Preliminary: Please do not cite without permission

Quantity-Based Price Discrimination Using Frequency Reward Programs

Wesley R. Hartmann
Assistant Professor
Stanford Graduate School of Business
518 Memorial Way
Stanford, CA 94305-5015
hartmann_wesley@gsb.stanford.edu
P. (650) 725-2311
F. (650) 725-7979

V. Brian Viard
Assistant Professor
Stanford Graduate School of Business
518 Memorial Way
Stanford, CA 94305-5015
viard_brian@gsb.stanford.edu
P. (650) 736-1098
F. (650) 725-0468

October 10, 2005

Abstract

This paper explores the quantity-based price discrimination and switching costs effects of reward programs by analyzing individual level choice data of participants of a frequent golfer program. Golf provides a setting to evaluate reward programs that are not designed to exploit a principal agent problem, as is likely the case in industries such as airlines and hotels. We estimate a dynamic structural model of consumer choice using a reward program database for a golf course. The expected value of participating in the program is shown to be negligible for low purchase probability golfers, but increasing monotonically with golfers' purchase probabilities. The marginal effects of the reward on purchase probabilities are also greater for high-volume consumers. If a firm limits the reward earning opportunities, however, these effects may be lower for high-volume individuals who earn all possible rewards without adjusting their behavior. Other counterfactual analyses demonstrate that the number of credits necessary to qualify for a reward and the time horizon for qualifying also influence the relationship between the program's effects on qualifying purchases and consumer types. We quantify the price discrimination effect of the program by calculating the uniform price for each individual that is necessary to achieve effects equivalent to that created by the reward program. Under certain program parameters, price discrimination can be virtually eliminated if it is an unintended consequence, as might be the case for airline and hotel programs. We also explore stylized examples of reward programs to compare the endogenous switching costs created by reward programs considered in the theoretical literature to those created by commonly implemented reward programs.

1 Introduction

Reward programs, also known as frequent-purchaser programs or loyalty programs, have become ubiquitous. Besides the well-known examples of airline frequent flyer programs, reward programs are employed in coffee shops, sandwich shops, grocery stores, hotels, and even shoe retailers.[1] Despite the prevalence of these programs, there is scarcely any empirical evidence evaluating their effectiveness. While previous academic work and many managers claim that reward programs increase loyalty and therefore provide a competitive advantage to the adopting firm or are an effective defensive maneuver against competing firms who already offer similar programs, other authors indicate that such programs offer little or no benefit and may not be sustainable.[2] In this paper we examine two common aspects of reward programs, price discrimination based on cumulative purchase quantity, and the creation of endogenous switching costs.

While firms may institute reward programs for a variety of reasons including scope economies, customer information, or exploiting principal-agent problems of business travelers, an inherent characteristic of these programs is quantity-based price discrimination. In some cases, this may be profitable and motivate use of the program. If higher volume customers have a relatively greater marginal elasticity of demand than lower volume customers, volume discounts may be profitable. In fact a reward program may be a reasonable substitute for a traditional volume discount in markets where volume accumulates across time, which makes it difficult for customers to reveal their types on their first purchase.[3] In other cases, the quantity-based price discrimination aspect may merely be a feature that the firm must manage while trying to obtain another objective with the program.

We distinguish between two types of discrimination created by reward programs: discriminatory effects on the expected values of the rewards and discriminatory effects on the purchase incentives of customers. Customers who, all else being equal, purchase greater quantities will have greater expected values of the reward because their likelihood of earning it is greater. However, the program can actually have a greater effect on the purchase incentives of lower-volume consumers relative to higher-volume customers, depending on the program parameters. For example, when a firm offers a single reward-earning opportunity, those customers that would purchase the qualifying amount absent the reward program will certainly receive the

---

[1] DSW (Discount Shoe Warehouse) offers a "Reward Your Style" program in which customers receive $25 when their cumulative purchases total $250 (see www.dswshoe.com/rys.jsp).

[2] See for example Borenstein (1996) who argues that the trading stamps employed by grocery stores in the 1950s and 1960s offered little, if any, value and eventually collapsed.

[3] A reward program may offer advantages over a traditional volume discount in implementing volume-based price discrimination. A traditional volume discount discriminates based on the quantity purchased at a given time. If this bundle includes future units of the good, varying marginal costs over time could lead to consumers selectively purchasing bundles with higher average costs. A reward program provides the discount ex-post, so that customers purchase all qualifying units at a price related to their marginal cost, then receive a discount on a single unit for which the firm can easily control the marginal cost.

reward, but their purchase incentives prior to earning the reward will be only marginally affected. Moreover, the effect will be very small relative to lower-volume consumers. Pre-reward purchase incentives are affected to the extent that the purchase is necessary to increase the customer's likelihood of receiving the reward. In choosing the number of reward earning opportunities, the length of the qualifying period, the number of credits required to earn a reward, the magnitude of the reward, and the length of time before a reward expires, the firm determines the impact on purchase incentives for each customer type.

Previous academic work has focused on the ability of reward programs to endogenously create switching costs. In a survey article of switching costs, Klemperer (1995) includes "discount coupons and similar devices" such as frequent flyer programs as a category of switching costs involving brand loyalty (page 517). The argument is that as consumers purchase more from a firm (i.e. have more credits) their purchase probabilities will increase. While this effect has been modeled before (see Caminal & Matutes (1990) and Fudenberg and Tirole (2000)) there has been relatively little empirical work to quantify these effects in actual programs. Moreover, these models have generally considered stylized reward programs such as discrete-choice, two-period models that result in price schedules that monotonically decrease with cumulative quantity purchased. Using simulated examples of reward programs we show that while the monotonic relationship between switching costs and number of credits holds in these stylized models, it does not necessarily hold for most observed reward programs, because these programs stipulate a non-monotonic price to credit relationship. In future work we plan to estimate the effects of the golf reward program on switching costs in our data.

To analyze the quantity-based price discrimination and switching costs aspects of reward programs, we estimate a dynamic model of consumer choice for a reward program offered by a golf course. From the estimated model, we can calculate the effective price discrimination of the reward program for different types of individuals under multiple sets of program parameters. We also outline a method for measuring the switching costs faced by individuals at various points in the program which will allow us to quantify switching costs faced by different consumer types and at different levels of accumulated credits.

We demonstrate that reward programs generate quantity-based price discrimination, whether desired by the firm implementing the program or not. These insights come from considering an empirical model of consumer demand under a firm's equilibrium choice of a reward program, rather than focusing on the competitive interactions of firms. We also demonstrate through simulations of stylized reward program examples that reward programs as commonly implemented result in zero switching costs at certain points in the program. This effect does not appear in simplified two-period, discrete-choice models that miss important aspects of reward programs by only considering a single reward-earning opportunity and implicitly assuming a monotonic relationship between price and number of purchases.

2 Related Literature

The majority of work on reward programs uses two-period theoretical models to analyze the competitive implications of the programs. This theoretical literature has primarily focused on the role of switching costs, though the relationship with price discrimination has been recognized. There is limited empirical investigation of reward programs and what little has been done primarily focuses on experiments designed to evaluate the psychological effects of reward programs on consumer behavior.

Much of the theoretical work followed Klemperer (1987a and 1987b) which analyzed exogenous switching costs but noted the role of switching costs in reward programs. Caminal and Matutes (1990) consider the case of endogenous switching costs and find that, in equilibrium, firms will commit to future prices for returning customers. The firms commit to a declining price path in which second-period prices are below marginal cost. Firms also set a second-period price for first time buyers. Competition in this model results in lower profits than would have been earned if firms could not commit to second period prices. This result contrasts with the finding by Klemperer that exogenous switching costs soften competition.

Fudenberg and Tirole (2000) positions Caminal and Matutes (1990) as a special case of models of customer poaching and brand switching. The authors consider various forms of price discrimination in which past purchases reveal demand characteristics of customers, such that firms can condition prices on previous purchases in order to segment customers. The model in Caminal and Matutes (1990) differs in that consumers are randomly relocated in the second period so that their first purchase reveals nothing about their second-period preferences. Price discrimination still plays a role, however because the prices guaranteed to repeat purchasers segment customers in the second period.

A limiting feature of two period models of reward programs is that only single reward-earning opportunities are possible and the prices charged with respect to cumulative purchases is monotonically declining. In most observed reward programs, on the other hand, the price path is non-monotonic (i.e. the price is typically constant with discrete drops when rewards are redeemed). We explore the implications of multiple reward-earning opportunities, non-monotonic relationships between price and purchases and varying program lengths to evaluate both the quantity-based price discrimination and switching cost aspects of the programs.

An important aspect introduced by the consideration of more than three periods is the role of marginal switching costs. Typically switching costs are thought of as a one-time investment that leads to a preference for a firm that is ex-ante homogeneous. However, subsequent purchases from the firm may increase or decrease the switching costs. While not an important characteristic for all switching cost environments, it is for reward programs. For most observed reward programs, marginal switching costs will be positive most of the time, but drop when a reward is redeemed, resulting in negative marginal switching costs. We demonstrate the pattern of marginal switching costs in such reward programs using a simple illustrative example in the following section.

There is a limited amount of empirical work on reward programs. Much of the work involves experiments designed to test psychological effects of reward programs on consumer behavior (e.g. Kivetz and Simonson (2002) and Kivetz et. al. (2005)). We instead focus on the economic incentives of the programs. The most closely related paper to ours is Lal and Bell (2003) who use actual grocery data. They investigate the effectiveness of a frequent shopper program in grocery retailing and find results consistent with our analysis of how reward program parameters influence quantity based-price discrimination. They find that a grocery chain's reward program increased purchases by its lower-spending customers by more than the increase in purchases among its higher-spending consumers. Although the program lost money among high-volume consumers after taking into account their higher redemption rates, increased profits among low-volume consumers more than offset these losses to net an overall gain.

This is exactly what one would expect from a reward program such as they analyze in which there is a single reward opportunity and a relatively low threshold for qualifying. Specifically, because customers qualify at low purchase levels, the marginal qualifying consumer is a relatively low-spending individual. As a result, the program discriminates within low-types, which increases profits, while limiting redemption losses to high types by limiting the number of rewards that can be earned. We describe such an effect with a counterfactual in Section 7.

Lal and Bell rationalize their findings in a model with two customer types, switchers (i.e. "cherry-pickers") and non-switchers. The model implies that the reward program will only influence the switchers' behavior. While such a result implies a switching cost explanation for reward programs, the model does not allow for the fact that had the grocery chain redefined the parameters of the reward program, it may have excluded all low-type consumers and included only high-type consumers. Their model will not accommodate this because high type consumers are assumed to be loyal non-switchers.

3 Illustrative Examples

Three main dimensions summarize a reward program: the number of credits required to qualify, the value of the discount rewarded, and the time horizon of the qualifying period. In the first part of this section we illustrate how these attributes can be manipulated to segment consumers. Figure 1 demonstrates how these dimensions can be manipulated to segment low- from high-volume consumers. The figure shows the expected value of a hypothetical reward program (given a $14.75 reward, 10 credits needed to earn it and varying time horizons) as a function of consumers' exogenous probabilities of purchase.[4] Increasing the number of credits required to qualify flattens this curve. For a given discount, setting this very low will cause most low-volume consumers to qualify and dilute the segmentation power of the program. Setting it very high, on the other hand, will cause only a small group of high-volume consumers to qualify. Similarly, decreasing the time horizon of the qualifying period flattens the curve. For a given discount, a lengthy time horizon results in most consumers qualifying while a very short time horizon results in only a few high-volume consumers qualifying.

Referring again to Figure 1, a reward program is only useful for purposes of price discrimination if high-volume consumers have a greater elasticity of demand than low types. Given a time horizon and number of credits required to qualify, the firm should set the discount optimally based on the distribution of types in the market and their relative elasticities.

Switching costs have been a commonly analyzed attribute of reward programs in theoretical models. However, there has been little investigation of these costs in models reflecting the long time horizons over which credits and rewards can be earned in actual programs. An important feature of most reward programs is that the relationship between price and the number of credits (i.e. purchases) is non-monotonic. That is, consumers must accumulate a certain number of credits before a reward is earned and the full impact of the discount is realized. This results in a non-monotonic relationship between switching costs and number of credits such that the costs rise until a reward is received and used and then drop substantially (typically back to zero).[5] One way to avoid this pattern is to specify a monotonically declining relationship between price and cumulative quantity purchased. While this avoids the non-monotonic switching costs pattern, it does require that after enough purchases, additional purchases will increase the magnitude of the switching costs only negligibly.

In the remainder of this section we demonstrate these patterns using simple illustrative examples of a non-monotonic and a monotonic reward program. We quantify the cumulative and marginal switching costs of these two program types. In the non-monotonic case, which corresponds to most observed programs, cumulative switching costs follow repeated patterns of increasing while a consumer qualifies for a reward and then falling when a reward is used, implying the existence of negative marginal switching costs. In the monotonic case, cumulative switching costs increase monotonically, while the marginal switching costs uniformly decrease toward zero over time.

[4] Exogenous in the sense that neither current nor future prices affect the purchase probability.
[5] Costs would not drop all the way to zero if the consumer earns an additional credit toward the next reward for the purchase in which they utilize the reward.

For demand in our examples, we specify a simple discrete choice model where the current period utilities from purchasing and not purchasing are:

$$u_1 = A - \gamma\left(p - d\left(C\right)\right) + \varepsilon_1$$
$$u_0 = \varepsilon_0.$$

A represents the utility of purchasing, while the utility of the outside option is normalized to zero. p is the price charged by the firm in every period. $d\left(C\right)$ is the discount as a function of the number of credits the customer has in the program and ε is an i.i.d. logit error.

In period t, the discounted present value of purchasing and not purchasing are:

$$V_{1t}\left(C\right) = A - \gamma\left(p - d\left(C\right)\right) + \beta E\left[\max\left\{V_{1t+1}\left(C+1\right), V_{0t+1}\left(C+1\right)\right\}\right]$$
$$V_{0t}\left(C\right) = 0 + \beta E\left[\max\left\{V_{1t+1}\left(C\right), V_{0t+1}\left(C\right)\right\}\right]$$

The above expressions are the choice-specific value functions, net of the logit errors. In the logit model, the probability of purchase depends on the difference between these value functions.

Switching costs exist when the relative value of purchasing to not purchasing is greater because of a previous purchase. We therefore specify switching costs when a customer has C credits as:

$$\begin{aligned} SC\left(C\right) &= \tfrac{1}{\gamma}\left[\left[V_{1t}\left(C\right) - V_{0t}\left(C\right)\right] - \left[V_{1t}\left(0\right) - V_{0t}\left(0\right)\right]\right] \\ &= d\left(C\right) - d\left(0\right) + \tfrac{\beta}{\gamma}\left[\left[E\max\left(C+1\right) - E\max\left(C\right)\right] - \left[E\max\left(1\right) - E\max\left(0\right)\right]\right] \end{aligned}$$

where $E\max(C) = E\left[\max\left\{V_{1t+1}\left(C\right), V_{0t+1}\left(C\right)\right\}\right]$. Thus, the switching costs are the difference between the relative value of purchasing and not purchasing at C credits in period t and the relative value of purchasing and not purchasing if the customer had never made a purchase as of period t. The switching costs decompose into two components. First, the magnitude of the discount offered at C credits relative to the discount having never purchased (usually zero). Second, the discounted difference in the change in expected utility from purchasing at C credits from the change in expected utility from purchasing at 0 credits.

3.1 Monotonic Reward Program

We first analyze the switching costs in a monotonic reward program. In this case the first component of the switching costs, $d(C)-d(0)$, monotonically increases over time. However, the second component is negative because $d(C)$ asymptotes to a maximum discount so that the difference in the change in expected utility is always negative. Specifically, the slope of the expected utility with respect to the number of credits is decreasing over time, because additional purchases obtain smaller and smaller discounts. The switching costs therefore asymptote to:

$$\max_{C}\left\{d(C)\right\}-\left[d(0)+\beta\big/\gamma\left[E\max(1)-E\max(0)\right]\right],$$

which, when there is no initial discount, is the maximum discount less the forward-looking incentive of the program when the individual has never purchased from the firm.

To illustrate the pattern of switching costs from a monotonic program, we specified the following parameters and calculated the resulting switching costs:

$$\begin{aligned}&A=4,\ \ p=50,\\&d(C)=0.25p\left(1-.99^{C-1}\right)\big/11,\\&\gamma=0.1,\ \ T=1{,}500,\ \ \beta=0.9999.\end{aligned}$$

We simulate a finite horizon problem because an infinite horizon would lead to an unbounded state space. The finite horizon implies that the switching costs will depend on how many periods remain. For illustration, we calculate the switching costs with 1,000 periods remaining. We find that the switching costs increase at a decreasing rate, reaching over \$0.08 if 250 purchases have been made. Because the magnitude of the switching costs is dependent on the parameters, the pattern of switching costs is most important. See Figure SC1, for an illustration of this pattern.

3.2 Non-Monotonic Reward Program

For the non-monotonic reward program, it is simplest to consider an infinitely lived program. Suppose the reward is only received every $\overline{C}$ purchases. The first component of the switching costs is zero except for every $\overline{C}^{\text{th}}$ purchase when it will equal the magnitude of the reward, R. In this case, $C\in\left\{0,1,2,\ldots,\overline{C}\right\}$, such that after a reward is received, the customer returns to 0 credits. The latter component of the switching costs will therefore have the same value, every $\overline{C}$ purchases. We should therefore expect switching costs as a function of the cumulative purchases to return to zero after every $\overline{C}^{\text{ th}}$ purchase. This decline generates negative marginal switching costs.

To illustrate the pattern of switching costs in the non-monotonic reward program, we use parameter values similar to those in the monotonic example described above. However, we use an infinite horizon, and change the function for the discount to $d(C) = p \times 0.25 \times \{C = 10\}$ where $\{\bullet\}$ is an indicator function. In Figure SC2, we see that the switching costs increase through the tenth purchase, after which they drop back to zero. This pattern repeats every eleventh purchase (i.e. for the 21st, 32nd, etc. purchases).

These patterns warrant a caveat about the role of loyalty in reward programs. Specifically, a rational customer is only loyal while being compensated for being loyal. Once a reward is cashed, the customer has no greater incentive to purchase from the firm than she did if she had never purchased from the firm. In other words, at the time of using the reward there are negative marginal switching costs equal to the size of the cumulative switching costs just before the reward was redeemed. Because of this a firm has an incentive to stagger the timing of redemptions across consumers so that not all consumers experience zero switching costs at the same time. A similar point has been made about staggering the expirations of other types of long-term contracts (see, for example, Shapiro and Varian (1999), page 168). A monotonic program continually accumulates switching costs; however this also implies that the most frequent purchasers will receive the best prices from the firm. As we point out in assessing the discriminatory aspects of these programs, this may not be in the best interests of the firm if high volume customers have less elastic demand.

4 Data and Application

To empirically evaluate frequency reward programs we use data from a frequent golfing program administered by a nationwide golf course management company, the American Golf Players Association (AGPA).

4.1 The Structure of the Reward Program

The AGPA rewarded golfers by giving them a green fee certificate after purchasing ten rounds of golf at AGPA courses. The green fee certificate entitled the golfer to a discount of 25%, 50% or 100% off the price of a round of golf, depending on the course. It could not be used on Fridays, Saturdays, or Sundays.

The program required a paid membership, but immediate benefits of the membership more than offset the monetary expense of signing up for the program. The membership lasted for one year and required a renewal to continue and retain credits earned. Our current analysis focuses on the golfers' first year in the program. We plan to expand it to consider subsequent years once we have reworked our estimation code to analyze longer time periods in an efficient amount of time.

4.2 Golf Details

Currently we are analyzing a course with a green fee certificate discount of 25%. Golfers at the course have the ability to purchase one of three types of rounds. An 18-hole round is the typical full round of golf with a price of about $59 ($79 on weekends and $69 on Fridays). Late in the day, a golfer may purchase a Twilight round of golf for $39 ($49 on Fridays, Saturdays and Sundays). This round typically involves somewhere between 9 and 18 holes, depending on the start time. Golfers can also purchase 9 or fewer rounds of golf for $20 ($25 on Fridays, Saturdays and Sundays) by golfing late in the day or on the back-9 in the morning.

The data currently in the analysis includes daily purchase decisions by each of the golfers. We observe where they are at in the reward program and when they purchase their qualifying round of golf to receive the reward. We have data on the exact date the reward is issued, but this is not yet included in the analysis. Currently we assume the reward is issued immediately. Another limitation of the data is that we do not observe when the golfer uses a reward. We assume they use it on their next eligible purchase. We also currently restrict the maximum number of rewards an individual can hold at a given time to one to keep the state space small. While these assumptions clearly would not fit well in other settings, such as those with a principal-agent problem where the traveler saves rewards for personal travel rather than using them on the next available trip regardless of purpose, they are reasonable in this empirical setting. While observing the use of a reward would provide price variation and help in identification, we describe below that identification does not rely on ever observing an individual qualifying for a reward. During pre-reward periods, changes in consumer expectations of receiving a reward over time provide sufficient variation to identify the price coefficient.

4.3 Summary Statistics

The current analysis considers 327 golfers that we observe for their entire first year in the program. The time period is currently restricted to the first year in the program, providing 327 * 365 = 119,355 observations. Summary statistics for the golfers are presented in Table 1. On average, the golfers played 10.78 times during the year. They earned between 0 and 7 rewards during this time period. Renewal rates in the program were generally increasing in number of rewards earned. The majority of rounds purchased were 18-hole rounds.

Notice that 185 of the golfers did not actually earn a reward during their first year in the program. The reward program was practically costless to the firm for these customers.[6] To the extent that some of these customers believed ex-ante that they

[6] The marginal costs of the program itself are likely negligible because the system is computerized.

might qualify for a reward and increased their play as a result, the course was able to increase the revenue from these customers without incurring any expense. Thirty-one percent of these golfers actually chose to renew their membership so they either believed they would earn a reward in the future, or they valued some other non-reward based benefits of renewal that we have not yet been able to quantify. These golfers played an average of 4.18 times during the year.

The remaining golfers earned between 1 and 7 rewards throughout their first year in the program. Based on Figure 1, these golfers likely placed a high expected value on the program. However, the relevant term is the effective price decrease of the program to these customers. We quantify this in the section analyzing counterfactuals and show how it differs for low- and high-volume customers.

We are currently entering weekly price data for this and other courses in the region to supplement our identification of the price parameters in the model defined in the following section. Currently the price parameter is identified from the value a golfer places on a reward. Highly negative price coefficients are consistent with individuals that are very responsive to the incentives of the program, while individuals with less negative price coefficients are less responsive to the program.

In order to understand how identification is possible even in the absence of price variation, it is helpful to consider a model with only an intercept and a price parameter in the choice equation. Individuals with different values of the coefficients may have the same static purchase probability. In the last period of a program, they will behave identically, but as soon as we consider earlier periods in the program, they will behave differently based on their price coefficients. An individual with a zero price coefficient would behave the same in all periods, while an individual with a very negative price coefficient would have an incentive to purchase a tenth round in the next to last period to get a reward in the last period.[7]

5 Demand Model

In this section, we define a dynamic model of demand that characterizes the purchase choices of consumers when they have 365 days to participate in a program that rewards them with a discount after every ten purchases. The demand model illustrates that the utility an individual receives from purchasing is composed of the current period utility plus the expected future utility from the purchase. This latter

[7] Israel (2005) considers a similar setting in which consumers receive a discount after three years with the same auto insurance company. He notes that identifying the price sensitivity from the distance to the discount is confounded by the fact that this is negatively correlated with tenure with the firm which has potential state-dependence impact of its own. In our setting this is only true for consumers who have never received a reward. Also, in our setting and even for those who have never purchased there is not a one-to-one correspondence between tenure and number of credits remained to earn a reward because consumers do not purchase every period and because consumers purchased even before they entered the program, Furthermore, tenure with the firm extends well before our data and the origin of the reward program itself.

component takes into account the incremental step toward a future reward. We first specify the current period utility and then define the finite horizon dynamic game to derive the discounted present value of expected future utility of choices. The demand model is similar to Hartmann (2004). The key differences are the incorporation of the reward into the pricing function, the specification of the additional states involving the reward, and the 365-day finite time horizon.

5.1 The Current Period Utility

The utility individual i receives each period conditional on the state, $S = \{H, C, D, R\}$, and preferences, γ, is:

$$u(S_{it}, y_{it}, p_{it}; \gamma) = \begin{cases} v_0 + \varepsilon_{0t} & \text{if } y_{it} = 0 \\ v_{ijt} + \varepsilon_{ijt} & \text{if } y_{it} = j \quad \forall \quad j \in \{1, ..., J\} \end{cases}$$

$$v_{ijt} = \gamma_{ij0} + \gamma_{i1} p_{jt}\left(D_t, R_{it}\right) + \gamma_{i2} \log\left(H_{it}\right) + \gamma_{i3}\left(H_{it} = 60\right) + \gamma_{i4}\left(D_t > 5\right) \tag{1}$$

where v_{ijt} is the utility associated with choice $j \in \{1, 2, 3\}$ where 1 indexes an 18-hole round, 2 a 9- to 18-hole round, and 3 a 9-hole or less round at time t net of an individual and time specific shock to preferences, ε_{ijt}. H_{it} is the number of days since the last purchase and enters the choice utility in the same manner as Hartmann (2004). We specify the log of H_{it} to additively enter the current period utility. There is also an indicator for $H = 60$. γ_{i4} is a taste for weekend golf relative to weekday golf. The price of product j depends on D_t, the day of the week, and R_{it}, an indicator for whether or not the individual has a reward. C is the number of credits an individual has toward a reward. v_0 represents the utility of the outside good and is normalized to 0.

The laws of motion of the state variables create the dynamics in the decision process. The length of time since the last round, H_{it}, increases by one whenever the outside good is chosen, and resets to 1 whenever the individual chooses one of the three types of rounds:

$$H_{it+1} = \left\{ \begin{array}{lll} H_{it} + 1 & \text{if} \quad y_{it} = 0 \quad \text{and} \quad H_{it} < 60 \\ H_{it} & \text{if} \quad y_{it} = 0 \quad \text{and} \quad H_{it} = 60 \\ 1 & \text{if} \quad y_{it} \in \{1, 2, 3\} \end{array} \right\}$$

H_{it} is bounded above at 60. A reward is earned once a tenth credit is earned:

$$R_{it} = \begin{cases} 1 & \text{if } C_{it-1} = 9 \text{ , } y_{it-1} > 0 \\ 0 & \text{if } y_{it-1} = 1, \mathrm{D}_{it-1} < 6, C_{it-1} < 9 \\ R_{it-1} & \text{otherwise} \end{cases}$$

To keep the state space small, we currently restrict the model such that individuals cannot hold more than one reward.

The number of credits evolves as follows:

$$C_{it} = \begin{cases} C_{it-1} + 1 & \text{if } y_{it-1} > 0 \text{ and } C_{it-1} < 9 \\ C_{it-1} & \text{if } y_{it-1} = 0 \\ 0 & \text{if } y_{it-1} > 0 \text{ and } C_{it-1} = 9 \end{cases}$$

The days of the week are numbered 1 to 7, beginning with Monday.

$$D_{t+1} = \begin{Bmatrix} D_t + 1 & \text{if} \quad D_t < 7 \\ 1 & \text{if} \quad D_t = 7 \end{Bmatrix}$$

5.2 Dynamic Optimization Problem

Dynamics arise in this model through the state variables H_{it} and C_{it}. Specifically, the consumer recognizes the implications the choice at time t has on H_{it+1}, which directly affects utility, and C_{it+1}, which affects future prices through R_{it}. We solve the model backwardly, starting at the terminal period (i.e. 365). We plan to extend estimation to incorporate the consumer's end-of-year renewal decision to more accurately represent the program. However, our current code needs to be modified to speed up the program to accommodate this improvement.

In the terminal period, individuals only receive their current period utility from golfing. The maximization problems solved is therefore:

$$\max_{y_{iT}} \left\{ v_{i0T}\left(S_{iT};\gamma_i\right) + \varepsilon_{i0T}, \ldots, v_{iJT}\left(S_{iT};\gamma_i\right) + \varepsilon_{iJT} \right\}$$

The model's dynamics are apparent in analyzing the choices in period $T-1$. Each choice will have a non-stochastic component of utility that is the sum of the non-stochastic portion of the current period utility specified above, v_{ijt}, plus the discounted present value of future utility from choosing choice j in state S_{iT-1}:

$$V_{ijT-1}\left(S_{iT-1};\gamma_i\right)=v_{ijT}+\beta E\left[\max_{y_{iT}}\left\{\begin{array}{l}v_{i0T}\left(S_{iT};\gamma_i\right)+\varepsilon_{i0T},\\ ...,\\ v_{iJT}\left(S_{iT};\gamma_i\right)+\varepsilon_{iJT}\end{array}\right\}\mid S_{iT-1},y_{iT-1}=j\right] \tag{2}$$

where the expectation above is taken only over the ε's because the state space evolves deterministically. $y_{it}\in\{0,1,...,J\}$ is the choice of individual i in period t. The choice-specific utilities for preceding periods can be specified as follows:

$$V_{ijt}\left(S_{it};\gamma_i\right)=v_{ijt}+\beta E\left[\max_{y_{iT}}\left\{\begin{array}{l}V_{i0t+1}\left(S_{it+1};\gamma_i\right)+\varepsilon_{i0t+1},\\ ...,\\ V_{iJt+1}\left(S_{it+1};\gamma_i\right)+\varepsilon_{iJt+1}\end{array}\right\}\mid S_{it},y_{it}=j\right] \tag{3}$$

We iteratively solve the above equation back to the first period, $t=1$. The maximization problem in all periods before T is therefore:

$$\max_{y_{it}}\left\{V_{i0t}\left(S_{it};\gamma_i\right)+\varepsilon_{i0t},...,V_{iJt}\left(S_{it};\gamma_i\right)+\varepsilon_{iJt}\right\} \tag{4}$$

In cases where j involves a purchase, the state changes such that the individual is one step closer to a reward (or receives a reward) and has one less day to finish the program. When an individual chooses not to purchase, she will be just as many credits away from a reward on the next day, but have one less day to actually earn the reward. In these cases the expected future value of the program should be decreasing.

6 Estimation

We specify the individual-specific parameters, γ_i s, as normally distributed random coefficients. We then estimate the model using Ackerberg's (2001) change of variables and importance sampling technique. The technique involves calculating the value of the likelihood at a wide range of candidate parameter values, then searching for the parameter vector that weights the likelihoods of the candidate parameters such that the value of the simulated likelihood function is maximized.

The choice probabilities are those of the typical discrete choice model, except that the choice specific value functions are inserted into the choice probability, rather than the current period utilities:

$$\Pr\left(y_{it} \mid S_{it};\gamma_i\right)=\sum_{j=0}^{J}\frac{\left\{y_{it}=j\right\}\exp\left(V_{ijt}\right)}{\sum_{k=0}^{J}\exp\left(V_{ikt}\right)} \tag{5}$$

Since we observe each of the individuals for 365 days, the likelihood function for an individual is:

$$L_i\left(S_{i1},\ldots S_{iT},y_{i1},\ldots,y_{iT};\gamma,\Sigma\right)=\int\prod_{t=1}^{T}\Pr\left(y_{it} \mid S_{it};\gamma_i\right)f\left(\eta_i\right)d\eta\,. \tag{6}$$

The random coefficients are specified as:

$$\gamma_i=\gamma+\Gamma\eta_i \tag{7}$$

where γ is the parameter vector representing the means of the random coefficients and Γ is the Cholesky decomposition of Σ, the variance-covariance matrix. η_i is a standard normally distributed vector that must be integrated over to evaluate the individual likelihood function. To limit the number of parameters we estimate, we restrict Σ such that the only correlations between random coefficients are between the intercepts and the price parameter[8].

The first step in the Ackerberg technique defines a change of variables. We use the random coefficient definition above to define the variable $u_i=\gamma+\Gamma\eta_i$. We therefore have the following likelihood function:

$$L_i\left(S_i,y_i;\gamma,\Sigma\right)=\int\prod_{t=1}^{T}\Pr\left(y_{it} \mid S_{it};u_i\right)f_U\left(u \mid \gamma,\Sigma\right)du \tag{8}$$

Next, we define an importance sampling distribution $g\left(u\right)$. This distribution must have a non-zero density over the support of u. The distribution is:

$$g\left(u\right)=f_U\left(u \mid \gamma_1,\Sigma_1\right) \tag{9}$$

Notice that the distribution is identical to the population distribution of the random coefficients; however it has a different mean and variance. In estimation, we hold the mean and variance of this importance sampling distribution fixed and solve for

[8] Current estimates of the model restrict all covariances to be zero; future estimates will allow for this correlation between intercepts and prices.

the mean and variance of the population distribution. Dividing and multiplying Equation (8) by $g(u)$, we get:

$$L_i(S_i, y_i; \gamma, \Sigma) = \int \prod_{t=1}^{T} \Pr(y_{it} \mid S_{it}; u_i) \frac{f_U(u \mid \gamma, \Sigma)}{g(u)} g(u)\, du \tag{10}$$

We take draws from $g(u)$ to simulate the above likelihood and then search for the set of parameters that maximize:

$$L = \prod_{i=1}^{N} L_i(S_i, y_i; \gamma, \Sigma). \tag{11}$$

The model estimates are reported in Table 2. The estimates themselves do not reveal much about the reward program because they describe the current period utility only. The price coefficient and state-dependence are both negative as expected. Golfers prefer 18-hole rounds over twilight rounds and the latter over 9-hole or less rounds. Golfers prefer to play on weekends relative to weekdays. There is significant heterogeneity in all of the random coefficients. This heterogeneity will be instrumental in evaluating how the firm's reward program discriminates between individuals. We evaluate the purchase probability distributions dictated by these parameters, the expected value of the reward program, and the program's effective pricing in counterfactual analyses in the following section.

7 Counterfactuals

In all of these counterfactuals we must assume that the other firms in the market maintain their same pricing strategies since we do not model the supply-side. We are unable to evaluate the profitability and solve for the optimal program structure for two reasons. First, we do not have the demand from all types of consumers, only those who enrolled in the reward program. Second, we do not know the firm's cost structure. It is clearly a capacity constrained firm. Marginal costs are likely negligible when the constraint is not binding. However, because we do not know the demand from the other types of consumers, we cannot calculate the appropriate opportunity costs, and hence marginal costs, of a round of golf when the constraint is binding. We therefore focus our analysis on demonstrating the price discrimination created by the program and evaluate the demand elasticities to determine whether this firm has an incentive to effectively lower its price to higher volume customers.

7.1 Demonstrating the Discount to High Types

To demonstrate the role the reward program plays in lowering prices to and raising demand from high type consumers, we compare the purchase probabilities of six types of golfers under the reward program and a counterfactual in which the reward program is removed. Under both scenarios, we assume that the golfer has played the day before, and under the reward program, we assume that the golfer has no credits, no rewards, and 365 days left before the end of the program, and it is a Monday. The top panel of Table 3 displays the purchase probabilities for six golfers that vary in their utilities for 18-hole rounds of golf.[9] The first golfer has a utility for an 18-hole round at the fifth percentile of the distribution of the random coefficient. Purchase probabilities for this golfer under both the reward program and uniform price appear identical. For the golfer with the mean value for an 18-hole round, the purchase probability is 1.27% under the reward program, and only 0.06 points lower under the uniform price. A golfer with a utility for an 18-hole round at the seventy-fifth percentile has a 2.16% purchase probability under the reward program and a 1.95% purchase probability under the uniform price. This 0.21 point difference in the purchase probabilities represents a 9.72% decline in the purchase probability of this golfer.

The last column of Table 3 computes the effective price decrease in a uniform price scenario that would mitigate these differences in purchase probabilities. The effective price decrease is zero for the fifth percentile consumer and increasing to $1.74 for the seventy-fifth percentile consumer. As shown in the bottom panel of Table 3, reducing the length of the reward program by half results in more distinct segmentation of the consumer distribution. The shorter program would affect only the top quartile of the distribution, with the largest effective price decrease to the individuals in the ninety-fifth percentile of the distribution.

7.2 Segmentation in Expected Value of Reward Program

Reward programs used for quantity-based price discrimination must offer a greater value to individuals likely to purchase more rounds. As we describe in regard to Figure 1 in Section 3, a firm can manipulate the length of the program or the number of credits to determine how individuals with different purchase probabilities will value the program. To assess the degree of segmentation created by the golf course's reward program, we compute the expected value of a single reward earning opportunity as a function of purchase probability and compare it with the underlying distribution of customers. The line with the dots (corresponding to the 5th, 25th, median, 75th, 95th and 99th percentiles of the distributions of purchase probabilities) in Figure 2 demonstrates that the expected value of the reward program cuts upward through the right-hand side of the distribution of customers. Individuals to the left of the mode of the distribution expect practically zero value from the program, while the 75th percentile customer expects about $5 from the program and the upper 5th percentile of customers expect more than $14 from the program. The next line to the right considers the expected values of the program if the golf course were to reduce

[9] The six consumer types are determined by adjusting their intercept utility up or down appropriately. The values of all other parameters are identical and equal to the mean value of the random coefficients depicted in Table 2.

the length of the program from one year to one half year. In this case, only the top 25 percent of customers receive any value. For these customers, the value increases from about \$0 to \$10.11 for the 99th percentile of customers. A further reduction in the program length to 90 days only yields a \$0.52 expected value to the 99th percentile of customers.

If a firm increases the scope of eligible purchases in its program, it generally shifts the distribution of purchase probabilities to the right. It may also raise the expected value of the reward if there is a greater probability a customer will use it. Just as many airlines have eliminated expirations on their rewards, suggesting they are not concerned with quantity-based price discrimination, recent agreements consolidating airline reward programs further decrease the segmentation in the value of the programs.

7.3 Effective Price Discrimination for Pre-Reward Purchases

While the preceding analysis demonstrated that the expected value of a reward program increases monotonically with the individual's probability of purchase, this does not imply that the effective price discount on pre-reward purchases is increasing monotonically with the probability of purchase. In fact, it actually begins to decrease for high-purchase probability individuals. The intuition is based on the opportunity cost of not purchasing from the firm and can be seen in a simple example. Consider a customer who normally purchases ten sandwiches per month from a sandwich shop that introduces a reward program offering one free sandwich per month after five purchases. This consumer has little opportunity cost of not purchasing one of her ten sandwiches from the firm with the reward program. On the other hand, an individual purchasing only five per month will lose the reward if it makes even one of its five purchases from another firm.

We consider an extreme case of only a single reward-earning opportunity in a year to demonstrate this phenomenon. Consider Figure 3b. We evaluate the effective price decrease of the reward program for the first purchase (i.e. the price decrease necessary to achieve the same purchase probability as the reward program). The figure illustrates that for a 365-day reward program, the effective price decrease is increasing up to the individual at the 75th percentile of the distribution, but then begins to fall such that the individual at the 99th percentile of the distribution only receives an effective price decrease of \$0.03. Note however that the cost to the firm (i.e. the expected value of the reward to this individual) is actually greatest for this 99th percentile customer, as can be seen in Figure 2. The upper line in Figure 3b illustrates what happens when you can earn an unlimited number of rewards during the year. A dip in the purchase probabilities still exists for the high types, but is not as exaggerated as in the single-reward opportunity. Figure 3a shows this curve and the comparable curve for a 182-day program. This illustrates that the firm should optimally choose the reward program duration and number of reward-earning opportunities based on the underlying distribution of demand elasticities of its customers.

7.4 Evaluating Elasticities

Currently the estimated elasticities from our model are decreasing in consumers' purchase probability. In this case, the firm does not have an incentive to use this reward program to price discriminate because it is lowering the price to those golfers for whom it should be raising the price. However, these elasticities are the result of a restriction in our estimation that we plan to relax. Currently the covariance between the intercepts for purchasing and the price coefficient is assumed to be zero. Therefore we have assumed that the distributions of price coefficients for golfers with different frequencies of play are all identical. The static logit elasticity equation implies that for a given price coefficient as shares increase elasticities must decrease. Allowing for correlation between the price coefficient and the intercepts will relax this restriction and allow us to determine the relationship between demand elasticities and player frequency.

7.5 Estimating Switching Costs

Using the results of our estimation, we can calculate the cumulative and marginal switching costs induced by the golf firm's reward program for a particular consumer at a given point in the program and with a given number of credits accumulated. We define the switching costs using the definition from Section 3 above. In the context of the model we defined, the switching costs for an 18-hole round of golf are:

$$SC(C_{it}) = 1/\gamma_{i1}\begin{bmatrix}\left[V_{i1t}(H_{it},C_{it},D_t,R_{it}) - V_{i0t}(H_{it},C_{it},D_t,R_{it})\right] - \\ \left[V_{i1t}(H_{it},0,D_t,0) - V_{i0t}(H_{it},0,D_t,0)\right]\end{bmatrix}$$

$$SC'(C_{it}) = 1/\gamma_{i1}\begin{bmatrix}\left[V_{i1t}(H_{it},C_{it},D_t,R_{it}) - V_{i0t}(H_{it},C_{it},D_t,R_{it})\right] - \\ \left[V_{i1t}(H_{it},C_{it}-1,D_t,R_{it}) - V_{i0t}(H_{it},C_{it}-1,D_t,R_{it})\right]\end{bmatrix}$$

where SC is the cumulative switching cost and SC' is the marginal switching cost.[10]

We plan to estimate the marginal and cumulative switching costs implied by our empirical data as future work. We will focus, in particular, on the effect of the non-monotonicities inherent in the program and compare the switching costs of low-versus high-frequency players

[10] Obviously, when we consider the switching costs at zero credits, the formula for the marginal switching costs will require a substitution of 9 for $C_{it}-1$, and perhaps a change in the number of rewards possessed.

8 Conclusions

Our empirical model considers an environment in which consumers are faced with uncertainty over their future choices. The estimates of this model allow us to perform counterfactual analyses that demonstrate how a reward program price discriminates based on the quantities individuals purchase. There are a variety of mechanisms in these programs that can affect the impact of the program across different consumer segments. Increasing (decreasing) the life of credits or the scope of the program, or decreasing (increasing) the number of credits necessary for a reward or the value of the reward itself, serves to reduce (augment) the ability of a reward program to segment customers based on quantity by making the program's effects more homogenous (heterogeneous) to infrequent and frequent consumers.

While many reward programs may have been introduced for other reasons (e.g. kickback schemes), the discriminatory pricing effects of the programs must be recognized. In fact, firms employing reward programs for other means may actually use the various mechanisms described in this paper to eliminate or reduce unintended quantity discounts. The elimination of expirations on miles by many airlines and agreements to consolidate programs reduce the quantity discount element of these programs.

We also outline a method for quantifying the switching costs created by a reward program and demonstrate in stylized examples the effects that non-monotonicities in the relationship between price and number of cumulative purchases has on these costs.

Bibliography

Ackerberg, D. (2001). "A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation," working paper.

Banerjee, A. and L. Summers (1987). "On Frequent Flyer Programs and Other Loyalty-Inducing Arrangements," working paper, Harvard University.

Borenstein, S. (1996). "Repeat-Buyer Programs in Network Industries," in *Networks, Infrastructure, and the New Task for Regulation*, The University of Michigan Press, Ann Arbor, Michigan, 1996.

Bulkley, G. (1992). "The Role of Loyalty Discounts When Consumers are Uncertain of the Value of Repeat Purchases," *International Journal of Industrial Organization*, 10, 91 – 101.

Cairns, R. D. and J. W. Galbraith (1990). "Artificial Compatibility, Barriers to Entry, and Frequent-Flyer Programs," *Canadian Journal of Economics*, 4, 807 – 816.

Caminal, R. and C. Matutes (1990). "Endogenous Switching Costs in a Duopoly Model," *International Journal of Industrial Organization*, 8, 353 – 373.

Fudenberg, D. and J. Tirole (2000). "Customer Poaching and Brand Switching," *RAND Journal of Economics*, 31, 634 – 657.

Hartmann, W. (2004). "Intertemporal Effects of Consumption and Their Implications for Demand Elasticity Estimates," working paper.

Israel, M. (2005). "Who Can See the Future? Information and Consumer Reactions to Future Price Discounts," working paper.

Katz, M. (1984). "Price Discrimination and Monopolistic Competition," *Econometrica*, 52, 1453 – 1471.

Kim, B., Shi, M. and K. Srinivasan (2002). "Reward Programs and Tacit Collusion," *Marketing Science*, 20, 99 – 120.

Kim, B., Shi, M. and K. Srinivasan (2004). "Managing Capacity Through Reward Programs," *Management Science*, 50, 503 – 520.

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

Kivetz, R., O. Urminsky and Y. Zheng. (2005). "The Goal-Gradient Hypothesis Resurrected: Purchase Acceleration, Illusionary Goal Progress, and Customer Retention," working paper.

Klemperer, P. (1987a). "Markets with Consumer Switching Costs," *The Quarterly Journal of Economics*, 102, 375 – 394.

Klemperer, P. (1987b). "The Competitiveness of Markets with Switching Costs," *RAND Journal of Economics*, 18, 138 – 150.

Klemperer, P. (1995). "Competition when Consumers Have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade," *Review of Economics Studies*, 62, 515 – 539.

Kopalle, P. K. and S. A. Neslin (2003). "The Economic Viability of Frequency Reward Programs in a Strategic Competitive Environment," *Review of Marketing Science*, 1, 1 – 39.

Kumar, N. and R. Rao (2003). "Effect of Loyalty Card Programs on Supermarket Prices," working paper.

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

Lederman, M. (2004). "Do Enhancements to Loyalty Programs Affect Demand? The Impact of International Frequent Flyer Partnerships on Domestic Airline Demand," working paper.

Shapiro, C. and H. Varian (1999). *Information Rules*, Harvard Business School Press, Boston, Massachusetts.

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

Table 1

Purchases and Rewards During First Year in Program

Total Rewards Earned	N	Average Number of Purchases: Total	18 Holes	9-18 Holes	9 Holes or less	Percent Renewed
0	185	4.18	3.18	0.89	0.11	31%
1	95	13.62	10.54	2.79	0.29	54%
2	27	23.26	15.19	6.74	1.33	63%
3	11	33.64	22.55	8.36	2.73	82%
4	5	45.20	23.20	21.40	0.60	100%
5	3	53.67	36.67	17.00	0.00	67%
6	0					
7	1	72.00	61.00	10.00	1.00	100%
Total	327	10.78	7.75	2.66	0.36	

Table 2

Model Estimates

	Random Coefficients: Mean	Standard Deviation
18-Hole Intercept	-0.833	0.353
	(0.035)	(0.026)
9-18 Holes Intercept	-3.553	0.910
	(0.074)	(0.058)
9-Hole or Less Intercept	-7.790	1.095
	(0.111)	(0.071)
Price Coefficient	-0.060	0.011
	(0.001)	(0.001)
Days Since Last Purchase	0.072	0.089
	(0.009)	(0.007)
60 Days Since Last	-0.546	0.646
	(0.063)	(0.048)
Weekend Indicator	1.167	0.580
	(0.050)	(0.038)

Table 3

Reward Program vs. Uniform Price
First Purchase Comparisons

Type of Customer*	18-Hole Round Purchase Probability: Reward Program	Uniform Price	Percent Difference	Effective Price Decrease of Reward Program
365-Day Program				
5th Percentile	0.37%	0.37%	0.00%	$0.00
10th Percentile	0.75%	0.74%	-1.33%	$0.11
Mean Consumer	1.27%	1.21%	-4.72%	$0.77
75th Percentile	2.16%	1.95%	-9.72%	$1.74
95th Percentile	4.22%	3.89%	-7.82%	$1.40
99th Percentile	6.61%	6.11%	-7.56%	$1.41
182-Day Program				
5th Percentile	0.37%	0.37%	0.00%	$0.00
10th Percentile	0.74%	0.74%	0.00%	$0.00
Mean Consumer	1.21%	1.21%	0.00%	$0.01
75th Percentile	1.95%	1.97%	1.03%	$0.15
95th Percentile	3.89%	4.28%	10.03%	$1.65
99th Percentile	6.11%	6.66%	9.00%	$1.53

* Customers analyzed have the mean value for all random coefficients except the intercept for 18-Hole rounds of golf. This coefficient is a value from a percentile of the distribution. All types of customers are assumed to have played the day before, a Monday. In the reward program scenario, purchase probabilities are for the first day in the program, with zero credits.

Figure 1
Effect of Reward Program Term on Per Credit Value to Consumer

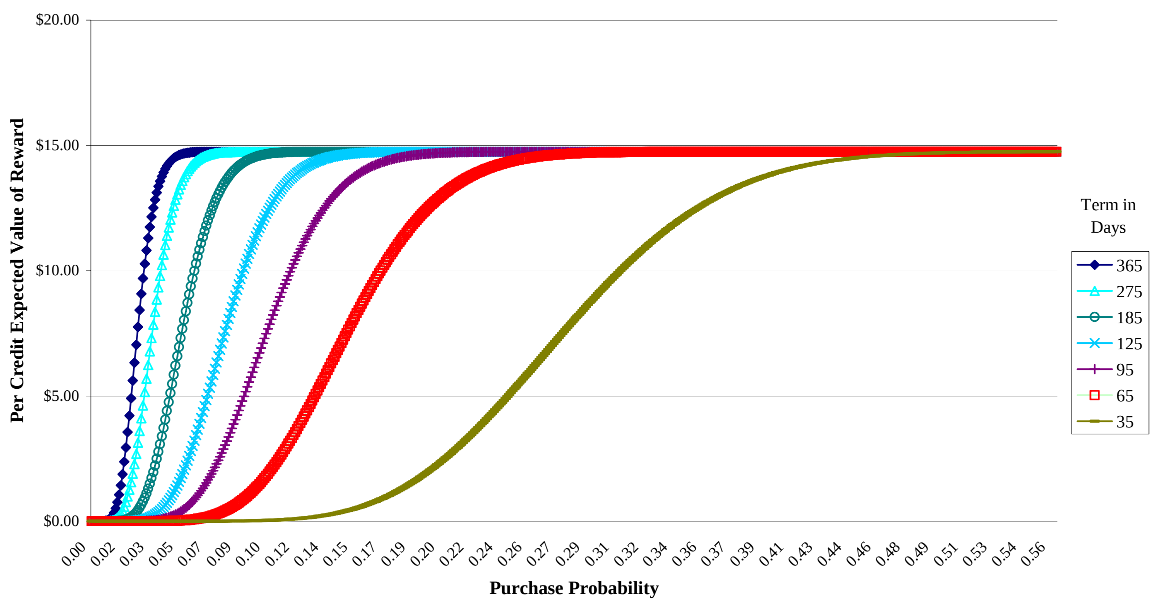

Expected values are calculated for a $14.75 reward (equivalent to that in the reward program in this paper). The probability of receiving the reward is 1 minus the cdf of a binomial distribution evaluated at 10 successes in the Term Length's number Bernouilli trials, where the probability of success is the purchase probability reported on the horizontal axis.

Figure 2
Expected Value of a Reward
by Program Length and Purchase Probability

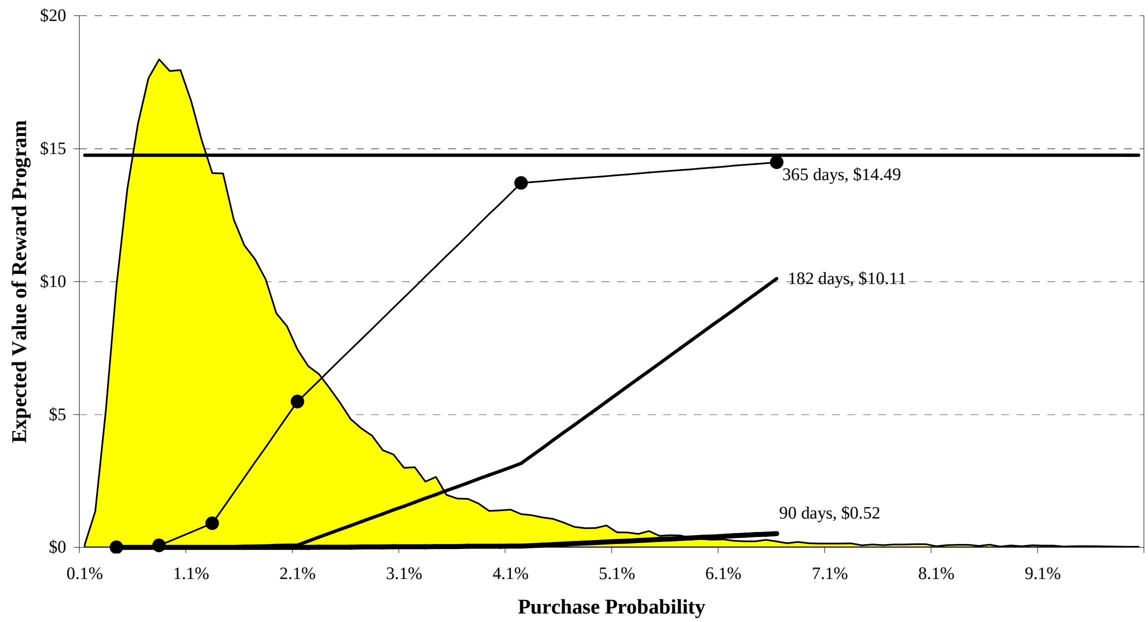

Figure 3a
Effective Price Decrease for First Purchase in Reward Program
365-Day Program vs. 182-Day Program

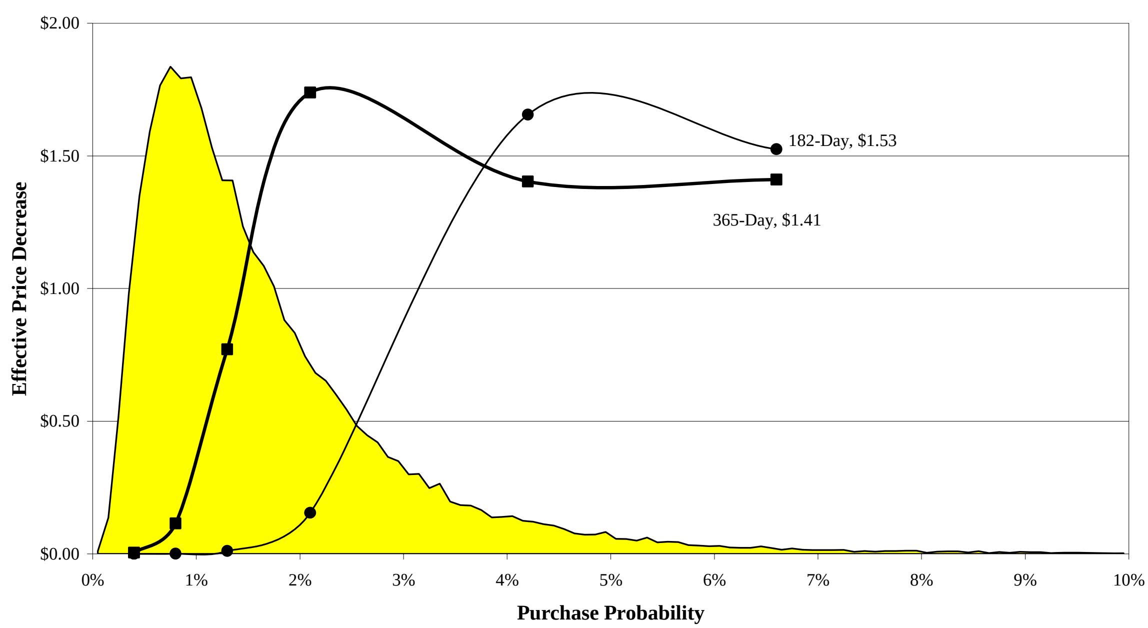

Figure 3b
Effective Price Decrease for First Purchase in 365-Day Reward Program
Unlimited Rewards vs. Single Reward Opportunity

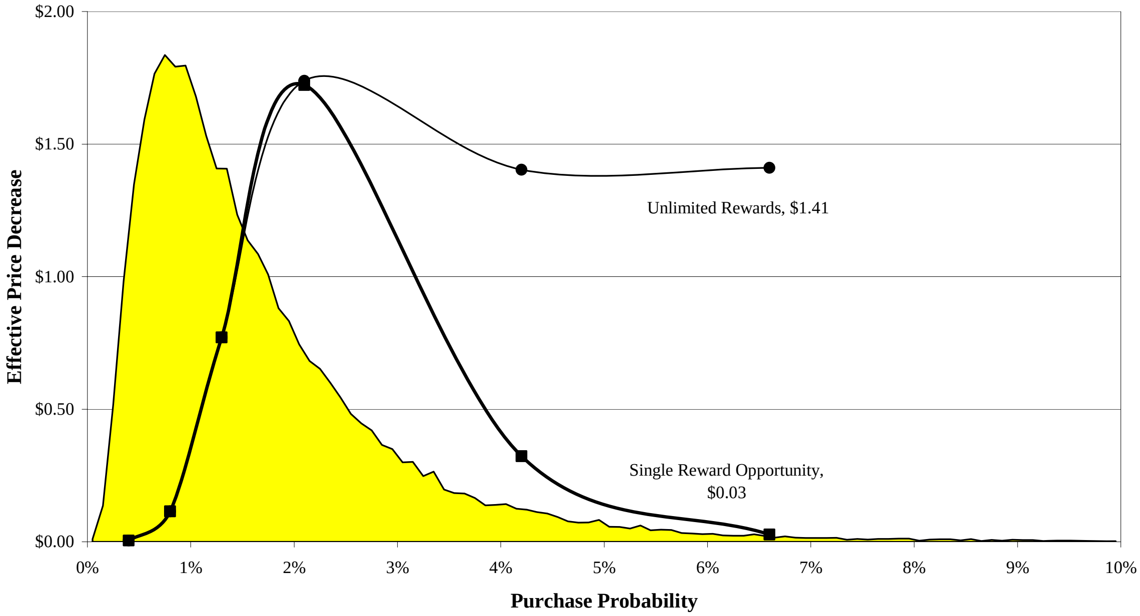

Figure SC1
Switching Costs in Monotonic Reward Program
500 periods in, 1000 remaining

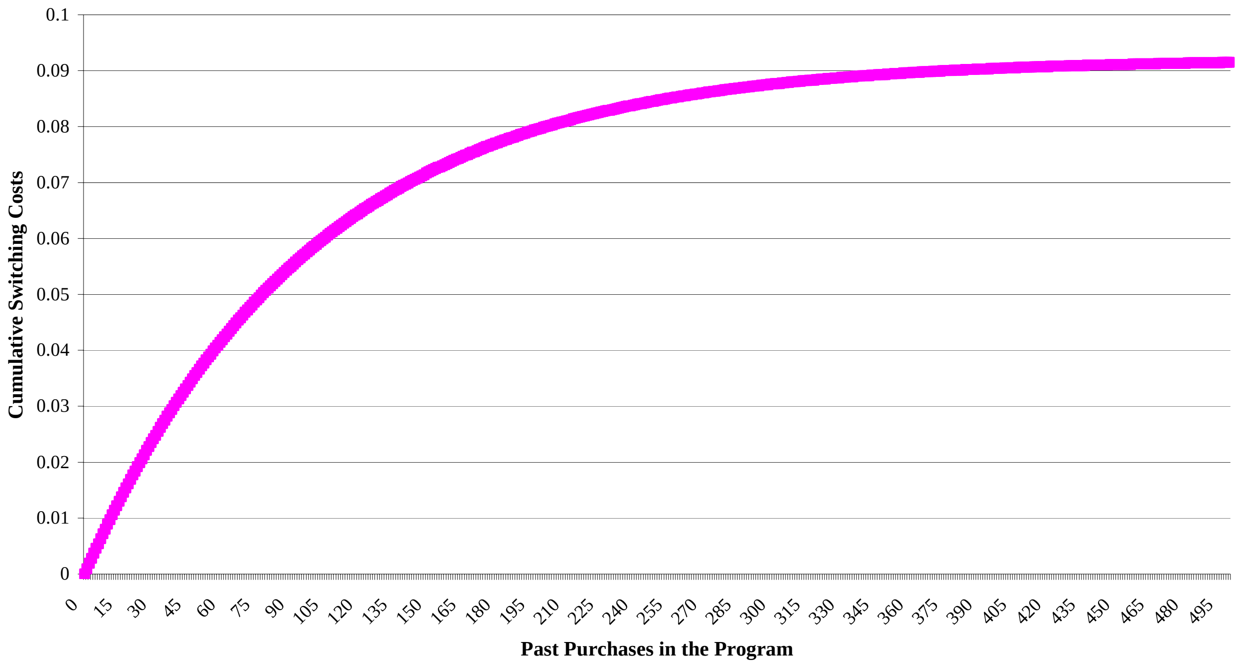

Figure SC2
Switching Costs in Non-Monotonic Reward Program