

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

Praveen K. Kopalle

Tuck School of Business at Dartmouth, Dartmouth College, Hanover, New Hampshire 03755,
praveen.kopalle@dartmouth.edu

Yacheng Sun

Leeds School of Business, University of Colorado, Boulder, Colorado 80309,
yacheng.sun@colorado.edu

Scott A. Neslin

Tuck School of Business at Dartmouth, Dartmouth College, Hanover, New Hampshire 03755,
scott.neslin@dartmouth.edu

Baohong Sun

Cheung Kong Graduate School of Business, New York, New York 10019, bhsun@ckgsb.edu.cn

Vanitha Swaminathan

Katz School of Business, University of Pittsburgh, Pittsburgh, Pennsylvania 15260,
vanitha@katz.pitt.edu

We estimate the joint impact of the frequency reward and customer tier components of a loyalty program on customer behavior and resultant sales. We provide an integrated analysis of a loyalty program incorporating customers' purchase and cash-in decisions, points pressure and rewarded behavior effects, heterogeneity, and forward-looking behavior. We focus on four key research questions: (1) How important is it to combine both components in one model? (2) Does points pressure exist in the context of a two-component loyalty program? (3) How is the market segmented in its response to the combined program? (4) Do the programs complement each other in terms of the incremental sales they produce?

Our most basic message is that the frequency reward and customer tier components of loyalty programs should be modeled jointly rather than in separate models. We find strong evidence for points pressure for both the customer tier and frequency reward components using both model-based and model-free evidence. We find a two-segment solution revealing a "service-oriented" segment that highly values cash-ins for room upgrades and staying in "luxury" hotels, and a "price-oriented" segment that is more price sensitive and highly values the frequency reward aspects of the loyalty program. Furthermore, we find that both components generate incremental sales. Also, there was slight synergy between the programs but not a huge amount. Overall, each component contributes to increased revenues and does not interfere with the other.

Key words: loyalty program; customer tier programs; frequency reward; database marketing; segmentation

History: Received: June 1, 2009; accepted: September 1, 2011; Eric Bradlow and then Preyas Desai served as the editor-in-chief and K. Sudhir served as associate editor for this article. Published online in *Articles in Advance* January 24, 2012.

1. Introduction

Loyalty programs, designed to maintain and enhance loyalty, have become "go-to" marketing programs for many companies (Deighton 2000, Lewis 2004, Liu 2007, Zeithaml et al. 2001). In recent years, various firms have initiated loyalty programs and have supported these programs with investments in creating and maintaining databases on loyalty program members. For instance, a recent survey conducted by research companies Ipsos Mori/The Logic Group suggests that almost two thirds (62%) of people say they belong to at least one loyalty program,

although only 26% agree that it leads to greater loyalty (Barnett 2010).

Thus, despite the growth of loyalty programs, their value to the firm is a source of debate. Advocates view loyalty programs as a means to soften price competition (e.g., Kim et al. 2001, Klemperer 1987) or build a customer database (Butscher 1998, Reynolds 1995), and as a dominant strategy when fixed costs are low (Kumar and Rao 2006). Because 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 et al. 2006, Hartmann and Viard 2008, 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 et al. 2008). Differentiating the relative effectiveness of these two components is important for evaluating and diagnosing the impact of any loyalty program. Frequency reward programs are of the form “buy X times, get something free.” These are the original trading-stamp programs. Customer tier programs are of the form “once you qualify for our diamond tier, we will provide you with special benefits and services.” Table 1 provides brief examples of loyalty programs, where we clearly see the frequency reward and customer tier components.

Both loyalty program components rely on accumulated customer sales to determine which customers qualify for which rewards. However, they differ in the nature of the reward as well as the means by which

customers attain it. First, a frequency reward is a one-shot affair—the customer redeems points for a free stay at a hotel, a free flight, a coupon, etc. In contrast, customer tier programs offer a steady stream of benefits as long as the customer is a member of that tier. Second, frequency reward programs typically require customers to proactively redeem their points, and customer tier programs dispense their reward automatically. Once customers qualify for a certain tier, they are notified and treated according to their tier status. Third, there is usually no expiration date for points inventory accumulated under frequency reward programs; in contrast, customer tier status does expire and needs to be re-earned.

Prior loyalty program research (e.g., Lewis 2004) has focused on the frequency reward component. This work (e.g., Bolton et al. 2000, Lal and Bell 2003, Roehm et al. 2002, Taylor and Neslin 2005) provides evidence of rewarded behavior effects; i.e., a reward earned in the last period increases the likelihood of repatronage in the next period. In a laboratory setting, Kivetz et al. (2006) and Nunes and Drèze (2006)

Table 1 Examples of Loyalty Programs Showing Frequency Reward and Customer Tier Components

Company	Frequency reward component		Customer tier component		
	Earning points	Cash-in rewards	Customer tiers	Tier eligibility requirements	Tier benefits
Starwood Hotels: Preferred Guest	Two “Starpoints” for every dollar spent at Starwood Hotel	Free night stay, free merchandise from “partners,” Banana Republic gift card	Preferred, Gold-Preferred, Platinum-Preferred	Preferred: automatic Gold: 10 stays (25 nights) in a calendar year Platinum: 25 stays (40 nights) in a calendar year	Gold: Late check-in, bonus points, best rate guarantee Platinum: Gold benefits + arrival amenity, best room guarantee, free Internet
Harrah’s Casino: Total Rewards	For example, one credit for every \$5 spent on slot machine, one credit for every \$10 spent on video poker	Merchandise, meals, golf	Gold Platinum Diamond	Gold: automatic Platinum: 4,000 tier credits in a calendar year Diamond: 11,000 tier credits in a calendar year	Gold: discounts at gift shops and travel Platinum: Gold benefits + birthday offer, exclusive gift, and free tournament entry Diamond: “Diamond lounges,” priority service, invitations to special events
Delta Airlines: SkyMiles	One mile flown = one “base mile,” bonus miles depending on class of ticket purchased	Free trip, seat upgrade; points required depends on destination	Silver Gold Platinum Diamond	Silver: 25K miles Gold: 50K miles Platinum: 75K miles Diamond: 125K miles Miles calculated starting January 1 of each year.	Silver: priority phone line, priority boarding Gold: higher priority phone line, Sky priority boarding Platinum: highest priority phone line, Sky priority boarding Diamond: VIP phone line, Sky priority boarding, Sky Club membership
Best Buy Reward Zone	\$1 spent = one point	Best Buy gift certificates	Regular Silver	Silver status: \$2500 spent in calendar year	Silver: Earn points faster, free standard shipping, extended return policy, no restocking fees, priority service support, free movie downloads

Sources. See the following for complete descriptions (all accessed September 22, 2010): Starwood Hotels: http://www.starwoodhotels.com/preferredguest/account/member_benefits/index.html, Harrah’s Casino: http://www.harrah.com/total/_rewards/overview/overview.jsp, Delta Skymiles: <http://www.delta.com/skymiles/index.jsp>, and Best Buy Reward Zone: <https://myrewardzone.bestbuy.com/>.

Note. The descriptions present summaries and examples.

have demonstrated the points pressure effect; i.e., customers increase their purchase frequency as they get closer to earning a frequency reward. Akin to Lewis (2004), we demonstrate the effect using actual purchase data. In addition, however, we investigate the extent to which points pressure effects exist separately for a customer tier program and for a frequency reward program.

With the exception of Drèze and Nunes (2008), very little has been learned about customer response to customer tier programs, particularly in a dynamic setting (see Blattberg et al. 2008). To our best knowledge, prior research has not studied frequency reward and customer tier programs in an integrated way. This does not sit well with the fact that many companies offer customer tier programs alongside the frequency reward component, as illustrated in Table 1. The key point is that in practice, both tier *and* frequency components are common elements of loyalty programs. As a result, an evaluation of a loyalty program requires an integrated analysis of both components. The purpose of this paper is to develop and estimate a model to provide such an evaluation, which allows us to answer the following questions:

- How important in terms of insight and model fit is it to combine the tier and frequency components in one integrated model?
- Do “points pressure” effects exist for both components? For example, do customers increase their purchase rate as they approach higher tier status?
- How is the market segmented in terms of response to the loyalty program? Is there a “frequency reward” segment and a “customer tier” segment?
- Are the frequency and tier components complementary in terms of incremental sales produced by the program?

To address these issues, the model must contain the following phenomena.

(a) *Forward-Looking Customers*—Both frequency reward (e.g., free hotel stay) and customer tier benefit (e.g., elite status) components 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, customer tier reward is delivered automatically.

(c) *Customer Heterogeneity*—This addresses the issue of market segmentation.

(d) *Rewarded Behavior*—The reward indeed is short term, but customer affect created by the reward can translate into an increase in loyalty.

We apply our model to the loyalty program instituted by a major hotel chain. The analyses yield interesting results: (1) Including frequency reward

and customer tier components in the same model improves model fit and yields additional insights. (2) There is a significant points pressure effect for both components. (3) There are rich customer segments where customers vary in their responses to the frequency and customer tier components as well as to price. (4) Both components produce incremental sales, with some complementarity but no cannibalization between the two components.

In summary, the contribution of this paper is threefold. First, we provide an integrated analysis of the impact of two important components (frequency reward and customer tier) of a loyalty program on customer behavior and corresponding sales. Second, we endogenize the cash-in decision for the frequency reward, a factor not considered in previous work. Third, we add to the empirical knowledge base on how loyalty programs work in terms of points pressure, rewarded behavior, and incremental sales impact.

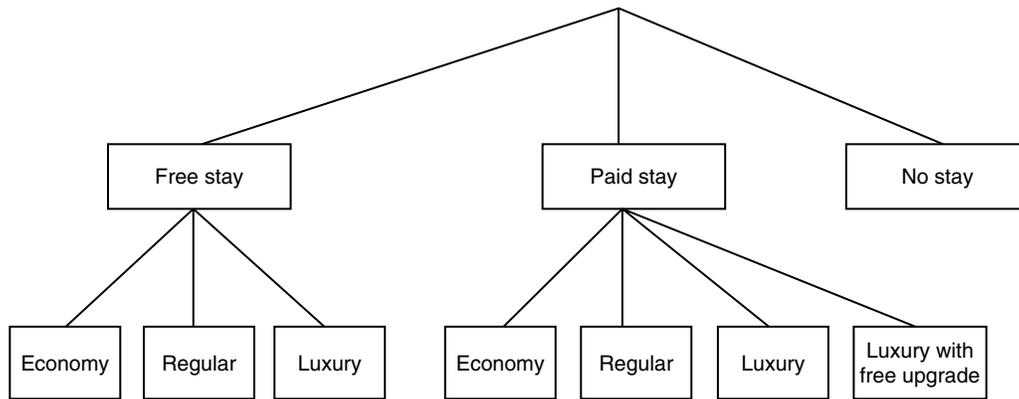
One novel aspect of our paper relative to previous work is the inclusion of explicit measures of competitive activity (competitive price and occupancy rates). We indeed find that competitive activity has a negative impact on sales, although that impact differs by market segment. Although this is an advance and shows the desirability of including competitive measures, our results must be seen as exploratory because we cannot differentiate competitive hotel stays from no stays, and our policy implications do not imply an equilibrium analysis that allows for competitive response. We will discuss this issue in more detail when we report competitive effects and policy simulations. But first we discuss the model, the data, and then the empirical results.

2. Model

2.1. Background

The model is generally applicable to loyalty programs comprising both frequency reward and customer tier components. To make the exposition clear, we describe the model in the context of our application to a hotel’s loyalty program. The hotel categorizes each of its properties as “economy,” “regular,” or “luxury.” All properties allow customers to cash in points for a free night, whereas cash-in for a room upgrade is only available at a luxury property. Points are accrued via paid stays, they do not expire, and no points are earned on free stays. Customers need to contact the hotel if they wish to cash in points. Customers could not cash in for two rewards at once, e.g., a free stay at a luxury hotel plus an upgrade. The tier component places customers in the Base, Platinum, or Diamond tier depending on how many

Figure 1 Decision Tree for Nested Logit Formulation



paid stays they had during the past calendar year.¹ 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

In formulating the model, we recognize that the customer faces two fundamental decisions—the “stay” decision and the “accommodation” decision. The stay decision is whether to stay at the hotel paying full price, paying no price (a free stay), or not staying at all. The accommodation decision entails the specific level of hotel—regular, economy, luxury, or upgraded room in a luxury hotel. We therefore will formulate a nested logit model, with the decisions organized as in Figure 1.

In addition to capturing the two types of decisions (stay versus accommodation), organizing the choices this way accounts for unobserved factors that systematically elevate some stay decisions over others. For example, a business trip will increase the utility of paid stay because a third party may be paying for the trip. The nested logit model will then capture the elevation of utilities for all four accommodation choices within paid stay. A leisure trip will increase the utility of a free stay, and similarly, the three choices within free stay. We now describe this formulation in detail.

We consider $i = 1, \dots, I$ customers who make choices at time t . Define C_{imt} to indicate which choice the customer makes regarding the stay decision. Specifically, $C_{imt} = 1$ if customer i chooses alternative m at time t and 0 otherwise, where

$$m = \begin{cases} 0 \rightarrow \text{no stay,} \\ 1 \rightarrow \text{paid stay,} \\ 2 \rightarrow \text{free stay.} \end{cases} \quad (1a)$$

We define D_{int} to indicate which choice the customer makes regarding the accommodation decision. Specifically, $D_{int} = 1$ if customer i chooses alternative n at time t and 0 otherwise, where

$$n = \begin{cases} 1 \rightarrow \text{economy,} \\ 2 \rightarrow \text{regular,} \\ 3 \rightarrow \text{luxury,} \\ 4 \rightarrow \text{luxury with free upgrade.} \end{cases} \quad (1b)$$

2.3. Customer Utility

We define U_{imnt} as the customer i 's utility for the choice combination (m, n) at time period t , which is part of an underlying dynamic model discussed in the next section, and where the allowable combinations of m and n are defined by the nested logit tree above. Then

$$U_{imnt} = U_{imt}^C + U_{int}^D + \mu_{imt} + \nu_{imnt}, \quad (2)$$

where

U_{imt}^C = utility determining choice among no stay, stay, and paid stay, because of factors observed by the researchers;

U_{int}^D = utility determining choice among economy, regular, luxury, and luxury with free upgrade accommodations, because of factors observed by the researcher;

μ_{imt} = contribution to utility because of unobserved factors determining choice among no stay, stay, and paid stay; and

ν_{imnt} = contribution to utility because of unobserved factors determining choice among all possible choice combinations (m, n) .

A key term in this formulation is the unobserved random effect μ_{imt} . After controlling for competitor price and occupancy rates, leisure travel might systematically tilt utility toward free stay. This would result in a relatively large value of μ_{i2t} . All accommodation types under free stay would share a boost in

¹ Note from Table 1, this reliance on “calendar year” purchases to assign tiers appears to be common.

utility. On the other hand, business travel might systematically tilt utility toward paid stay (μ_{i1t} would be large) and boost utility for all accommodation types under paid stay. Using this formulation, we therefore can account for a factor we are unable to observe—business or leisure travel.

The last term in Equation (2), ν_{imnt} , represents unobserved time-specific determinants of customer i 's utility for choice combination (m, n) . We assume the ν_{imnt} s are independent, identically distributed extreme value random variables. This treatment is consistent with prior research (Ben-Akiva and Lerman 1985, Chib et al. 2004, Chintagunta 2002, Erdem et al. 2003, Sun 2005).

We will now specify the model for observed factors affecting type of stay (no stay, paid stay, free stay):

$$U_{imnt}^C = \begin{cases} \beta_{i0}PC_t + \beta_{i1}OR_t & \text{for } m = 0, \\ \beta_{i2}BHU_i + \beta_{i3}WDS_i + \delta_{i1}C_{i2t-1} & \\ + \delta_{i2} \sum_{d=1}^2 C_{idt-1} + \sum_{s=1}^2 \lambda_{is}E_{ist} & \text{for } m = 1, 2. \end{cases} \quad (3)$$

The outside option ($m = 0$) corresponds to either (1) the customer does not stay at any hotel at all or (2) the customer stays at a competitive hotel. Unfortunately, we do not have the data to distinguish between these possibilities. This is similar to the situation faced by previous researchers, e.g., Gönül and Srinivasan (1996) and Lewis (2004). Nevertheless, we can address this to some extent by having the utility of not staying in the hotel depend on competitive prices (PC_t) and competitor occupation rate (OR_t , the number of rooms sold divided by number of rooms available). Thus, β_{i0} and β_{i1} can be interpreted as the extent to which change in competitors' price and vacancy affects the relative attractiveness of staying with the focal hotel. We expect both β_{i0} and β_{i1} to be negative because high prices and occupancy rates make competitors unattractive.

The utility of staying in the hotel, be it a free or paid stay, is influenced by five observed factors: (1) Base hotel usage rate measured by the number of stays in the initialization period (BHU_i).² We expect $\beta_{i2} > 0$. (2) The customer's propensity to stay on weekdays versus weekends, measured by the number of stays that start on a weekday in the initialization period (WDS_i). The sign of β_{i3} could be positive or negative, depending on whether weekday stayers are more likely to visit the hotel ($\beta_{i3} > 0$) or less likely ($\beta_{i3} < 0$). A positive β_{i3} thus suggests the hotel is more

geared toward business travelers, because they are more likely to have a large number of weekday stays in the initialization period (a negative β_{i3} would suggest non-weekday stayers; e.g., leisure travelers are the more likely hotel visitors).³ (3) δ_{i1} represents the strength of rewarded behavior, i.e., the impact on utility in period t if the customer received a free stay in the previous period ($C_{i2t-1} = 1$). Based on previous research, we expect $\delta_{i1} > 0$ because receiving a reward increases customer "delight" (Rust and Oliver 2000) and therefore increases subsequent utility.⁴ (4) δ_{i2} captures state dependence, i.e., the impact on utility in period t if the customer stayed in the focal hotel in the previous period ($C_{i1t-1} = 1$ or $C_{i2t-1} = 1$). As in most previous choice model research, we expect the state dependence effect to be positive ($\delta_{i2} > 0$). (5) λ_{is} denotes the utilities of membership in customer tier s ($E_{i1t} = 1$ if the customer is in the Platinum tier and 0 otherwise; $E_{i2t} = 1$ if customer is in the Diamond tier and 0 otherwise). Because tier membership provides extra benefits and the Diamond tier provides the most benefits, we expect $\lambda_{i2} > \lambda_{i1} > 0$.

Note that price of the focal hotel does not appear in the above formulation. Price will be included in the various accommodation types and impact the choice of paid versus free versus no stay via the inclusive value in the nested choice probabilities (Equation (15)). We now specify the utility models that determine choice among accommodation types:

$$U_{int}^D = \begin{cases} 0 & \text{for } m = 0, \\ \alpha_{in} + \beta_{i4}P_{nt} & \text{for } m = 1, n = 1, 2, 3, \\ \gamma_{i0} + \alpha_{i3} + \beta_{i4}P_{3t} & \text{for } m = 1, n = 4, \\ \gamma_{i1} + \alpha_{in} & \text{for } m = 2, n = 1, 2, 3. \end{cases} \quad (4)$$

For no stay ($m = 0$), there is no accommodation decision, and we set utility equal to zero for scaling purposes. For a paid stay ($m = 1$), we define alternative-specific constants α for economy ($n = 1$), regular ($n = 2$), and luxury ($n = 3$), and we expect $\alpha_{i3} > \alpha_{i2} > \alpha_{i1}$. We express the alternative-specific utility for the luxury upgrade ($n = 4$) as α_{i3} (staying in the luxury hotel) plus γ_{i0} , which captures the additional utility as a result of the upgrade. We expect $\gamma_{i0} > 0$. Note also that price for economy, regular, or luxury

³ As a simplification, we assume the β s do not differ by the stay branches of the nested logit tree. This means that, e.g., the impact of base hotel usage rate is the same on the utility of both paid and free stays. However, this does not mean these two utilities are equal. First, the unobserved factor μ_{imnt} will cause them to differ, and the inclusive value term from the lower part of the nest will cause these utilities to differ.

⁴ Note, however, an alternative possibility is that δ_{i1} would be negative because of purchase acceleration; that is, customers move forward purchases in time and therefore purchase at a lower rate after they cash in, i.e., a "post-promotion dip" (Neslin 2002, pp. 23–24).

² The first 52 weeks of the data are used for initialization. We believe this is long enough to smooth out the variations in factors such as price promotion and seasonality. Previous research (e.g., Bucklin et al. 1998) also used a similar length for the initialization period.

hotels now enter the utility function (P_{nt}). For a free stay ($m = 2$), we need not include price because the price is zero. However, we have the same alternative-specific constants for each hotel type, plus a term (γ_{i1}) indicating utility for a free stay.⁵ Intuition suggests $\gamma_{i1} > 0$ because of the “transaction utility” of a great “deal” (Thaler 1985). However, if the hassle costs are high for cashing in points, it is possible that $\gamma_{i1} < 0$.

2.4. Points Accumulation

We now introduce notation to represent the points requirements of the frequency reward and customer tier requirements. Denote R_1 as the points inventory requirement for an upgrade, and let R_2 , R_3 , and R_4 denote the points requirements for a free stay at economy, regular, and luxury properties, respectively. Let T_s denote the required number of paid stays for tier s . The structure of a two-component loyalty program in our case can thus be summarized by the vector $(R_1, R_2, R_3, R_4, T_1, T_2)$.

Customers must accrue enough points to choose either an upgrade or a free stay. Points “inventory” accrues as follows:

$$INV_{it+1} = INV_{it} + PointsAccrued_{it} - PointsCashedIn_{it}, \quad (5)$$

where

INV_{it+1} = customer i 's cumulative points inventory at the beginning of period $t + 1$,

$PointsAccrued_{it}$ = points earned by customer i in period t ,⁶ and

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

Notice that the points inventory determines the choice set of the customer. Not staying in the hotel and a paid stay are always options for the customer. If customer i has enough points for an upgrade, then the choice set would include ($m = 1, n = 4$). If she is eligible for a free stay in an economy hotel, then ($m = 2, n = 1$) is included in the choice set, etc. In essence, we have an economic model of consumer choice under cutoffs (see Swait 2001 for a detailed discussion).

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

$$INVP_{it+1} = INVP_{it} + PStay_{it}, \quad (6)$$

⁵Note that the *incremental* utility of a free stay relative to a paid stay at property n for customer i is $\gamma_{i1} - \beta_{in}P_{nt}$, and this increases with higher prices ($\beta_{in} < 0$).

⁶The hotel we study allowed customers to purchase a very limited number of points. However, we find no evidence of this in our data, perhaps because of the generally high level of points inventories carried by customers.

where

$PStay_{it}$ = number of paid stays added to customer i 's inventory in period t .

Membership in a customer tier, s , is determined as follows:

$$E_{ist} = \begin{cases} 1 & \text{if } T_s \leq \min(INVP_{it}, INVP_{i, \text{LastYear}}) \text{ and} \\ & \max(INVP_{it}, INVP_{i, \text{LastYear}}) < T_{s+1}, \\ 0 & \text{otherwise,} \end{cases} \quad (7)$$

where

E_{ist} = indicator variable signifying that customer i is in customer tier s at period t ,⁷

T_s = required number of paid stays for tier s ,

$INVP_{it}$ = number of paid stays by customer i at period t in the current calendar year, and

$INVP_{i, \text{LastYear}}$ = number of paid stays by customer i during the last calendar year.

Note our model captures the typical situation (see Table 1) that the requirements for the frequency and tier rewards are often based on different accumulation rules. First, whereas $INVP_{it}$ will expire at the end of a calendar year, INV_{it} will never expire. Second, when the customer chooses a paid stay, $INVP_{it}$ will always increase by one. On the other hand, the amount by which INV_{it} increases depends on the monetary expenditure of the customer, which depends on her utilities for different hotel levels. Third, when customers make a cash-in decision, INV_{it} will decrease, and furthermore, the exact amount of decrement depends on what reward the customer redeems for; on the other hand, $INVP_{it}$ will not decrease as the result of a cash-in decision.

2.5. 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. However, there is an interesting dynamic interaction between the frequency reward and customer tier components: for example, a forward-looking customer may not wish to cash in for a free stay because free stays do not count toward 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

$$\max_{C_{imt}, D_{int}, m, n \in N_{it}} E \left\{ \sum_{t=1}^{\infty} \sum_{m, n \in N_{it}} [U_{imnt} \cdot C_{imt} \cdot D_{int}] \delta^{t-1} \right\}, \quad (8)$$

⁷We know from the data what tier the customer was in initially, so we have an initial value for E_{ist} .

where δ is the per period discount factor and N_{it} is customer i 's choice set at time t . The operator $E[\cdot]$ (we suppress the subscripts i and 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

$$\begin{aligned} V_{it}(INV_{it}, INVP_{it}, F_{it}, C_{imt-1}) \\ &= \max_{C_{imt}, D_{imt}, m, n \in N_{it}} V_{imnt} \\ &= \max_{C_{imt}, D_{imt}, m, n \in N_{it}} \{U_{imnt} + \delta E[V_{i,t+1}(INV_{i,t+1}, \\ &\quad INVP_{i,t+1}, F_{it+1}, C_{imt})]\}, \quad (9) \end{aligned}$$

where F_{it} denotes the current calendar year period for customer i indicated by period t . This is a state variable because it reflects seasonality (see §2.6 for how seasonality enters the model) and the time remaining before customer i 's current tier status expires. The decision rule is to choose the decision (m, n) that maximizes the choice-specific value function V_{imnt} :

$$V_{imnt} = U_{imnt} + \delta E[V_{i,t+1}(INV_{i,t+1}, INVP_{i,t+1}, F_{it+1}, C_{imt})]. \quad (10)$$

The state variables are as follows: inventory of total points (INV_{it+1}), given by Equation (5); inventory of paid stays ($INVP_{it+1}$), given by Equation (6); and past decision (C_{imt}) and the calendar period (F_{it+1}). Given the complexity of the dynamic programming problem, we adopt linear interpolation (Keane and Wolpin 1994) to estimate the model. See the appendix for a description of our solution approach for parameter estimation, dynamic programming algorithm, and linear interpolation.

2.6. Customer Expectations of Future Demand and Prices

The customer's decision is derived from a dynamic program. Therefore, when evaluating expected future utility, customers need to incorporate their expectations regarding (1) future prices at different types of properties of the focal hotel chain, (2) price of the competitors, and (3) length of stay with the hotel. These, in turn, determine their expected future points.

We assume that consumer i has rational expectations for the prices of the focal company's accommodation type n (P_{int} , $n = 1, 2, 3$):

$$\ln P_{int} = \lambda_{1in} + \sum_{k=1}^3 \lambda_{1+k,i} S_{kt} + \eta_{int}, \quad n = 1, 2, 3. \quad (11)$$

P_{int} is individual specific and depends on seasonality, captured by dummy variables S_{kt} , $k = 1, 2, 3$, where η_{int} is the error term and

$$\begin{aligned} S_{1t} &= 1\{t \text{ is either in December, January, or February}\}, \\ &\quad 0 \text{ otherwise;} \\ S_{2t} &= 1\{t \text{ is either in March, April, or May}\}, \\ &\quad 0 \text{ otherwise;} \quad \text{and} \quad (12) \\ S_{3t} &= 1\{t \text{ is either in June, July, or August}\}, \\ &\quad 0 \text{ otherwise.} \end{aligned}$$

Consumers' expectations for the competitive price and occupancy rate are modeled similarly:

$$\ln PC_{it} = \lambda_{1ic} + \sum_{k=1}^3 \lambda_{1+k,ic} S_{kt} + \eta_{ict}, \quad (13.1)$$

$$\ln OR_{it} = \lambda_{1io} + \sum_{k=1}^3 \lambda_{1+k,io} S_{kt} + \eta_{iot}. \quad (13.2)$$

Thus, λ_{1in} ($n = 1, 2, 3$), $\lambda_{2-4,i}$, λ_{1ic} , $\lambda_{2-4,ic}$, λ_{1io} , and $\lambda_{2-4,io}$ characterize consumers' expectations. We estimate these coefficients by assembling the focal price, competitive price, and occupancy rate data in a time series and then estimating each equation using ordinary least squares. We follow a two-step approach for compiling the customer-specific prices and competitive occupancy rates needed for these regressions. The first step is to determine the location of the city relevant to the customer; the second step is to assemble the prices and occupancy rates for that location (see the appendix for further details). If the customer stayed in the focal hotel in period t , we know directly from the data the price for that accommodation and the occupancy rate (and price) of competitors in that period. If the customer did not stay in the focal hotel in period t , we still observe competitive data, and we impute prices at the focal hotel using the average price paid by other consumers for the same accommodation during period t in that city. This method is similar to the imputation of missing price information in packaged goods categories (e.g., Erdem et al. 2008). The average R^2 (across customers) for these equations were 0.42 (P_{1t}), 0.35 (P_{2t}), 0.36 (P_{3t}), 0.91 (PC_{it}), and 0.55 (OR_{ct}).

Customers also have expectations regarding their length of stay, denoted LOS_{it} , where $LOS_{it} \sim$ left-truncated Normal($\overline{LOS}_i, \sigma_{LOS,i}^2$). We determine the respective means (\overline{LOS}_i) and standard deviations ($\sigma_{LOS,i}^2$) of these distributions from each customer's transaction history during the initialization period. These distributions are then used in computing future value function, where we draw from corresponding normal distributions with the estimated means and variances (see the appendix for further details).

Note that these expectations, and the resultant choices and points accrual that emerge from them, are not explicitly included in the utility function. However, they play an important role in solving the dynamic program and hence in our predictions and estimates, because the inventory of points is a state variable (Equation (9)). Therefore, the customer maximizes utility, taking into account the likely points accumulation scenario that will result from his or her choices.

2.7. Heterogeneity and Estimation

A long history of work in the choice models literature has established the importance of incorporating cross-consumer heterogeneity. Two fundamental ways of modeling heterogeneity are continuous and latent class. We adopt latent class because one of our goals is to identify and interpret segments, and latent class naturally reveals those segments (Kamakura and Russell 1989). Let $\omega_l = \alpha_{jl}, \gamma_{1l}, \gamma_{2l}, \beta_{0,l}, \dots, \beta_{4,l}, \delta_{1l}, \delta_{2l}, \lambda_{sl}, j = 1, 2, 3; s = 1, 2$ be the vector of coefficients to be estimated for each latent class l . Then, the likelihood function of the sequence of choices of all customers is

$$L(\omega) = \sum_{l=1}^L \prod_{i=1}^I \prod_{t=1}^T \sum_{m=0}^2 \sum_{n=1}^4 [C_{imt} \cdot D_{int} \cdot \Pr(C_{imt} = 1, D_{int} = 1 | INV_{it}, INVP_{it}, S_{it}, C_{im,t-1}, D_{in,t-1}, \omega_l, i \in l) \cdot \Pr(i \in l)], \quad (14)$$

where l indexes latent classes and $\Pr(i \in l)$, the probability customer i is in latent class l , is estimated. Given the extreme value distribution of the error term, the probability customer i in latent class l makes decision (m, n) at time t is

$$\Pr(C_{imt} = 1, D_{int} = 1 | INV_{it}, INVP_{it}, S_{it}, C_{im,t-1}, D_{in,t-1}, \omega_l, i \in l) = \left(\frac{e^{U_{imt}^C + \gamma_m \ln \sum_n e^{(U_{imt}^D + \delta \cdot EV_{imnt})}}}{\sum_j e^{U_{ijt}^C + \gamma_j \ln (\sum_h e^{(U_{ijt}^D + \delta \cdot EV_{ijht})})}} \right) \cdot \left(\frac{e^{(U_{imt}^D + \delta \cdot EV_{imnt})}}{\sum_h e^{(U_{imht}^D + \delta \cdot EV_{imht})}} \right), \quad (15)$$

where U_{imt}^C and U_{imt}^D are defined in Equations (3) and (4), respectively; δ is the discount factor; EV_{imnt} is the integrated value function on the pair $(\mu_{imt+1}, \nu_{imnt+1})$; and $\ln(\sum_h e^{(U_{imht}^D + \delta \cdot EV_{imht})})$ is the inclusive value for alternative m . This inclusive value is 0 for $m = 0$. Finally, $\gamma_j, j = 1, 2$ are the log-sum coefficients.

The estimation (see the appendix) is performed using one year of data (26 biweekly periods), assuming a finite, forward-looking horizon that goes one year beyond the end of the data (total of two years). In an earlier version of the model, when we did not use initialization variables, we utilized 52 biweeks

for our estimation, and the results regarding tier and reward response were quite similar to the results we obtain now with 26 biweekly periods. We use customer biweek as the unit of analysis because that period is long enough to cover the full length of all stays yet short enough so there are no multiple stays within a period. We use a finite horizon for the following reasons.

1. This is an approximation that is computationally attractive relative to the infinite horizon problem. What is important in the context of our interpolation is that the parameters of the regression interpolation should stabilize before the last data period is reached (Erdem et al. 2008); indeed, they do stabilize in our estimation. Furthermore, in our algorithm, the value function calculation is for a total of two years; i.e., customers are looking ahead one year beyond the last period of our data. Allowing customers to look ahead an additional year outside of data not only approximates an infinite horizon problem but also ensures that the last period in our data holds the same importance regardless of the period a customer is in at a given point in time. This thus allows the customer to consider the benefits of reaching a certain tier in year 1 and reaping the benefits in year 2.

2. Using a finite horizon is consistent with previous dynamic structural models (e.g., Erdem and Keane 1996; Erdem et al. 2003, 2008; Sun et al. 2003; Sun 2005 in a brand choice context; and Lewis 2004 in a frequency reward program context).

3. The limitation of the finite horizon is that at the end of that horizon there is no salvage value. Our approach limits the amount of “end game” behavior on the part of the customer in the last period of the data, because the customer is considering the implications for the following year. Consequently, customers are taking into consideration the additional one-year benefit of achieving higher tier status or of having a higher points inventory.

3. Data

Our data for the hotel application are from the loyalty program of a major hotel chain, covering 3,907 customers who were members of the hotel’s loyalty program⁸ over two years (January 1, 2002–December 31, 2003). Included are the dates of stay by customer, type of property stayed in (economy, regular, or luxury), price paid,⁹ points earned, etc. Each trip includes

⁸ This suggests that our results apply strictly to the population of focal hotel loyalty card holders. However, (1) this is a very important group for hotel managers, (2) our sample is not unusually loyal, averaging a reasonable 3.0 stays per year, and (3) our later findings of significant competitive effects suggest that they are not single-brand loyal.

⁹ Note that because we observe price paid but not the list price of all accommodations, we must impute the price available for a

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. We integrated these data with the competitor pricing and occupancy rates, which were purchased from Smith Travel Research, a syndicated data firm that tracks pricing and occupancy information for the hotel industry on a monthly basis. These data covered hotels in each of the 14 submarkets (e.g., Los Angeles Airport vicinity, San Diego–La Jolla area, San Francisco–Market Street area, San Francisco–Nob Hill/Wharf area) in which our customers exclusively stayed. The competitive pricing and occupancy data were combined with the customer stay information from the focal hotel chain based on the time variable.

Our focal chain’s policy does not allow points to be accumulated from other sources, and our data include all point transactions. An analysis of our data revealed that all customer stays involved hotel points 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. There were no cases where significant hotel points, unrelated to a paid stay, were added to a customer’s 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.

The hotel’s frequency reward program involves four categories of frequency reward: free stay at an economy (5,000 points), regular (8,000 points), or luxury (12,000 points) property, and a free upgrade at a luxury property (3,000 points). Five points are earned per dollar spent, while Platinum and Diamond tier members earn an additional 15% and 30%, respectively, per dollar spent.¹⁰ 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, $YR - 1$), will be reassigned to a lower tier in year $YR + 1$ if the number of *paid* stays in year YR do not meet the requirements for the higher tiers. During a year, the average customer paid \$121.60 per night, accumulated approximately 2,395 points per stay, stayed 3.0 times, with 15.8% of stays at an economy property and 48.9% and 35.3% stays at regular and luxury properties. About 18.9% and 2.4% qualified for the Platinum and Diamond tiers, respectively.

given customer at a particular time period based on the prices other customers paid for the various alternatives. This is a relatively standard procedure with panel data analyses of scanner data.

¹⁰ Bonus points for elite customers are common. See Table 1.

Table 2 Model Fit Statistics and Parameter Estimates

	Estimation sample	Holdout sample
Number of observations	10,400	10,400
Myopic model with 2 segments:		
LL	-3,746.0	-3,960.5
AIC	7,554.0	7,983.0
BIC	7,616.5	8,045.5
Dynamic model with 1 segment:		
LL	-4,024.3	-4,266.4
AIC	8,080.6	8,564.8
BIC	8,112.9	8,597.1
Dynamic model with 2 segments:		
LL	-3,414.0	-3,664.8
AIC	6,890.0	7,391.6
BIC	6,952.5	7,454.1
Dynamic model with 3 segments:		
LL	-3,401.1	-3,650.9
AIC	6,894.2	7,393.8
BIC	6,987.0	7,486.6

Notes. The numbers in bold denote best model fit. LL, log likelihood.

For estimating the model, we chose a random sample of 400 customers.¹¹ We used the first year of data as an initialization period to compute the base hotel usage rate (*BHU*) and weekday stay (*WDS*) variables. Thus, there were 26 decision periods for each customer during the one-year time frame used in the estimation. In our data, 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) a myopic model with two latent segments, (2) a dynamic model with one segment, (3) a dynamic model with two segments, and (4) a dynamic model with three segments. The discount factor δ in Equation (9) is set to 0 in the myopic model, whereas it is set at a biweekly equivalent of 0.995 in the dynamic model (Lewis 2004, Sun 2005). Fit statistics for the four models are provided in Table 2. The two-segment dynamic model provides the best fit according to Akaike information criterion (AIC) and Bayesian information criterion (BIC).¹²

The superiority of the two-segment dynamic model is supported by its performance in a holdout sample. We applied the model to 400 customers not used

¹¹ There is nothing about the algorithm that would preclude a larger sample size, but with 400, the model took six to seven days to converge. This sample size is similar to other research using dynamic models (Erdem et al. 2008, 2005; Osborne 2007).

¹² The within-group correlation for the paid stay nest was 0.35, and the correlation for the free-stay nest was 0.38, suggesting it is important to include the μ_{int} terms in the model (Equation (2)).

in the estimation. To assign holdout customers to segments, following Kamakura and Russell (1989), we use the estimated proportions of the two segments as our prior probability of segment membership (73.3% for the first segment and 26.7% for the second segment). We then update this probability for each customer by incorporating the likelihood of his or her decision sequence. Based on Cramer (1999), who advocates using the prior probability as the cutoff, we assigned customers to the segment for which their posterior probability of being in that segment is larger than the prior probability. This means we used a 73.3% prior probability as the cutoff for classifying someone in the price segment. The rest are then assigned to the second segment. The corresponding fit statistics are shown in Table 2. The two-segment dynamic model again performs best.

The superiority of the dynamic to the static model is also reflected in the number of significant coefficients (21 for the dynamic model, 15 for the static model) as well as the diagnostics from the coefficients that are significant. For example, there is no rewarded behavior effect for the myopic model, yet there is one for the price-oriented segment in the dynamic model. This makes sense and is consistent with previous research that has found rewarded behavior effects.

4.2. The Two-Segment Solution

Table 3 reports the parameter estimates and their *t*-values for the two-segment model. Interestingly, Segment 1 derives significant positive utility from a free stay and shows no significant preference for an upgrade, whereas Segment 2 prefers the upgrade and no significant preference for a free stay. This is consistent with differences in price sensitivity: Segment 1 is more price sensitive, especially in response to the focal brand's prices. Segment 2's preference for a luxury property is positive and significant, whereas for Segment 1, it is not significantly different from 0. Members of Segment 2 have significant positive coefficients for both levels of the customer tier program, whereas the coefficient for only the Diamond tier is significant for Segment 1.

Given the above observations, we label the first segment "price oriented" in that its members are more price sensitive and are more attracted to free stays. We label the second segment "service oriented" because its members derive positive value for both customer tiers, favor upgrades, and prefer a luxury hotel. Consistent with this interpretation, the service-oriented segment exhibits a positive coefficient for the weekday stay variable, suggesting that its members are more likely to be business travelers, whereas the price-oriented segment has a significant negative weekday stay coefficient, suggesting that its members

Table 3 Parameter Estimates

Parameter estimates (dynamic model with 2 segments)				
Variable	Segment 1 (73.3%): Price oriented		Segment 2 (26.7%): Service oriented	
	Parameter estimate	<i>t</i> -value ^a	Parameter estimate	<i>t</i> -value ^a
Competitor price (β_0)	-0.0333	-6.39	-0.0260	-9.75
Competitor utilization rate (β_1)	-0.0252	-2.94	-0.026	-0.23
Base hotel usage rate (β_2)	0.201	4.28	0.115	4.09
Weekday stay (β_3)	-0.007	-6.70	0.025	2.25
Focal hotel chain price (β_4)	-0.094	-4.60	-0.008	-12.07
Inertia of cash-in (rewarded behavior, δ_1)	0.005	2.92	0.001	0.44
Inertia of staying (state dependence, δ_2)	0.032	12.74	0.027	0.56
Belonging to the Platinum tier (λ_1)	0.035	0.17	0.025	1.97
Belonging to the Diamond tier (λ_2)	0.254	5.74	0.277	9.49
Cash-in for upgrade (γ_0)	0.45	1.13	0.658	3.48
Cash-in for free nights (γ_1)	0.784	2.01	0.56	1.80
Preference for economy hotel (α_1)	-3.959	-4.81	-3.518	-19.8
Preference for regular hotel (α_2)	-3.247	-5.88	-1.443	-32.5
Preference for luxury hotel (α_3)	-10.741	-0.61	0.108	3.63

^aSignificant ($p < 0.05$) estimates are indicated in bold.

are more likely to be leisure travelers.¹³ This is consistent with the finding that Segment 1 is more price sensitive. The price-oriented segment is in the majority, 73.3%, of the sample, and the service-oriented segment is 26.7%. Furthermore, there is support for the rewarded behavior effect for the price-oriented segment, because δ_2 is positive and significant. However, the rewarded behavior effect is not significant for the service-oriented segment.

We provide further information on the price-oriented versus service-oriented segments in a holdout sample (see Table 4). As expected, the service-oriented segment pays a higher price (even when controlling for accommodation type), stays at a luxury hotel, and often qualifies for Diamond and Platinum tiers. Consistent with our observations above, the price-oriented segment is more likely to stay over the weekend.

Although not the focus of the paper, we close this section with a brief discussion of the competitive variable results. Table 3 shows that competitive price is

¹³ Note we considered an alternative definition of the weekday variable, allowing it to include Thanksgiving as well as Christmas as "weekend." The results were invariant to this robustness check.

Table 4 Characteristics of Segments in Holdout Sample

Summary characteristics of two segments		
Characteristics	Segment 1: Price oriented	Segment 2: Service oriented
	Mean (Std. error)	Mean (Std. error)
Average price per night	\$91.2 (2.45)	\$155.90 (3.30)
Average price paid at economy hotel	\$82.9 (7.3)	\$96.4 (7.9)
Average price paid at regular hotel	\$91.4 (5.4)	\$144.6 (5.7)
Average price paid at luxury hotel	\$126.2 (24.2)	\$177.3 (6.5)
Average paid stays	1.25 (0.31)	3.02 (0.77)
Average percent stays at economy hotel	28.8 (1.9)	9.79 (2.0)
Average percent stays at regular hotel	64.9 (2.0)	41.5 (3.3)
Average percent stays at luxury hotel	6.29 (1.2)	48.7 (3.4)
Average percent weeknight stays (Sun–Thu)	36.6 (2.3)	54.5 (3.5)
Average per period reward cash-ins (%)	1.90 (0.6)	2.12 (1.0)
Percent qualifying for Platinum tier	12.2 (1.4)	26.5 (3.0)
Percent qualifying for Diamond tier	1.4 (0.5)	3.49 (1.1)

Note. Four hundred customers; 0.733 cutoff value.

significant in both segments, and occupation rate is significant for the price-oriented segment, with the expected signs. This suggests the value of including competitive data in the model. However, as noted earlier, these results must be taken as exploratory because our competitive variables are aggregates, and our “no-stay” option includes competitive stays as well as “non-stays.” These cross effects may therefore be underestimated.

4.3. Points Pressure Because of Loyalty Program Components

Points pressure refers to the buildup 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 2(a)–2(c) and Figures 3(a)–3(c), where we use simulations of the model as well as the raw data to calculate the probability of a paid stay (i.e., choosing $m = 1$ and $n = 1, 2, 3,$ or 4) as the customer gets closer to earning the reward.

In Figures 2(a) and 2(b), the simulation begins with customers assigned to one of the two segments. We then assume each of them is five stays away from the next tier and estimate the probability of a paid stay. We do this for 5, 4, . . . , 0 stays away. We then plot the averaged simulated probability of a paid stay versus the closeness of the customer to the next tier

level. Figures 2(a) and 2(b) show that both segments increase their likelihood of paid stays as they get closer to earning membership in the next tier.¹⁴

These results are driven by the positive estimated coefficient ($\lambda > 0$) for customer tier membership. The points pressure behavior can be explained by the dynamic decision process: when the accumulated inventory is getting close to the next threshold, the rewards of higher tier membership can be attained sooner and hence are less discounted. As a result, total long-term utility of purchasing is higher, and customers purchase more.

Although tier classification is based on *previous* purchase levels, an alternative explanation for the finding of points pressure is that tier membership is a proxy for high usage. Our base hotel usage rate variable controls for this because it reflects the number of stays in the initialization period. However, given the importance of the issue, we include Figure 2(c), model-free evidence of the points pressure effect. To create this figure, we divide customers into heavy versus light users and calculate using the actual data the percentage 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 suggests that tier effects are real and not a reflection of the individual usage rate.¹⁵

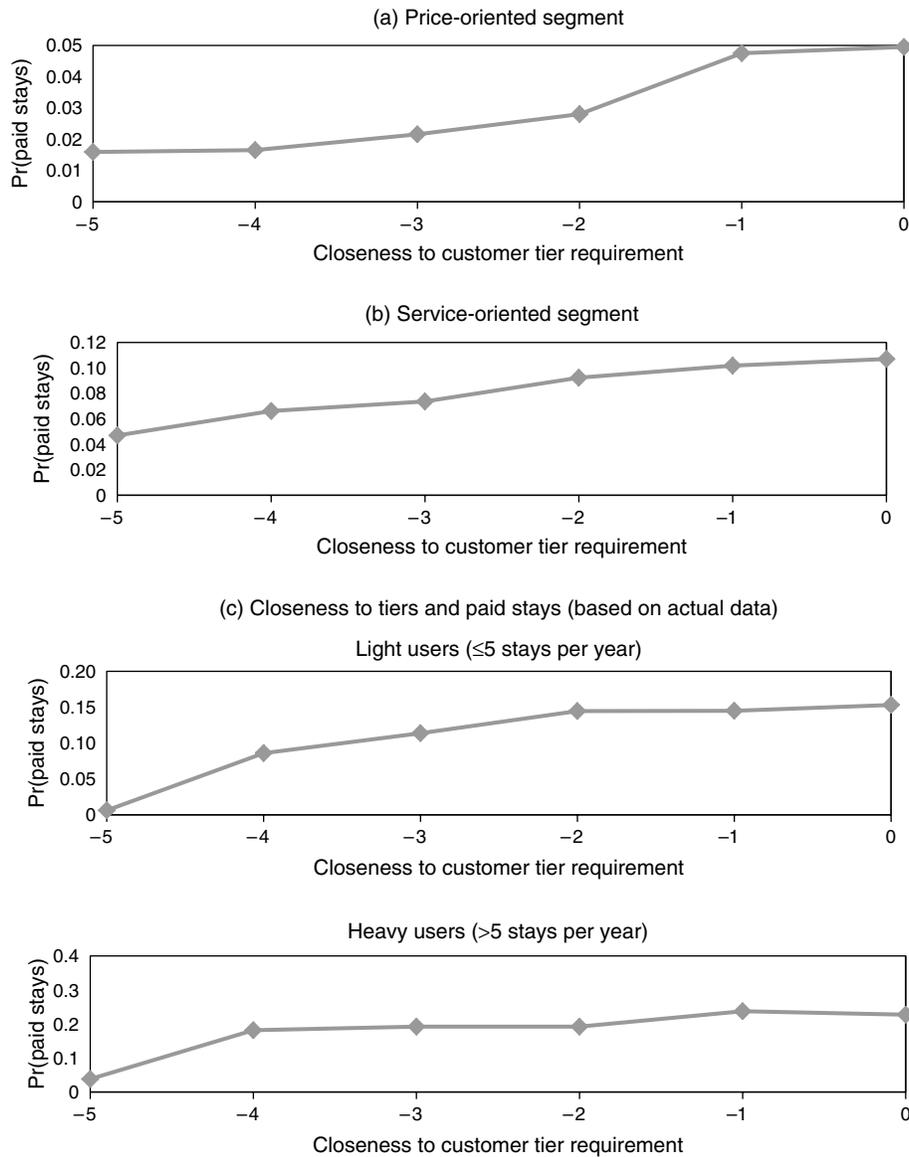
Using the same method as for the customer tier graphs, we investigated points pressure effects for the frequency reward program. Figures 3(a) and 3(b) are based on model-calculated probabilities of paid stay. Per the parameter estimates, we conclude that the points pressure for the price-oriented segment is driven by the free cash-in stays, and upgrades drive it for the service-oriented segment. Figure 3(c) is based on actual data, and akin to Figure 2(c), it provides model-free evidence of points pressure because of the frequency reward program. As in the case of customer tier, the probability of paid stays increases when reaching the reward level, now because of the high utility of a free upgrade as well as higher future utility.

Although points pressure has been demonstrated by previous research, that evidence was with respect

¹⁴ Note that when customers reach the next tier (i.e., zero stays away), their stay probability further increases because of the following: there is an increase in their utility because of the added benefits, and they continue to accumulate points for the next higher tier.

¹⁵ In this calculation, it is possible that different customers are included in different closeness states. To investigate this, we also performed the calculation for customers who were in each of the closeness states at least once during the data. We calculated the average probability of paid stay for each customer in each state and then averaged those averages across customers and plotted closeness to the next tier versus probability of paid stay. The graphs showed the same increasing pattern as in Figure 2, as well as in Figure 3, for the frequency reward component.

Figure 2 The Points Pressure Effect Because of the Customer Tier Component



Note. Panels (a) and (b) are based on simulations assuming no frequency reward program.

to frequency rewards. Our finding is that points pressure exists for *both* frequency reward and customer tier components of a loyalty program.

4.4. Comparison to Two Separate Models for Frequency Reward and Customer Tier

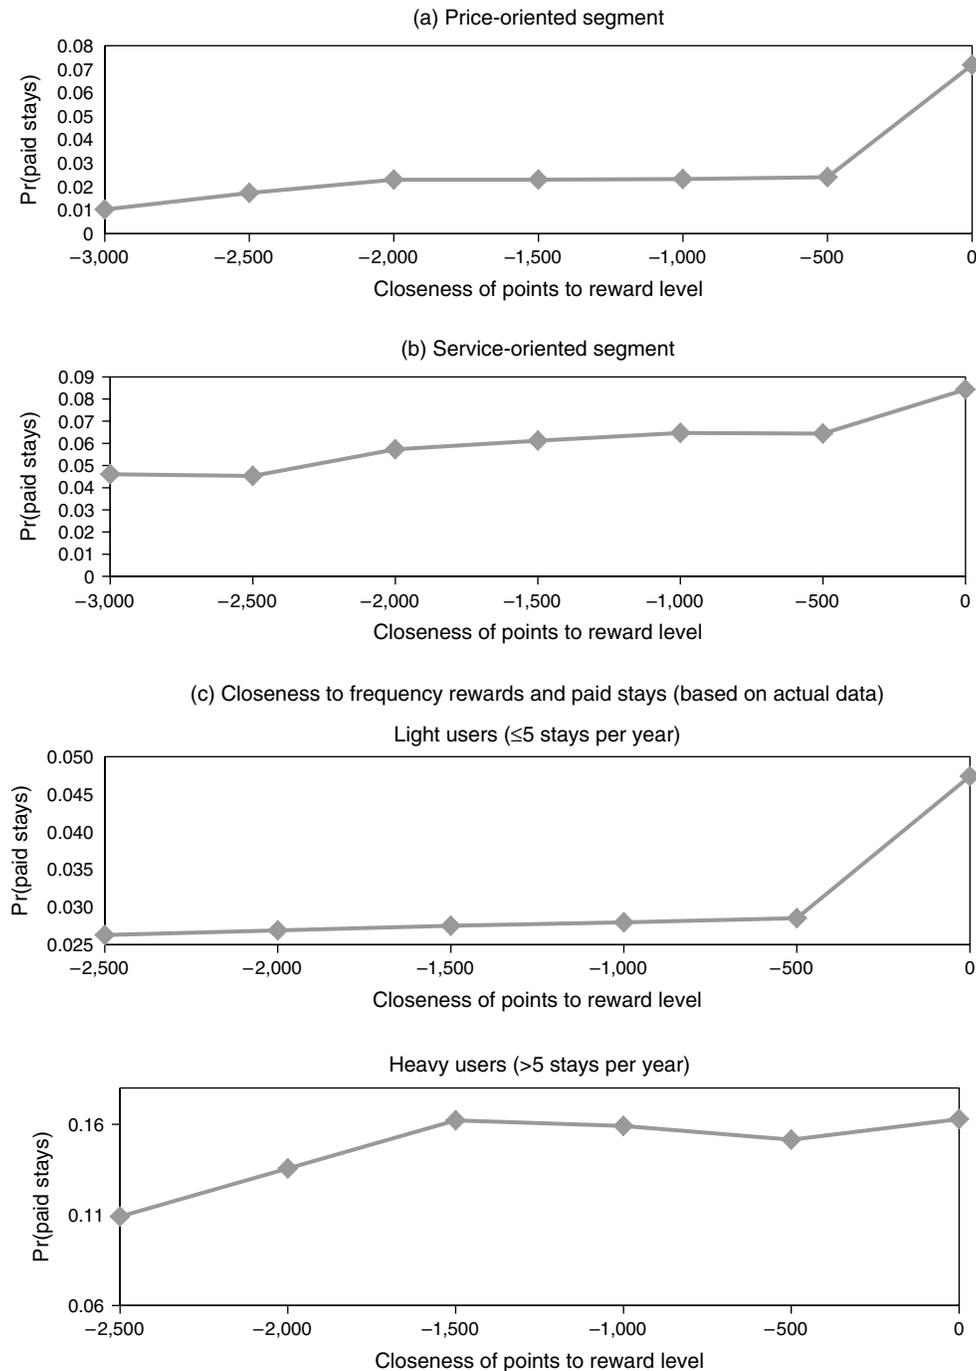
One of the important research questions we aim to address with this research is the value of jointly modeling the frequency reward and customer tier components of loyalty programs, as opposed to modeling them separately. One obvious advantage of the joint approach is that we can predict changes in demand depending on simultaneous changes in the points requirement for both components. We will discuss these analyses in §5. However, to further address this issue, we estimated two additional models.

1. *Only frequency reward program:* In this case, we removed the two customer tier variables (Platinum and Diamond) from the model and reestimated it.

2. *Only customer tier program:* In this case, we removed the frequency reward variables (free night, free upgrade, and inertia of cash-in) and reestimated the model.

We first find that the fit of the combined model (accounting for the number of parameters) is much better than the fit of either the frequency reward-alone model or customer tier-alone model (BIC = 6,953 for combined model, 8,134 for customer tier-alone model, 8,941 for frequency reward-alone model). Clearly, if a research is modeling just one of the components, there is ample explanatory power to be gained by including the other as well.

Figure 3 The Points Pressure Effect Because of the Frequency Reward Component



Note. Panels (a) and (b) are based on simulations assuming no customer tier program.

In addition, the diagnostics from the various models are different. For example, the competitive price variable is significant for only one segment among the four total segments in the customer tier-alone and frequency reward-alone models, whereas this variable is significant for both segments in the combined model. The combined model sensibly suggests competitive prices are important for both segments. As another example, in the frequency

reward-alone model, the coefficients pertaining to the frequency reward program—*inertia of cash-in*, *cash-in for upgrade*, and *cash-in for free nights*—are all larger than their corresponding coefficients in the combined model, suggesting they may be overestimated in the separate models. As Table 4 shows, there is a positive association between membership in higher tiers and *cash-in*. Omitting tier membership therefore (mistakenly) ascribes more importance to *cash-in*.

In summary, modeling the two components together produces better explanatory power and more sensible diagnostics compared to modeling the components separately. In the next section, we show how the combined model can be used to make predictions as we change the points requirement for both components.

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 3 to simulate four scenarios: (1) the hotel employs neither the frequency reward nor the customer tier component, (2) the hotel employs both frequency reward and customer tier components, (3) the hotel employs the frequency reward but not the customer tier component, and (4) the hotel employs the customer tier but not the frequency reward component. For each scenario, we simulate customer behavior for a year and record stays and revenues per customer. The number of stays without cashing in is calculated as the average, over customers and time periods, of the probability of choosing $m = 1$ and $n = 1, 2$, or 3 , i.e., stay without cashing in. Similarly, we calculate the expected number of upgrades ($m = 1$ and $n = 4$) and free stays ($m = 2$ and $n = 1, 2$, or 3). The results are presented in Table 5 and yield several conclusions.

First, both components generate incremental paid stays. This is particularly important for the frequency component, where one might conjecture that free stays would substitute for paid stays. However, this does not occur because the reward generates points

pressure in anticipation of the reward and then a rewarded behavior carryover after the cash-in. Both these factors increase paid stays. Focusing on the price-oriented segment under the frequency reward-alone scenario, the annual number of paid stays per customer with no cash-in increases from 1.11 to 1.37, and the number of cash-in upgrades increases from 0.00 to 0.10. Thus the total increase is 0.36. The same calculation for the service-oriented segment reveals the total increase is 0.68.

Second, there is no cannibalization between the programs; i.e., they do not cannibalize sales that would have occurred anyway. In fact, there may be some synergy, which can occur because both segments gain positive utilities from cashing in and elite status, and purchase probability is nonlinear in these utilities. Note that we are focusing on the revenue-generating aspects of these programs and do not take cost into consideration because of data limitations. For the service-oriented segment, the number of paid stays increase by 0.68 ($2.88 - 2.20$) when the frequency program is offered alone and by 0.11 ($2.31 - 2.20$) when the customer tier program is offered alone. If there was neither substitution nor complementarity, with both programs, the total number of paid stays would be $2.20 + 0.68 + 0.11 = 2.99$. However, the number is a bit higher ($3.05 = 2.50 + 0.55$). The corresponding numbers for the price-oriented segment are 1.50 and 1.55. These synergy effects are not huge but certainly show that the components do not cannibalize each other and show how the model can be used to measure these interactions.

Third, both segments respond to the customer tier program. This can be seen by comparing the base case scenario to the customer tier-alone scenario. For the price-oriented segment, the simulated number of

Table 5 Simulated Annual Stays, Cash-ins, and Upgrades per Customer Under Different Program Scenarios

Scenario	Revenue (\$)	No. of total stays	No. of stays with no cash-in ($m = 1, n = 1, 2, 3$)	No. of cash-in upgrades ($m = 1, n = 4$)	No. of cash-in free stays ($m = 2, n = 1, 2, 3$)
Price-oriented segment					
1. Neither component	343.7	1.11	1.11	0.00	0.00
2. Both	443.8	2.43	1.45	0.10	0.88
3. Frequency reward component only	420.6	2.31	1.37	0.10	0.84
4. Customer tier component only	372.0	1.14	1.14	0.00	0.00
Service-oriented segment					
1. Neither component	1,189.3	2.20	2.20	0.00	0.00
2. Both	1,598.6	3.83	2.50	0.55	0.92
3. Frequency reward component only	1,511.8	3.75	2.36	0.52	0.86
4. Customer tier component only	1,251	2.31	2.31	0.00	0.00

total stays per customer annually increases from 1.11 to 1.14. For the service-oriented segment, the number of stays increases from 2.20 to 2.31.¹⁶ 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. Fourth, both segments respond to the frequency reward program but for different reasons. The price-oriented segment is directly motivated to achieve free stays because of its large, significant coefficient. The service-oriented segment is more motivated by upgrades. However, for an upgrade to happen, the utility for an upgrade has to overcome the negative price effect of a purchase, whereas no such price effect exists for a cash-in free stay. So we find the service-oriented segment compiles both free stays and upgrades.

Fifth, although both components increase paid stays, the frequency component has a stronger impact than the customer tier program. This is especially true for the price-oriented segment, but even for the service-oriented segment. As we noted before, this segment will still cash in for free stays because although it is not as price sensitive as the price-oriented segment, it still is somewhat price sensitive, and a free stay saves a lot of money.

In summary, we find that *both* components generate incremental sales and do not cannibalize each other. If anything, there is some synergy between them. These results of course are a function of our specific estimates and invite further research to generalize. We will take one step in that direction with our second application. The key point is that the dynamic structural model can be used to evaluate crucial issues involving the impact of these components.

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

Here, we explore how a firm could use the model to adjust the requirements of its loyalty program to increase revenues. Two important notes before proceeding: First, the forthcoming scenarios do not consider competitive response. This is certainly a limitation, and although the data requirements would be daunting, our hope is that future research can model consumer response to competitive loyalty programs. Second, one could raise the interesting question of why there should be a need to do this. We are assuming the customer is a sophisticated dynamic optimizer—surely, the firm should be as sophisticated. But our view is that in order for the firm to

improve the effectiveness of its marketing mix, it must first understand how its customers respond to loyalty programs. Then the firm can adjust. In this sense, the firm is a first mover—it tunes the loyalty program requirements taking into account the optimal customer response to these requirements. Note also that we are considering firm revenues, not profits.

Revenues can be increased by setting program requirements correctly, but this is a nontrivial task because there are many complexities and trade-offs. 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.

We conduct a grid search and simulations to determine the frequency reward and customer tier requirements that would most improve revenue per customer. We varied the free upgrade requirement from 1,000 points to 7,000, the free stay requirement at economy, regular, and luxury properties from 3,000, 5,000, and 7,000 points, respectively, to 15,000, all in steps of 1,000. Furthermore, we also varied the paid stay requirement for Platinum and Diamond tiers from 1 and 4, respectively, to 20 in steps of 1. In each cell, we simulated the behavior of 200 consumers. Table 6 presents the results for a subset of the design cells.

We find that although lower requirements encouraged more staying, this reduces revenue, primarily because of free cash-in stays. For example, when the point requirements for an upgrade and free stays at the three types of properties are 1,000, 3,000, 5,000, and 9,000, respectively, the annual revenue per customer is only \$544.70. On the other hand, the revenue increases to \$558.10 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 \$608.30) is 6,000, 8,000, 10,000, and 14,000 for the reward program and 3 and 6 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, whereas it is less stringent for the Platinum and Diamond tiers (because of significant positive utility for segments and because there is less direct cannibalization of paid

¹⁶ Note the baseline sales level for the service-oriented segment (2.20) is greater than that for the price-oriented segment (1.11). This is due to lower alternative-specific constants for the price-oriented segment.

Table 6 Simulated Revenue Projections Under Various Loyalty Program Designs

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
	1,000	3,000	5,000	9,000	7	15	544.7
	1,000	4,000	6,000	13,000	3	15	557.9
	1,000	5,000	11,000	15,000	3	13	555.1
	1,000	6,000	8,000	11,000	8	14	555.3
	2,000	4,000	10,000	13,000	8	11	563.4
	2,000	5,000	7,000	11,000	4	7	559.4
Current	3,000	5,000	8,000	12,000	5	12	558.1
	4,000	7,000	10,000	14,000	3	17	576.8
	4,000	6,000	9,000	13,000	4	8	580.6
	6,000	8,000	10,000	12,000	4	8	591.9
	6,000	8,000	10,000	13,000	8	17	589.8
Optimal	6,000	8,000	10,000	14,000	3	6	608.3
	6,000	8,000	11,000	15,000	3	7	599.8
	6,000	8,000	11,000	13,000	8	16	570.8
	6,000	8,000	11,000	15,000	3	19	576.4
	6,000	8,000	11,000	15,000	4	12	577.7
	6,000	8,000	11,000	15,000	5	9	575.1

stays with free stays in the customer tier program). Of course, our recommendation might not maximize profits. The costs of providing extended services may actually be convex in the number of customers in a given tier. However, data on this cost function could easily be incorporated in the analysis. The example shows how the model can be used to set reward requirements to increase revenues for a loyalty program, taking into account both the frequency reward and customer tier components.

6. 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 incorporating customers' purchase and cash-in decisions, points pressure and rewarded behavior effects, heterogeneity, and forward-looking behavior. Importantly, we model the customer's cash-in decision for free stays or upgrades endogenously—cash-in is a decision the customer makes taking into account the immediate utility of the rewards as well as the long-term implications, such as the impact on future points inventory levels.

We focused on four key research questions: (1) How important is it to combine both components in one model? (2) Does points pressure exist in the context of a two-component loyalty program? (3) How is the market segmented in its response to the combined program? and (4) Do the programs complement each other in terms of the incremental sales they produce? The following paragraphs describe our findings with respect to these questions.

Our basic message is that it is important to jointly model frequency reward and customer tier components of loyalty programs rather than in separate models. We demonstrate this in four ways. (1) We show how the combined model can be used to predict the impact of simultaneous changes in requirements in both components. (2) We show how the combined model can investigate potential complementarity between the components. (3) We find that the combined model adds explanatory power to a model that captures only one of the components. (4) We find that the combined model yields more sensible diagnostics compared to the single-component models. For example, the frequency component-alone model may attribute too much impact to the frequency reward variables because these can be correlated with tier membership. Overall, the conclusion is simple—a marketing mix model is better to the extent that it includes more elements of the marketing mix. For example, modeling price effects but not advertising effects results in a less precise model. Similarly, modeling just one component of the loyalty program results in a loss of information relative to a model incorporating both components.

We find that the phenomenon of "points pressure," previously demonstrated with respect to frequency reward programs, also exists with respect to customer tiers. We find that consumers increase their purchase rate as they get closer to a reward, whether that reward be a free hotel stay or a higher customer tier.

The segmentation results are quite interesting. We find that customers are heterogeneous with respect to their response to these programs. We find a price-oriented segment that strongly values the frequency reward and reacts more strongly to prices,

especially of the focal brand. This is in contrast to the service-oriented segment that highly values upgrades, is less price sensitive, and gains more utility especially from luxury accommodations. Both segments highly value the customer tier component. Reinforcing these findings are post hoc analyses that show the service-oriented segment pays higher prices per stay, has more paid stays (especially in luxury accommodations), and is more likely to quality for the Platinum or Diamond tiers.

We investigated the extent to which the frequency and tier components generate incremental sales, as opposed to producing stays that substitute for purchases that would have occurred anyway. We found that both components generate incremental sales. We also investigated whether employing both programs resulted in complementing each other's contribution to incremental sales. We found there was a slight synergy between the programs but not a huge amount. Overall, each component contributes to increased revenues and does not interfere with the other. This is somewhat a function of the segments; for example, Segment 2 especially valued the higher tiers whereas Segment 1 placed less value on them. However, the finding of little synergy may be the result of the fundamentally different designs of the components—frequency reward is a one-off “gift” when a customer makes a proactive effort to cash in; customer tier provides a long-term increase in ongoing service.

Our results have several managerial 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. These effects are important—the firm is “giving away” a free night or a free flight, but points pressure generates significant increases in paid stays as the customer builds up toward the reward and benefits from a positive halo after the reward is cashed in. Consistent with our joint revenue maximization analysis, which shows that the firm should capitalize on these effects by increasing reward thresholds, 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; see Higgins 2006) while emphasizing the benefits of a customer tier program (Elliott 2006).

For researchers, the implications are as follows. (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. 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 than do myopic models. Our results reinforce the value of dynamic structural models to evaluate customer response to future-oriented incentives, a theme similar to recent pioneering work on salesperson incentive plans (Chung et al. 2009, Misra and Nair 2011).

There are several avenues for extending this research. First, our model may be applied to other industries such as airlines to generalize the substantive results. Second, although we incorporated measures of competitive price and occupancy rate and found indeed that they exerted a significant impact on customers' stay decisions, these findings must be taken as exploratory because the competitive variables are aggregated across competitors, and our no-stay option includes competitive stays as well as “non-stays.” However, in previous versions of the model when we did not include competitive data, our main findings—e.g., the existence and interpretation of two segments and the “non-finding” with regard to component complementarity—were very similar, and the fact that we found significant effects despite measurement error suggests that future work on examining competitive effects in more depth for loyalty programs would be worthwhile. Third, future work should delve into the factors that decrease perceived costs and increase perceived value of the frequency reward component (see Kivetz and Simonson 2002, 2003; Kwong and Soman 2006; Roehm et al. 2002; and Drèze and Nunes 2008 for insights on the value side of this equation) as well as pinpoint the exact source of these incremental sales as a result of points pressure. Fourth, we need to better understand when firms should offer free products versus other types of rewards (e.g., Kim et al. 2001) or combine currencies for cash-in (e.g., Drèze and Nunes 2004). Fifth, it would be interesting to examine the role of third-party payers in a reward program context. There are many cases, especially in the travel industry where a traveler's business is paying for the trip, although there may be some internal pressures put on traveler's to conserve funds. Along these lines, we note that free stays are disproportionately less likely to begin during the week (71.6% of paid stays begin during the week, whereas 58.9% of free stays begin

during the week), consistent with the notion that travelers use free stay for leisure travel, presumably using the points they accumulate on business travel. Sixth is the consideration of price endogeneity. There is plenty of evidence that current prices are often determined by previous prices (e.g., Nijs et al. 2007). Furthermore, the inclusion of endogeneity in dynamic structural choice models is not common (e.g., see Erdem and Keane 1996; Erdem et al. 2003, 2008; Hendel and Nevo 2006; Sun 2005; as well as Lewis 2004). Therefore, future research should consider this issue. Seventh would be to examine the case where the rewards from the frequency reward and customer tier program overlap—e.g., when one can obtain an upgrade by either cashing in points or belonging to the Platinum tier. To accommodate this would probably require the choice set to change depending on tier status, a conceptually “doable” modification but one that would make the model more complex. In conclusion, although 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 fully understand this ever-popular marketing activity.

Acknowledgments

The authors thank the editor-in-chief, associate editor, and two anonymous reviewers for their helpful comments and suggestions on earlier drafts of this manuscript. The authors also thank the attendees at the Marketing Science Conference, Marketing Dynamics Conference, Association for Consumer Research Asia-Pacific Conference, North American Society for Marketing Education in India International Conference, and Yale Conference on Customer Insights as well as seminar participants at the Indian School of Business, Rensselaer Polytechnic Institute, University of Maryland, and University of Groningen for their helpful suggestions on this research. Part of this research was conducted at the Indian School of Business and Rensselaer Polytechnic Institute when the first author was a visiting scholar.

Appendix. Solution Approach for Parameter Estimation

We embed a dynamic programming backward induction method with a finite horizon, 1 to $T = 26$ periods, within a maximum likelihood estimation procedure. The state space is the Cartesian product $S \times C \times F \times Y \times V$, where

- $S = \{0, 1\}$: 1/0 denotes that consumer stays/does not stay at the focal hotel,
- $C = \{0, 1\}$: 1/0 denotes that consumer redeems for reward/does not redeem for reward,
- $F = \{1, 2, \dots, 26\}$: $f \in F$ denotes that the current period is the f th fortnight of the year,
- $Y = \{1, \dots, 26\}$: $y \in Y$ denotes that consumer has accumulated y paid stays, and
- $V = \{0, 1, \dots, 200\}$: $v \in V$ denotes that consumer has accumulated a points inventory of 200 v .

For the ease of exposition, we use m to denote the customer’s stay decision (Equation (1a)). Note that SC are related to the consumer’s choice of m such that $S = 1$ ($m \geq 1$) and $C = 1$ ($m = 2$).

Maximum Likelihood Estimation Procedure

Begin iteration for an arbitrary set of parameter values

For customer $i = 1$ to 400

For period $t = 1$ to $T = 26$

Determine the levels of state variables for each customer

For each of the seven possible combinations of (m, n)

Compute customer’s current period utility (Equations (3) and (4))

Recall future value function determined as a function of customer i ’s state variables

Compute probability of choice using nested 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

Dynamic Programming Algorithm for Computing Value Function

For customer $i = 1$ to 400

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

For a randomly set of selected state vector (S, C, F, Y, V) ¹⁷

For each of the seven possible combinations of (m, n)

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

Compute customer’s (1) expectation of price,

(2) expectation of nights stay, and

(3) customer tier status (see state transition map for customer points inventory $(v_{i,t+1})$ for computation details)

Compute the option-specific current period utility (Equations (3) and (4))

Compute the option-specific value function (Equation (9)) 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

¹⁷We restrict Y to be no greater than F to reflect the fact that the paid stay inventory expires at the end of a calendar year.

Increment to next state vector in the selected set
 Having obtained the value for each state vector in the selected set, we use simple linear interpolation to compute the value function for each possible value of the state vector (S, C, E, Y, V).
 Decrement to next t
 Increment to next customer

State Transition Map

$$s_{i,t+1} = \begin{cases} 1 & \text{if } m = 1, 2, \\ 0 & \text{if } m = 0, \end{cases}$$

$$c_{i,t+1} = \begin{cases} 1 & \text{if } m = 2, \\ 0 & \text{if } m = 0, 1; \end{cases}$$

$$f_{t+1} = \begin{cases} f_t + 1 & \text{if } f_t < 26 \text{ (or between 27 and 52),} \\ 0 & \text{if } f_t = 26 \text{ (or 52);} \end{cases}$$

$$y_{i,t+1} = \begin{cases} y_{it} + 1 & \text{if } f_t < 26 \text{ (or between 27 and 52) and} \\ & m_{it} = 1, \\ y_{it} & \text{if } m_{it} = 0, 2, \\ 0 & \text{if } f_t = 26 \text{ (or 52);} \end{cases}$$

$$v_{i,t+1} = \begin{cases} v_{it} + W^{-1} \cdot \sum_{w=1}^W \text{prate}_{it} \cdot p_{in,w} \cdot \text{LOS}_{i,w} & \text{if } m_{it} = 1, \\ v_{it} & \text{if } m_{it} = 0 \\ v_{it} - R_1 & \text{if } m_{it} = 1 \text{ and } n_{it} = 4, \\ v_{it} - R_{(n_{it}+1)} & \text{if } m_{it} = 2; \end{cases}$$

where 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), and Platinum and Diamond members will earn an extra 15% and 30% of their points, respectively, for each dollar they spend; thus,

$$\text{prate}_{it} = \begin{cases} 5 & \text{if } y_{it} < 5, \\ 5.75 & \text{if } y_{it} \geq 5 \text{ and } y_{it} < 12, \\ 6.5 & \text{if } y_{it} \geq 12. \end{cases}$$

The variables $p_{in,w}$ and $\text{LOS}_{i,w}$ 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 normal (LTN):

$$p_{in,w} \sim \text{LTN}(\overline{p}_{in}, \sigma_{p_{in}}^2) \text{ and } \text{LOS}_{i,w} \sim \text{LTN}(\overline{\text{LOS}}_i, \sigma_{\text{LOS}_i}^2).$$

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 estimate the expectations Equations (11) and (13), we follow a two-step approach for compiling customer-specific prices and occupancy rates. The first step is to determine the location of the city relevant to the customer; the second is to assemble the prices and occupancy rates. In the first step, we know for each customer the location he or she stayed in when staying in the focal hotel. For other periods, we impute a location based on the observed choice

history of the consumer. We assume that in future periods, customers are more likely to visit the same cities they had visited recently. In the second step, note that the competitive occupancy rates and prices in these markets are observed. For the focal hotel prices, we use the observed price if the customer stayed in the hotel in that week; otherwise, we use the average price paid by other customers staying in the focal hotel in that week. For customers who never stayed at a particular type of accommodation type in a relevant city, the time series is constructed using prices paid by other customers for that accommodation type in that city. Our approach is similar to the imputation of missing price information in packaged goods categories (e.g., Erdem et al. 2008).

Interpolation Technique (Based on Keane and Wolpin 1994)

The dynamic programming problem in Equation (9), i.e., the Bellman equation for customer i at time t , has five state variables and all are integers: INV (inventory of points), $INVP$ (inventory of paid stays), F (calendar period of the year), C (past purchase/cash-in decision), and D (past accommodation decision). We chose ranges for these state variables based on the data and consider discrete points within the ranges where the problem is evaluated. Whereas the respective ranges of $INVP, F, C,$ and D are small, INV 's range is in hundreds of thousands. Even though we discretize the range of INV into 8,000 distinct points, the overall state space is still very large and thus suffers from the "curse of dimensionality." To alleviate the severe computation burden, we use the linear interpolation technique of Keane and Wolpin (1994). We follow the backward induction, and starting from the last period, we follow three steps for each period: (1) compute the value functions at a subset of points in the state space for that period, (2) use the value functions at these state points to fit a regression line with the value function as the dependent variable and the state variables as the independent variables, and (3) based on the regression results, impute the value functions for all remaining state points in that period. For each period, the computation of value function in step (1) is based on 500 draws, and the interpolation in step (3) is based on about 2% (randomly chosen) of all possible points in the state space.

References

Barnett, M. 2010. How brands can play the loyalty card. *Marketing Week* 33(41) 28–30.
 Ben-Akiva, M., S. R. Lerman. 1985. *Discrete Choice Analysis*. MIT Press, Cambridge, MA.
 Blattberg, R. C., B. Kim, S. A. Neslin. 2008. *Database Marketing: Analyzing and Managing Customers*. Springer, New York.
 Bolton, R. N., P. K. Kannan, M. D. Bramlett. 2000. Implications of loyalty program membership and service experiences for customer retention and value. *J. Acad. Marketing Sci.* 28(Winter) 95–108.
 Bucklin, R. E., S. Gupta, S. Siddarth. 1998. Determining segmentation in sales response across consumer purchase behaviors. *J. Marketing Res.* 35(2) 189–197.

- 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. *J. Marketing Res.* 41(May) 184–196.
- Chintagunta, P. 2002. Investigating category pricing behavior in a retail chain. *J. Marketing Res.* 39(May) 141–154.
- Chung, D., T. Steenburgh, K. Sudhir. 2009. Do bonuses enhance sales productivity? A dynamic structural analysis of bonus-based compensation plans. Working paper, School of Management, Yale University, New Haven, CT.
- Cramer, J. S. 1999. Predictive performance of the binary logit model in unbalanced samples. *J. Roy. Statist. Soc. Ser. D* 48(1) 85–94.
- Deighton, J. A. 2000. Frequency programs in service industries. D. Iacobucci, T. A. Swartz, eds. *Handbook of Services Marketing and Management*. Sage Publications, Thousand Oaks, CA, 401–408.
- Dowling, G. R., M. Uncles. 1997. Do customer loyalty programs really work? *Sloan Management Rev.* 38(Summer) 71–82.
- Drèze, X., J. C. Nunes. 2004. Using combined-currency prices to lower consumers' perceived cost. *J. Marketing Res.* 41(February) 59–72.
- Drèze, X., J. C. Nunes. 2008. Feeling superior: The impact of loyalty program structure on consumers' perceptions of status. *J. Consumer Res.* 35(6) 890–905.
- Elliott, C. 2006. I'm elite and you're not. *New York Times* (December 26) <http://www.nytimes.com/2006/12/26/business/26elite.html>.
- Erdem, T., M. P. Keane. 1996. Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Sci.* 15(1) 1–20.
- Erdem, T., S. Imai, M. P. Keane. 2003. Brand and quantity choice dynamics under price uncertainty. *Quant. Marketing Econom.* 1(March) 5–64.
- Erdem, T., M. P. Keane, B. Sun. 2008. A dynamic model of brand choice when price and advertising signal product quality. *Marketing Sci.* 27(6) 1111–1125.
- Erdem, T., M. P. Keane, T. S. Oncu, J. Strebler. 2005. Learning about computers: An analysis of information search and technology choice. *Quant. Marketing Econom.* 3(3) 207–246.
- Gönül, F., K. Srinivasan. 1996. Estimating the impact of consumer expectations of coupons on purchase behavior: A dynamic structural model. *Marketing Sci.* 15(3) 262–279.
- Gupta, S., D. R. Lehmann, J. A. Stuart. 2004. Valuing customers. *J. Marketing Res.* 41(February) 7–18.
- Hartmann, W. 2010. Demand estimation with social interactions and the implications for targeted marketing. *Marketing Sci.* 29(4) 585–601.
- Hartmann, W. R., V. B. Viard. 2008. Do frequency reward programs create switching costs? *Quant. Marketing Econom.* 6(2) 109–237.
- Hendel, I. A. Nevo. 2006. Measuring the implications of sales and consumer stockpiling behavior. *Econometrica* 74(6) 1637–1673.
- Higgins, M. 2006. Hotel reward programs: That "free" hotel stay gets less free. *New York Times* (December 24) <http://travel.nytimes.com/2006/12/24/travel/24prac.html>.
- Kamakura, W. A., G. J. Russell. 1989. A probabilistic choice model for market segmentation and elasticity structure. *J. Marketing Res.* 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. *Rev. Econom. Statist.* 76(4) 648–672.
- Keiningham, T. L., T. G. Vavra, L. Aksoy. 2006. Managing through rose-colored glasses. *Sloan Management Rev.* 48(Fall) 15–18.
- Kim, B.-D., M. Shi, K. Srinivasan. 2001. Reward programs and tacit price collusion. *Marketing Sci.* 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. *J. Marketing Res.* 39(2) 155–170.
- Kivetz, R., I. Simonson. 2003. The idiosyncratic fit heuristic: Effort advantage as a determinant of consumer response to loyalty programs. *J. Marketing Res.* 40(May) 454–467.
- Kivetz, R., O. Urminsky, Y. Zheng. 2006. The goal-gradient hypothesis resurrected: Purchase acceleration, illusionary goal progress, and customer retention. *J. Marketing Res.* 43(February) 39–58.
- Klemperer, P. D. 1987. Markets with consumer switching costs. *Quart. J. Econom.* 102(May) 375–394.
- Kopalle, P., S. A. Neslin. 2003. The economic viability of frequency reward programs in a strategic competitive environment. *Rev. Marketing Sci.* 1(1) Article 1.
- Kumar, N., R. Rao. 2006. Using basket composition data for intelligent supermarket pricing. *Marketing Sci.* 25(2) 188–199.
- Kwong, J. Y., D. Soman. 2006. The role of fluency in the decision to spend loyalty program points. Working paper, University of Toronto, Toronto.
- Lal, R., D. E. Bell. 2003. The impact of frequent shopper programs in grocery retailing. *Quant. Marketing Econom.* 1(2) 179–202.
- Lewis, M. 2004. The influence of loyalty programs and short-term promotions on customer retention. *J. Marketing Res.* 41(August) 281–292.
- Liu, Y. 2007. The long-term impact of loyalty programs on consumer purchase behavior and loyalty. *J. Marketing* 71(4) 19–35.
- Misra, S., H. Nair. 2011. A structural model of sales-force compensation dynamics: Estimation and field implementation. *Quant. Marketing Econom.* 9(3) 211–257.
- Neslin, S. A. 2002. *Sales Promotion*. Marketing Science Institute, Cambridge, MA.
- Nijs, V. R., S. Srinivasan, K. Pauwels. 2007. Retail price drivers and retailer profits. *Marketing Sci.* 26(4) 473–487.
- Nunes, J. C., X. Drèze. 2006. The endowed progress effect: How artificial advancement increases effort. *J. Consumer Res.* 32(March) 504–512.
- Osborne, M. 2007. Consumer learning, switching costs and heterogeneity: A structural estimation. *Quant. Marketing Econom.* 9(1) 25–70.
- Peppers, D., M. Rogers. 1997. *Enterprise One to One*. Doubleday, New York.
- Reynolds, J. 1995. Database marketing and customer loyalty: Examining the evidence. *Eur. Retail Digest* 7(July) 31–38.
- Roehm, M. L., E. B. Pullins, H. A. Roehm Jr. 2002. Designing loyalty-building programs for packaged goods brands. *J. Marketing Res.* 39(May) 202–213.
- Rust, R. T., R. L. Oliver. 2000. Should we delight the customer? *J. Acad. Marketing Sci.* 28(Winter) 86–94.
- Sharp, B., A. Sharp. 1997. Loyalty programs and their impact on repeat-purchase loyalty patterns. *Internat. J. Res. Marketing* 14(5) 473–486.
- Shugan, S. M. 2005. Brand loyalty programs: Are they shams? *Marketing Sci.* 24(2) 185–193.
- Sun, B. 2005. Promotion effect on endogenous consumption. *Marketing Sci.* 24(3) 430–443.
- Sun, B., S. Neslin, K. Srinivasan. 2003. Measuring the impact of promotions on brand switching under rational consumer behavior. *J. Marketing Res.* 40(4) 389–405.
- Swait, J. 2001. A non-compensatory choice model incorporating attribute cutoffs. *Transportation Res. Part B* 35(10) 903–928.
- Taylor, G. A., S. A. Neslin. 2005. The current and future impact of a retail frequency reward program. *J. Retailing* 81(4) 293–205.
- Thaler, R. 1985. Mental accounting and consumer choice. *Marketing Sci.* 4(3) 199–214.
- Zeithaml, V. A., R. T. Rust, K. N. Lemon. 2001. The customer pyramid: Creating and serving profitable customers. *Calif. Management Rev.* 43(Summer) 118–142.