Can Non-Tiered Frequency Reward Programs be Profitable?

Zhenling Jiang
Yulia Nevskaya
Raphael Thomadsen
Washington University in St. Louis

February 29, 2016

Abstract

We examine the effectiveness of a customer loyalty program with a non-tiered reward structure. We use a unique data set consisting of all transactions at a chain of hair salons from both before and after the implementation of the loyalty program. We find that the loyalty program is profitable despite the approximate linear structure of the reward that provides a 5% discount to members who regularly return. Importantly, the increased profitability would be hard to pinpoint if a researcher did not have data from both before and after the reward program introduction. The increased profits come from increasing the frequency of visits, and more importantly, from reducing the attrition rate. We find that the *de facto* discount received by members who redeem rewards is only about 2.7% (the average discounts to all members is only 0.7%). These redeeming members increase the frequency of visits by 1.9%. Further, the introduction of the loyalty program leads to an approximately 3% increase in customer lifetime value across all consumers. With some basic assumption, this reflects a 6% increase in the value of the program's members, meaning that the program has a significant net benefit to the firm.

1. Introduction

Customer loyalty reward programs are commonly implemented across a wide range of businesses. In 2013, the average US household belonged to 21.9 loyalty reward programs and actively participated in 9.5 of them (Berry 2013). Nevertheless, there is no consensus in the academic literature on the effectiveness of loyalty programs in influencing customer behavior, especially given the wide variety of reward designs that are used.

The goal of this paper is to measure the effectiveness of a non-tiered reward program.[1] We define a non-tiered reward program as one where the reward per dollar spent does not increase as the consumer makes additional purchases from the firm. The key characteristic of such program is the absence of tiered benefits, i.e. a customer does not receive progressively larger rewards as her cumulative spending with the firm grows over a period of time. In the program we study, which is a reward program from a chain of Hair Salons, consumers gain a $5 off coupon for every $100 they spend. This program is similar to many other non-tiered programs such as "buy 10 get 1 free" programs that are often employed by small retail or service businesses such as frozen yogurt shops, sandwich shops, or golf courses. These non-tiered programs contrast with tiered programs, such as most airlines' frequent flier programs, where consumers earn elite statuses with increasing perks if they fly enough. The benefits of tiered programs accrue through many mechanisms which are absent in non-tiered programs, including lock-in due to increasing benefits, as well as consumer self-signalling (e.g., Drèze and Nunes, 2009).

Most of the literature on non-tiered rewards programs finds that such programs are not effective or have only a small level of effectiveness before accounting for the loss of revenue from rewards programs. Sharp and Sharp (1997) analyze the Fly Buys program, a non-tiered rewards program for a group of retailers. They consider whether the rewards program affects the venues in which consumers shop, as well as whether shoppers adjust their purchase frequency or become more loyal to stores. They find mixed results, and are unable to conclude that loyalty programs have a substantial impact on repeat-purchase behavior. Hartmann and Viard (2008) tests whether a 'buy 10 get 1 free' rewards program creates switching costs for members. They examine a particular manifestation of switching costs by looking at whether customers accelerate their purchases as they move closer to achieving a reward. They find that there is limited acceleration, and that switching costs are not a significant feature of the rewards program because frequent

[1] This type of program is also referred to as a linear rewards program or a frequency rewards program.

customers will earn a reward soon anyway, and low-involvement customers rarely get close to a reward. That said, Lewis (2004) evaluates a frequency rewards program for a retailer and estimates that the presence of the rewards program increases revenues by approximately 2%. Similarly, Leenheer et al. (2007) find that membership in a loyalty program increases a store's share-of-wallet by about 4%. Vana et. al. (2016) study a cash back program for an online vendor, and find that if the company delays the cash back this can lead to higher retention and spending. Bolton et al. (2000) finds that membership in a non-tiered credit card rewards program has a large impact on retention and spending rates, but this paper does not control for a key selection effect: members may be selectively pre-disposed to valuing the card more than non-members. Consistent with previous papers (van Heerde and Bijmolt, 2005; Leenheer et al., 2007), we show that such a selection effect is large, and that members and non-members exhibit very different purchase patterns even in the absence of a loyalty program.

There are also several papers that study a different type of non-tiered rewards program, where customers who spend enough money in a fixed time period receive a non-monetary reward, such as a ham or a turkey from a grocery store. Lal and Bell (2003) find that these programs lead to shifts in spending at the store with the give-away. However, this paper potentially suffers from selection effects, since the main analysis is whether the customers who redeemed for a ham increased their spending more than customers who did not redeem, conditional on the customer's previous-year's spending category. It would be expected that consumers who end up spending more serendipitously would also be more likely to qualify for and receive a ham, so it is hard to assess whether the measured effect fully reflects the impact of the promotion. Taylor and Neslin (2005) also find that a similar program increased sales by approximately 6% during the program's 8-week period, and by 1.8% over the 7 weeks after the promotional period. While both of these papers suggest that non-tiered rewards programs could be profitable, the limited time nature of these programs may affect the size of the measured impact.

One key issue that comes up in trying to assess the effectiveness of the loyalty programs is that of how to obtain an unbiased measure of the program's impact. Many papers consider only the behavior of consumers during a time period after the program has started (Bolton et al., 2000; Verhoef, 2003). Other papers consider only the behavior of consumers who are in the program (Hartmann and Viard, 2008; Kopalle et al., 2012; Stourm et al., 2015), or even a subset of members consisting of the most-involved members (Lewis, 2004; Liu, 2007). In general, one cannot measure the impact of a rewards program by comparing the behavior of members and non-members because the choice to join a program is likely to be correlated with the customer's expected level of engagement with the company. This means that one would expect members to spend more, shop more frequently, and to have higher levels of retention even if the company did not offer a rewards program. Leenheer et al. (2007) address such selection issues through the use of an instrumental variables approach, but such a solution depends on having strong and valid instruments, which are hard to find.

Some papers that analyze the impact of rewards programs based on data from members of the program during a time period where the program exists use models to identify the effect of the program. These papers capture the modeled effect of the rewards program (based on rational economic behavior), but may not capture the entire impact of the rewards program because a consumer's responses to a loyalty program can go beyond those based on pure "economic utility" (Henderson et al., 2011). For example, Hartmann and Viard look at the current reward program participants and use structural modeling to analyze how the timing between golf games differs as consumers get closer to earning a free game. While this approach will capture part of the value of the rewards program – that part of the customer's response that is based on rational behavior – it is possible that the presence of a program can change the behavior of the consumers through other psychological ways that might not lead consumers to shift their behavior as they get closer to a goal. For instance, the presence of the rewards program can make the customer feel an emotional connection to the particular golf course, which leads to them frequenting the course more often. In a similar vein, consumers do not necessarily equate rewards with the same value as the cash back or discount they earn. Nevertheless, in some cases – generally those with tiered loyalty programs – researchers have been able to show that the programs have a positive impact even when only measuring the modeled effects (e.g., Kopalle et al., 2012).

In contrast, we measure the impact of the rewards program by using a census of data from before and after the introduction of a rewards program. We can then see how customer behavior changes before and after the program is introduced. We consider three dimensions of behavior that can change as a result of the program: the amount that customers spend on any particular visit, the frequency with which the customers patronize the firm, and the probability that the customer is retained by the firm. For the first two measures, we are able to control for the selection issues of who joins a loyalty program by employing a fixed effects regression, so the impact of the rewards program is measured by how spending and frequency change for a particular individual because of the presence of the program. Measuring customer attrition is more challenging because consumers generally leave the company only once, precluding the use of a fixed effects regression. Further, there is significant consumer heterogeneity and this heterogeneity is very likely to be correlated with the choice to join a reward program, the tenure that a customer has with the company, and even the speed at which the customer accumulates points in the program. Given the above, we measure the impact of the program by comparing the attrition rate across *all* consumers before and after the introduction of the reward program.

Ultimately, we find that the program has only a negligible impact on spending per visit but increases the frequency of visits among redeeming members by 1.9%. However, the program increases retention, leading to an approximately 3% increase in customer lifetime value across all consumers. Because the cost of the program is small – even among redeeming members, the cost is only 2.7% of revenues – the program has a net benefit to the firm.

The rest of the paper is organized as follows: In the next section, we describe the data and the specifics of the loyalty program. Then we analyze the impact of the loyalty program in three areas separately: spending behavior, frequency of visit and customer attrition. Finally we summarize our findings.

2. Data

We use a rich panel dataset obtained from a small chain of hair salons that cater to men. For each customer visit, the data state the date of the visit, the dollar amount spent, the services and products purchased, any applied discounts, and a unique customer identification number. The data cover 39 months, out of which we use 30 months for analysis.[2] Importantly, our data contains the universe of transactions for the covered time period. There are 807,527 unique transactions from 87,417 unique customers. After the first 15 months in our data, the company introduced a customer loyalty reward program. Therefore, we have observations of individual customers both before and after the start of the loyalty program, which provides an opportunity to observe the changes in individual behavior that can be attributed to the effect of the loyalty program.

[2] The company introduced new services around 4 months into the data, and it introduced a new type of discount unrelated to the loyalty program in the last 5 months of our data. To cleanly measure the effect of the loyalty program, we omit from analysis the first 4 and the last 5 months of the data.

The reward component consists of distributing a $5 off coupon to members via email for every $100 spent on services or hair care products. The coupon is valid for 90 days from the date of issue and cannot be combined with other discounts. Further, customers must bring the coupon to the store to use it, although showing the coupon email on a cell phone is generally accepted. Either because the customers did not bring or remember the coupon, or for other unknown reasons, some coupons go unredeemed even among members who have earned rewards and should be redeeming them.

The loyalty program is free to join. To become a member the customer needs to provide their full name (which is requested at the time of checking with the receptionist and when making an appointment) and email, and agree to receive electronic marketing messages, which are generally sent less than once a month. Members who join the program may also get a reminder to come in for a service if they do not return to the hair salon for an extended period of time. We show in Appendix D that these reminders do not have much of an effect on customer behavior.

Approximately 40% of customers in our dataset are members of the loyalty program. There are two types of members, which we term as "automatic" and "non-automatic." Some customers voluntarily provided their email addresses in order to receive marketing messages before the start of the loyalty program. The company automatically enrolled these 15,115 customers into the program on the date that it was launched, hence we call these customers "automatic" members. Note that we do not observe any purchases after the start of the loyalty program for some automatic members, reflecting that they churned before the program's start. Automatic members make up approximately half of all loyalty-program members in our dataset. The other half of the members joined the loyalty reward program by providing an email address during one of their visits after the program was launched; we call these members "non-automatic" members.

Out of 30,004 members, only 3,476 redeemed a reward during the time of our analysis. We term these members "redeeming members." Note that only these customers entail a cost to the company.

Table 1 shows the composition of the company's customer base. Members make up 40% of customers, but 52% of transactions. Similarly, automatic members and redeeming members engage in a disproportionate number of transactions. Table 2 provides summary statistics of the variables that are the focus of our study. We again observe different average behavior by each of the groups.

Table 1: Customer base composition

	Number of customers	Share of customer base	Average number of transactions	Share of all transactions
All Customers	75,553	100%	8.24	100%
Members	30,004	40%	10.87	52%
Non-Members	45,549	60%	6.51	48%
Automatic Members	15,115	20%	11.72	28%
Redeeming Members	3,476	5%	23.83	13%

Table 2: Customer behavior

	Average spending per visit ($)	Average frequency (Days between visits)	Average attrition rate (%)
All Customers	21.0	40.1	15.9
Members	21.6	38.0	10.7
Non-Members	20.4	42.7	21.6
Automatic Members	21.8	38.5	11.8
Redeeming Members	22.3	32.3	1.8

One complicating factor in measuring the impact of the loyalty program using our data is that the price for many of the core services increased by $1 one month after the launch of the loyalty program, which reflects an approximately 5% increase. In the analyses below, we use a variety of tactics to account for the impact of this price increase. Because one of the variables we consider is the change in spending, we calculate spending using the post-increase prices. Thus, a customer who buys the same haircut each time they go to the salon would be viewed as spending the same amount per visit before and after the price change. Our findings are the loyalty program increases the frequency that members visit the store and increases the retention rates at the stores. In general, we would expect the price increases to reduce the frequency of visits and the retention rate, if there is any effect. Thus, our findings should be viewed as a lower bound on the size of the impact of the rewards program. We note that the size of the price increase is approximately the same as the maximum discount from the loyalty program. This aids our interpretation of the finding that the loyalty program increases frequency and retention. In particular, it demonstrates that our results must come from the effect of having a loyalty program, and not just from the consumers' purely rational response to a net reduced price from the program's discounts.

3. Analysis of Loyalty Program Effectiveness

We measure the effectiveness of the loyalty program along the following dimensions: spending per visit, frequency of visits, and customer attrition. We find that the rewards program has a minimal impact on spending per visit, while the rewards discount decrease revenues per visit for redeeming members by approximately 2.7%. The frequency of visits for these redeeming members increases by approximately 1.9%. We find that the loyalty program reduces customer attrition, which results in a 3% higher customer lifetime value. In the next three sections, we describe in detail the methodology and results of the analysis.

3.1. Spending Per Visit

We begin by quantifying the impact of the loyalty program on the amount of money a customer spends during a visit. Figure 1 shows the monthly averages of per-visit spending for members and non-members, before and after the start of the loyalty program (month 12).[3] For this figure, we define a member as a customer who eventually joins the program, no matter when he does so. For members, we show the numbers both with and without accounting for the reward discount. The dotted lines represent the average spending for the period before and after the introduction of the program. We observe lower spending per visit on months 6, 18 and 30 because of a promotion the company runs every year on the same month.

Figure 1 demonstrates that, on average, members spend more per visit than non-members, both before and after the introduction of the loyalty program. Therefore, we observe strong customer self-selection into the rewards program. This demonstrates that comparing the behavior between members and non-members would overstate the impact of the rewards behavior.[4]

The average spending for both members and non-members increases slightly after the program is implemented, although the changes are small. For members, the average spending per visit before the loyalty program is $21.73. After the start of the program, it becomes $21.94 and translates to $21.84 when accounting for the reward discounts used. The spending per visit for non-members are $20.46 and $20.69 before and after the program, respectively.

---

[3] As mentioned above, the prices for most of the basic services went up in month 13, one month after the introduction of the program. The calculations for Figure 1 and the analyses to follow are done using new prices in order to eliminate the effect of price inflation.

[4] Potentially offsetting this overstating effect is the fact that some customers who never gave their email but churned before the loyalty program might have joined the loyalty program once it was offered (and thus would have been counted as members if the timing of their patronage was different).

Figure 1: Monthly averages for spending per visit before and after the start of the loyalty program

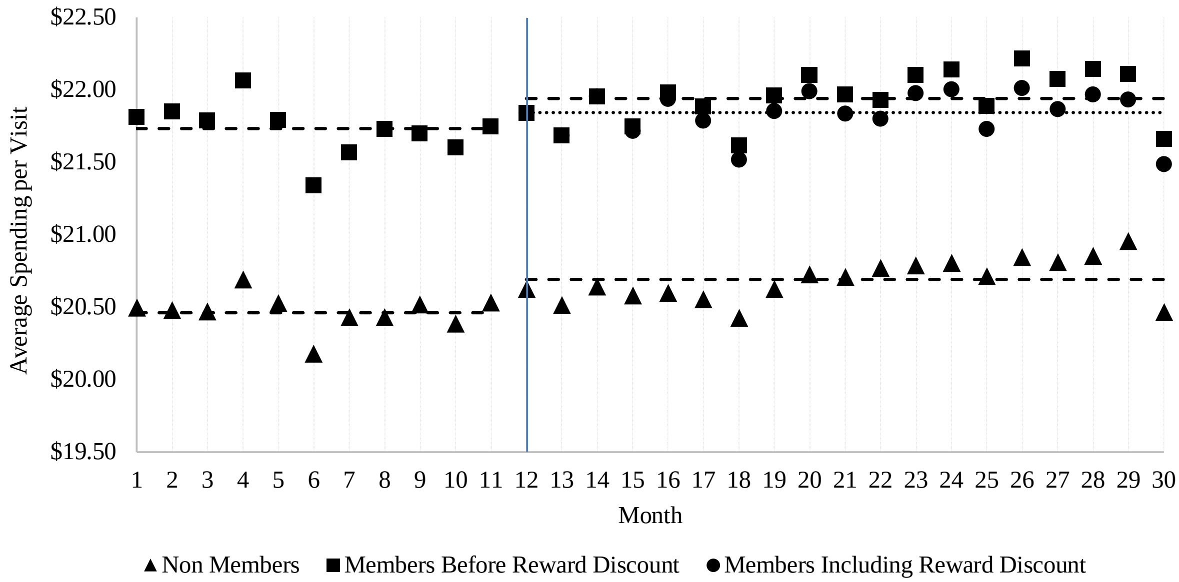

To correct for selection effects, we specify a fixed-effects model of per-visit spending:

$$\mathrm{S}_{ik} = \beta_{ap}^{m} \cdot AP_{ik} + \beta_M \cdot M_k + \alpha_i + \epsilon_{ik}, \tag{1}$$

where S_{ik} is customer i's ($i = 1, .., N$) dollar amount spent on visit k ($k = 1, .., K_i$), AP_{ik} is an indicator that takes on the value of 1 for visits after customer i becomes a member (for non-members, *AP* = 1 indicates visits after the start of the program), M_k is a vector of month dummies (with one of the dummies omitted) to control for seasonality, α_i is an individual fixed effect, and ϵ_{ik} is an error term.

For the regression in equation (1) we pool individuals who ultimately became members and those who remained non-members, but allow for different coefficients on the "after program" indicator variable AP_{ik} for these two groups of customers, β_{ap}^{m}, where index m indicates the ultimate membership status. Note that the definitions of the AP_{ik} variable, and therefore the interpretation of its coefficients, are different for members and non-members. To create an AP_{ik} indicator for members, we use the actual point of time that a member joined the program, but we use the date that the rewards program started to create that indicator for the non-members. As a result, for the member group the coefficient $\beta_{ap}^{m=1}$ measures the average effect of the loyalty program, while for non-members the coefficient $\beta_{ap}^{m=0}$ measures the effect of unobserved (to the researcher) factors occurring at the time of the program. However, there are only 6122 customers that join the rewards program at a date that does not coincide with the first visit where the rewards program exists (out of 34,215 members), so the effect of this approach to variable construction is not large. The estimates from equation (1) appear in Table 3.

Table 3: Estimation results – spending per visit

	Non-Members	Members (w/o discount)	Members (w/ discount)	Redeeming Members (w/o discount)	Redeeming Members (w/ discount)
	(1)	(2)	(3)	(4)	(5)
β_{ap}^{m}: AP (After Program)	0.262***	0.334***	0.163***	0.378***	-0.224***
	(0.025)	(0.029)	(0.029)	(0.064)	(0.064)
Month Dummies Included	Yes	Yes	Yes	Yes	Yes
Mean Fixed Effect	20.18	21.17	21.23	22.22	22.24
Note:					*p<0.1; **p<0.05; ***p<0.01

We find that members (before accounting for the program's rewards, Column 2) increase their spending slightly more than non-members (column 1) after loyalty program was implemented, with this effect being statistically significant but economically small. The smaller increase may be a result of the industry, where generally the only way to spend more per visit is to upgrade to a premium service such as a scalp massage or neck shave or to purchase hair care products. We also observe that the spending after the loyalty program for members once we account for the reward discounts (column 3) are about 17¢ lower than the spending without the discounts.[5] Thus, the average discount among members is approximately 0.8%. This is much lower than the 5% theoretical cost of the program because only a subset of all members earn and redeem a discount.

It might be reasonable to hypothesize that the impact of the loyalty program is larger for members who are active enough to redeem rewards. Further, these members are also the only members for which the firm incurs a marginal cost of servicing the program. Therefore, we also run the regression separately for those individuals who redeem at least one reward discount during the period of our data. Column 4 shows that the spending increase by redeeming members cannot be statistically distinguished from either the non-member or the member spending (columns 1 and 2), although its magnitude is similar to the magnitude of the increase by members as a whole. The impact of the discount is larger for redeeming members (column 5), however. For these customers,

[5] While the monthly dummies are the same for columns 1 and 2, they are estimated separately for columns 3, 4 and 5.

their spending after the introduction of the rewards program decreases by 60¢, or approximately 2.7%, which is again lower than the 5% potential maximum discount.

We now turn to checking the robustness of our findings. One can argue that if the timing of when to join the loyalty program is at the customer's discretion then even putting in an individual-level fixed effect does not fully account for the endogeneity in the decision to join the program. Therefore, as a robustness check, we run the analysis for "automatic" members, i.e. customers who gave their email address to the company sometime before the start of the program and were automatically enrolled right after it was launched. This group of customers is valuable for our analysis since the timing of their joining the program is exogenous, which allows us to claim the causal effect of the loyalty program. The regression results with automatic members are presented in Appendix A. The estimated impact of the program shows a similar pattern as the estimates in Table 3, although the already small impacts of the loyalty program on spending are even smaller (and not statistically significant). Thus, we conclude that the loyalty program has only a small effect on spending. This may not be very surprising in the industry we choose because while there are some upgrades available, most customers just consider getting the basic haircut and do not value much the optional upgrades.

3.1.1. Impact of Rewards on Upgraded Service and Hair Product Purchases

While the magnitude of the impact on overall spending per visit is not statistically significant, the point estimates do suggest a small increase. One reason that the impact could be small is that it is sometimes hypothesized that the benefits of the program accrue only once the customer earns a reward (Drèze and Nunes 2011, Wang et. al. 2016). This delayed benefit may come either because the reward may make the benefit of the program more salient, or because the reward might be viewed as a treat, which should then be spent on something nice. We examine the role the reward plays by examining the choices of redeeming members to opt for upgraded services or buy products at different stages of the reward cycle. More specifically, we compare a redeeming member's behavior in these four stages:

(1) Before Program: visits that happen before joining the loyalty program;

(2) Before Reward: visits **after** loyalty program but before using the **first** reward discount;

(3) Reward: visits where the member redeems a reward discount;

(4) After Reward: all visits after redeeming the first reward, except for visits where the member redeems a reward.

We run the following fixed effects regression separately for service upgrades and product purchases:

$$g_{ik} = \beta R_{ik} + \alpha_i + \epsilon_{ik}, \quad (2)$$

where g_{ik} is an indicator variable for customer i's service upgrade or product purchase during visit k, R_{ik} is a vector of dummies indicating the stages of the reward cycle, α_i is an individual fixed effect, and ϵ_{ik} is an error term. We chose to use a linear probability model for ease of direct interpretation, where the coefficients can be interpreted as approximating the increase in choice probability relative to one before joining the program, which is the base category. The estimates are shown in Table 4. We find that the probability of a service upgrade and a product purchase both increase after joining the program, which is consistent with the after program increase for all groups in Table 3.

Table 4: Regression results - upgraded service and product purchase by redeeming members

	Upgraded Service	Product Purchase
Before Reward	0.030^{***}	0.009^{***}
	(0.004)	(0.002)
Reward	0.057^{***}	0.035^{***}
	(0.005)	(0.004)
After Reward	0.046^{***}	0.010^{***}
	(0.005)	(0.002)
Mean Fixed Effects	0.17	0.054

Note: *p<0.1; **p<0.05; ***p<0.01

What is especially interesting is that we see a statistically significant increase in both types of purchases during visits where a reward is redeemed. Further, it appears that customers may learn that they enjoy the upgraded service – a significant number of the consumers continue choosing the upgraded service after the reward redemption. There is no such bump in the probability for purchasing products, perhaps reflecting the possibility that a customer does not learn about the value of a product so saliently or perhaps because a customer who has bought product might not need to buy product every visit.

3.2. Frequency of Visits

In this section, we assess the impact of the loyalty program on customers' frequency of visits, as measured by the number of days that pass between each pair of consecutive visits conditional on customers not churning. Our approach is to compare the average number of days between visits before and after the start of the loyalty program.

We put an upper limit of 91 days on the number of days between visits, and consider anyone who did not return within 91 days to have "churned." The last 91 days of the data are used only to determine the number of days since the previous visit. In choosing the cap of 91 days we follow the company's definition of a churned customer.[6] For robustness, we have also run a version of analysis with the cap of 182 days, which gives similar results (see Appendix B).

We need to use the above cap because our data is right-censored in terms of time between visits. The nature of the potential bias is as follows: Suppose that a customer that has a long duration between visits. For visits that occur early in our data, we observe the customer returning and include the number of days between visits in our calculations. Now, consider a customer with the same duration between visits but who shows up late in our data. We would not observe this customer's return, so the long duration between visitors would not be counted. Of course, the time period before the loyalty program occurs early in our data, while the time period later in the data reflects the post-program period. If we did not cap the number of days used to measure the duration between visits, we would measure a longer inter-purchase period before the loyalty program than after the program, even if the program had no impact on frequency. Once we impose the limit on the number of days between visits (and use the last period of the data to measure returning customers) we no longer have this issue.

In Figure 2, we plot the average number of days between visits for members and non-members for each month. As in the spending analysis section, members are defined as anyone who eventually joins the rewards program. We again observe strong self-selection among customers: customers who visit more frequently before the start of the program are more likely to become members. The average number of days between visits decreases slightly for members after the start of the loyalty program, shrinking from 38.1 days to 38.0 days. Just the opposite occurs for non-members, with the average number of days between visits increasing from 42.6 to 42.9 days. We see a larger effect for redeeming members: the reduction is from the average of 32.7 to 32.2 days.

[6] 91 days is a multiple of 7 days, reflecting that some customers cut their hair on a particular day of the week.

However, as before, the customer composition effect may mask the impact of the loyalty program in this initial analysis. In particular, customers with moderate patronage levels might not provide the business with an email address before the rewards program and therefore be observed as non-members if they churn before the program, but these same types of customers might join the program if they patronize the firm after the program has started. We control for this effect in our econometric analysis.

Figure 2: Monthly averages for number of days between visits

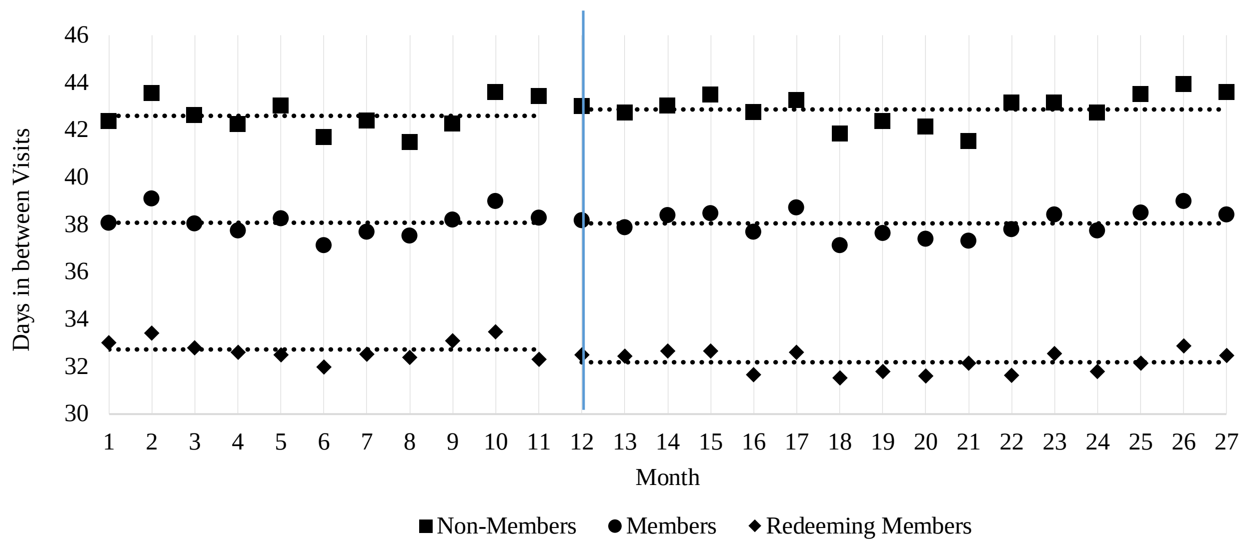

We run a fixed-effects regression to pin down the impact of the loyalty program on the frequency of visits by customers who chose to become members. We specify a regression as follows:

$$days_{ik} = \beta_{ap}^{m} * AP_{ik} + \beta_m \cdot M_k + \alpha_i + \epsilon_{it}, \tag{3}$$

where $days_{ik}$ is a number of days between the k^{th} and $(k+1)^{th}$ visits for member i, AP_{ik} is an indicator variable that equals 1 for members after they join the loyalty program (for customers who never joined the loyalty program, $AP_{ik} = 1$ for visits after the start of the loyalty program), M_k is a vector of month dummies to control for seasonality, and α_i are individual fixed effects, which captures individual heterogeneity in average visit frequency.

We run a pooled regression for all members and non-members. As was the case with the spending regression, we interact the "after-program" variable with a membership indicator, so members and non-members have separate coefficients, but the seasonality coefficients are common across the two groups. We also run a separate regression for redeeming members.

Table 5 reports the results of the regression analysis of visit frequency. Consistent with Figure 2, we observe a small increase in the number of days between visits after the program for non-members. This may reflect a general trend that customers tend to come back somewhat less frequently over time. Among members, we observe a small but statistically insignificant decrease in the number of days between visits for members after the program. As with spending, we might expect that the rewards program might have most of its effects on redeeming members. Columns 3 and 4 confirm that is the case. Redeeming members return 0.725 days earlier after joining the loyalty program, which represents an increase in the frequency of visits of 1.9%.[7]

Table 5: Estimates for frequency of visits regression

	Non-Members	All Members	Redeeming Members	Non-Redeeming Members
	(1)	(2)	(3)	(4)
β_{ap}^{m}: AP (After Program)	0.160^{**}	-0.041	-0.725^{***}	0.283^{***}
	(0.076)	(0.064)	(0.098)	(0.083)
Month Dummies Included	Yes	Yes	Yes	Yes
Mean Fixed Effect	49.04	45.17	37.56	46.42
Note:			$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$	

The company increased the price for most of its core services by $1 one month after the start of the program. The price increase might be expected to increase the time between visits, and may explain the reason why non-members have an increase in the number of days between visits after the program is implemented. The effect of price on frequency of visits might be present for redeeming members, too. If such an effect occurs, it means that the measured impact of the loyalty program on redeeming members in Table 5 represents a lower bound of the loyalty program effect on the frequency of visits.

Regression results for "automatic" members are presented in Appendix B. The estimates for this group has a similar magnitude and statistical significance as those presented in Table 5.

[7] In Appendix B we show that the increase in the frequency of visits for redeeming members is 2.9% if we use a 182-day cap on the number of days between visits.

We next consider how reward redemption affects the frequency of visits. The results are presented in Table 6. We see that the frequency of visits is higher for redeeming members in all stages after the customer joins the program compared to before the program, but that the effect is largest for the visit where customers redeem a reward. For this visit, the time between visits is over 1 day shorter than it is before the loyalty program (an increase in frequency of 3.4%). Even before redeeming their first reward discount, members return 0.59 days earlier (1.7% increase in frequency) than before the program. After redeeming the first rewards coupon, there is a long-term effect of a 0.89-day reduction (2.5% increase in frequency) in the time between visits, possibly reflecting that the benefit of the loyalty program is more salient after the first reward is earned.

Table 6: Frequency of visit regression – reward effect

	Days between Visits
Before Reward	-0.594 ***
	(0.114)
Reward	-1.236 ***
	(0.182)
After Reward	-0.893 ***
	(0.141)
Mean FE	35.959
Note:	*p<0.1; **p<0.05; ***p<0.01

3.2.1. Does Loyalty Program Lead to Visit Acceleration or Reduced Visit Skipping?

Customers reduce the average time between visits after joining the loyalty program. Is this change due to customers coming back more quickly for their next visit or reducing the number of skipped visits, when they presumably get a haircut elsewhere. We show that both effects occur in economically meaningful ways.

We begin by examining acceleration. To do this, we run the regression described in Eq. 3 on redeeming members but using only observations where the number of days between visits is less than 1.5 times the individual's fixed effect. The assumption we are making is that if a customer skipped a visit by getting a haircut elsewhere, they should return to the store at a time that is approximately twice the length of their usual time between visits, which is measured by the estimated individual-level fixed effects. Thus, we assume that the visits that are less than 1.5 times the fixed effects generally contain visits where the customer did not skip a haircut. If we still see a decrease of the average time between visits, it is attributed to a likely effect of increased frequency. The results are reported in Table 7. The results indicate that the loyalty program reduces the time between visits for this restricted set by approximately 0.365 days. This reduction of 0.365 days can then be attributed to the lower bound on how much of the effect is from acceleration (vs. skipping), where this is a lower bound because the effect is also mechanically somewhat smaller due to the fact that we limit the analysis to observations smaller time intervals between visits.

Table 7. Frequency of visits – acceleration by redeeming members

Estimate	Program Effect	S.E.
Original Sample	-0.725***	0.098
Limited Sample ($days_{ik} < 1.5\alpha_i$)	-0.365***	0.078

We now turn to assessing whether this effect could also come from reduced visit skipping. We infer the incidence of customers skipping a visit, and measure the effect of the program by comparing the incidence of skipping visits before and after program. We compute the incidence of skipping visits among program members before and after joining the program using the multiplier of 1.8 in the above definition of a skipped visit. In other words, if a customer comes back after 1.8 times their usual cycle, then they are likely to have skipped a visit. The result is presented in Table 8. For robustness, we also do the same calculation where we use 1.5 times the individual's fixed effect for the cutoff.

We observe that the incidence of visit skipping does decrease after program, and the amount of decreased skipping is substantial in terms of reducing the remaining skipping that occurs. Thus, we conclude that the effect of the program both accelerates purchases and reduces instances of skipping.

Table 8. Skipping visits

	Incidence before program	Incidence after program
Days > 1.8 *α_i	1.58%	1.03%
Days > 1.5*α_i	3.66%	2.61%

3.3. Customer Attrition

We now consider how the loyalty program impacts customer attrition. Retaining existing customers is crucial to companies like the one in this paper. New customer acquisition costs are high and the company relies a lot on repeat customers, which make up approximately 84% of the salon's traffic.

The impact of the loyalty program on customer retention is an important component of the loyalty program's effectiveness. However, many previous studies that assess the effectiveness of loyalty programs ignore retention, largely due to data limitations. If a researcher only has data on consumers while they are enrolled in the program, it is not possible to assess how the customer's propensity to leave the program depends on the program's presence; It is possible to examine how a customer's choice to leave depends on how close the customer is to reaching a reward threshold, but this only provides a partial picture of the impact. Similarly, if one has data on members and non-members, one cannot just compare the attrition rates between these groups because of self-selection in the decision to join the loyalty program. For example, we observe that the attrition rate is 11% for members and 22% for non-members. However, we measure that the impact of the rewards program on attrition is much smaller than one would infer from such a "naïve" comparison. Another strategy one might consider would be to examine customers who provided email addresses before the start of the program, and compare the attrition rates of these customers before vs. after they were automatically enrolled in the loyalty program. However, it is well-known that retention rates increase with consumer tenure (and we observe this in our data as well). Therefore such a comparison would overstate the extent that the program increases retention since customers who remain in the "after program" period must also have a longer tenure at the salon.

Our strategy is to identify the loyalty program's impact on retention from the change in the attrition rate across all customers – not just members – after the program introduction. We look at the change in attrition across all customers in order to avoid self-selection issues, such as the ones described above. This strategy leverages the fact that we have a census of data before and after the loyalty program introduction. Behind this analysis, we make two key assumptions: 1) the most significant change in the business during the observation period is the introduction of the loyalty program, and 2) the distribution of the tenure consumers have with the firm is at a steady-state.

The first assumption should be valid with one significant exception: the price for most of the core services went up by $1 one month after the start of the loyalty program. We would expect the price increase to cause higher attrition, if there is any effect at all, so what we measure as the impact of the loyalty program on attrition should be a lower bound on the program's true impact. We also examine the change in the attrition rate during the month right after the program starts but before the price increase occurs and obtain a similar effect.

To verify the second assumption, Figure 3 plots the percentage of new customers in each month.[8] We observe that the percentage of new customers remains relatively steady over time, with the exception of holiday periods (which might reflect an influx of out-of-town visitors). The company has been in business for a substantial time before the start of the observation window, so we have every reason to believe that the business is in a relative steady state.

Figure 3. New customer traffic by month, % of customer base

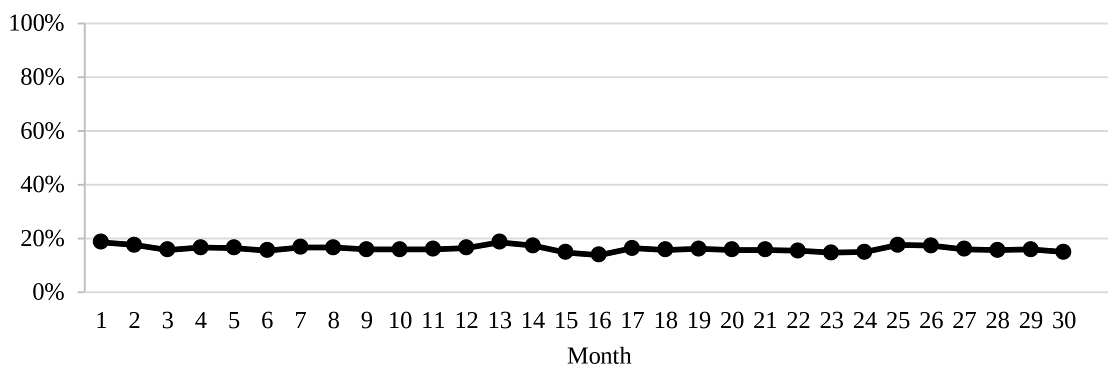

Figure 4 shows the attrition rate in each month, where the dotted lines show the mean attrition rates before and after the program. The loyalty program was introduced at the beginning of month 12. We observe that the average attrition rate after the loyalty program introduction is lower than the average attrition rate before the program was implemented. Note that the prices were increased in month 13, and that the attrition rates in months 13 and month 14 are much higher than the average attrition rate for the rest of the post-program period. This suggests that some consumers leave the salon after they find out about the higher prices, suggesting that the estimates we calculate for how much the rewards program reduced attrition will underestimate the true effect.

[8] For internal consistency, the new customer is defined as someone who did not visit the salon chain in the previous 91 days. If a customer returns after 91 days, then upon returning, they are considered to be a new customer.

Figure 4. Overall Attrition Rate over Time

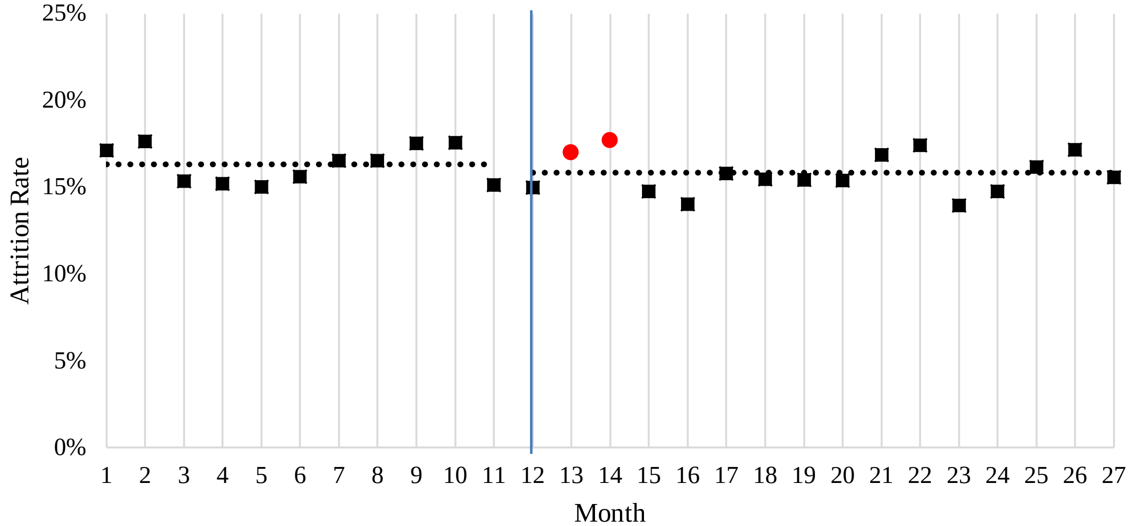

We run a logit regression to estimate how the loyalty program influences a customer's probability of churning. As noted above, we define a churned customer as one who does not return within 91 days after their previous visit.[9] We include an indicator variable for all visits after the introduction of the loyalty program as well as a vector of month dummies to account for the seasonality effect.

$$Prob(churn = 1) = \frac{\exp(\beta_0 + \beta_{ap} AP + \beta_m M)}{1 + \exp(\beta_0 + \beta_{ap} AP + \beta_m M)} \tag{4}$$

The results are presented in Table 9. The parameter estimates, which have the expected signs and are statistically significant, indicate that the loyalty program leads to lower attrition rate.

Table 9: Attrition Regression Estimates

	Attrition
Program Effect, β_{ap}	-0.037***
	(0.008)
Constant, β_0	-1.693***
	(0.013)
Month Dummies Included	Yes
Observations	558,743
Note:	*p<0.1 **p<0.05 ***p<0.01

[9] Results using a 182-day cap appear in the appendix, and yield similar results.

We immediately observe that the impact of the program on the overall attrition rate is negative (meaning that more customers stay with the firm) and statistically significant. In Table 10, we convert these estimates into average attrition rates before and after the program's introduction. The overall customer attrition rate goes down from 16.2% to 15.7%, reflecting a 3% (s.e., 0.6%) reduction in the level of attrition.

Table 10. Program effect on the probability of attrition

	Before Program	After Program
Attrition Rate	16.2%	15.7%
(s.e.)	(0.08%)	(0.06%)

We use the customer lifetime value as a metric to link the reduction in attrition to the profitability of the company (Kumar and Shah, 2004). We calculate the impact of this lower attrition rate by assuming that consumers visit once every 40 days, and use an annual discount rate of 5%, 10% and 15%, which may span the reasonable costs of capital for business ventures, and calculate the customer lifetime multiplies, which is $m = \frac{1+i}{1+i-r}$, where *r* is the retention rate and *i* is a 40-day discount rate (Fader and Hardie, 2012). The results are presented in Table 11. We observe that the program increases the average customer lifetime value by 3%.[10] Note that this effect is measured across all customers, whereas the only customers for whom the company bears a cost of servicing in the program. If one multiplies the 2.7% cost of providing the discount to the redeeming members by the 13% of the transactions that are done by redeeming members, the average cost of the program become 0.35% overall, or 0.7% among members. Thus, we see that the program is largely profitable from the increase in customer lifetime value alone. Alternatively, one might believe that the program only increases retention among members. In such a case, the impact of the loyalty program on customer lifetime value of members would be at least 6%.[11]

[10] This effect would be even larger if one assumed annual discount rates below 10%.

[11] The 6% figure is based on using the population's average attrition rate as a baseline, and calculating the change in the multiplier from that rate. However, members have lower-than-average attrition rates. If we subtract (1/0.52) times the coefficient in Table 9 from the attrition rates for members, we obtain an even higher number. Thus, the 6% is a conservative estimate.

Table 11: Change in customer lifetime value

Annual Discount Rate	5%	10%	15%
Change in CLV Multiplier	3.1% (0.6%)	3.0% (0.6%)	2.9% (0.6%)

4. Conclusion

We use a unique dataset with all customer purchases both before and after the introduction of a non-tiered loyalty program to show that such a program can be effective. We measure how the program changed the spending per visit, the frequency of visits, and the retention rates. We find that because of imperfect redemption, the cost of the program is 2.7% of revenues among redeeming members, or 0.7% across all members. We find that there is a small and not statistically significant impact of the loyalty program on spending (before applying the loyalty discount). We find that customers that redeem rewards increase the frequency of their visits by 1.9%. Finally, we show that among all customers, the customer lifetime value increases by 3% after the implementation of the program. If we were to assume that this gain came only from the program's members, we would then conclude that the program increases the customer lifetime value of these members by over 6%.

Looking closely at the behavior of redeeming customers at different stages of the redemption cycle provides us with further insights of how the rewards program affects customer behavior. More specifically, we observe that customers are more likely to buy a service upgrade or make a product purchase during a visit where they redeem a reward. For the case of service upgrades, this behavior appears to continue even after the reward is redeemed, perhaps reflecting customer learning about the value of these enhanced services. Frequency also increases after the customer joins the rewards program, but we again see that the number of days between visits is smallest by quite a lot on the visit when the customer redeems the reward.

References

Berry, Jeff (2013). The 2013 COLLOQUY Loyalty Census. COLLOQUY Industry Report. http://www.colloquy.com/files/2013-COLLOQUY-Census-Talk-White-Paper.pdf

Bolton, Ruth N., P. K. Kannan and Matthew D. Bramlett (2000), "Implications of Loyalty Program Membership and Service Experiences for Customer Retention and Value," *Journal of the Academy of Marketing Science*, 28(1), 95-108.

Breugelmans, Els., Tammo H. A. Bijmolt, Jie Zhang, Leonardo J. Basso, Matilda Dorotic, Praveen Kopalle, Alec Minnema, Willem Jan Mijnlieff and Nancy V. Wünderlich (2014), "Advancing research on loyalty programs: a future research agenda," *Marketing Letters*, 26(2), 127-139.

Drèze, Xavier and Stephen J Hoch (1998), "Exploiting the Installed Base Using Cross-merchandising and Category Destination Programs," *International Journal of Research in Marketing*, 15(5), 459-471.

Drèze, Xavier and Joseph C Nunes (2009), "Feeling Superior: The Impact of Loyalty Program Structure on Consumers' Perceptions of Status," *Journal of Consumer Research*, 35(6), 890-905.

Drèze, Xavier and Joseph C. Nunes (2011), "Recurring Goals and Learning: The Impact of Successful Reward Attainment on Purchase Behavior," *Journal of Marketing Research*, 48(2), 268-281.

Fader, Peter S. and Bruce G. S. Hardie (2012), "Reconsiling and Clarifying CLV Formulas," Note. http://brucehardie.com/notes/024/

Hartmann, Wesley R. and V. Brian Viard (2008), "Do frequency reward programs create switching costs? A dynamic structural analysis of demand in a reward program," *Quantitative Marketing and Economics*, 6(2), 109-137.

Henderson, Conor M., Joshua T. Beck and Robert W. Palmatier (2011), "Review of the theoretical underpinnings of loyalty programs," *Journal of Consumer Psychology*, 21(3), 256–276.

Kopalle, Praveen K., Yacheng Sun, Scott A. Neslin, Baohong Sun, Vanitha Swaminathan (2012), "The Joint Sales Impact of Frequency Reward and Customer Tier Components of Loyalty Programs," *Marketing Science*, 31(2), 216-235.

Kumar, V. and Denish Shah (2004), "Building and sustaining profitable customer loyalty for the 21st century," *Journal of Retailing*, 80(4), 317-329.

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

Leenheer, Jorna, Harald J. van Heerde, Tammo H.A. Bijmoltc and Ale Smidtsd (2007), "Do loyalty programs really enhance behavioral loyalty? An empirical analysis accounting for self-selecting members," *International Journal of Research in Marketing*, 24(1), 31-47.

Lewis, Michael (2004), “The Influence of Loyalty Programs and Short-Term Promotions on Customer Retention,” *Journal of Marketing Research*, 41(3), 281-292.

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

Meyer-Waarden, Lars and Christophe Benavent (2006), “The Impact of Loyalty Programmes on Repeat Purchase Behavior,” *Journal of Marketing Management*, 21(1-2), 61-88.

Meyer-Waarden, Lars and Christophe Benavent (2009), “Grocery Retail Loyalty Program Effects: Self-selection or Purchase Behavior Change?,” *Journal of the Academy of Marketing Science*, 37(3), 345-358.

Sharp, Byron and Anne Sharp (1997), “Loyalty programs and their impact on repeat-purchase loyalty patterns,” *International Journal of Research in Marketing*, 14(5), 473-486.

Stourm, Valeria, Eric T. Bradlow, and Peter S. Fader (2015), “Stockpiling Points in Linear Loyalty Programs,” *Journal of Marketing Research*, 52(2), 253-267.

Taylor, Gail Ayala and Scott A. Neslin (2005), “The Current and Future Sales Impact of a Retail Frequency Reward Program,” *Journal of Retailing*, 81(4), 293-305.

Vana, Prasad, Anja Lambrecht and Marco Bertini (2016), “Cashback is Cash Forward: Delaying a Discount to Increase Future Spending,” *working paper*.

Verhoef, Peter C. (2003), “Understanding the Effect of Customer Relationship Management Efforts on Customer Retention and Customer Share Development,” *Journal of Marketing*, 67(4), 30-45.

van Heerde, Harald J. and Tammo H.A. Bijmolt (2005), “Decomposing the Promotional Revenue Bump for Loyalty Program Revenue Members Versus Nonmembers,” *Journal of Marketing Research*, 42(4), 443-457.

Wang, Yanwen, Michael Lewis, Cynthia Cryder and Jim Sprigg. Enduring Effects of Goal Achievement and Failure Within Customer Loyalty Programs: A Large-Scale Field Experiment. *Marketing Science*, forthcoming.

Appendix A – Spending Regressions for Automatic Members

Table A1: Estimation results – spending per visit for "automatic" members

	Non-Members	Members (w/o discount)	Members (w/ discount)	Redeeming Members (w/o discount)	Redeeming Members (w/ discount)
	(1)	(2)	(3)	(4)	(5)
AP (After Program)	0.262***	0.292***	0.164***	0.341***	-0.176*
	(0.025)	(0.036)	(0.037)	(0.090)	(0.092)
Month Dummies Included	Yes	Yes	Yes	Yes	Yes
Mean FE	20.18	21.36	21.38	22.46	22.49
Note:				*p<0.1; **p<0.05;***p<0.01	

Appendix B – Frequency Regression Robustness Checks

Table B1: The Estimates from Table 5 Using 182 Days Between Visits

Dep. Variable: Days between Visits				
	Non-Members	All Members	Redeeming Members	Non-Redeeming Members
	(1)	(2)	(3)	(4)
AP (After Program)	0.282**	-0.082	-1.164***	0.389**
	(0.125)	(0.105)	(0.137)	(0.138)
Month Dummy Included	Yes	Yes	Yes	Yes
Mean FE	63.01	54.42	40.37	56.68

Note: *p<0.1; **p<0.05; ***p<0.01

Table B2: Frequency of visit regression – automatic members using 91-day cap

	Non-Members	All Automatic Members	Redeeming Automatic Members	Non-Redeeming Automatic Members
	(1)	(2)	(3)	(4)
AP (After Program)	0.160**	-0.047	-0.875***	0.269***
	(0.076)	(0.082)	(0.131)	(0.102)
Month Dummies Included	Yes	Yes	Yes	Yes
Mean FE	49.04	45.56	37.87	46.69

Appendix C – Customer Attrition Robustness Checks

Table C1. Attrition regression with 182-day cap - estimates

	Estimate (S.E.)
Program Effect, β_1	-0.041*** (0.010)
Constant, β_0	-2.278*** (0.018)
Month Dummies Included	Yes
Observations	495,423

Note: *p<0.1 **p<0.05 ***p<0.01

Table C2. Program effect on the probability of attrition with 182-day cap

	Before Program	After Program (All Visits)
Attrition Rate (S.E.)	10.0% (0.06%)	9.7% (0.06%)

Table C3. Change in customer lifetime value with 182-day cap

Annual Discount Rate	5%	10%	15%
CLV Multiplier – 1	2.9% (0.8%)	2.8% (0.8%)	2.6% (0.7%)

Appendix D – Effect of Reminders on Customer Attrition

There are two main components of the loyalty program. Besides the rewards component, the company also sends out service reminders by email when members do not return for their next service after a certain number of days. Customer retention can be affected by both components of the program. To demonstrate that the email communication component of the program we study had minimal if any effect on customer retention, we take advantage of the variation in the reminder email policy. This variation comes from the experiments with the reminder policies that the company ran. Table D1 describes the timing and policies used in the experiments.

Table D1. Reminder email policies

		←Start of Loyalty Program		End of Data Observation→
Reminder Emails		Period 1	Period 2	Period 3
Short Term	Days after Visit	45 days	35 days	45 days
	Email Target	***New*** customers (1st visit)		***Repeat*** customers (4+ visits)
Medium Term	Days after Visit	None	70 days	65 days
	Email Target		***New*** customers (1st visit)	***Repeat*** customers (4+ visits)
Long Term	Days after Visit	90 days		
	Email Target	***New*** customers (1st visit) & ***Repeat*** customers (4+ visits)		***Repeat*** customers (4+ visits)
Defector	Days after Visit			225 days
	Email Target	None		***Repeat*** customers (4+ visits)

In periods 1 and 2, new customers get reminder emails if they do not return within certain number of days, while in period 3 new customers do not get reminder emails. The opposite is true for repeat customers with four or more visits: they get reminder emails in period 3 but not in periods 1 and 2.

If reminder emails contribute to reducing attrition rate, we would expect to see new customers having a higher attrition rate in period 3 than in period 1 and 2, and repeat customers having a higher one in period 1 and 2 than in period 3. We start the analysis one month after loyalty program was introduced, so that customers in different reminder conditions face the same prices. The attrition rates for different experimental conditions are presented in Tables D2 and D3. Table D2 shows that the new customer attrition rate is the same with and without the reminder emails. Table D3 shows similar results for repeat customers – attrition rate for this group does not change when the customers receive reminder emails. Therefore, we conclude that reminder email component of the loyalty program does not drive the reduction in attrition rate for new customers or repeat customers. As such, we attribute the decrease in the attrition rate to the effect of the rewards of the loyalty program.

Table D2. Attrition rate in the two experimental conditions – new customers

Period 1 & 2 (Reminder Emails)	Period 3 (No Reminder Emails)
54.8% (0.23%)	54.8% (0.34%)

Table D3. Attrition rate in the two experimental conditions – repeat customers with 4+ visits

Period 1 & 2 (No Reminder Emails)	Period 3 (Reminder Emails)
5.0% (0.05%)	5.0% (0.07%)