IO
7/20/06
3:45 PM

Do Frequency Reward Programs Create Switching Costs?[1]

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

July 10, 2006

Abstract

This paper examines a common assertion that customers in reward programs become "locked in" as they accumulate credits toward earning a reward. We define a measure of switching costs and use a dynamic structural model of demand in a reward program to illustrate that frequent customers' incentives to purchase are practically invariant to the number of credits. In our empirical example, these customers comprise over eighty percent of all rewards and over two-thirds of all purchases. Less frequent customers may face substantial switching costs when close to a reward, but rarely reach this state.

Keywords: switching costs, reward programs, dynamic programming, discrete-choice.

[1] The authors would like to thank Dan Ackerberg, Eyal Biyalogorsky, Mara Lederman, Phillip Leslie, Marc Rysman, Andy Skrzypacz, Raphael Thomadsen, conference participants at SICS, IIOC, and the Berkeley/Santa Clara/Stanford Joint Marketing Seminar, and seminar participants at Yale University, UCLA, University of Toronto, University of Colorado at Boulder, Washington University, St. Louis, University of Chicago GSB, Harvard, MIT and Stanford GSB for helpful comments. We would also like to thank Andrew Nigrinis and Liang Qiao for valuable research assistance. Any errors are our own.

1 Introduction

Switching costs are one of the most commonly cited effects of reward programs. If firms are able to "lock-in" customers as they progress through these programs, the switching costs may reduce welfare by leading to inefficient switching or reduced competition. Despite these concerns, there has been no work defining and estimating a measure of switching costs in reward programs.

We define switching costs in a reward program as a function of the opportunity cost of not purchasing, and hence not progressing further in the program. The opportunity cost can involve the lost opportunity to use a reward, a decrease in the expected number of rewards that will be earned, or a delay in the earning and redemption of future rewards. Switching costs depend on how much lower the opportunity cost is when a customer has zero credits than when the customer has earned credits toward a reward.

Frequent customers face nearly identical opportunity costs at all levels of credits, indicating negligible switching costs. As an example, consider a "buy ten get one free" program. Because the next reward is never far away, purchasing at zero credits has a similar effect on the timing of the first reward as purchasing at nine credits has on the timing of the second reward.

Less frequent customers realize smaller opportunity costs of not purchasing at zero credits and may therefore incur economically significant switching costs when close to earning a reward. At zero credits, these customers either have a small probability of earning a reward or will earn it so far in the future that its value will be heavily discounted. At higher credit levels, rewards are more imminent or likely so opportunity costs are greater.

Illustrating this effect of demand intensity on switching costs requires a demand model with both dynamics and heterogeneity. We develop a structural model with these features and estimate it using data from a reward program offered by a golf course. Our estimates confirm that switching costs are negligible for frequent golfers and economically significant for less frequent golfers when they are close to earning a reward. The importance of switching costs therefore depends on the demand comprised by each of these types and how often infrequent customers are close to earning a reward.

Frequent customers purchase more and represent a greater fraction of demand. In our case they comprise more than two-thirds of all purchases and more than eighty percent of all rewards earned. Less frequent customers rarely progress to credits with higher switching costs, but if they do, they progress out of these states faster than those with little or no switching costs. In our case, most infrequent golfers exited the program before ever reaching half the necessary credits for a reward.

A way of assessing the implications of any existing switching costs is to evaluate the program's effect on demand elasticities. We find that the elasticity under the reward program is generally no lower than it would be if the firm had not offered the program and instead lowered the uniform price by an amount equal to the per-purchase value of the reward.

The following section describes how switching costs can arise in reward programs and defines an equation to measure them. Section 3 defines a random utility model of demand in a reward program. Section 4 describes our specific application and data. Section 5 describes the empirical implementation of the demand model in the context of our data. Section 6 analyzes our estimated model to evaluate the switching cost effects of the reward program. Section 7 concludes.

2 Switching Costs and Reward Programs

Previous academic work has focused on the ability of reward programs to endogenously create switching costs or lock-in.[2] The notion is that as customers accumulate credits toward a reward, they will be less likely to choose other alternatives.

2.1 A Measure of Switching Costs

To define switching costs in a reward program it is useful to start with exogenous switching costs as commonly modeled in a two-period setup. Klemperer (1987a) treats switching costs as a "start-up" cost. In the first period, a customer incurs this cost regardless of the firm chosen such that it does not affect the relative valuation of the firms. In the second period, if the customer chooses the same firm, there is no start-up cost. If it chooses a different firm, it must pay the start-up cost again, thereby incurring a switching cost.

While Klemperer (1987a) models switching costs as an additional cost of choosing a different firm (e.g., the sunk cost of opening a new account or negotiating a contract with a supplier), the two-period model is analytically identical if there is a benefit to choosing the same firm a second time (e.g., non-portability of telephone numbers increases the value of purchasing from the same telephone company again as in Viard (2005)). In summary, the switching costs are a value, s, either added to the costs of the new firm or to the benefits of the previously chosen firm.

Switching costs are defined by the opportunity cost of not repeating a choice made earlier. To illustrate this, we define switching costs incurred by a customer in period 2, given that the customer chose option 1 in period 1, as:

[2] Klemperer (1987b and 1995) cite frequent flier programs as an example of switching costs.

$$SC_{2|y_1=1} = \{V_{12}(S=1) - V_{02}(S=1)\} - \{V_{12}(S=\bullet) - V_{02}(S=\bullet)\}, \ (1)$$

where V_{j2} is the value of option j in period 2, y_1 is the choice in the first period and $S \in \{\bullet, 0, 1\}$ denotes whether the customer is not committed, committed to option 0, or committed to option 1. The first bracketed expression is the forgone value from not choosing option 1 again, given that the customer is committed to option 1. The second bracketed expression is the forgone value from not choosing option 1, given that the customer is uncommitted. The switching cost is the difference between these two opportunity costs.

Applying this formula to the Klemperer (1987a) model of switching costs where there is a setup cost to choosing a firm for the first time yields the switching costs, s:

$$SC = \begin{bmatrix} \{(r - tx - p_{12}) - (r - t(1-x) - p_{02} - s)\} \\ -\{(r - tx - p_{12} - s) - (r - t(1-x) - p_{02} - s)\} \end{bmatrix} = s, \ (2)$$

where r represents customers' reservation value, t represents transport costs, x represents distance from the left-hand side firm in a Hotelling model, and p_{j2} is the price of firm j in period 2. In Equation (2), not purchasing from firm 1 in the terminal period involves forgoing the savings of setup costs s.

In the case of a benefit to choosing the same firm a second time, the formula also yields switching costs of s:

$$SC = \begin{bmatrix} \{(r - tx - p_{12} + s) - (r - t(1-x) - p_{02})\} \\ -\{(r - tx - p_{12}) - (r - t(1-x) - p_{02})\} \end{bmatrix} = s. \ (3)$$

In Equation (3), not purchasing from firm 1 in the terminal period involves forgoing the benefit s of purchasing from firm 1 again.

2.2 Switching Costs in Reward Programs

In the theoretical literature on reward programs, the switching costs, s, is a price discount for a returning customer. That is, it is an added benefit of choosing the same firm as in the first period. In Caminal and Matutes (1990), new customers of firm 1 in the second period pay p_{12}, while returning customers to firm 1 in the second period pay $p_{12} - s$.[3] This corresponds to the application of our formula in Equation (3) above, since $\left(r - tx - p_{12} + s\right) = \left(r - tx - \left(p_{12} - s\right)\right)$.

Most actual reward programs do not conform to a two-period model. The additional time periods have two important implications for program design and the switching costs created. First, most rewards, s, in the model, can be redeemed in any one of many periods after they are earned. In this case, V_{02} in the first bracketed expression in Equation (1) (i.e., the value of choosing a different alternative in period 2) includes the discounted present value of using the reward in a later period. This is an option value $Es' < s$ such that:

$$SC = \begin{bmatrix} \left\{\left(r - tx - \left(p_{12} - s\right)\right) - \left(r + Es' - t\left(1 - x\right) - p_{02}\right)\right\} \\ -\left\{\left(r - tx - p_{12}\right) - \left(r - t\left(1 - x\right) - p_{02}\right)\right\} \end{bmatrix} = s - Es'. \quad (4)$$

Although this is a slight abuse of notation because we remain in a two-period setting and allow an option value beyond that, it makes clear that the option value of a reward decreases the switching costs. Two-period models force the option value to zero because a customer must use a reward in the second period or lose it, even if the other choice is somewhat preferred in the second period.[4] When a customer has a choice of when to redeem a reward, he may wait to redeem it until a period in which he prefers the firm for which he has a reward. The option to wait reduces the switching costs in the immediate period.

The second implication of a longer time horizon is that multiple purchases may be required to obtain the reward, s. This introduces intermediate time periods in which a customer has some credits toward a reward, but has not yet earned it. This can potentially lead to switching costs even when a customer does not have a reward. When a customer in period t has C_t credits toward a reward, but has not yet earned the reward, the switching costs, following Equation (1) above, are:

$$SC_t\left(C_t, R_t = 0\right) = \begin{pmatrix} \left\{V_{1t}\left(C_t, R_t = 0\right) - V_{0t}\left(C_t, R_t = 0\right)\right\} - \\ \left\{V_{1t}\left(C_t = 0, R_t = 0\right) - V_{0t}\left(C_t = 0, R_t = 0\right)\right\} \end{pmatrix}, \quad (5)$$

where R_t is an indicator for whether an individual has a reward in period t. The uncommitted state analogous to $\left(S = \bullet\right)$ in Equation (1) is $\left(C_t = 0, R_t = 0\right)$. Consider

[3] Caminal and Matutes (1990) solve the model two ways. In one case the firms choose a second period price and discount, where the discount is the switching cost. In the other case, the firms choose a second period price for new customers and inherit a pre-committed price to returning customers, where the difference is the switching cost.

[4] Most theoretical models of reward programs and exogenous switching costs "relocate" customers' preferences each period such that their preferred firm may change.

the common "buy ten, get one free" reward program design. In the intermediate nine time periods this equation will describe a customer's switching costs. When he possesses a reward, the first bracketed expression will change such that $R_t > 0$.

As described earlier, the switching cost in Equation (5) measure what an individual gives up by not purchasing (and hence not progressing in the program) at C credits relative to if the customer were uncommitted (had zero credits). Opportunity costs will be large when a customer's decision not to purchase significantly delays or reduces the probability of receiving (and hence redeeming) a reward. Opportunity costs will be low when not purchasing has little effect on the timing or likelihood of receiving a reward.

2.3 Why Customers Experience Negligible Switching Costs

In this section we describe why switching costs are always small for frequent customers and only significant for infrequent customers on rare occasions. The effects differ between these types because their timing and likelihood of receiving rewards differ. Very frequent purchasers routinely earn rewards and therefore have similar opportunity costs at all credit levels.[5] To illustrate this, consider a "buy ten get one free" program. The effect of not purchasing at zero credits on receipt of the first reward is quite similar to the effect of not purchasing at nine credits on receipt of the second reward. Because there are reasonably short intervals between receiving rewards, neither of these cases will involve significant discounting. The similarity of opportunity costs across credits implies that frequent customers will incur negligible switching costs.

Very infrequent purchasers do not place much value on the firm's good and therefore not on the in-kind reward either. This implies low opportunity costs of not purchasing and negligible switching costs for these types.

In between these two there are some moderate-frequency purchasers whose opportunity costs change significantly as they progress through the program. At zero credits, these customers will face a long, heavily discounted, horizon before earning a reward, if ever. Therefore the opportunity cost of not purchasing for these customers is negligible when uncommitted. However, when they obtain enough credits that a reward is much closer, or more likely, the opportunity cost of not purchasing becomes much greater. This gap between ex-ante and ex-post opportunity costs creates switching costs.

While these moderate frequency customers can face switching costs when close to earning a reward, they will rarely be at such a point. If the program has a finite

[5] Some reward programs do not provide an opportunity to earn multiple rewards. In these cases, very frequent customers face negligible opportunity costs of not purchasing at all credit levels because they expect to earn the reward without adjusting their purchase frequency. We evaluate such a program in the Appendix.

horizon or a costly renewal, they may never get close to earning a reward. If the program never ends (has an infinite horizon), they will visit all states with some frequency but the switching costs will actually lead them to depart switching cost states faster than non-switching costs states (i.e., the steady-state fraction of the time in low switching costs states will be greater).

The purchase frequency also affects the switching costs when a customer has a reward. As described above, customers have an option value of waiting to use a reward until they have a preference for the firm that gave the reward. If a frequent customer opts not to use a reward, there will likely be an imminent period when the customer will want to use the reward. This implies a high option value and hence decreased switching costs when possessing a reward. On the other hand, if a less frequent customer opts not to use a reward, there is unlikely to be an imminent period in which he will want to use the reward. This implies a lower option value and therefore a negligible reduction in switching costs.

In the sections that follow, we show these effects in our empirical setting using simple descriptive regressions and analysis of our estimated structural model. The descriptive analysis allows us to confirm that moderate-frequency customers have switching costs and others do not, but does not allow us to measure their magnitude. Analysis of the estimated structural model allows us to quantify the switching costs by customer type and confirms the implications above.

If reward programs are designed to create switching costs, the implications of this analysis, and its confirmation by our empirical results, is that these programs are not aimed at frequent purchasers. While earning a reward, these customers face negligible switching costs because they realize much of the gains of the program when making their first purchase. When holding a reward, high-volume customers typically have greater option values of rewards implying small switching costs. If these programs are aimed at creating switching costs for less frequent purchasers, any gains must be weighed against the effects of giving discounts to high-volume customers that may represent a large portion of overall sales.[6]

2.4 "Behavioral" Sources of Switching Costs

It is possible that switching costs could arise from other, behavioral explanations. For instance, a business traveler may be a high-volume customer but have a low option value for a reward because he rarely assumes the role of leisure traveler. This explanation is a consequence of the separation of principal and agent, which is a likely motivation for these programs in the travel industry (see Borenstein (1996) and Cairns and Galbraith (1990)). Analyzing the role of the principal-agent problem

[6] We discuss the implications of reward programs for quantity-based price discrimination in detail in Hartmann and Viard (2006).

in reward programs is beyond the scope of this paper and not a prevalent feature of our empirical example.

Other behavioral factors in reward programs documented by the psychology and marketing literature could potentially increase switching costs. For example, if the receipt of a reward builds goodwill with a customer, he may be more likely to purchase after earning a reward. This would result from the customer failing to recognize that the reward program is merely a specific pricing schedule. Kivetz et.al. (2005) has noted that customers may care about how far they have progressed toward a reward. This could also add to switching costs in pre-reward periods. While we ignore these effects in favor of focusing on the pecuniary incentives of reward programs, our model could be used to test for these behaviors.

The rational economic approach to reward programs that we take matches that of the existing theoretical literature, which follows Klemperer (1987a) and Caminal and Matutes (1990). Our approach generalizes the demand side of their models and shows that this generalization diminishes the extent of switching costs in reward programs. It does not, however, eliminate their role and therefore quantifying them remains relevant. Due to the complexity of our demand model and data constraints, exploring the strategic implications addressed by this theoretical literature is not currently feasible in our analysis.[7] However, our generalization of demand in a reward program and the resulting implications for switching costs suggests an avenue for future theoretical work on reward programs.

3 A Model of Demand in a Reward Program

In this section, we define a dynamic demand model that characterizes customers' purchase choices under a reward program. We first develop the random utility model generally and then tailor the model to the specifics of our empirical setting in Section 5. In this section, we specify the model conditional on a given customer type. As motivated in earlier sections, the switching costs and purchase frequency vary across customers' innate preferences, so we specify a distribution over these types in Section 5.

In a reward program, the utility an individual receives from purchasing is composed of the current period utility plus the expected future utility from the purchase, which includes the expectation of earning a reward in the future. We specify customers' current period utilities, define the dynamic game to derive their discounted present value of expected future utility, and end with a discussion of identification.

[7] Other papers considering the strategic implications of reward programs include Banerjee and Summers (1987) and Kim et.al. (2002).

3.1 Current Period Utility

The current period utility individual i receives each period t conditional on the state, S, choice $y_{it} \in \{0,1,2,...,J\}$, and preferences, γ, is:

$$u(S_{it}, y_{it}; \gamma) = \begin{cases} v_{i0t} + \varepsilon_{i0t} & \text{if } y_{it} = 0 \\ v_{ijt}\left(p_{ijt}\left(R_{it}\right)\right) + \varepsilon_{ijt} & \text{if } y_{it} = j \;\; \forall \;\; j \in \{1,2,...,J\} \end{cases}, \quad (6)$$

where v_{ijt} is customer i's utility associated with choice j at time t net of an individual- and time-specific shock to preferences, ε_{ijt}. v_0 represents the utility of the outside good.

The outside option captures substitution toward either no purchase or purchase of a competing firm's product. In most contexts, the researcher will not have data for competing firms so the outside alternative will need to include competitors. This prevents consideration of competitive effects of reward programs, but it does not prevent estimating switching costs generated by the program. Switching costs occur when a customer has an increased incentive to purchase from a firm. This implies a decreased incentive both to purchase from another firm and to not purchase at all. An example of switching costs primarily decreasing the likelihood of no purchase is cigarette addiction.

We assume that the ε_{ijt}'s are independently and identically distributed Type-1 extreme-value errors, conditional on the customer's state. The uncertainty faced by customers is represented by their future ε_{ijt}'s. Customers observe the current value of their error but know only the distribution of future values, while the econometrician knows only the distribution for all periods.

In Equation (6), price, p, is a function of the stock of rewards the individual possesses, R. This is typical of most reward programs, which provide an "in-kind" reward, but there are exceptions in which the reward is a cash payment or not in-kind. Cash payments are easily accommodated in the model because their value can be measured relative to the marginal utility of income, which is a parameter contained in γ. Rewards of goods in other markets are more difficult to incorporate because either the demand model must include demand in the other market, or it must incorporate the reward's cash value.[8]

[8] Lewis (2004) estimates demand in a program which rewarded frequent flyer miles for credits earned through the purchase of grocery and drugstore items. Lewis neither incorporates airline demand nor the cash value of the miles in his model. The value of rewards is therefore mis-specified. Instead, he values them by the additional willingness to pay for groceries on the day a customer receives a reward. As analysis of our model will show, this is an inappropriate measure since a customer may highly value a reward, but it may have negligible impact on his purchase decision relative to periods when he does not earn a reward.

In addition to R, the other essential state variable in a model of demand in a reward program is C, the number of credits toward a reward. C does not appear in Equation (6) because credits do not provide utility, but rather only matter in their effect on future values of R. While this assumption is consistent with rational utility maximization, there is behavioral research indicating that customers may value how many credits they have earned.[9]

Estimation of demand in some programs requires a state variable, say W, to track whether or not a customer has renewed their membership. In our application, renewal is required to retain credits. In other cases, membership may not expire but credits may. In Southwest Airlines' *Rapid Rewards* program, credits expire if a reward has not been earned within a certain amount of time. When credits do expire it is necessary to keep track of the time remaining for each credit. This drastically expands the state-space, which adds significant computational time.

The laws of motion for these state variables create the model's dynamics. These depend on the actual program design but typically involve C increasing with each purchase and R as a function of C. These transition equations can also capture nonlinear, tiered reward structures (e.g., twice the number of credits required for one reward may yield a reward worth more than twice as much). The pace of credit accumulation could also be modeled to depend on previous rewards earned (e.g., Platinum status). Multiple- (versus single-) reward earning opportunities can be accommodated by allowing (or not allowing) the credits to be re-accumulated once a reward is earned. In fact, subject to the state space not growing too large, virtually any relationship between purchases and rewards could be accommodated. We next describe the dynamics.

3.2 Dynamic Optimization Problem

The solution of the dynamics in the model depends on the relevant time horizon. If the program's duration is infinite and customers are automatically renewed then one can solve for a fixed point of the infinite-horizon game. If the program is finite in duration the values in the model will be indexed by time and solved backwardly from the terminal period, T. When T is large enough, early values in the finite horizon will resemble those in the infinite horizon. To be general enough to cover both potential horizons, we therefore proceed by describing a finite horizon.

[9] Kivetz et.al. (2005) illustrate this in an experiment manipulating the total number of credits required to earn a reward, but holding fixed the number of remaining credits left to earn a reward by giving free credits. While their results suggest that customers value credits, there are two caveats. First, their analysis does not consider that the V's in Equation (5) will differ for these programs if customers expect they will have the opportunity to earn future rewards. Second, these biases may not persist through multiple opportunities to earn rewards as customers learn the program's true nature.

In the terminal period, individuals receive only current period utility. The customer's maximization problem is:

$$\max_{y_{iT}}\left\{v_{i0T}\left(S_{iT};\gamma_i\right)+\varepsilon_{i0T},...,v_{i3T}\left(S_{iT};\gamma_i\right)+\varepsilon_{i3T}\right\}. \quad (7)$$

The primary dynamics in the model arise through the state variable C_{it}. A customer recognizes the implication of a choice at time t on C_{it+1}, which affects future prices through R. We specify the dynamics of all state variables, S, following Rust (1987).

Beginning with the penultimate period, the utility from each choice, j, has a non-stochastic component that equals the sum of the non-stochastic portion of current period utility plus the discounted (by factor β) expected future utility from choice j in state S_{iT-1}:

$$V_{ijT-1}\left(S_{iT-1};\gamma_i\right)=v_{ijT-1}+\beta E\left[\max_{y_{iT}}\left\{\begin{matrix}v_{i0T}\left(S_{iT};\gamma_i\right)+\varepsilon_{i0T},\\ ...,\\ v_{i3T}\left(S_{iT};\gamma_i\right)+\varepsilon_{i3T}\end{matrix}\right\}\mid S_{iT-1},y_{iT-1}=j\right], \quad (8)$$

where the expectation is taken over only the ε's because the state space evolves deterministically. The choice-specific utilities for preceding periods are:

$$V_{ijt}\left(S_{it};\gamma_i\right)=v_{ijt}+\beta E\left[\max_{y_{it+1}}\left\{\begin{matrix}V_{i0t+1}\left(S_{it+1};\gamma_i\right)+\varepsilon_{i0t+1},\\ ...,\\ V_{i3t+1}\left(S_{it+1};\gamma_i\right)+\varepsilon_{i3t+1}\end{matrix}\right\}\mid S_{it},y_{it}=j\right]. \quad (9)$$

This equation is solved iteratively back to the first period, $t=1$. The maximization problem in all periods before T is:

$$\max_{y_{it}}\left\{V_{i0t}\left(S_{it};\gamma_i\right)+\varepsilon_{i0t},...,V_{i3t}\left(S_{it};\gamma_i\right)+\varepsilon_{i3t}\right\}. \quad (10)$$

Equation (9) reveals the reward program dynamics. If an individual makes an eligible purchase, the state changes to reflect that the individual's credits have increased, moving him closer to earning a reward. In addition, the progression of time will reflect that he has one less day remaining in the program before either exiting or renewing. If an individual does not purchase, the state changes to reflect an equal number of credits but one less day to earn the reward without renewing, decreasing the individual's expected future utility.

3.3 Identification through Dynamics

A convenient feature of the model is that the dynamics help identify the model parameters. Specifically, there is a unique combination of the discount factor, marginal utility of income, and preferences for the alternatives that produce a given trajectory of choices over time. If the econometrician is willing to assume that customers discount the future at a rate derived from current savings or borrowing interest rates, the full set of parameters can be identified without any exogenous price changes.

The price parameter (negative marginal utility of income) in our model is identified from the value a customer places on a current or future reward. Highly negative price coefficients are consistent with individuals that are very responsive to program incentives, while small negative price coefficients are consistent with individuals that are less responsive. While a customer receives a reward (a price reduction) only at certain points in his purchase history, a reward program affects incentives at every purchase. These incentives differ depending on whether the customer is close to or far away from earning a reward.

To understand how identification is possible in the absence of price variation, consider a model with a single customer and only an intercept and a price variable in the choice equation (the argument easily extends to a model with heterogeneous consumers and additional control variables). In a static setting, the lack of price variation would prevent us from separately identifying the intercept and price coefficient.

In a dynamic setting, identification becomes apparent by considering the extreme case of a customer in the penultimate period of a finite-horizon program and one credit away from earning a reward to use in the last period. If the customer has a zero price coefficient there would be no effect of the possibility of qualifying for the reward. On the other hand, if the price coefficient were very negative there would be an added incentive to purchase in the penultimate period to earn the reward for use in the last period.

A further implication of this argument is that the model can be identified using data before a customer ever earns a reward. This is particularly useful in our setting because we do not observe the exact timing of customers receiving rewards. Thus, our estimates will be less sensitive to the assumptions we make about this timing.

We take advantage of this identification approach in our setting by assuming a value for the discount factor and evaluating a course with fixed prices over time. We choose such a course for two reasons. First, if a firm varies its price over time, the expectations in Equations (8) and (9) must include future prices, greatly expanding the state-space. This must be weighed against the importance of including other variables in the state-space (e.g., other forms of state-dependence that could be correlated with C or R). Second, this allows us to avoid unnecessary complications arising from competitive price dynamics. By focusing our analysis on a municipal course with fixed prices, the prices charged by the firm are not a function of the prices chosen by competitors.

A caveat about this identification argument is that it requires dynamic variation in responsiveness to the program. Specifically, if a customer has a high enough innate preference for the firm such that his closeness to earning a reward does not matter, the model cannot distinguish between a slightly greater intercept and a slightly less negative price coefficient. Similarly, if a customer has a low enough innate preference for the firm such that he is unaffected by the program the model cannot separately identify a slightly lower intercept from a slightly more negative price coefficient.

In the next section we demonstrate that in our empirical setting there is a segment of customers that alter their behavior as they approach a reward. This implies that the price coefficient in our model will be identified by these customers. The price coefficient for other types of golfers will be "extrapolated" based on the form of correlation assumed in the heterogeneity distribution.[10] That said, our empirical application contains additional variation that helps identify the price coefficient, which we describe in Section 5.

Israel (2005) considers identification in a similar setting in which customers 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 its negative correlation with tenure with the firm, which has a potential state-dependence impact of its own. In a reward program setting with multiple reward-earning opportunities this is only true for customers who have never received a reward. Even for those who have never received a reward or for participants in a single-reward program there is not a one-to-one correspondence between tenure and number of credits remaining to earn a reward because customers do not purchase every period and because customers may have a history with the firm before joining the program.

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.

[10] While this assumption is restrictive, it will not affect inference of switching costs. Even if intercepts and price coefficients are incorrectly extrapolated, our model will identify customers not responsive to the number of credits as not having switching costs. Making this assumption results in a simpler model to understand the deeper question of exactly how the incentives of these programs affect demand.

4.1 The Structure of the Reward Program

The program rewarded golfers by giving them a green fee certificate after purchasing ten rounds of golf at member 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. Credits toward the reward could be earned any day of the week, but the reward could not be used on Fridays, Saturdays, or Sundays.

The golf management company designed most dimensions of the reward program after Southwest Airlines' *Rapid Rewards* program, one of the biggest and most successful airline frequent flyer programs. The most important common features relate to how credits and rewards are earned. Like Southwest's program, a purchase of any type of round (whether cheap or expensive) yielded the same credit toward a reward. In addition, once a reward was earned, the customer could begin earning a new reward, rather than saving credits for a reward of greater value.

The program required a paid membership, but immediate benefits of the membership roughly offset the monetary expense of signing up. Membership cost $29.95, but entitled the golfer to an immediate discount of $16.50 off the price of a Monday through Thursday game, $10 worth of balls at the driving range, and other smaller promotions. The membership lasted for one year and required a renewal within the sixty following days to continue and retain credits earned. Though the membership fee had to be repaid for renewal, we assume away any pecuniary cost by the same logic as the initial membership payment. This does not, however, assume costless renewal. Those golfers that purchased less frequently would have found it more costly to renew as it might require an extra trip to the course. Our analysis focuses on the golfers' first year in the program, leading up to their first renewal decision.

We analyze a municipal course located in southern California with a reward discount of 100% (i.e., a free round). The course is open to the public (does not require a membership fee) and golfers do not need to belong to the reward program to play on the course. We do not observe golfers outside the program.

There are four other eighteen-hole courses within a five-mile radius which might be considered its potential competitors. Three of these courses are priced at over $50 on weekdays; the fourth has a price of over $20 on weekdays. The equivalent price at the course we analyze is about 25 percent less than the cheapest of these courses. Although we do not observe purchases at any of these courses, their pricing did not affect the pricing at the course we study because the local government set its prices.

4.2 Golf Details

Golfers can purchase one of three types of rounds. An 18-hole round is the typical round with a price of $16.50 ($20 on Saturdays and Sunday). Late in the day, a golfer may purchase a Twilight round for $10.50 ($12.50 on Saturdays and

Sundays), which typically involves between 9 and 18 holes depending on the golfer's start time. Golfers can also purchase 9 or fewer rounds for $9.50 ($11.50 on Saturdays and Sundays) late in the day or on the back-9 in the morning.

Our data includes daily purchase decisions by each golfer between January 1, 2000 and December 31, 2001. We focus on golfers that joined and finished their first year of the program during this period. Each golfer is therefore observed for a period of 365 days, although the exact calendar days differ across golfers according to when they joined the program.

For each golfer, we observe the number of credits earned and when they qualify for a reward. We do not observe the exact date the golfer receives the reward, so we assume the reward is issued immediately. We also do not observe when the golfer uses a reward so we assume it is used when making the next eligible purchase and we restrict golfers to hold at most one reward to keep the state space small. While these would not be good assumptions for some other settings, such as those with a principal-agent problem where the traveler saves rewards for personal travel rather than use them on the next available trip, it is reasonable here. While observing the use of a reward provides price variation and aids in identification, it is not necessary. As explained in the previous section, intertemporal changes in customer expectations of earning a reward during pre-reward periods provide sufficient variation.

4.3 Summary Statistics

The analysis considers 531 golfers that we observe for their first year in the program. This provides 193,815 observations. Summary statistics for the golfers are presented in Table 1. On average, the golfers played 11.55 times and earned between 0 and 8 rewards. The majority of rounds purchased were 18-hole rounds. Renewal rates in the program generally increase in the number of rewards earned.

Three hundred, thirty-one of the golfers did not earn a reward during their first year. The reward program was practically costless to the firm for these customers.[11] To the extent that some of these customers believed ex-ante that they 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. Ninety-five percent of these customers did not renew, consistent with more costly renewal for less frequent players. There are at least two reasons why five percent of these golfers might renew their membership despite not having earned a reward. They may have valued non-reward benefits of renewal[12] or they may have believed it was worth it to retain their credits to earn a reward in the future. In fact, of the golfers who never earned a reward, those who renewed had an average of 6.8 (median of 7.5) credits in their pocket at the end of the year while those who did not renew had an average of 3.7

[11] The marginal cost of the program is negligible because the system is computerized.

[12] An example of a non-reward benefit is that members of the program had access to Twilight and Super-Twilight rates one hour earlier than non-members.

(median of 3). This pattern also holds for golfers who earned rewards. Those with more credits at the end of the year were much more likely to renew.[13]

4.4 Descriptive Analysis

Our switching costs discussion implies that the reward program should have the greatest effect on moderate-frequency purchasers relative to low- and high-frequency purchasers. We can confirm this in a descriptive analysis by determining if the number of credits a customer holds is related to their purchase probability and how this varies by type of customer. One way to do this is to relate the time between purchases to the number of credits held and allow this effect to differ across consumer types. To do so, we run a fixed-effects regression of logged time between purchases on logged number of credits held and split the sample among low-, moderate-, and high-frequency players.

We define the customer type by the maximum number of days between purchases for each customer over their year in the program.[14] We define low-frequency purchasers as those with a maximum time between purchases of ninety days or more (24% of the sample), high-frequency purchasers as those with a maximum time between play of fewer than forty days (27% of the sample) and the remaining customers in between as moderate-frequency purchasers (48% of the sample). These findings are robust to changing the cutoffs by up to two weeks in either direction.

Specifically, we estimate the following regression on the three subsamples:

$$\log\left(timebw_{ir}\right) = \lambda \log\left(credits_{ir}\right) + \varepsilon_{ir}\ , \quad (11)$$

where $timebw_{ir}$ is the time between games r and $r-1$ and $credits_{ir}$ is the number of credits after game $r-1$ for golfer i.[15] In performing the analysis, we demeaned all of the variables within each golfer so this is equivalent to a regression with fixed effects at the golfer level. Results for the low-, moderate-, and high-frequency golfers are shown in Columns 1, 2 and 3 respectively of Table 2.

---

[13] To confirm this we ran the following logit regression:

$$\underset{(0.400)}{\underset{-3.70}{renew_i = \alpha}} + \underset{(0.159)}{\underset{0.802}{\beta_1 rewards_i}} + \underset{(0.068)}{\underset{0.275}{\beta_2 credits_i}} + \underset{(0.032)}{\underset{-0.039}{\beta_3 rewards_i * credits_i}} + \varepsilon_i$$

where $renew_i$ is an indicator for whether golfer i renewed or not at the end of the year, $rewards_i$ is the number of rewards golfer i has at the end of the year, and $credits_i$ is the number of credits golfer i has at the end of the year. Each additional credit increases the odds of renewing by 2.6% which is a large effect given that the average probability of renewal is 15% and each additional reward increases the probability of renewal by 7.6%.

[14] We focus on the maximum time between purchases because it separates out those with extended layoffs due to injury or some other persistent positive shock to the opportunity cost of golfing.

[15] Since we do not observe when golfers first began playing, we have two less observations than the total number of games played for each golfer.

The results are consistent with the implications from our switching costs definition. Number of credits earned has a significant and negative effect on the time between play for the moderate-frequency players, with an elasticity of eight percent. Thus, the program appears to be effective in creating lock-in for these customers. Number of credits does not have a significant effect on the high-frequency players, consistent with their opportunity costs of not purchasing remaining constant as they progress through the program. Number of credits actually has a positive and significant effect on the low types. Holding more credits increases time between play with an elasticity of 13%. Given our definition of low types as those with a maximum time between purchases greater than ninety days, it may be more accurate to describe these golfers as those experiencing a significant layoff at some point in the program.

Extended layoffs in the program could occur from injury, moving, or some other persistent positive shock to the opportunity cost of golfing. Because such a shock is more likely to have occurred later in the program, this is reflected by a positive relationship between duration and credits. Identifying switching costs for these individuals is therefore not possible in this simple specification.

A persistent shock to the outside alternative also raises a potential concern about our i.i.d. assumption of the logit errors. However, as we will show in the next section, including a parameter that allows for an increase in the value of the outside good when a customer has been inactive for more than sixty days helps accommodate these occurrences.[16] By accounting for these layoffs, this parameter allows us to actually identify whether switching costs exist for these individuals, which is not possible in the simple specification.

5 Model Estimation

In this section we discuss the specific application of the model to our setting and discuss its estimation. In our setting, $J=3$ with $j=1$ corresponding to an 18-hole round, $j=2$ to a 9- to 18-hole round, and $j=3$ to a 9-hole or less round. Utility net of the logit errors for the inside and outside goods is parameterized as follows:

$$\begin{aligned} v_{ijt} &= \gamma_{ij0} + \gamma_{i1} p_{ijt}\left(D_t, R_{it}\right) \\ v_{i0t} &= \gamma_{i2} \log\left(H_{it}\right) + \gamma_{i3}\left(H_{it} \geq 60\right) + \gamma_{i4}\left(D_t > 5\right) \end{aligned}. \quad (12)$$

As is common, we specify the non-stochastic portions of current period utility to be additively separable. As in Section 3, p_{ijt} is a function of R_{it}, but also the day of the week, D_t, to account for the fact that the course offers a lower price on weekdays.

[16] The indicator only accounts for the shock once the customer has been inactive for 60 days. To account for inactivity before 60 days an unobserved binary shock with a Markov transition matrix could be included in the model.

D_t must also enter the state space because customers' decisions will depend on future prices.

While we primarily rely on the model dynamics to identify the price coefficient, price variation across days of the week also plays a role in determining the price coefficient. In a static sense, the relative prices of the choices differing across weekdays and weekends can determine the price coefficient because we do not estimate a separate weekend indicator for each type of round. Furthermore, in a dynamic sense, different weekday demands either leading up to or following higher weekend prices can also affect the price coefficient.

The other primary difference from the model specified in Section 3 is the inclusion of H_{it}, the number of days since customer i last purchased. Following Hartmann (2006), this allows for the marginal utility of golf to fall after consumption. We also assume the log of H_{it} affects the current period utility additively and cap H_{it} at sixty days to capture the effect of golfers' layoffs. γ_{i2} captures the marginal affect of waiting an additional day before consuming and γ_{i3} is an indicator to account for the cap at 60 days.

γ_{i1} captures the golfer's price sensitivity (or negative marginal utility of income). γ_{ij0} captures customer i's taste for good j, while γ_{i4} captures customer i's taste for weekend relative to weekday golf. We expect a golfer's utility from the outside good to be lower on the weekend than on weekdays. The complete set of state variables in our application is $S = \{H, C, D, R, W\}$.

5.1 Laws of Motion for State Variables

The number of credits increments with any purchase, remains the same with no purchase and resets to zero at the tenth credit when a reward is earned:

$$C_{it} = \begin{cases} C_{it-1} + 1 & \text{if } y_{it-1} > 0 \text{ and } C_{it-1} < 9 \\ C_{it-1} & \text{if } y_{it-1} = 0 \\ 0 & \text{if } y_{it-1} > 0 \text{ and } C_{it-1} = 9 \end{cases} . \quad (13)$$

A reward is earned after making ten purchases of any type of round and is used on the next 18-hole round purchased on a Monday through Thursday. To keep the state space small, we restrict individuals to hold at most one reward. This is a reasonable assumption given that a golfer would lose the time value of money by holding onto a free game. One would want to relax this assumption if modeling an airline or hotel reward program in which the customer earns rewards as a business traveler and consumes them as a leisure traveler, which would lead to stockpiling of rewards. Thus, for our empirical setting the transition equation for rewards is:

$$R_{it} = \begin{cases} 1 & \text{if } \{C_{it-1} = 9\,, y_{it} > 0\} \\ 0 & \text{if } y_{it-1} = 1, \mathrm{D}_{it-1} < 6, C_{it-1} < 9 \\ R_{it-1} & \text{otherwise} \end{cases} . \quad (14)$$

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} is bounded above by 60:

$$H_{it+1} = \left\{ \begin{array}{lll} H_{it} + 1 & \text{if} \quad y_{it} = 0 & \text{and} \quad H_{it} < 60 \\ H_{it} & \text{if} \quad y_{it} = 0 & \text{and} \quad H_{it} = 60 \\ 1 & \text{if} \quad y_{it} \in \{1,2,3\} & \end{array} \right\} . \quad (15)$$

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

$$D_{t+1} = \left\{ \begin{array}{ll} D_t + 1 & \text{if} \quad D_t < 7 \\ 1 & \text{if} \quad D_t = 7 \end{array} \right\} . \quad (16)$$

All golfers in our data set have initially enrolled in the reward program for 365 days at the end of which they have sixty days to renew their membership and retain any earned credits. There is a renewal fee, but the immediate benefits approximately offset this. The renewal decision thus depends on the desire for future play. The variable W indicates whether or not the customer has renewed their membership to retain credits for another year. Its law of motion is:

$$W_{it} = \begin{cases} 1 & \text{if } y_{it-1} > 0 \text{ and } 365 < t - 1 < 425 \\ W_{it-1} & \text{if } y_{it-1} = 0 \text{ and } t - 1 \neq 365 \\ 0 & \text{if } t - 1 = 365 \end{cases} . \quad (17)$$

This specification assumes that the renewal and purchase choice during the sixty-day window for a non-renewed member are the same. We therefore avoid modeling a separate renewal choice.

The presence of the renewal decision requires us to account for the time until renewal. This increases the size of the state space enough that it is too computationally intensive to estimate an infinite-horizon problem. We therefore solve the model for a two-year horizon where the second year determines the value of renewing, which affects decisions in the first year. We then use the solution of the model for the first year to estimate the likelihood of the data for golfers’ first year in the program.

5.2 Heterogeneity Specification

Up to now, we have specified the model conditional on a given individual's parameters, γ_i. Heterogeneity is essential for two reasons. First, as described in Section 2, the level and prevalence of switching costs vary by customer type. Second, a reward program makes demand state-dependent and it is commonly recognized that extensive heterogeneity is necessary to properly identify state-dependence.

To take the model to data and estimate the mean and variance of these parameters, we assume that they are normally-distributed, random coefficients. We estimate the model using simulated maximum likelihood and Ackerberg's (2001) importance sampling technique. This involves calculating the likelihood at a wide range of candidate parameter values, then searching for the parameter vector that weights these to maximize the likelihood.

The choice probabilities have the typical logit form with the choice specific value functions instead of the current period utilities:

$$\Pr\left(y_{it} \mid S_{it};\gamma_i\right) = \sum_{j=0}^{3} \frac{\left\{y_{it} = j\right\} \exp\left(V_{ijt}\right)}{\sum_{k=0}^{3} \exp\left(V_{ikt}\right)} . \ (18)$$

Since we observe each individual for T days, the individual's likelihood function is:

$$L_i\left(S_{i1},...S_{iT}, y_{i1},..., y_{iT};\gamma,\Sigma\right) = \int \prod_{t=1}^{T} \Pr\left(y_{it} \mid S_{it};\gamma_i\right) f\left(\eta_i\right) d\eta . \ (19)$$

The random coefficients are:

$$\gamma_i = \gamma + \Gamma \eta_i , \ (20)$$

where the γ vector contains the mean parameters of the random coefficients and Γ is the Cholesky decomposition of Σ, the variance-covariance matrix of the random coefficients. η_i is a standard normally distributed vector that we must integrate out to evaluate the individual likelihood function.

The joint likelihood is the product of the individual likelihoods is:

$$L = \prod_{i=1}^{N} L_i\left(S_i, y_i;\gamma,\Sigma\right) . \ (21)$$

Table 3 reports the model estimates. The estimates themselves reveal little about the reward program because they describe only the current period utility. The price coefficient and state-dependence are both negative as expected. Golfers prefer 18-hole rounds to twilight rounds and the latter to 9-hole or less rounds. Golfers prefer to play on weekends relative to weekdays. Golfers who have not played in over 60 days are less likely to play, consistent with them experiencing layoffs. There is significant heterogeneity in all of the random coefficients. In the next section, we evaluate the switching costs effects implied by these parameters.

6 Measuring Switching Costs in Reward Programs

We quantify the switching costs created by the reward program using various counterfactuals. We evaluate these under two scenarios: the observed program and a program not requiring renewal (henceforth referred to as a continuous program). We use the continuous program as a baseline case because the prospect of renewal distorts the pattern of switching costs from this baseline. As a result, it is easiest to describe the observed program by referencing how it deviates from the continuous.

To measure the magnitude of switching costs and how they evolve as a customer moves through the reward program we use our model and estimated preference parameters. Given the customer's state, his switching costs at C credits are:

$$SC\left(H_{it}, C_{it}, D_t, R_{it}\right) = {}^{1}\!/_{\gamma_{i1}} \left[\begin{array}{l} \left\{ \log\left[\exp\left(V_{i1t}^{C}\right) + \exp\left(V_{i2t}^{C}\right) + \exp\left(V_{i3t}^{C}\right) \right] - V_{i0t}^{C} \right\} - \\ \left\{ \log\left[\exp\left(V_{i1t}^{0}\right) + \exp\left(V_{i2t}^{0}\right) + \exp\left(V_{i3t}^{0}\right) \right] - V_{i0t}^{0} \right\} \end{array} \right], \quad (22)$$

where $V_{ijt}^{C} = V_{ijt}\left(H_{it}, C_{it}, D_t, R_{it}\right)$ and $V_{ijt}^{0} = V_{ijt}\left(H_{it}, 0, D_t, 0\right)$. This definition modifies Equation (5) to include four choices and the logit errors specified in Section 5. Utility is converted to dollars by dividing the bracketed expression by the price coefficient (which is assumed to equal one in the simple Hotelling-type model considered in Section 2). Given the model specification, Equation (22) is also the common log-sum expression of a customer's willingness to pay for a policy change in a logit demand model.[17]

After measuring the switching costs in these two scenarios, we consider how the reward program (and any resulting switching costs) affects demand elasticities.

[17] This expression can also be derived using Shum (2004)'s definition of setting the switching costs equal to the dollar value necessary to equalize purchase probabilities in non-committed and committed states.

6.1 Continuous Program

We analyze the switching costs at zero to nineteen past purchases for golfers with varying preferences for 18-hole rounds of golf at the course. Specifically, we consider customers from the 5th, 25th, median, 75th and 95th percentiles of the distribution of customers' innate preferences.[18] We measure their switching costs on a Monday when the golfer just purchased the day before (i.e., $H = 1$).[19] Following the discussion in Section 2, we describe the switching costs separately for cases when a customer does not possess a reward and has just earned a reward.

Figure 1 illustrates the switching costs in a continuous program. There are four important characteristics of switching costs when customers do not hold a reward (i.e., at zero to nine and eleven to nineteen past purchases).

First, switching costs monotonically increase (almost linearly) as a customer earns additional credits toward a reward. As he does so, the next and future rewards become more imminent, increasing the opportunity cost of not purchasing. Second, a customer's switching costs return to their initial level after he cashes a reward.[20] This highlights a caveat about the role of lock-in in most reward programs: customers are routinely "released" such that they have little or no lock-in.

Third, frequent players have negligible switching costs when not holding a reward. This arises because the reward program has a similar effect on their opportunity cost of not purchasing for any number of credits they may possess (including before they ever purchase). This is most easily seen in Figure 2, which depicts the value of the additional purchase incentives generated by the program at different credit levels. The flat slope in Figure 2 for the top quartile indicates that they essentially view the program as a uniform price decrease of just less than $1.65 (the per-purchase value of a reward).

Fourth, less frequent customers have non-negligible switching costs when close to earning a reward as seen in Figure 1. This is consistent with the steep slope in their purchase incentives shown in Figure 2. For example, purchase incentives for the 5th percentile golfer are valued at $0.70 at zero credits but at $1.29 just before earning a reward leading to switching costs as high as $0.59 just before earning a reward.

---

[18] We determine these types by adjusting each of their 18-hole intercept and price coefficients such that their utility from playing an 18-hole round places them in the appropriate percentile of the play frequency distribution under the uniform price regime. In doing so, we account for the correlation between a golfer's intercept and price coefficient. We set the value of all other parameters to their mean values in Table 3.

[19] While we could have chose the average value of H, a value of 1 is the only value that every golfer certainly achieves at some point in the program, given that all purchased at least once.

[20] In the program we analyze, a customer earns a credit even when cashing a reward so switching costs never return to zero. In programs where customers do not earn a credit upon using a reward, switching costs drop to zero after reward use.

Although switching costs are significant for less-frequent customers just before earning a reward, these customers spend very little time in these states. The switching costs will lead to greater purchase probabilities and consequently faster exit from these states than those without switching costs.

We now turn to switching costs when golfers have a reward (i.e. past purchases equal ten in Figure 1). In this case, switching costs remain small for the highest types but are as large as $1.72 for the fifth percentile golfer. Frequent customers have greater option values from waiting to use a reward because they are more likely to want to purchase again soon. This lowers the opportunity costs and consequently the switching costs.[21]

6.2 Observed Program

The observed program differs from the continuous program in that after 365 days in the program, the customer has sixty days to make a purchase and renew in order to retain their credits. Frequent purchasers will be relatively unaffected by renewal because they are very likely to renew or purchase frequently enough that they can earn a reward before expiration without altering their behavior much. Less frequent purchasers, on the other hand, are unlikely to renew. Depending on the prospects of earning the reward before expiration, a customer may either accelerate purchases to get a reward before expiration, or give up on the prospect of earning a reward.

Figure 3 depicts the switching costs by number of credits for each type of golfer in the observed reward program. All states used to measure the switching costs are identical to those used in the continuous case except we assume golfers are ninety days into the program. This makes it possible for all types to reach any of the ten possible credit levels.

When customers do not hold a reward, there are two primary differences from the continuous case. First, the highest switching costs, those faced by the 5^{th} percentile golfer, are more than 3.5 times greater than in the continuous program, reaching a high of just over two dollars. This is because the opportunity costs when close to earning a reward are exaggerated by the need to qualify before expiration and the opportunity costs when far from a reward are diminished because there is little prospect of ever qualifying. Customers in the top quartile do not face the same dilemma when close to earning a reward because they are likely to renew and retain their credits. These customers face negligible switching costs as they did in the continuous program.

Second, switching costs can decrease with the number of credits, rather than rising in a nearly linear fashion, as in the continuous program. For example, switching costs

[21] This applies more generally to any type of coupon or promised future rebates. Its effect on the customer's purchase propensity is lower when this option value exists.

for the median golfer decrease slightly between eight and nine credits. Holding the time until expiration fixed, additional credits can either increase the incentive to purchase, if a customer previously faced little prospect of earning a reward, or decrease the incentive to purchase, if a customer no longer needs to accelerate his purchases in order to ensure he will earn a reward. These effects are amplified in a finite-horizon, single-reward opportunity program, which we consider in the Appendix. The observed renewable program is a hybrid between such a finite horizon program (which can be viewed as having an infinite renewal cost) and the continuous program (which can be viewed as having costless renewal).

The overall message from Figure 3 is that switching costs generally do not arise. Customers in the top quartile essentially face no switching costs, but represent over eighty percent of rewards earned and more than two-thirds of purchases. Lower-frequency golfers can have switching costs but, as shown in Table 1, the 331 golfers that never earned a reward only purchased 3.86 times on average and most lost their credits because they did not renew. These golfers therefore rarely reached states in which switching costs are significant. If the firm were to have used a continuous program, all of these customers would have eventually reached these states, but they would have spent most of their time in states with low switching costs.[22]

6.3 Elasticity Implications of Switching Costs

Since we do not observe demand by all the firm's customers we cannot adequately measure the firm's marginal (opportunity) costs and evaluate the program's profitability. However, to evaluate the role of these switching costs we measure the firm's demand elasticity with and without a reward program. If the reward program generates significant switching costs, demand should be less elastic when the firm has a reward program. Figure 4 provides the elasticities under the reward program and two uniform price scenarios: the undiscounted price at the course and a uniform price lowered by $1.65 (the per-purchase value of an earned reward).

Evaluating the elasticities by type and credit in the program, we find that if a customer has substantial switching costs (i.e., is below the median and is close to earning a reward) their demand is less elastic. When these customers realize practically no incentive from the program (e.g., zero credits for the 5th percentile golfer), elasticities are similar to those if the program did not exist. When the program begins to affect the customer's incentives (e.g., three through five credits for the 5th percentile golfer), demand becomes more elastic than would be the case if the program were not offered. This is likely because these customers realize further purchases will lock them into the firm.

[22] If one were to solve for the ergodic distribution, greater purchase probabilities at higher credit levels would imply more time spent at low credit levels in the steady-state.

For high types, the demand elasticity under the reward program at all credit levels is almost exactly what it would be under a uniform price program with a price decrease of $1.65, consistent with them having negligible switching costs. Overall, switching costs generally do not reduce the demand elasticity. Demand is only less elastic on the rare occasions when lower-frequency customers have switching costs. This comes at the cost of more elastic demand before these customers become locked-in.

7 Conclusion

Our analysis suggests that switching costs are not an important feature of reward programs. The primary insight is that customers who highly value a firm's reward program before participating will face negligible switching costs because their incentive to start purchasing in the program is as strong as their incentive to continue purchasing once in the program. Customers who highly value the program are those who highly value the product. These customers purchase with the greatest frequency and therefore comprise the greatest fraction of demand.

The programs can lock-in those who place a lower value on the program when they are close to earning a reward because they have little incentive to purchase before earning their first credit but a large incentive when the prospect of a reward is near. However, customers who place little value on the program up-front are those who place a low value on the product. These customers rarely purchase frequently enough to get close to a reward.

The key to assessing switching costs in reward programs is comparing a customer's opportunity costs of not purchasing as they progress through the program to those before entering the program. We demonstrate this in a program with a simple "buy ten get one free" program. Analysis of switching costs in more complicated programs should proceed by considering how the opportunity costs will differ from those in such a simple program.

For example, consider a traveler who earns credits in an airline frequent flyer plan primarily on business flights and uses the rewards primarily for leisure flights. The fact that the traveler's employer pays for the qualifying tickets lowers the cost of purchasing at all levels of accumulated mileage. This includes the first qualifying flight, so switching costs are unlikely to be higher. When the traveler holds a reward, the fact that the traveler will wait to use the reward for a leisure trip decreases the option value of using the reward later. Since this does not have a corresponding effect on their opportunity cost on their first purchase, this increases the traveler's switching costs when he holds a reward.

Our switching costs analysis also suggests at least two reasons to extend theoretical models of reward programs beyond two periods. First, customers in reward programs typically can wait to use an earned reward until they favor the firm more highly. A two-period model imposes that a customer use or lose the reward in the second period, overstating switching costs. Second, typical reward programs require multiple qualifying purchases, which allows for the possibility that a firm can build switching costs before customers are guaranteed a discount. These factors may change the competitive dynamics that have been the focus of this literature. An empirical analysis of the competitive equilibrium is currently infeasible because modeling both forward-looking firms and customers is still computationally intractable and it is unlikely a researcher could gain access to individual-level data on the credits earned in multiple programs.

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.

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.

Hartmann, W.R. (2006). "Intertemporal Effects of Consumption and Their Implications for Demand Elasticity Estimates," *Quantitative Marketing and Economics*, forthcoming.

Hartmann, W.R. and V.B. Viard (2006). "Quantity-Based Price Discrimination Using Frequency Reward Programs," working paper.

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

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

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). "The Competitiveness of Markets with Switching Costs," *RAND Journal of Economics*, 18, 138-150.

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

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.

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

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

Rust, J. (1987). “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55(5), 999 – 1033.

Shum, M. (2004). “Does Advertising Overcome Brand Loyalty? Evidence from the Breakfast-Cereals Market,” *Journal of Economics & Management Strategy*, 13(2), 241 – 272.

Viard, V. B. (2005). “Do Switching Costs Make Markets More or Less Competitive?: The Case of 800-Number Portability,” forthcoming in *RAND Journal of Economics*.

Appendix: Switching Costs in an Expiring Program

In this appendix, we consider the switching costs effects of a finite-horizon program. While switching costs from the continuous program are always positive, this is not necessarily true for a finite-horizon program. The intuition is clearest in a finite-horizon, single-reward program. Figure A1 depicts switching costs in such a program for the types of golfers considered in the counterfactuals in Section 6.

We analyze this program using the same states and customer types as those used in the observed case. In this case, the switching costs of customers below the median are similar to the observed program. However, the switching costs for the most frequent purchasers differ. Both the 75th and 95th percentile golfers experience negative marginal switching costs (switching costs decline as customers gain an extra credit holding all else equal). These decreasing switching costs eventually lead to negative switching costs for both types. At nine credits, switching costs are \$-0.06 and \$-0.34 for the 75th and 95th percentile customers respectively. At zero credits, both of these types have some incentive to purchase and earn the reward before the end of the program. However, when fewer credits remain to earn a reward, holding all else equal, they need not accelerate purchases to qualify. To take an extreme case, a very frequent purchaser with nine credits and 360 days left to earn the last credit for the only possible reward has no reason to purchase faster than if the program did not exist.

When customers have a reward, switching costs still reflect the option value effect. After redeeming a reward, the customer has no prospect of earning a reward and is therefore strictly worse off than before beginning to purchase from the firm. This is reflected by the negative values from 11 to 19 credits. The absolute values of these numbers mirror the purchase incentives of the program ex-ante.

Although we have illustrated negative switching costs using a single-reward program, they can also occur in unlimited reward programs when there is a finite horizon or costly renewal. The intuition is that as a customer approaches the program's end or a prohibitively costly renewal decision, the customer has limited time to earn a single reward. In these cases, scenarios similar to Figure A1 arise.

Table 1

Purchases and Rewards During First Year in Program

Total Rewards Earned	N	Average Number of Purchases				Percent Renewed
		Total	18 Holes	9-18 Holes	9 Holes or less	
0	331	3.86	1.50	1.42	0.94	5%
1	95	13.98	5.49	3.63	4.85	22%
2	55	24.64	10.95	6.45	7.24	38%
3	21	33.95	15.43	7.43	11.10	29%
4	16	42.50	17.25	11.88	13.38	31%
5	7	54.71	39.43	3.86	11.43	71%
6	5	62.80	36.20	15.60	11.00	60%
7	0					
8	1	82.00	21.00	3.00	58.00	0%
Total	531	11.55	5.08	3.06	3.41	

Descriptive statistics for 531 golfers in the sample during their first year in the program. Sample consists of all golfers who joined and finished at least one year in the reward program between January 1, 2000 and December 31, 2001.

Table 2

Fixed Effects Regressions of Log(Time Between Play) for Different Frequency Golfers

	Low Freq.		Medium Freq.		High Freq.	
	1		**2**		**3**	
Log(Credits)	0.128	**	-0.079	***	0.002	
	(0.061)		(0.030)		(0.034)	
Intercept	2.577	***	2.314	***	1.679	***
	(0.092)		(0.048)		(0.055)	
N	1,365		2,700		1,539	
R^2	0.635		0.376		0.231	

Player fixed effects included in all models. Model 1 includes all golfers who had a maximum time between play of 90 days or more, Model 2 includes all golfers who had a maximum time between play of 40 or more days but less than 90 days, and Model 3 includes all golfers who had a maximum time between play of 40 days or less. Standard errors in parentheses. * = 10% significance, ** = 5% significance, *** = 1% significance.

Table 3

Model Estimates

	Mean	Random Coefficients: Variance-Covariance Matrix						
Golfing								
18-Hole Intercept	-2.993	**4.021**	2.399	2.079	-0.206	0.000	0.038	0.000
	(0.139)							
9-18 Holes Intercept	-4.015	2.399	**2.175**	1.288	-0.140	0.018	0.019	-0.262
	(0.104)							
9-Hole or Less Intercept	-4.473	2.079	1.288	**1.973**	-0.124	0.007	-0.001	-0.103
	(0.102)							
Price Coefficient	-0.123	-0.206	-0.140	-0.124	**0.013**	0.000	-0.002	0.011
	(0.008)							
Outside Alternative (Not Golfing)								
Days Since Last Purchase	-0.024	0.000	0.018	0.007	0.000	**0.002**	-0.002	-0.018
	(0.003)							
60 Days Since Last	0.068	0.038	0.019	-0.001	-0.002	-0.002	**0.006**	0.001
	(0.005)							
Weekend Indicator	-1.207	0.000	-0.262	-0.103	0.011	-0.018	0.001	**0.439**
	(0.044)							

Estimates from the dynamic demand model of customers in the golf reward program. Data consists of daily purchase decisions of the 531 golfers in the sample over their first year in the program. Standard errors are in parentheses. Variances of the random coefficients are in bold. Off-diagonal elements are the covariances of the random coefficients.

Figure 1
Switching Costs by Type of Golfer
Continuous Program

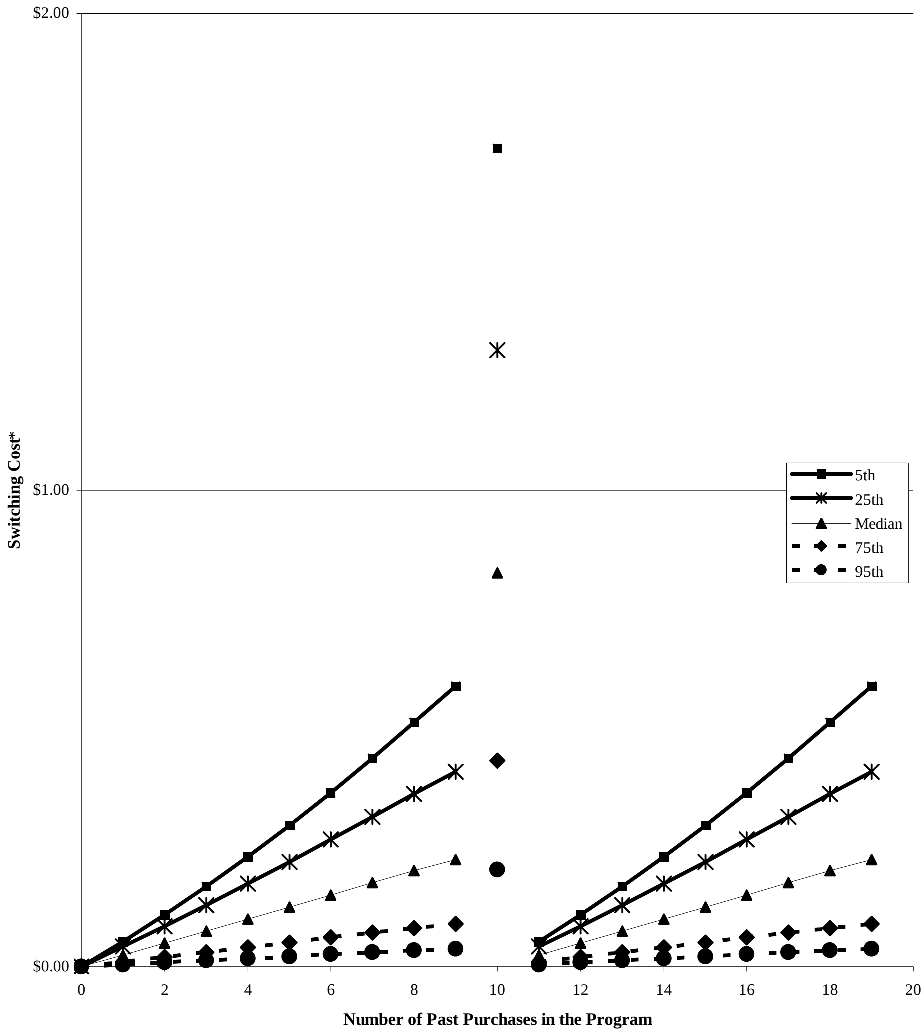

Switching costs in a continuous program as a function of number of credits earned for five different customer types. Switching costs are measured as defined in Section 6 of the text and calculated assuming it is a Monday and the golfer has played the day before. At 0 to 9 and 11 to 19 past purchases, the switching costs are measured when the customer does not possess a reward. At 10 past purchases, the customer is assumed to have a reward. Customer percentiles are defined by adjusting their 18-hole intercept and price coefficients so that their utility from playing an 18-hole round places them in the appropriate percentile of the play frequency distribution under the uniform price regime. In doing so, we account for the correlation between a golfer's intercept and price coefficients. We fix all other coefficients at their mean value.

Figure 2
Value of Purchase Incentives in Continuous Reward Program
by Credits and Type of Golfer

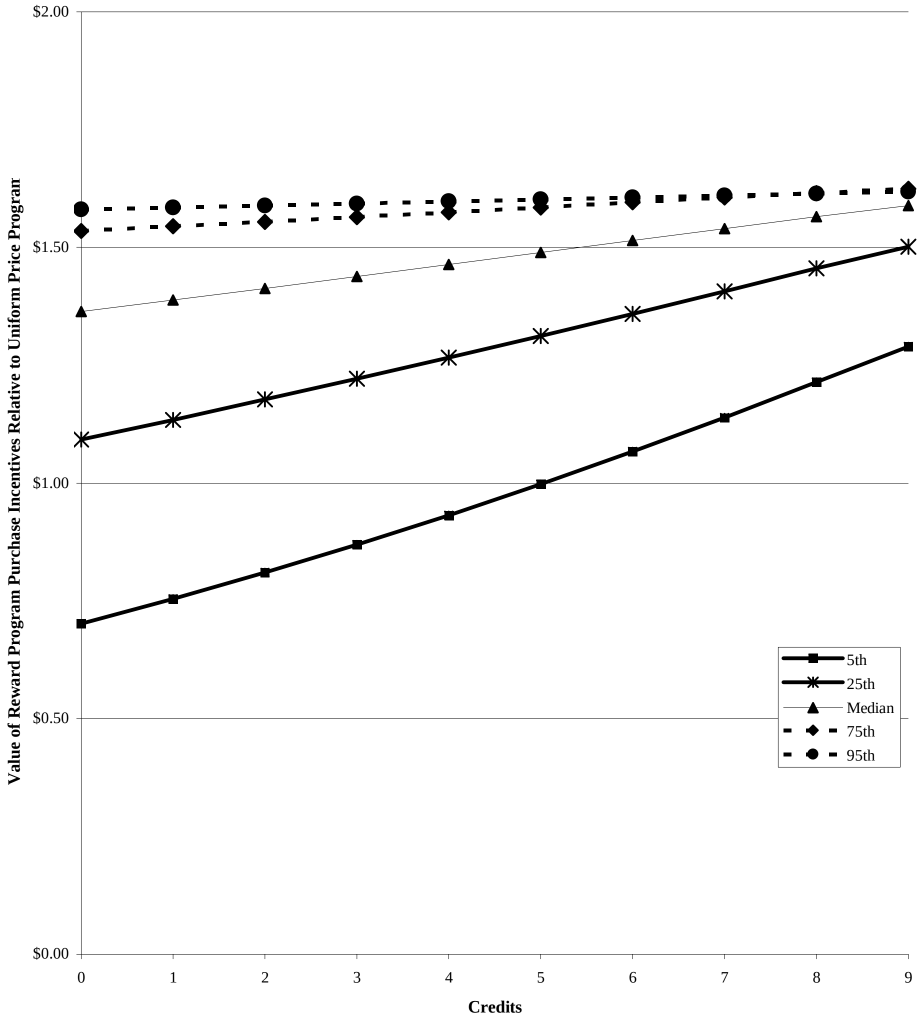

Value of purchase incentives in continuous reward program relative to a uniform price scenario as a function of number of credits earned for five different customer types. Program evaluated assuming it is a Monday, the golfer has played the day before, and under the reward program, the golfer has no rewards. Customer percentiles are defined by adjusting their 18-hole intercept and price coefficients so that their utility from playing an 18-hole round places them in the appropriate percentile of the play frequency distribution under the uniform price regime. In doing so, we account for the correlation between a golfer's intercept and price coefficients. We fix all other coefficients at their mean value.

Figure 3
Switching Costs by Type of Golfer
Observed Program

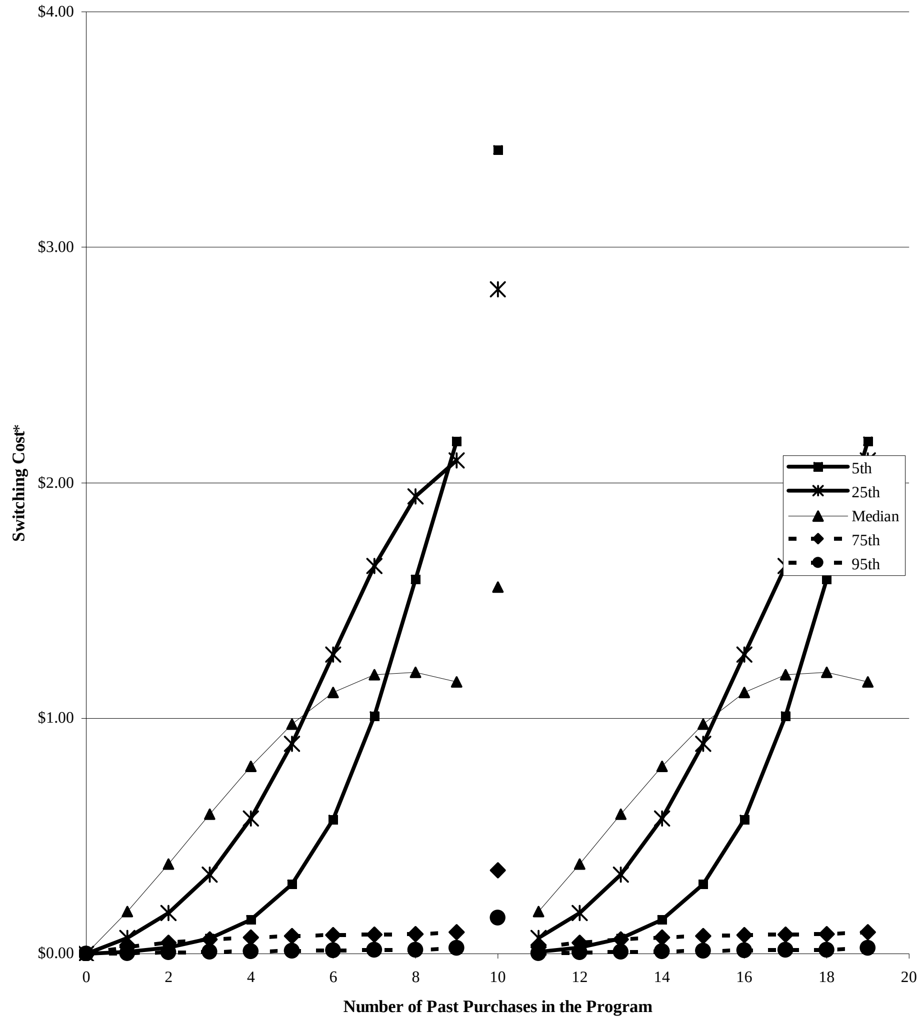

Switching costs in the observed reward program as a function of number of past purchases for five different customer types. Switching costs are measured as defined in Section 6 of the text and calculated assuming it is a Monday, the golfer has played the day before, and is 90 days into the program. At 0 to 9 and 11 to 19 past purchases, the switching costs are measured when the customer does not possess a reward. At 10 past purchases, the customer is assumed to have a reward. Customer percentiles are defined by adjusting their 18-hole intercept and price coefficients so that their utility from playing an 18-hole round places them in the appropriate percentile of the play frequency distribution under the uniform price regime. In doing so, we account for the correlation between a golfer's intercept and price coefficients. We fix all other coefficients at their mean value.

Figure 4
Elasticity of 18-Hole Weekday Round of Golf
by Number of Credits or Uniform Price

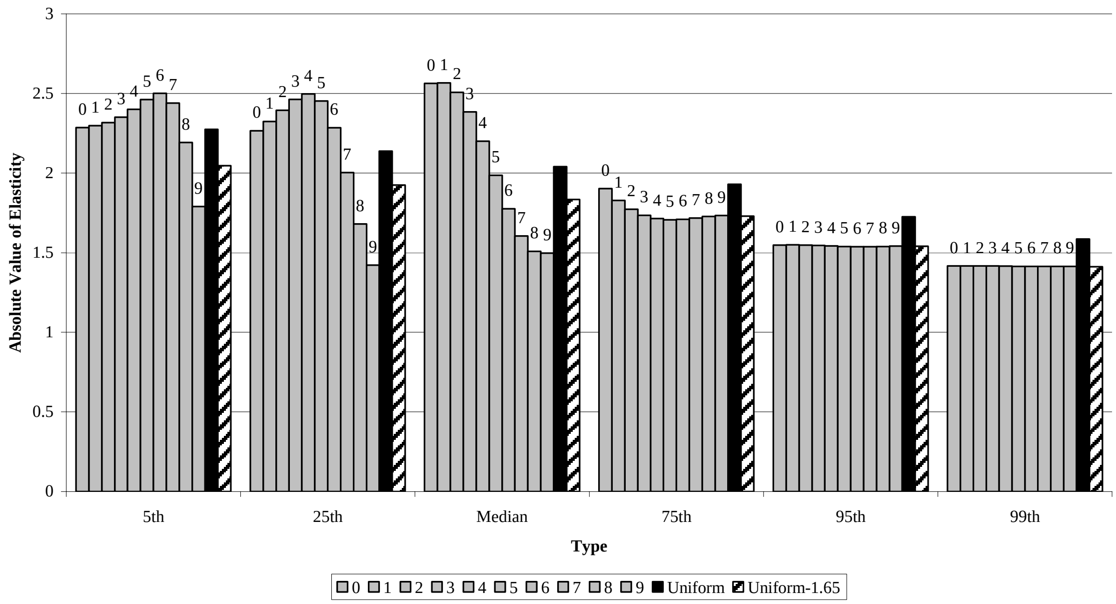

Demand elasticity of 18-hole weekday round of golf for six different customer types. Elasticity is calculated at different credit levels under the reward program, under the firm's posted uniform price, and under the posted uniform price less $1.65 (the per-purchase value of an earned reward). Elasticities are calculated for a Monday, assuming the golfer has played the day before and is 90 days into the program. Customer percentiles are defined by adjusting their 18-hole intercept and price coefficients so that their utility from playing an 18-hole round places them in the appropriate percentile of the play frequency distribution under the uniform price regime. In doing so, we account for the correlation between a golfer's intercept and price coefficients. We fix all other coefficients at their mean value.

Figure A1
Switching Costs by Type of Golfer
Single-RewardOpportunity, 365-Day Program

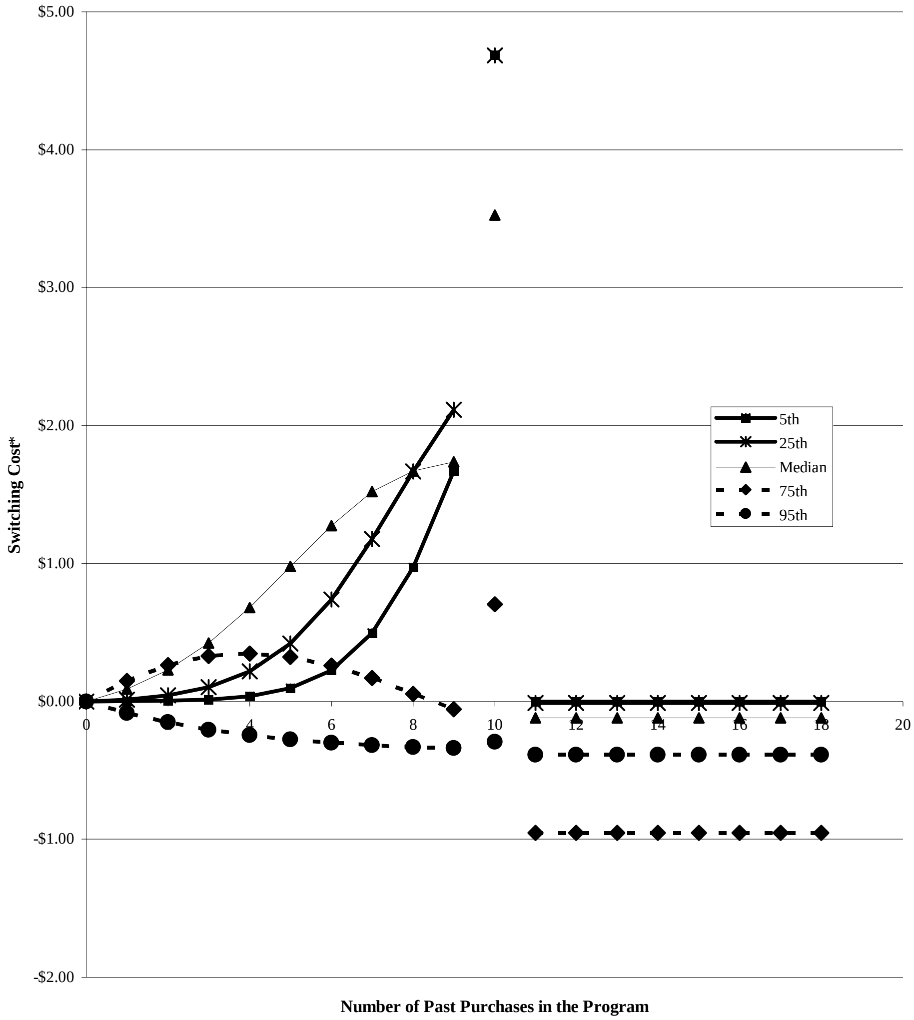

Switching costs in a reward program offering a single reward over a 365-day time horizon as a function of number of credits earned for five different customer types. Switching costs are measured as defined in Section 6 of the text and calculated assuming it is a Monday and the golfer has played the day before, and began the program 90 days before. At 0 to 9 and 11 to 19 past purchases, the switching costs are measured when the customer does not possess a reward. At 10 past purchases, the customer is assumed to have a reward. Customer percentiles are defined by adjusting their 18-hole intercept and price coefficients so that their utility from playing an 18-hole round places them in the appropriate percentile of the play frequency distribution under the uniform price regime. In doing so, we account for the correlation between a golfer's intercept and price coefficients. We fix all other coefficients at their mean value.