample of 800 customers (.727 cut-off value)

Summary Characteristics of Two Segments		
Characteristics	**Segment 1: "Price-Oriented"**	**Segment 2: "Service-Oriented"**
	Mean (Std. Error)	Mean (Std. Error)
Average price per night	$87.0 (2.29)	$158.90 (3.28)
Average paid stays	1.20 (0.31)	3.12 (0.83)
Average percent stays at economy hotel	30.4% (1.9%)	9.79% (2.0%)
Average percent stays at Regular hotel	63.6% (2.0%)	41.5% (3.3%)
Average percent stays at Luxury hotel	5.99% (1.0%)	48.7% (3.4%)
Average percent weeknight stays (Sun-Thu)	37.2% (2.0%)	58.7% (3.3%)
Average per period reward cash-ins	1.95% (0.6%)	2.10% (1.0%)
Percent qualifying for Platinum tier	13.4% (1.4%)	25.5% (3.0%)
Percent qualifying for Diamond tier	1.48% (0.5%)	2.89% (1.1%)

Table 3. Annual Impact of Frequency Reward and Customer Tier Components Per Customer

	"Price-Oriented" Segment				
	Simulated Revenue	**Simulated # of Total Stays**	**Simulated # of Stays With No Cash-in (k = 1,2,3)**	**Simulated # of Cash-In Upgrades (k = 4)**	**Simulated # of Cash-in Free Stays (k = 5,6,7)**
1. Neither Component	$251	1.21	1.21	0.00	0.00
2. Both	$278	2.31	1.31	0.03	0.97
3.Frequency Reward Component Only	$196	2.22	1.24	0.03	0.95
4. Customer Tier Component Only	$259	1.25	1.25	0.00	0.00

	"Service-Oriented" Segment				
	Simulated Revenue	**Simulated # of Total Stays**	**Simulated # of Stays With No Cash-in (k = 1,2,3)**	**Simulated # of Cash-In Upgrades (k = 4)**	**Simulated # of Cash-in Free Stays (k = 5,6,7)**
1. Neither Component	$1066	2.56	2.56	0.00	0.00
2. Both	$1306	4.01	2.64	0.52	0.85
3. Frequency Reward Component Only	$1294	3.97	2.61	0.51	0.85
4. Customer Tier Component Only	$1074	2.59	2.59	0.00	0.00

Table 4. Revenue Maximizing Loyalty Program Design (a Subset of the Design Cells)

Design	Points for Upgrade	Free stay points: Economy hotel	Free stay points: Regular hotel	Free stay points: Luxury hotel	Number of Stays: Platinum Tier	Number of Stays: Diamond Tier	Simulated Annual revenue ($) per customer
	1000	3000	5000	7000	1	4	305.07
	1000	4000	6000	13000	7	16	358.08
	1000	5000	12000	14000	3	10	335.81
	1000	6000	8000	11000	4	7	429.84
	2000	4000	10000	13000	4	11	459.69
Current	**3000**	**5000**	**8000**	**12000**	**5**	**12**	**560.00**
	3000	10000	12000	14000	4	8	557.04
	5000	8000	11000	15000	3	15	677.73
	5000	9000	13000	15000	3	11	707.29
	6000	8000	10000	12000	5	8	705.35
	6000	8000	10000	12000	8	18	680.72
Optimal	**6000**	**8000**	**10000**	**13000**	**3**	**14**	**734.14**
	6000	8000	10000	15000	3	10	682.83
	6000	8000	11000	13000	8	13	699.15
	6000	9000	11000	14000	7	19	682.26
	6000	9000	11000	15000	7	12	690.69
	6000	9000	12000	15000	4	14	682.87
	6000	9000	13000	15000	7	14	690.16
	6000	11000	13000	15000	8	16	684.57
	7000	9000	11000	13000	5	17	673.84
	7000	9000	11000	14000	4	7	724.83

FIGURE 1. Demonstrating Points Pressure Effect for Customer Tier: The Relationship Between Closeness to Customer Tier and Predicted Paid Stay Probability

a. Price-Oriented Segment

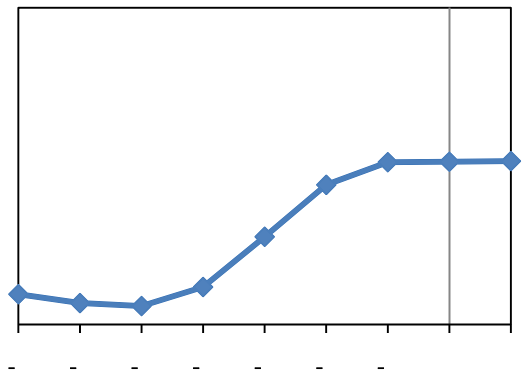

b. Service-Oriented Segment

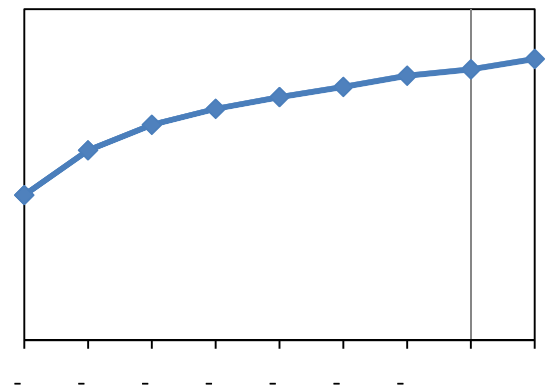

c. Closeness to Tiers and Paid Stays (Based on Actual Data)

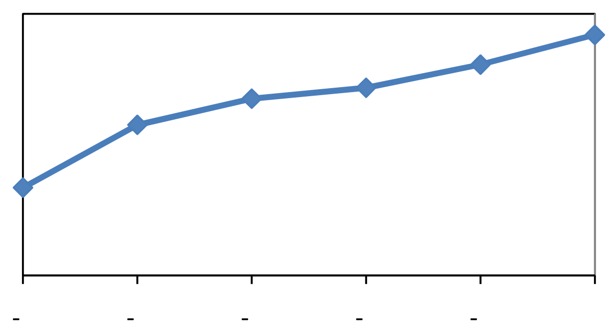

Note: Figures 1a-1b are based on simulations assuming no frequency reward program; Figure 1c is based on actual data

FIGURE 2. Demonstrating Points Pressure Effect for Frequency Reward: The Relationship Between Closeness to Points Required for a Reward and Predicted Paid Stay Probability

a. Price-Oriented Segment

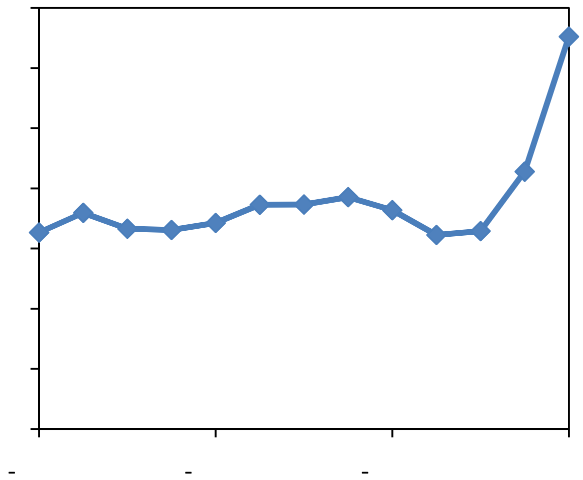

b. Service-Oriented Segment

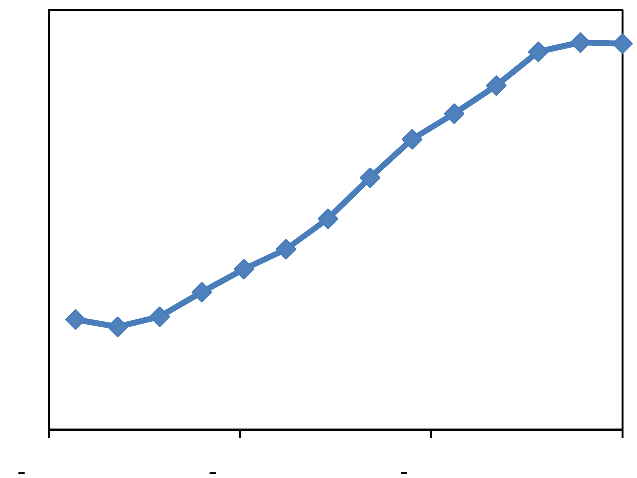

c. Closeness to Frequency Rewards and Percentage Paid Stays

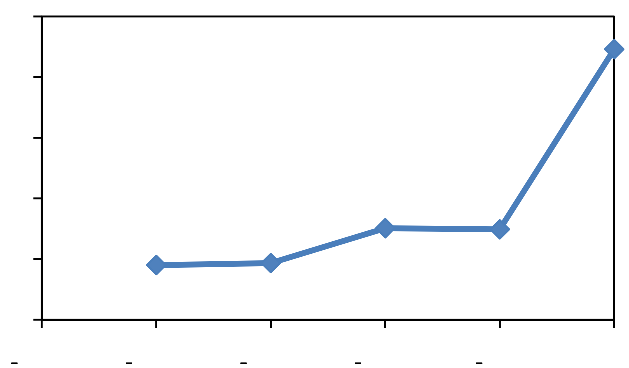

Note: Figures 2a-2b are based on simulations assuming no customer tier program; Figure 2c is based on actual data

FIGURE 3. The Rewarded Behavior Effect: The Relationship Between Reward Cash-In and Subsequent Predicted Probability of Staying

(Based on simulations assuming no customer tier program)

a. Price-Oriented Segment

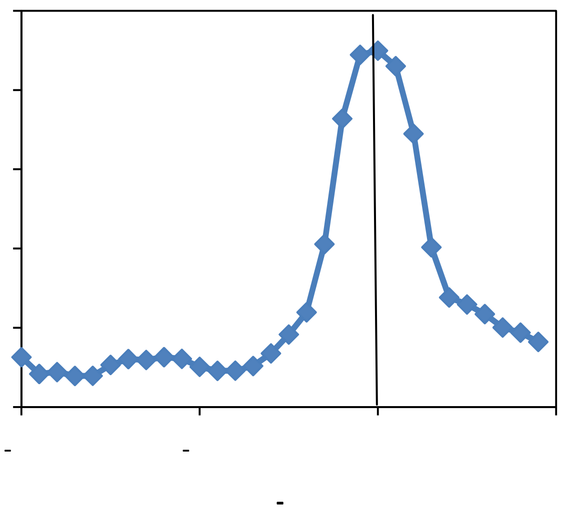

b. Service-Oriented Segment

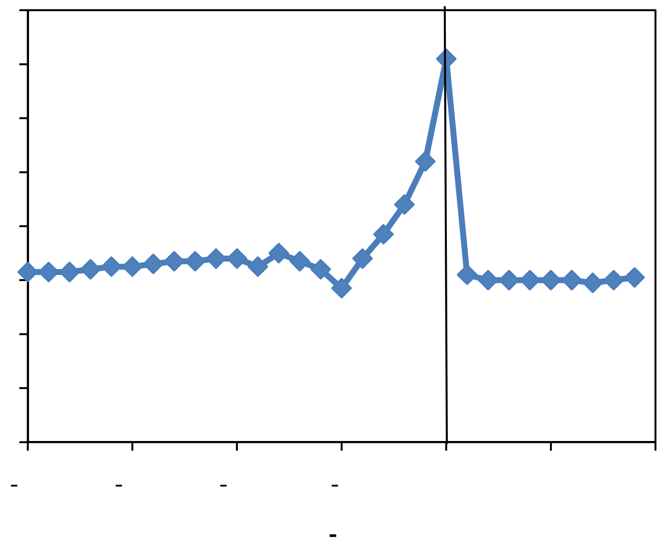

APPENDIX: Solution Approach for Parameter Estimation

We embed a dynamic programming backward induction method with a finite horizon, 1 to T=52, within a maximum likelihood estimation procedure. The state space is the Cartesian product where

Maximum Likelihood Estimation Procedure:

Begin iteration for an arbitrary set of parameter values

- For customer i = 1 to 200
  - For period t = 1 to T = 52
  - Determine the levels of state variables for each customer
    - For each option k = 0 to 7
      - Compute customer's current period utility (Equation 5)
      - Recall future Value function determined as a function of customer i's state variables
      - Compute probability of choice using Multinomial Logit formula
    - Increment to next option
    - Identify the log-likelihood of the observed choice
  - Increment to next period
- Increment to next customer
- Sum the log likelihoods over all customers and t
- Use Newton-Rhapson method to find the next set of parameter values

Continue iteration till parameters that maximize the overall log-likelihood are arrived at

<u>Dynamic Programming Algorithm for Computing Value Function</u>:

For Customer i = 1 to 200

For $t = T = 52$ to 1 (where t is a time index for the t^{th} future period from now)

For each possible value of the state vector (S, C, F, Y, V)[1]

For each option k = 0 to 7

Compute the option-specific current period utility (Equation 5)

Determine future option-specific expectations of prices and length of stay by averaging over draws from respective customer-specific distributions (both fitted as left-truncated normal)

Compute the expectation of points based on customer's (1) expectation of price, (2) expectation of nights stay, and (3) customer tier status (see state transition map for customer point inventory ($v_{i,t+1}$) below for computation details)

Compute the option-specific Value Function (Equation 7) as a function of state variables

Increment to next option

Determine the optimal option with the maximal value function by averaging across draws to arrive at the expected value function

Store the expected optimal value as functions of the current set of state variables

Increment to next state vector

Decrement to next t

Increment to next customer

<u>State Transition Map</u>

$$s_{i,t+1} = \begin{cases} 1 \text{ if } k_{it} = 1, 2, \ldots 7 \\ 0 \text{ otherwise} \end{cases}$$

$$c_{i,t+1} = \begin{cases} 1 \text{ if } k_{it} = 4, 5, 6, 7 \\ 0 \text{ otherwise} \end{cases}$$

[1] We restrict that Y to be no greater than F to reflect the fact that the paid stay inventory expires at the end of a calendar year.

$$v_{i,t+1} = \begin{cases} v_{it} + W^{-1} \cdot \sum_{w=1}^{W} prate_{it} \cdot p_{ik,w} \cdot n_{ik,w} & \text{if } k_{it} = 1,2,3 \\ v_{it} & \text{if } k_{it} = 0 \\ v_{it} \text{ - } R_{(k_{it}-3)} & \text{if } k_{it} = 4,5,6,7 \end{cases}$$

$prate_{it}$ is the ratio of earned points to the eligible dollar spent. This ratio is 5 for typical customers (whose number of paid stays is less than 5), while platinum and diamond members will respectively earn an extra of 15% and 30% points for each dollar they spend, thus,

---

and are the w^{th} draw of the prices and number of nights at time t drawn from the customer-and-category specific price and customer specific night distributions, both assumed to be left-truncated normals:

x and x

The product between the per-night price and the number of nights stayed, when averaged over a total of W draws gives us the expected number of dollars spent by customer i; further multiplying by the point-to-dollar rate ($prate_{it}$) gives us the expected number of hotel points earned. To impute the mean and variances of customers' expectation about their prices, we distinguish between two cases:

If the consumer i has stayed in a category k hotel for at least once, then we compute and using the average and variance of the prices paid by consumer i. Otherwise, we impute and using all observed prices paid for a category k hotel. Consumers' expectations of nights stay are computed similarly.