Evaluating Loyalty Programs with Endogenous Redemption

Mihaela Alina Nastasoiu, *The University of Western Ontario*

Supervisor: Mark Vandenbosch, *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Business

Abstract

Evaluating the capacity of consumer loyalty programs to generate additional sales is essential for marketers who run such programs. However, customers' self selection into the loyalty programs makes this evaluation difficult. This is the case especially in set-ups where the reward is not granted automatically upon achieving a certain number of points. In the case of automatic rewards, based on temporal discounting and the diminishing perceived distance to the goal, marketing theory predicts that points accumulation accelerates as consumers approach the threshold of necessary points for the reward, and is also boosted after the redemption, in what is called 'the rewarded behavior effect'. These mechanisms generate sales that are attributed solely to the reward program.

In this thesis I use these insights to develop two models for evaluating loyalty programs where the rewards are not granted automatically, but consumers accumulate points and are free to decide the moment and the size of the redemption. These types of programs are expanding in sectors such as retail and finance, so being able to evaluate them becomes increasingly important. The first model applies to programs where consumers use the accumulated points like cash, for groceries and other day-to-day expenses, while the second applies to programs where consumers use the points for non-ordinary expenses or treats, which on average are much larger. Both are models of consumer behavior, which retailers need to understand in order to learn how to improve the profitability of their loyalty programs. The difference between these two set ups is important because consumers plan for redemptions in different ways in these two contexts. I estimate the parameters of both models using data provided by AIR MILES, Canada's largest coalition loyalty program. I show how sample heterogeneity and the non-random timing of the reward cash-in can be confounded with true loyalty program effects and I tease apart these effects to obtain non-biased estimates of program profitability. I use the model insights to suggest ways in which the loyalty program owner can change the program to further boost its profitability and discuss how the recommendations are contingent on retailers' contribution margins. The dissertation advances the literature on loyalty program evaluation by developing structural models for set ups where retailers do not impose automatic redemptions upon consumers.

Keywords: consumer loyalty programs, endogenous redemption, structural models, points

pressure, rewarded behavior effect

Contents

List of Figures

List of Tables

List of Appendices

Chapter 1

Introduction

Cultivating customers' loyalty is key for any business. Besides the product's intrinsic qualities and promotion and advertising campaigns, one tool that companies use increasingly to build customers' loyalty is to offer reward programs (RPs). These programs offer consumers loyalty points (LPs) or miles proportional to their purchases (e.g. 1 mile per $10 spent) and allow them to use the amassed LPs to redeem different types of rewards. As RPs impose costs on companies, one of the critical questions that managers need to answer is how effective the RP is in generating sales that would have not been realized in the absence of the RP.

Evaluating the effectiveness of RPs is difficult, given consumers' self selection into the programs. Two main strategies have been followed: experimental or quasi-experimental approaches and structural approaches. The first approach aims to compare RP members to non-members (after accounting for self-selection) and draw conclusions about the overall effectiveness of the programs based on this comparison. Structural approaches are more analytical, in the sense that they focus on identifying the theoretical mechanisms that make RPs effective and measuring the impact of those specific mechanisms. Points pressure (the tendency to accelerate LPs collection as one approaches the redemption moment) and the rewarded behavior effect (the tendency to only gradually decrease collection after a redemption) are the prime mechanisms that have been identified to enhance the effectiveness of the RPs. Chapter 2 presents in detail both the experimental and the structural approaches and reviews their strengths and challenges. Also, this chapter presents the empirical evidence on RP's effectiveness.

Looking back at past research in this area, it is apparent that most work is done on RPs where the reward is pre-set by the merchant and granted automatically to participants - like the automatic status upgrades that frequent fliers receive upon accumulating a certain number of flown miles. However, loyalty reward programs are rapidly expanding in areas such as retail and financial services, where the main attraction is represented by the prizes that consumers can redeem with their LPs, not the status upgrade. In these programs consumers typically accumulate points which they can later use either to reduce their bills or redeem rewards from a given set offered by the company that runs the program. The time and the size of the redemption are usually entirely at consumers' discretion.

Since the existing models are not well equipped to deal with this endogenous redemption decision, I address this deficiency. In this thesis I propose two models that focus jointly on how consumers collect and spend their LPs. The first model is applicable to RPs where prizes are discounts offered to consumers on their bill (this will be called the 'cash' setup), while the second is applicable to RPs that offer more substantial prizes, like flights, cruises or merchandise (the 'dream' setup). The difference in the prize structure is important because it dictates how consumers think about their LPs and organize their redemptions. I argue that in the 'cash' setup consumers plan to accumulate and spend LPs in a continuous way; they derive utility from the rewards they receive for their points and they permanently struggle between giving in to the temptation of redeeming something now and hoarding points in order to be able to make large redemptions which seem to bring more utility per mile redeemed than smaller redemptions. No single redemption is a goal in itself, but part of the collection-redemption cycle that consumers pursue in order to maximize their utility. Chapter 3 presents a model of the collection-redemption cycle that is applicable in this context. The parameters of the model are estimated using a sample of AIR MILES cash collectors (participants who chose to use their AIR MILES only to obtain discounts at different retailers). Based on these parameters, policy experiments are conducted to determine the sales generated exclusively due the RP and ways in which the RP can be modified in order to further boost sales. The results show that different program set ups are optimal depending on retailers' contribution margins.

In contrast to RPs that offer cash discounts as prizes, in the 'dream' setup any individual redemption stands out because it is larger in value and it usually represents a treat that is not purchased habitually. At any point in time consumers may or may not have a specific target (or redemption goal), but once they form such goals (e.g., 'I'm using my points to redeem gifts for Christmas' or 'to book flights for the family's summer holiday'), they focus entirely on that goal and disregard goals that they might form into the future, after achieving (or failing to achieve) the current goal. This is not to say that that consumers in the 'dream' setup are shortsighted, but that they care about the future only up to a certain point - i.e., they bracket more narrowly (Read et al., 1999). This is because each redemption is a different event in itself, while in the case of cash discounts, it is easier for consumers to think about the redemptions collectively, as they are more homogeneous. Chapter 4 presents a model where consumers are assumed to form specific goal redemptions and work towards attaining their goals. The parameters of this model are estimated using a sample of AIR MILES collectors who opted to accumulate points that can be used to purchase 'dream rewards' such as flights, cruises or merchandise. Similarly to the 'cash' model, the estimates are used to conduct policy experiments that tease out the share of points accumulated due solely to the program vs. the share of points corresponding to sales that would have been attained even in the absence of the RP. Also, the estimated parameters are used to gain insights into how the efficiency of the program would improve if AIR MILES managed to persuade the program participants to shorten their inter-redemption spells.

This thesis advances research on loyalty reward programs in several ways. First, there are only two structural models that make the redemption moment endogenous: Kopalle et al. (2007, 2012). These models are applicable to RPs in industries such as hospitality and travel where consumers face the trade-off between redeeming rewards and amassing points in order to gain status. However, this trade-off is not present in industries such as retail or the financial sector; in these areas sacrificing a redemption in order to build balance for a status upgrade is not a relevant consideration. So in this thesis I expand the set of models with endogenous redemption to RPs in the retail and financial sectors, which have come to jointly account for 60% of RP memberships in the US (Berry, 2013). Furthermore, focusing on the timing of redemptions is important because it is redemptions that actually makes RPs successful, by setting in motion mechanisms such as points pressure and the rewarded behavior effect. Unredeemed LPs can apparently make a loyalty program seem more profitable, but redemptions are the engine that power reward programs. For most customer loyalty programs, LP issuance is a liability; earnings are recognized only when a redemption is made and a portion of that liability is converted to cash (Konsewicz, 2007). So in the short-run, it is the redemption, not the issuance, that improves RPs books and allows them to reduce their liabilities and enhance their financial situation. Redemption is also important for the long run success, especially in the case of frequent-flyer programs that sell their LPs or miles to credit card companies (The Wall Street Journal, 2010) and coalition programs administered by a third-party such as AIR MILES, Nectar or FlyBuys. To work, a reward loyalty program should actually reward customers.

Secondly, this is the first attempt to use a structural approach in evaluating an *on-going* retail RP (as opposed to an occasional retail program such the Thanksgiving turkey program studied by Taylor and Neslin (2005)), where consumers accumulate LPs and are free to decide when and how many of their points to use for redeeming rewards. Given consumers' self-selection into RPs, I believe that this is a more sound method of evaluation, compared to quasi-experimental approaches. (Why structural approaches are more sound will become clearer later in this dissertation.) Moreover, this method can be applied to coalition RPs, where consumers can acquire points at dozens of sponsors in very different sectors, so finding adequate controls for a quasi-experimental investigation would be even more challenging in respect of a coalition program. With a structural approach I am able to evaluate how strong the points pressure and rewarded behavior effects are and back out the effectiveness of the RP from these estimates.

Thirdly, I use the results of the structural models to emphasize how the sample heterogeneity and the fact that program participants select the timing of their redemptions can be confounded with program effects. In other words, these two mechanisms give rise to patterns of collection that are similar to the patterns of collection that one would expect if the RP changes participants' behavior. Therefore, in order to accurately determine the effect of the RP in driving sales, it is essential that these concurrent effects are isolated and teased apart. In Chapter 3 I show that failing to account for heterogeneity and the selection of the redemption timing produces upward biased estimates of the efficiency of the program. Somewhat surprisingly, it appears that heterogeniety plays a less significant role in the case of 'dream' collectors (Chapter 4), as the model applied on the entire, heterogeneous sample produces similar results to a specification where I apply the model separately to 3 a priori identified segments and then I aggregate the results.

Fourthly, I build a rich structural model that captures consumers' psychological processes

(Chintagunta et al., 2006), by modeling fallible cognitive process such as limited foresight, impulse purchases and fixed transaction costs for making a redemption. Theoretical models that capture consumer thought processes are valuable because they bolster our confidence in the results they produce and the recommendations that are based on those results. However, these more complex models usually suffer from not being amenable to a closed from solution for estimation. I overcome this problem by relying on simulation techniques. Both models are estimated through indirect inference (Gourieroux et al., 1993; Smith, 2008), a class of GMM estimators (Hansen, 1982), where the parameters of the structural model are pinned down by matching selected moments observed in the data to the corresponding moments of a simulated data set. My models are therefore an excellent managerial tool for evaluating retail RPs and testing different ideas on how these RPs can be improved. The model can be used in deciding, for example whether the amount of rewards given out for a given amount of points accumulated should be increased or decreased.

The next chapter presents the existing work on loyalty programs, differentiating the wider concept of customer loyalty from the 'engineered' loyalty that RPs aim to create. In this literature review, I focus both on the theoretical mechanisms through which RPs can increase sales as well as on the empirical findings so far, emphasizing the different approaches of experimental, quasi-experimental and structural approaches. Chapters 3 and 4 present two distinct models of LPs acquisition and redemption. The main difference between the two models is the set ups in which they apply. Chapter 3 is suitable to analyze RPs where the rewards are small, in-store discounts and, conditional on having enough LPs for a reward, consumers decide in every period whether they cash in points. In contrast, Chapter 4 is suitable to analyze RPs that offer more significant rewards - like flights or hotel stays. In these programs, consumers are assumed to form stable redemption goals and work towards attaining their goals. Both chapters start by exposing the theoretical models, presenting the data used to estimate them and then presenting the estimated parameters. I assess the implications for the break-even margin (the margin above which running a RP contributes positively to retailers' profits) and the share of sales attributable to the program. Furthermore, the estimated parameters are used to conduct 'what-if' analysis (also called counterfactuals) - that is evaluate how profitable the programs would be if retailers implemented certain changes (for example if they increased the amount of rewards per mile redeemed). Chapter 5 re-iterates the differences between the models presented in Chapters 3 and 4 and elaborates on the reasons for which the ‘dream’ rewards program appears to be more profitable than the ‘cash’ rewards program. In this chapter I also re-state the contribution to the field of loyalty programs that I make in this dissertation and I outline fruitful areas for further research.

Chapter 2

Loyalty reward programs

Loyal customers are an asset for any company or brand. As intuition would suggest that, all else equal, more loyal customers are less likely to be swayed by competition's attractive offers, more likely to spread positive word-of-mouth, more forgiving, and also more likely to pay higher prices for the product/ brand towards which they are loyal. It is no surprise then that loyalty occupies such a central place both in academic marketing research, as will become apparent in the literature review that I am conducting, as well as on the agenda of many practitioners (O'Brien and Jones, 1995; Reichheld, 1996).

A literature review conducted by Jacoby and Chestnut (1978) covering more than 50 years of research (1923-1976) reveals that early attempts to operationalize the concept of brand loyalty focused almost exclusively on behavioral measures such as share of purchases or sequence of purchases. However, in the 60's, researchers started to suggest that besides the behavioral component, brand loyalty should also include attitudinal measures that tap into consumers' commitment to the brand. Day (1969) and Dick and Basu (1994) clearly argue that, besides the behavioral aspect of repeated purchase, the customer loyalty construct should also comprise an attitudinal element; thus true loyalty exists wherever there is *both* repeat patronage and highly positive attitude for the focal brand relative to other brands. Chaudhuri and Holbrook (2001) also explicitly differentiate between purchase loyalty and attitudinal loyalty, both types of loyalty stemming from brand affect and brand trust. More recently Liu-Thompkins and Tam (2013) differentiate between attitudinal loyalty and habit.

However, adding the attitudinal measure to the concept of brand loyalty doesn't conclude the discussion around it. Positive attitudes towards brands can stem from their intrinsic qualities or from the marketing programs that surround them. Researchers in this area have been careful to emphasize the concept's property of being different from brand qualities relative to competition or contextual factors, such as lack of alternatives (Day, 1969; Aaker, 1991; Dick and Basu, 1994; Oliver, 1999). For example, Oliver (1999) defines loyalty as 'a deeply held commitment to rebuy or repatronize a preferred product/ service in the future (...) *despite* situational influences and marketing efforts having the potential to cause switching behavior'.

Lemon et al. (2001) provide what I think is the most clean and rigorous definition of loyalty. For them programs meant to develop customer loyalty contribute to customer relationship equity which they define as 'the tendency of the customer to stick with the brand, *above and beyond customers' objective* and subjective *assessment of the brand*' (emphasis added). Thus, in this definition, even emotional ties that customers have with a certain product or brand, or the attitudes they have towards it (the subjective assessment) are separated from brand loyalty, which is thus conceptualized as a 'stickiness' of the relationship that cannot be explained either by the value proposition of the brand (the more objective attributes) or by consumers' emotional ties to the brand. This contrasts with Day (1969)'s tentative conclusion that 'loyalty is based on a rational decision made after the evaluation of the benefits of competing brands' which is later reinforced by the observation that 'commitment is never total, the decision is reviewed when competitive or other circumstances change'.

If the early literature in this area, as outlined so far, was more concerned to clarify conceptually what customer loyalty is, recently the literature seems to have agreed on a definition of loyalty that is closely aligned to the one provided by Lemon et al. (2001). This refinement of the concept has evolved in step with the proliferation of loyalty programs offered by companies for their customers. Through these programs companies try to enhance and prolong the relationship with their consumers. Customers acquire points (sometimes called miles after the successful programs launched in the American airline industry) based on the amounts they spend with the company. As they collect enough points, consumers can then use them as a sort of currency to redeem different rewards. Also, consumers who collect enough points within given time frames are assigned to superior tier statuses (for example gold or platinum) that offer prestige, different perks and/ or improved service. Conceptually, loyalty has been restricted to refer purely to the stickiness of the relationship, which gives a one to one relationship between this concept and loyalty programs.

The questions that both academics and managers are now interested in, is whether, and how, these sort of programs bring value to the brand. In other words the interest lies in untangling between, on the one hand, the equity that is generated through the brand's objective qualities or through the emotional ties it fosters with consumers, and on the other hand, by the loyalty program. The shift can be easily noticed in the managerial recommendations that the academic loyalty literature provides. Early papers emphasized the role of loyalty antecedents such as trust, relational benefits and value (Sirdeshmukh et al., 2002) trust and brand affect (Chaudhuri and Holbrook, 2001), cognitive antecedents (accessibility, confidence, etc.), affective antecedents (emotion, satisfaction etc.) and conative antecedents (switching cost, sunk cost and expectations) (Dick and Basu, 1994) or quality and perceived associations (Aaker, 1991). These authors follow a logic where these antecedents lead (somewhat mechanically) to loyalty which in turn leads to outcomes of interest (market share, willingness to pay, word of mouth, etc.). The focus on these antecedents of loyalty seems to be in contradiction with the definition provided, which emphasized its property of being different from brand intrinsic subjective or objective qualities as assessed by consumers. Loyalty seems to be rather the outcome of good subjective and objective value provided by the company and not something that can be fostered exogenously. In contrast, more recent articles focus on evaluating the return of programs that that are strictly meant to increase loyalty, independently of programs meant to increase the objective value provided by the brand or the emotional response it generates (Rust et al., 2004).

In this research, I adopt a view similar to that of Lemon et al. (2001); Rust et al. (2004) and refer to loyalty as a mechanism that leads to re-purchase or re-patronizing above and beyond product quality or favorable emotional ties to the product. Specifically, I am interested in loyalty as a cause of brand or firm superior performance, not as its effect (Reichheld, 1996). I do not contest the fact that product quality or its capacity to generate positive attitudes are important drivers of commitment towards, and re-purchase of the product, and that these drivers can create what in lay terms would be called 'loyalty'. But in this research I am interested in loyalty as a driver (tool) that is separated from these other drivers, as loyalty programs represent a different managerial lever. Therefore, I study loyalty programs that reward consumers based on the volume of the previous purchase and have the capacity to create sunk and switching costs for consumers. This sort of programs may create repeated purchases, or loyalty, the mechanism being completely orthogonal to product quality and advertising strategy.

In contrast to Kumar and Shah (2004), who are interested in developing programs that 'proactively reward customers "today" for their "future" shopping', I am specifically interested in programs that try to change current behavior by offering future rewards. In order to follow Kumar and Shah (2004)'s suggestion companies would need to know who the customers with high future shopping are, so that it can reward them. But if those customers are highly likely to patronize the company in the future, why reward them? By rewarding them there is a risk of creating a ceiling effect where their purchases cannot be further increased instead of offering incentives to those customers who under the current state of affairs are not likely to buy that much into the future, but who can be easily persuaded to change their purchase patterns by offering them a reward? My emphasis on RPs as a *managerial tool* strongly favors the second approach.

The history of modern loyalty (or frequency reward) programs can be traced back to the trade stamps (Davis, 1959), that have been used by retailers as a means of increasing customers' loyalty since the beginning of the 20$^{\text{th}}$ century. In the 80's, this business model was adopted by airlines in their frequent flyer programs and financial service and credit card companies; with the advent of scanners in retailing and services, LPs have become common in hotel chains, car rental companies, supermarkets, music and chemist stores (Financial Times, 2006). A recent study (Nielsen, 2013), which surveyed 29,000 Internet respondents in 58 countries, revealed that almost 60% said that loyalty programs were available where they shopped; 84% of these reported that they are more likely to visit those retailers. The wide spread of such loyalty programs warrants a special focus on them.

The next subsection presents the business model for loyalty programs, distinguishing between LPs that re run by a third party and those that are run 'in-house' by a single retailer. The next two subsections present and systematize the research on loyalty programs: section 2.2 presents different theoretical mechanisms through which loyalty programs could generate extra-sales and section 2.3 reviews the empirical evidence on this topic.

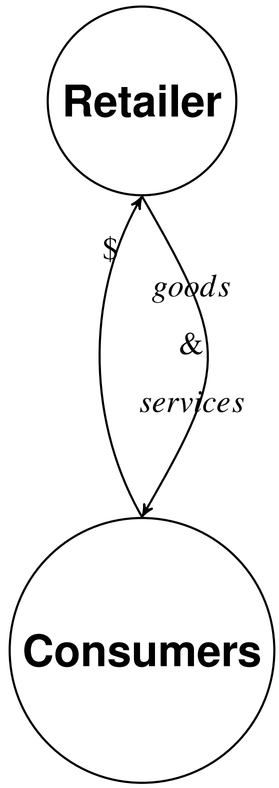

Figure 2.1: The business model for a retailer that does not offer a loyalty program

2.1 The business model

In this subsection I explain how loyalty programs work. I build the exposition from the simple case of a single vendor with no LP to the most complex one of an LP run by a third party. A merchant who does not use a loyalty program makes their profit from the difference between the sales they generate and their costs - the cost of goods sold and any other administrative and selling costs they may incur. This case is presented in Figure 2.1.

In Figure 2.2, I look at the set-up where the retailer runs an 'in-house' LP. In this set-up, consumers purchase goods and services from retailers and pay for them with money, exactly as in the previous case. But now, the retailer awards consumers with loyalty program points, usually proportional with the amount of money they spend. Some customers sign up for the LP program, some don't. For those who don't sign up for the program the business model is the same as the one presented in Figure 2.1. In Figure 2.2, they are the non-LP consumers in the left side of the figure.

The consumers who sign up for the program (the LP consumers on the right side in Figure 2.2) now receive points for the goods and services they purchase. For example, a Shoppers Drug Mart customer who spends $42 in store, receives 420 points (according to the rule that each dollar spent by the consumer is rewarded with 10 points). In offering these rewards, the retailer hopes that the quantity of goods and services purchased will increase compared to the scenario without LP (Figure 2.1). This uncertain additional spending is captured in Figure 2.2 by the dotted line labeled *incremental $*. I used the dotted line precisely to emphasize that when offering an LP, retailers cannot know for sure that the LP will generate incremental sales, or that the margin on the incremental sales is large enough to cover the cost of the LP program. Once consumers have accumulated enough points, they redeem these points for rewards (the gray full arrows in Figure 2.2. A redemption means that consumers' balance of points is decreased and they receive a reward in exchange. The reward can be an in-store discount or a gift that the consumer chooses. The retailer's *profit from the loyalty program* is the margin from the *incremental sales* due to the loyalty program minus the cost of running the LP. In turn, the cost of running the LP is given by the sum between the cost of the rewards and the administrative cost of running the loyalty program. The total profit is the profit from the loyalty program plus the profit as calculated for the retailer in the case of no loyalty program, as described in the previous paragraph.

Figure 2.3 shows the more complex case where the loyalty program is administered by a third party. This third party issues the loyalty points, and sells them to the retailers. Like in the previous set-up, the retailers award the miles to the consumers who participate in the loyalty program, proportional with the amount of money that the consumers spend to acquire goods and services from the retailers. As in the case of the in-house LP, retailers hope that the points they award will generate incremental sales (captured in Figure 2.3 bu the dotted gray line).

Once the consumers have accumulated enough points, they go the party that issued the loyalty points (the Reward Program, as labeled in the figure) and trade their points for rewards. Note that the rewards can be of broad types. The Reward Program can procure them either from the same retailers to which it sells the points, either from other vendors, but this choice is not consequential for the models that I am developing, so this this part of the chain is not shown either in Figure 2.2 or in Figure 2.3. The Reward Program's profit is given by the money

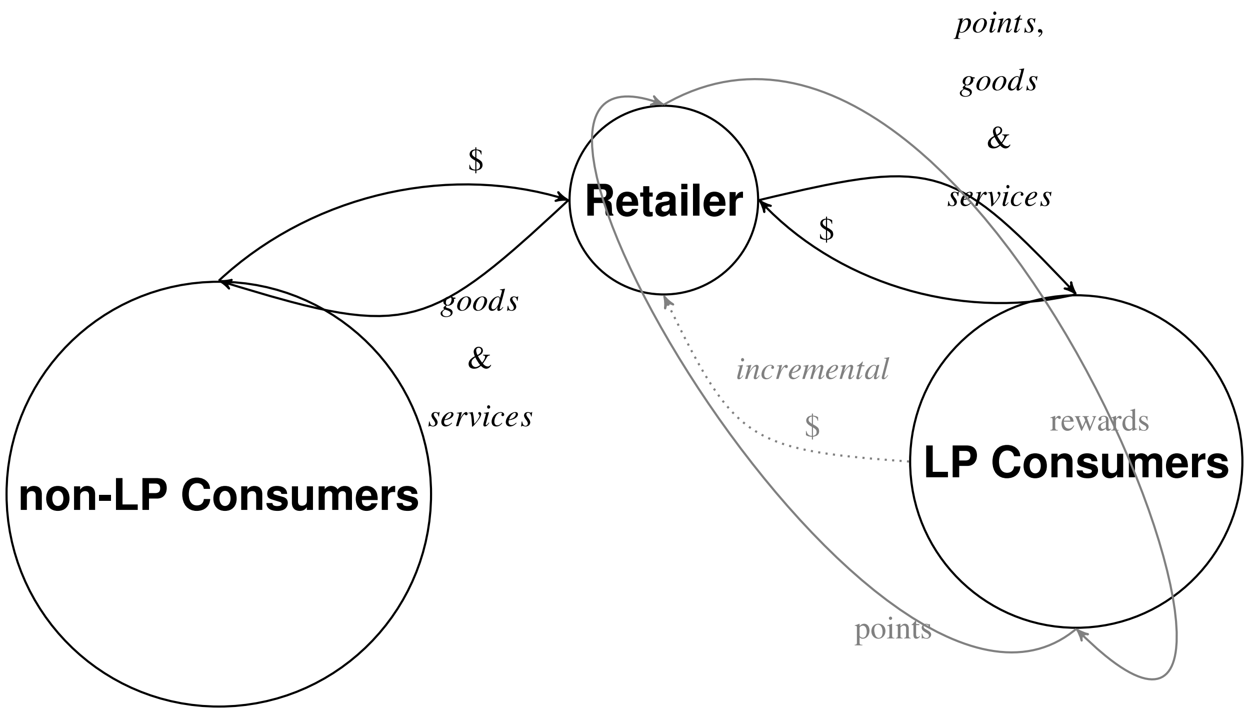

Figure 2.2: The business model for loyalty programs run ‘in-house’ by a single retailer

it receives from the retailers for the points it sells minus the sum of the cost of the rewards and administration costs. For example if the Reward Program retailers 12 cents for each loyalty point and the average cost of rewards is 9 cents, the contribution per point for the Reward Program is 3 cents.

Beyond this hypothetical example, it is difficult to know how profitable third party administered programs are, mainly because these programs are run by larger companies, that also run marketing analytics and retail solutions divisions and who report their consolidated earnings and profits. For example, in Canada, AIR MILES is the largest third-party run loyalty program. But the program is run by LoyaltyOne, which according to Hoover's [1] reported earnings of $44 million. Other third-party run programs are Aeroplan (in Canada) and Nectar (in the UK and Italy). Both these programs are run by Aimia which reported revenues of almost $600 million[2]. These numbers, while not specific for the LP divisions, still suggest that there is a thriving market for third-party run LPs.

In the case where the program is run by a third party the retailers no longer incur a direct cost of the reward program (the administrative expenses and the cost of the rewards), but they simply pay for the loyalty points they hand out to their customers. Retailers' *profit from the loyalty program* is given by the difference between the margin from the *incremental sales due to the loyalty program* and the cost of the points. For example if a certain retailer gives out on average 10,000 loyalty points per day, the retailer's LP cost is $1,200 per day. If the LP generates incremental sales of say 5,000 per day and the margin on those sales is 30%, then the daily profit from the LP for the retailer $300 (30%*$5,000-$1,200). Retailers' total profit is the profit from the loyalty program plus the profit as calculated for the retailer in the case of no loyalty program, as described in the case of Figure 2.1.

So far I have used the term *profit from the loyalty program*. However, it is not clear whether, or to what extent loyalty programs are profitable. This is actually the question that I am trying to answer in this dissertation. While the costs of the LP are easily observable and measurable, the *incremental sales* cannot be directly measured. If the margin from the incremental sales is

[1] *http://www.hoovers.com/company-information/company-search.html?term=loyaltyone* accessed July 28th 2016

[2] *https://www.aimia.com/content/dam/aimiawebsite/financial_reports/2016/First%20Quarter/AIMIA_Q1%202016_Highligh* accessed July 28th 2016

less than the cost of the program, then there will be a *loss from the loyalty program*.

The two models that I am proposing in this dissertation can be applied to both set-ups described in Figures 2.2 and 2.3. Since I am interested in how the rewards influence the consumers' purchase patterns, whether the program is administered in house by the retailer, or run by a third party, is, in itself, irrelevant for how the model plays out.

2.2 In theory, can loyalty reward programs change behavior?

A loyalty program is considered successful when it generates sales that otherwise would have not been realized and when the cost of generating these sales doesn't exceed the additional profit that is garnered. This relatively simple definition of success however does not engender an equally simple operationalization. The main difficulty, especially with secondary data provided by loyalty programs, is, as in all social sciences that try to gauge causal effects of different programs or polices (Imbens and Wooldridge, 2009), assessing the counterfactual, i.e., assessing how much consumers would have purchased from the focal firm in the absence of the loyalty program. Given that consumers cannot be randomly placed in conditions that either enroll or don't enroll them in a loyalty program, self-selection is always a concern in empirical studies that try to draw conclusions about the efficacy of a loyalty program based on comparisons between members and non-members.

In this review I will discuss how previous research has tried to circumvent this problem. But before moving in this direction I will discuss some theoretical reasons that previous researchers have put forth either to support or to argue against loyalty programs. First of all, reward loyalty programs are built on the idea that by promising and eventually offering rewards to customers whose purchases from the target company are over a certain threshold, or making the reward contingent on the amount spent, companies transform a spot market to a multi-period market where consumers take into account not only the present but also the future (Beggs and Klemperer, 1992; Dowling and Uncles, 1997). This change can alter consumers' decision process and make them allocate a larger share of wallet than they would have done in the absence of

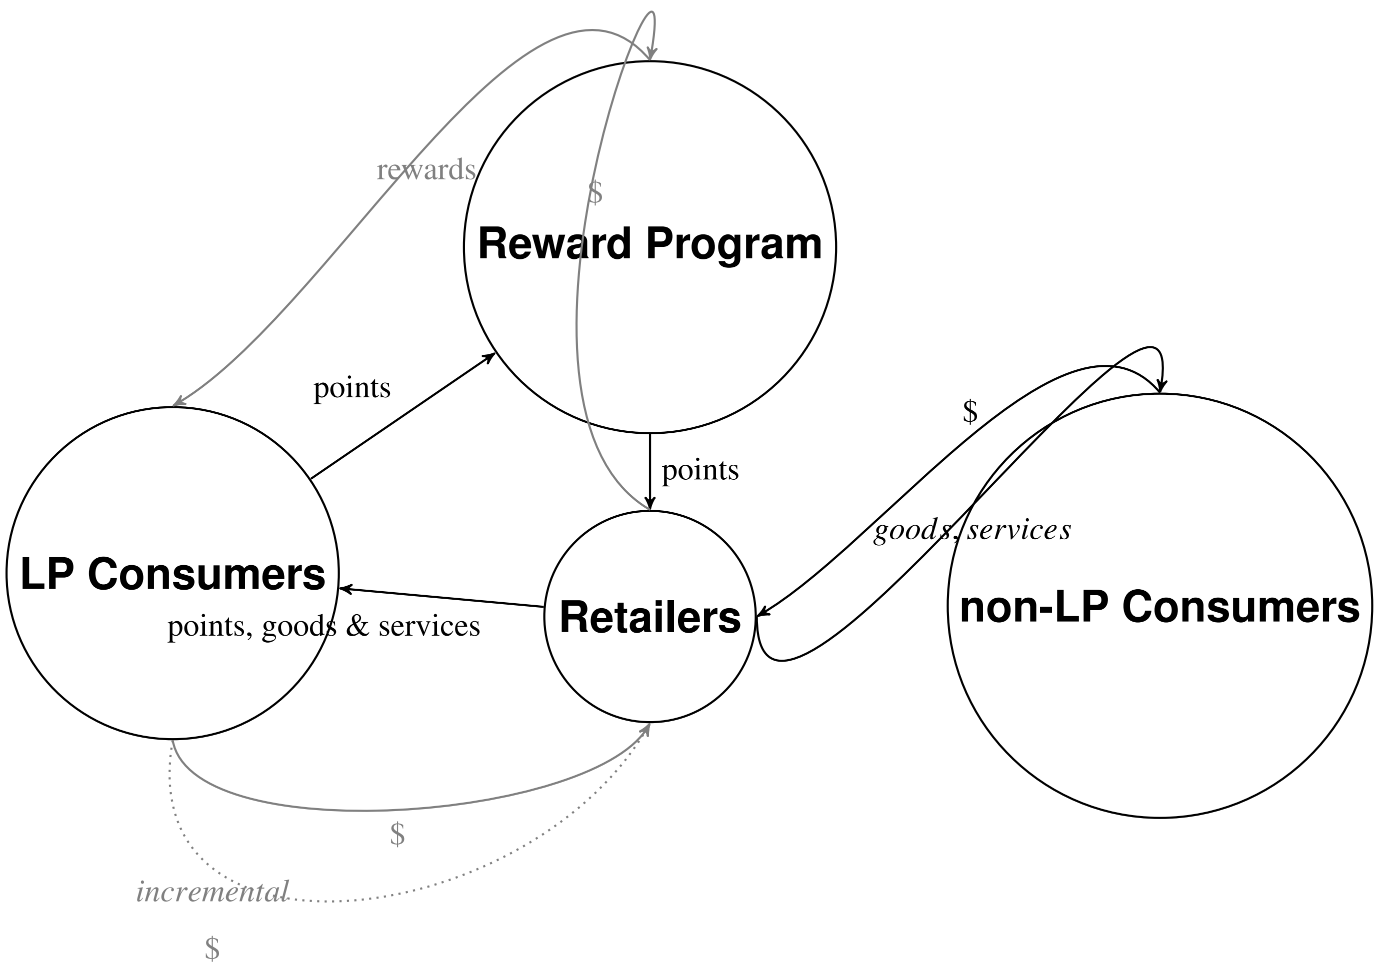

Figure 2.3: The business model for loyalty programs run by third parties

the program and also increase retention.

2.2.1 Switching costs

Loyalty programs have the ability to lock in consumers into a relationship and thus prevent them from switching to a competitor. The idea here is that once a consumer has already accumulated a number of points that can be used to obtain a reward from a certain competitor in the market, switching away from that competitor may entail the loss of the already accumulated points.

In the two-period theoretical model presented by Kim et al. (2001) the heavy users (i.e. those who consume the product in both periods) do not switch providers in the second period because they don't want to lose the reward offered to those who buy in two consequent periods from the same provider. However, Hartman and Viard (2004), point to the fact that two-period models overstate switching costs. In a two-period model the switching cost is equal to the reward value, but in a multi-period consumers can redeem a reward in any of the many future periods, so failing to re-patronize the target retailer in a certain period does not necessarily thwart consumers' chances of gaining the reward.

To the extent that they generate switching costs, loyalty programs have a potential to lock-in consumers in a relationship with a provider even in a model where consumers realize that firms have an incentive to increase prices in the second period precisely because of the switching costs. Based on an analytic model, Kim et al. (2001) find that in most cases, in a duopoly, firms are better off with than without a reward program. The same conclusion is drawn in a model where two firms compete in an infinite number of periods (Beggs and Klemperer, 1992). However, Dowling and Uncles (1997) notice that in the cases of the British grocery market and the American air carriers, after the competitors started introducing loyalty programs, market shares remained steady; even more germane is the observation that in Britain, Asda, the only chain not offering a loyalty program among the top 4, was the fastest growing one (Passingham, 1999). This may appear to contradict the theoretical findings exposed above, but in the absence of data on how the loyalty programs affected the profitability of each competitor, it is hard to draw a firm conclusion.

Kopalle and Neslin (2003) suggest that the major airlines introduced reward programs to counteract a stronger outside category - i.e. low-cost new entrants that were allowed to enter the market after the 1978 Airline Deregulation Act. This idea is derived from their analytical model which shows that when competitors have a potential to expand the market, they are better off by offering reward programs. Also, in another analytic model of two retailers located on a Hotelling line, Lal and Bell (2003) show that when only one of the competitors (or both) introduces a frequent shopper program, customers' cherry picking is reduced or eliminated and the profits of both competitors are increased relative to the situation where no retailer offers a loyalty program. However, their analytic model shows that with heterogeneous consumers, competing loyalty programs do not always result in higher profits.

In markets where competitors offer more or less similar or homogeneous products, loyalty programs act as a tool that help to create differentiation, which in turn makes demand more inelastic and thus can give rise to a non cooperative equilibrium that looks like a collusive equilibrium (Klemperer, 1987). However, the incentives that companies offer to consumers in the early stages, in the hope that this initial commitment will keep them captive in the later stages, increase competition and may leave firms worse-off than in a world without reward programs (Klemperer, 1987). The extra-returns anticipated from the loyalty programs may never materialize or be dwarfed by the costs. Authors who are pessimistic abut loyalty programs' value and the likelihood that they increase profitability (Dowling and Uncles, 1997; Shugan, 2005) put more emphasis on these costs rather than on the uncertain and difficult to measure benefits; in this sense reward programs are seen as producing liabilities that are not compensated by the increase in purchases. Dowling and Uncles (1997) argue that such programs are easy-to-replicate ad-ons (i.e. they don't provide durable differentiation), especially those programs with non relation specific benefits (Deighton, 2000), that can be easily emulated by competition. Moreover, these programs, by offering delayed rewards, do not even create such a strong responses as an immediate reward that is given to consumers through a price cut (Dowling and Uncles, 1997). This observation is qualified by the findings of Zhang et al. (2000). In comparing immediate versus delayed rewards (or front-loaded vs. real-loaded promotion), they find that the latter generate higher profits in a market where consumers seek variety, because in this case front-loaded rewards subsidize purchases that would have occurred even without the subsidy, while rear-loaded promotions retain customers that would have otherwise switched.

2.2.2 Points pressure

Blattberg and Neslin (2008) suggest that there are three specific mechanisms through which firms can use loyalty programs to increase sales: points pressure, rewarded behavior and personalized marketing. In this subsection I am focusing on the first one, which is best illustrated in Figure 2.4 through the black ascending portions that lead to redemption. Additionally, this figure illustrates the rewarded behavior mechanism, which is explained in section 2.2.3.

In this figure the X axis shows the time periods and Y axis the number of loyalty points collected per period. There are spikes in collection in the periods right before a redemption takes place. This pattern can be explained by future discounting. For example, if someone needs 1,000 points to redeem a reward, and is in a situation where they have only 10 points into their account, exerting some effort to collect more points doesn't pay off because the cost of the effort is immediate, while the reward is obtained with delay, when enough points will be accumulated. The discounted reward is simply not large enough to compensate for the cost of effort. However, if the same person is in a situation where they already have, say, 950 points into their account, a small effort to expedite the moment when one can claim the 1000 points worth prize, seems justified. Consequently consumers will accelerate collection as they near the requirements for a reward (the black portions of the line in the figure).

The idea of cost of effort behind the points pressure mechanism is perfectly aligned with the definition of loyalty programs that stresses the fact that they can drive *additional* sales beyond those driven by product qualities, be they objective and measurable or subjective emotions felt by consumers for that product. Without any loyalty program, consumers choose products that they find the best based on their own criteria and goals (Bettman et al., 1981). The loyalty program alters these decision criteria. The additional sales generated by the loyalty program are obtained at the expense of other products or by encouraging consumers to purchase something that they would have not necessarily purchased. Therefore they represent departures from the optimal or normal course of action that the consumer would have taken, so the effort and cost are relative to the course of action where a loyalty program wouldn't have been available. This isn't to say that consumers who incur these costs are overall worse-off. They accept to incur the costs precisely because they are compensated by the reward.

Kivetz et al. (2006) depart from the strictly calculative arguments (current effort vs. discounted reward) and rely on the research in the behaviorist and motivation literatures to build a goal distance model that explains the acceleration in collection through the fact that achievement motivation increases with smaller goal distance. They call this idea the *goal gradient* hypothesis. This explanation doesn't contradict the first account, but rather emphasizes the fact that what matters is the perceived, rather than the real progress towards the goal.

Dorotic et al. (2014) propose a new mechanism related to points pressure which they call *redemption momentum*. In their conceptualization redemption momentum motivates consumers to collect more points in the periods leading to redemption even when consumers aim to make a redemption for which they already have the sufficient number of points. In my models I do not specifically refer to this effect, but I capture it in two ways: 1) by allowing consumers to have uncertainty around the size of the redemption that they will make (Chapter 3) and 2) by allowing consumers to be forward looking: after any given redemption the program participants are likely to extract more utility from future redemptions when they are left with a higher balance which allows them to make higher redemptions in the future (Chapters 3 and 4).

The points pressure and the switching costs are closely related mechanisms. Both effects increase as the consumers approach the requirements that qualify them for the reward. They both express the idea that RP entrenchment happens progressively: at low levels of accumulated points the program is less relevant in altering choice than at high level of accumulated points. One difference between these two constructs can be understood in terms of continuity. With switching costs, the focus is on the discrete choice of whether to buy or not from the focal retailer that offers the RP; with points pressure the focus is on how much to buy from the focal retailer.

The other, more important difference between switching costs and points pressure stems only from their optics on competition. Switching costs clearly emphasize the idea of a competitor; Hartman and Viard (2004) outline them as difference-in-differences measure: supposing that A offers the RP and B is the outside option, switching costs are defined as the value of choosing A over B when A offers the RP minus the value of choosing A over B when A doesn't

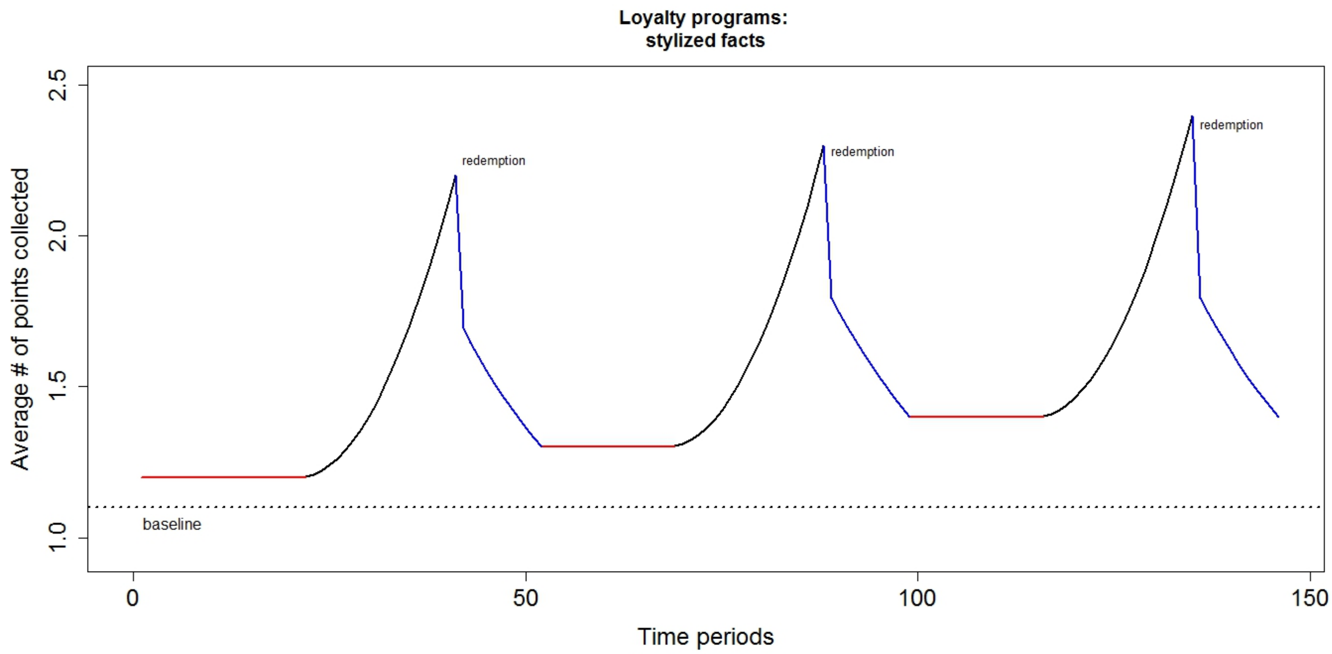

Figure 2.4: Stylized facts on the dynamics of loyalty points collection

offer the RP. In contrast, the points pressure mechanism may be exerted even in the absence of a competitor; additional sales realized by the firm offering the program may stem either from gain share from competition, from advancing future purchases or from, category expansion, the mechanism being agnostic with respect to which of these sources is at play.

The distinction between different sources of additional sales becomes important in the analytic model of Kopalle and Neslin (2003). In an environment with two competitors and one outside option, they find that RPs are more likely to be adopted when the category is highly expandable (i.e. when consumers have a high utility from the outside good). The logic is that when the outside option is unattractive (the market is served at full potential), if a company introduces a PR, most of the additional sales come from competition, which thus has an incentive to respond and wash away any benefits. In contrast, when the market is expandable, by introducing an RP, a company is less likely ruffle its competitor; both can enjoy higher profits by introducing RPs and expanding the market.

Finally, besides switching costs and points pressure another related mechanism which can contribute to increased sales as a result of the introduction of reward program is the zero-cost or effortless switching. For example if a consumer is indifferent between the product or retailer which offers the reward and another product or retailer that doesn't adopt such a program, before the introduction of the program the consumer might choose randomly between the two. After the introduction he or she only buys from the provider with the RP without making any sacrifice. While this effect may be significant, unfortunately without data on individual purchases at all competitors both before and after the introduction of the reward program, it cannot be identified. In Figure 2.4 this effect is represented by the difference between the dotted line (the baseline) and the leftmost red line.

2.2.3 Rewarded behavior

The second mechanism emphasized by Blattberg and Neslin (2008) - rewarded behavior - is illustrated in Figure 2.4 through increased collection rates *after* the redemption moment (the blue portions): the drop from the pre-redemption highest point is not sudden and then the collection flattens out a level that is superior to the previous level. This effect follows from behavioral learning theory which states that 'behavior that is positively reinforced is more likely to recur' (Rothschild and Gaidis, 1981). Simply put, the reward obtained by consumers represent an *unconditioned stimulus* - something that that generates positive feelings and is desired in itself. The product or brand that offers the reward is the *conditioned stimulus* that, through its association with the unconditioned stimulus elicits the same positive feelings (Blattberg and Neslin, 1990, p. 22). This association becomes more apparent after a redemption, so the post-redemption increased collection stems from the goodwill that is projected over the brand through its association with the reward.

This theory, inspired by classical conditioning seems to denude human decision making of any higher order motivation or cognition, relegating it 'to the same level as pigeons pecking levers to obtain their food' (Blattberg and Neslin, 1990, p. 27). While this critique seems valid, Blattberg and Neslin (1990) argue that there is little cognitive activity in low-involvement decision making, so in these type of situations consumers can indeed become susceptible to behavioral conditioning.

Another theoretical explanation for sustained collection effort after a reward redemption is provided by Drèze and Nunes (2011). They argue that being able to achieve challenging goals (for example earning enough miles to be eligible for superior status tier in an airline frequent flyer program), 'can lead to self-learning, resulting in an enhanced sense of proficiency' (p.270) which, in turn, leads to an increased base level of effort.

While both the rewarded behavior (Blattberg and Neslin, 1990) and the partial post-reward reset hypotheses (Drèze and Nunes, 2011) aim to explain the post-redemption increased efforts towards collection, their specific predictions are different. First, the rewarded behavior mechanism seems more likely to be applicable for rewards that are smaller in size, where consumers are not highly involved either with the category, the product or the reward. In contrast, the partial post-reward reset seems more fitted for high-involvement categories. Secondly, while the rewarded behavior hypothesis predicts a gradual decrease in effort after redemption, as the association between the conditional and unconditional stimuli starts to fade away, (the blue portions in figure 2.4) the partial post-reward reset predicts that after a successful goal achievement, the baseline collection is higher than the baseline collection before the reward was achieved (the difference between the red segments in figure 2.4).

Finally, the rewarded behavior effect can also manifest through the reciprocity norm (Kumar and Shah, 2004) - customers feeling that they need to make up to the company for the reward they have received.

Even though the continuous line in Figure 2.4 shows points collected, it is easily apparent that sales follow the same pattern. All the departures from the dotted baseline represent thus sales that would have not been realized in the absence of the loyalty program. In my models I will focus on the points pressure and rewarded behavior mechanisms, since the impact of effortless switching can only be evaluated in contexts where consumers' purchase patterns are available across competitors both before and after the introduction of the RP. But before I conclude this section on the theoretical mechanisms through which customer reward programs can benefit companies, in the next subsection I discuss the third mechanism proposed by Blattberg and Neslin (2008).

2.2.4 Building customer relationships

The last mechanism through which loyalty programs can increase sales of the three described by Blattberg and Neslin (2008) refers to the provision of personalized marketing. This mechanism comprises actions like targeted promotions, cross-selling or personalized customer service to meet individual customers' needs better (Sopanen, 1996). When marketers have access to customer's purchase patterns, they can provide them private offers, that aren't visible to competitors ('stealth marketing') and thus can't be easily copied (Deighton, 2000; Kumar and Shah, 2004).

It seems that this mechanism is in the early stages of theoretical development as the current theory can't answer specific questions such as 'how should targeted promotions be implemented in order to maximize revenues?' or 'how can personalized marketing be used to stimulate cross-selling better than traditional methods?'. So far the answers have been vague, such as 'the answer lies in sophisticated data collection and analytical techniques' (Kumar and Shah, 2004).

Another theoretical argument for loyalty programs is that they allow companies to collect large amounts of data on their customers, which they can subsequently use to better analyze and diagnose their market and eventually implement changes that would increase their profitability. According to Deighton (2000), retailers can sell the data collected from their shoppers to manufacturers, who can thus contact specific buyers with targeted offers.

However, this data is only partial; it doesn't allow us to observe purchases from competitors and moreover is made up of self-selected customers (Dowling and Uncles, 1997; Passingham, 1999) so the data may not be representative. While these draw backs are not fatal, it is worth mentioning that the data collected through a loyalty program can't be relied on as the only source of market intelligence (Kohli and Jaworski, 1990).

Those loyalty programs that assign customers to tiers can make them feel special by assigning them to an elite tier, that provides improved service to the customer, but doesn't cost the company too much - for example priority boarding in the airline industry or annual expenditure reports for heavy users of credit cards (Deighton, 2000). Experimental research by Drèze and Nunes (2009) showed that three-tier programs (e.g. gold, silver and no status) is more satisfying to all involved than a two-tier program (gold and no status). However, whether the preference that respondents have shown for three-tier programs is translated into behavioral measures (such as incidence of re-purchase) remains a still unaddressed question. Also, RPs may induce consumers 'a feeling of participation, the anticipation of future rewards, and a sense of belonging' (Dowling and Uncles, 1997), but how these benefits translate in sales is again an under-research area.

Thus, from a theoretical point of view, loyalty programs can bring value to companies through several mechanisms as summarized in Table 2.1. However, for most of these mechanisms, there are factors that impinge on their effectiveness. Given that strong theoretical arguments have been laid both for and against the effectiveness of loyalty programs, next, I will review the empirical evidence.

2.3 Measuring loyalty programs' effectiveness

In this subsection, I review empirical studies that try to asses the effects of loyalty programs on outcomes of interest such as purchases (defined either in absolute value per period of time or as share of wallet) or retention. I focus on studies that measure this impact either by observing transaction data or consumer diaries and exclude those that rely solely on survey data (Liberman, 1999). Also, I exclude those that focus on estimating brand loyalty rather than its effect on other outcomes (Che and Seetharaman, 2009), as well as those studies that focus on other aspects of loyalty programs, such as speed of adoption (Allaway et al., 2003).

Two main approaches have been used in order to identify such effects: experimental or quasi-experimental and structural. In the first type of approach, the main challenge is finding the appropriate counter factual, or the appropriate comparison term against which to evaluate the changes observed with a loyalty program treatment (Heckman et al., 1999). Experimental studies with random assignment are considered best at illuminating causal inference (Shadish et al., 2002, p.18), but having customers randomly assigned to either loyalty program or non-loyalty program regime is not how companies run their businesses, so my review of the literature revealed just two such field experiments. Moreover, as I will show, even random assignment may not always rule out all the alternative explanations. Quasi-experimental designs

Positive effects	Negative effects or weaknesses of the positive effects
• create switching costs for consumers (Beggs and Klemperer, 1992; Kim et al., 2001)	• the incentives provided are costly (Dowling and Uncles, 1997; Shugan, 2005)
• create differentiation (Klemperer, 1987)	• this differentiation is easily replicated (Dowling and Uncles, 1997)
• points pressure effect (Blattberg and Neslin, 2008; Kivetz et al., 2006)	• points pressure sets in motion only at high levels of accumulated points (Dowling and Uncles, 1997)
• rewarded behavior effect (Blattberg and Neslin, 2008)	
• partial post-reward resetting (Drèze and Nunes, 2011)	
• possibility of personalized marketing (Blattberg and Neslin, 2008)	
• make data collection easy (Dowling and Uncles, 1997)	• gathered data is limited (Dowling and Uncles, 1997)
• customer tiers can make customers feel special (Drèze and Nunes, 2009)	• may not translate into increased purchases
• the anticipation of future rewards and a sense of belonging (Dowling and Uncles, 1997)	• may not translate into increased purchases

Table 2.1: Summary: Theoretical mechanisms through which loyalty programs can affect firm performance

require even stronger assumptions that allow one to interpret the differences between the control and the treatment group as attributable to the loyalty reward program. Quasi-experimental approaches may appear to be interested only in the overall results, by trying to assess whether there a difference between the two treatments, and neglecting the process or the explanation that drives the difference. However, by looking for differences among the RP regime and non-

RP regime in different theoretically driven contexts (why measure something unless having an *apriori* expectation, or an expectation deducted from some theory?), this approach implicitly builds in theory in its development.

The second approach (structural modeling) places more accent up front on the process, or the theoretical mechanisms outlined in section 2.2. This approach prioritizes theory (typically derived from an optimizing behavior) as the guiding source for empirical specification (Chintagunta et al., 2006). While these models have been criticized for relying on numerous assumptions, their defenders (Keane, 2010) argue that data cannot simply 'speak' and assumptions are always necessary in order to learn anything from the data (Shugan, 2005). Moreover, the data itself cannot speak on unobservable constructs that only exist in theory. Here Keane (2010) lists concepts such as the opportunity cost of time; for loyalty programs such constructs would be reward valuation, switching cost or cost of effort to increase loyalty points collection.

It is important to note that the division between the two approaches is not as clear cut as the short descriptions above may seem to imply. As I will underline in the literature review, some studies exhibit the features of both approaches, with the advantage that the obtained results are more robust. In addition, both approaches agree on the necessity of being clear on the assumptions that allow inferring causal effects of the RPs and conducting sensitivity analysis to see how results change with changes in assumptions and on the necessity of running out-of-sample validations.

2.3.1 Experimental and quasi-experimental studies

In two field experiments Kivetz et al. (2006) were able to assign participants randomly to either a loyalty program or non-loyalty program condition. Both studies involved a campus café. In the first study, willing participants were randomly assigned to either a 'reward program' or a 'purchase-habit' condition. Participants in the first condition were given cards that allowed collecting stamps with the purchase date on 10 purchase occasions. After completing the 10 purchases, they would receive a free coffee. In the 'purchase-habit' condition, participants also received a stampable card and were asked to collect a stamp each time they purchase coffee. As a reward they were offered $5 when receiving the card and another $15 when returning the card at the end of the observation period, no matter how many stamps they collected.

In line with the points pressure hypothesis, the authors found that participants in the 'reward program' condition showed decreased inter-purchase times as they approached the goal (i.e. the 10^{th} purchase). In contrast, those in the 'purchase-habit' condition showed deceleration, i.e. the inter-purchase times in their case increased as they approached the completion of the card. Kivetz et al. (2006) were able to rule out alternative explanations such as time effects by using controls, or habituation, by analyzing the data of respondents in the'reward program' condition who after receiving the first reward, were reengaged in the program for a second reward. These participants showed longer inter-purchase times at the beginning of their second streak than at the end of their first streak.

What makes this design robust is fact that the effect of the program is emphasized though a difference-in-differences approach: inter-purchase times when far vs. close to the goal between the randomly assigned treatment and control conditions. The authors thus don't rely solely on comparing the treated to the non-treated (the experimental approach), but also on checking whether, in line with the theory, the inter-purchase time decreases with decreasing distance to goal for those on the 'reward program' condition (the more structural approach). Testing specific theoretical mechanisms (such as the points pressure, as in this case), is, I think, one way in which empirical research in this area can overcome the major problem of self-selection into loyalty programs.

But even in this felicitous case where selection into the treated and non-treated group is random, internal validity is not guaranteed. We are interested in the effects of the 'reward' and 'no reward' treatments on purchases. However, both treatments are included in more complex packages, that include elements extraneous to our interest - in Shadish et al. (2002)'s conceptualization, the causes in the field experiment are *molar*, not *molecular*. Thus, the reward treatment might have implicitly brought in a higher motivation to record purchases accurately (the reward being contingent on the purchases) than the no reward treatment, where there was no such contingency. Thus, it is plausible that the participants in the second condition might have made un-recorded purchases (either they didn't have the card with them or were unmotivated to pull it out of their wallet), so their recordings might be negatively biased. Moreover, if participants in this condition were more likely to neglect using their card towards the end of the study than at the beginning (perhaps researchers [posing as café employees]' instructions were followed more closely at the beginning), then the observed difference-in-differences approach can no longer be entirely attributable to the reward treatment. By exposing this argument, I don't want to imply that Kivetz et al. (2006)'s results should be discarded, but merely to show that the golden standard of random assignment does not necessarily take care of all the problems in causal inferences.

In second field study reported by Kivetz et al. (2006), participants in the café reward program were again randomly assigned either to a condition where they received a card that required 10 stamps for a free coffee, either to a condition that required 12 stamps, but had the first two already marked, presumably as a bonus. So even if in both conditions participants needed to make 10 purchases to obtain the reward, the illusion of progress was manipulated such that participants in the second condition perceived that they had a smaller distance proportion to the goal ($\frac{10}{10}$ vs $\frac{10}{12}$). Consistent with the goal gradient hypothesis, on average, participants in the second condition reached their goal faster than participants in the first condition.

Finally, Kivetz et al. (2006) report the results of an analysis performed on secondary data provided by a site that rewards participants for rating songs: $25 Amazon gift certificates are provided for participants who rate 51 songs. They find that participants visit the site more frequently and increase the number of songs rated per session as they approach the reward. Even though this set-up is somewhat different from retailing contexts where consumers are rewarded for previous purchases, the results lend further support to the goal gradient hypothesis.

While in their interpretation of the results Kivetz et al. (2006) emphasized only the points pressure mechanism, Drèze and Nunes (2011) note that these results are also consistent with post-reward partial reset effect, in the sense that the observed inter-purchase times after re-engagement are slightly lower than the inter-purchase times observed in the beginning, when participants first joined the rewards program. Drèze and Nunes (2011) use data provided by the loyalty program of an air carrier that grants tier 1, 2 or 3 statuses to customers who fly more than a number of miles in each year. Consistent with the points pressure mechanism, the authors show that customers reduce their inter-flight times as they get closer to the threshold that grants them a superior status. Moreover, they show that after successfully achieving status, customers work harder for the next year's status - which is consistent with their partial post-reward reset hypothesis. The alternative explanation that customers who are able to achieve status already have a higher propensity to fly with the company than those who don't is ruled out by using a regression discontinuity design. Drèze and Nunes (2011) also provide experimental evidence to show that only non-trivial rewards are motivating and elicit partial post-reward rest and to support the proposed mechanism (success increases self-efficacy).

Using transaction data provided by a grocery retailer, augmented by matched survey data, Taylor and Neslin (2005) also test both for the points pressure and a rewarded behavior hypotheses. Customers who spent enough during the period in which the program lasted were eligible to receive a turkey. By comparing the purchase levels in the pre-program period to those during the program period, they noted that the program had a positive impact on sales (the points pressure mechanism). Furthermore, by comparing the difference between the post-program and pre-program purchases among redeemers and non-redeemers, Taylor and Neslin (2005) conclude that a rewarded behavior effect is also supported, even though this effect is smaller than the points pressure effect. In order to account for the fact that there is self-section among reward redeemers, they use a switching regression model, where the probability of redeeming is estimated from customer characteristics, collected through the survey data.

Even though Taylor and Neslin (2005); Drèze and Nunes (2011); Kivetz et al. (2006) (study 3) did not use random assignment into treatment and control conditions, their evaluation of the effects of the reward programs they study are relevant, as they relied on theoretical insights to create meaningful controls.

A less usual approach to build a comparison term (or a counterfactual) is followed by Sharp and Sharp (1997): they rely on somewhat intricate 'secondary source contrasts' (Shadish et al., 2002, p.127-128). They exploit the observed empirical generality of positive correlation between brand penetration (i.e. the share of those buying the brand at least once out of the total number of potential customers) and share of customer requirements (SCR) where SCR represents the average share of purchases of the focal brand by those who purchased at least once, out of their total purchases in the category, being therefore a measure of brand loyalty. Their logic is that the effect of the introduction of a loyalty program should be visible in altering this empirical generality, as a loyalty program is expected to increase SCR, but not penetration. More precisely, those who don't buy at all in a category won't be swayed by a loyalty program, but those who already buy in the category are more likely to purchase more from the brand that offers the loyalty program, that is to increase the SCR for that brand

In fact, Sharp and Sharp (1997) focus on a coalition program, namely Fly Buys. This coalition program was launched in Australia in 1994 and it allowed consumers to collect points at different retailers or gas stations or by using a specific credit card. They could then redeem the points for air travel or accommodation. The authors survey a panel of consumers in the weeks following the launch of the program and used the data to estimate a descriptive Dirichlet model of repeated purchase patters (Ehrenberg et al., 2004). The model can be fit using data on category penetration, penetration of one particular brand, average frequency of buying the category per category buyer and the average frequency of buying the particular brand. It captures the well-established empirical generality of positive correlation between a brand's penetration and its share of purchases for the customer who already makes purchases in the category. So their strategy is to use the Dirichlet prediction for penetration and SCR as baselines and compare these predictions with the observed measures. Upward deviations from the model predictions in SCR and downward deviations from the model predictions for penetration are interpreted as evidence that the loyalty program is effective. Their finding is that 'there is a trend of excess loyalty for Fly Buys brands, though for most it is disappointingly small' (p. 480).

The other way in which the authors exploit the predictions of the Dirichlet model is by looking at the observed and theoretical percentages of the customers who bought from two coalition members. The logic here is that the loyalty program could potentially partition the market and increase the duplication of purchase among coalition members relative to the Dirichlet prediction. The finding here is again that the deviations are very minor. However, the conclusions of this study rely on the implicit assumption that the studied markets should obey the empirical generalization of double jeopardy captured by the Dirichlet model, which states that there should be a positive correlation between a brand's penetration and its SOR (Bendle et al., 2016). This relationship has been observed in a series of markets, but it is an inductive, a-theoretical observation that may not be entirely accurate, so the conclusion drawn based on it may not be valid, either.

In loyalty programs, measuring customers propensity to purchase from the target company before the introduction of the program is most of the times impossible, because purchases are tracked using precisely the loyalty cards that are part of the treatment. Also, the data collected through loyalty programs offers information only on members, i.e. only on the 'treated' group, so obtaining purchase information on the (yet) non-treated is often impossible. However, there several ways in which researchers have tried to overcome this challenge: considering purchases in the first month after joining the program as the 'untreated' outcome (Liu, 2007); using data from industries where data is collected for all customers, not only loyalty program members (Bolton et al., 2000; Verhoef, 2003); collecting panel data on people's purchases at all competing retailers (Nako, 1992; Mägi, 2003; Leenheer et al., 2007; Meyer-Waarden, 2008), using data provided by special promotions ran though the loyalty program (Lal and Bell, 2003; Taylor and Neslin, 2005), or promotions where the RP is implemented only for a category of goods (Drèze and Hoch, 1998). While each of these approaches comes with strings attached, in the sense that additional assumptions of non-selection need to hold in order to be able to identify the effect of the loyalty program on purchases, they still represent worthwhile endeavors.

Industries such as insurance or financial services record information about their customers irrespective of whether they are or not members of the loyalty program. This feature provides researchers with an easily available comparison group, even though the customers self-select to participate into the program. Bolton et al. (2000) use data provided by a financial services company, while Verhoef (2003) use data from an insurance company. In the absence of instruments, or even control variables to account for self-selection into these programs, the positive estimates of loyalty program membership on outcomes such as retention, account usage or share development (defined as the change in the share of type of insurances typically purchased by consumers from the target provider), are not informative with respect to the causal effect of the loyalty programs.

Drèze and Hoch (1998) implement a quasi-experimental design, where consumers are awarded loyalty points only for purchases made in a certain category (baby products, in their case). Once consumers accumulate 100 'Baby Bucks', they can use them to redeem a $10 coupon for other store purchases. The results of the program are evaluated using a difference in differences approach with two sources of contrasts: sales in the baby category during the program compared to one year earlier at the target chain vs. the same difference in the baby products market; and also sales in the baby category during the program compared to one year earlier vs. the same difference for grocery sales at the target chain.

Liu (2007) observes the behavior of a sample of customers who have just joined a loyalty program offered by a convenience store. The data is truncated to include only customers who made at least two purchases during the observation period. She assigns members to three groups (light, medium and heavy) based on their spending in the first month of the program.The first month is considered the baseline. She then compares the change in inter-purchase frequency and the transaction size over time for each group. While the heavy buyers maintain their behavior, the other two groups show a decrease in inter-purchase frequency and an increase in transaction size. Excluding explanations such as learning effects or overall positive trends in the market, she concludes that the loyalty program made the light and medium buyers more profitable and it hasn't affected the heavy buyers. This conclusion relies on the assumption that the behavior during the first month after joining the program isn't affected by the participation in the program and therefore the effect of the loyalty program on the heavy buyers may be underestimated. However, excluding customers with less than two purchases is likely to overestimate the impact of the program. Moreover, including a control on whether the current month is the last month in which the customer makes a transaction at the store (effectively using future information to construct an explanatory variable) is also likely to introduce positive bias in the estimated effects of the program. The reason is that this dummy variable partials out all the decreases in transaction size and increases in inter-purchase frequency and attributes them to the 'last month effect' which works against the effects of the RP.

By using data from a consumer panel, Meyer-Waarden and Benavent (2009), are in a better position to draw a causal inference from a pre-test \post-test contrast, as they observe the purchases of a group of customers who join the loyalty program both before and after they sign up (group 1), as well as the behavior of a group of customers who don't join the program during the observation period (group 2). Thus, they also have access to a control group. However, their modeling efforts are concentrated towards self selection (hazard model for the time to sign up), mean comparisons between group 1 and group 2 over time, but without controls. They also use a vector-autoregressive model and its associated impulse-response functions to test the relationship between the number of cards distributed and outcomes of interest for the store, such as share of wallet, frequency of purchase and mean basket.

Lal and Bell (2003) also find that reward programs affect light and moderate users more than heavy users. Their conclusion is based on data provided by a retailer who offers rewards to customers whose purchases are above given thresholds within certain periods of time (for example, spend more than $475 in a five-week period and receive through mail a coupon for a full ham). Comparing the changes in spending between the the promo period and the period before the promo among redeemers (i.e. presumably those who are interested and involved in the program) and non-redeemers within each category (light, medium, heavy), they note that these differences are largest for the light buyers and smallest for the heavy buyers.

Another approach (Mägi, 2003; Meyer-Waarden, 2008) is to simply compare loyalty program members and non-members, without measuring the differences between the two groups before the introduction of the program. In order to conduct this comparison, researchers use a panel of consumers whose purchases at any of the grocery stores in their area are recorded. Some of the stores offer loyalty programs. Mägi (2003) finds that having a loyalty card is associated with increased purchases at the focal chain, while Meyer-Waarden (2008) conclude that the members of the loyalty programs purchase more frequently, have larger basket sizes and make purchases in more categories from the store whose loyalty program members they are. However, given these data sets and the methodology, it is impossible to assess whether these effects are attributable to the reward programs themselves or are driven by a self-selection effect. If customers who already have a preference for a certain store are more likely to join that store's loyalty program, then the observed positive outcomes are certainly driven by self-selection. In this context, an unbiased estimation of the loyalty program's contribution to increased sales becomes impossible.

Nako (1992) also observes the purchases made by a panel of customers in the airline industry. As expected, similarly to the grocery examples above, members of a given airline's loyalty program are more likely to choose that airline vs. a competitor. But again, selection into the program is endogenous. However, in his data there are additional sources of variation, such as fares and service characteristics (average travel time, total flights - a proxy for convenience, percent of direct service and on-time performance for each airline). So he includes in the choice model interactions between frequent-flyer membership status and these variables; none of them turns out to be significant, but the sign of the coefficients suggests that for example the members of the frequent-flyers are less sensitive to price. Assuming that those who chose to join the frequency program were not less price sensitive to begin with, these estimates can better answer the question of how the program influenced sales.

In order to account for this self-selection problem, Leenheer et al. (2007), use instrumental variables to separate the effect of the loyalty program from that of the selection into the program. The Dutch households panel that they use contains information on all the supermarket purchases done by the members of the panel over two years, as well as loyalty program membership. A distinctive feature of the programs offered by the studied grocery chain is that they have both a discount feature (that offers instant discounts to members) and a saving feature (that allows accumulating points that can be later exchanged for rewards). Moreover, the authors have access to survey data on panel member's attitude towards loyalty programs in general (such as perceived costs and benefits and privacy concerns) and membership in gasoline loyalty programs. These variables serve as instrumental variables since they affect participation in loyalty programs, but don't affect the preference for a specify grocery chain. Their estimates show that on average a loyalty program membership yields € 240 additional revenue per customer per year.

Dorotic et al. (2014) analyze a Dutch coalition program where consumers accumulate points at the participating retailers and then choose themselves when and how much to redeem. This is exactly the type of program that I study, but my methodological approach is structural, while they rely on a reduced form Bayesian model which describes jointly the redemption and purchase patterns of program participants. They present point estimates which quantify the effect of different explanatory variables on the purchase incidence and amount and on the redemption incidence and amount. However, they do not use these point estimates to quantify the amount of sales that are due to the program and the minimum margin at which the program is profitable. Their conclusion is centered around the idea that redemption is the key mechanism that drives profitability for RPs. Therefore, I think that the structural approach that I develop in this dissertation can provide additional insights and complement this study.

2.3.2 Structural studies

Structural models lay out explicit equations capturing that utility that the consumer obtains from choosing the good that offers the reward, as well as the utility from the reward itself. The choice of whether (or how much) to buy represents the solution of a dynamic model in which consumers are assumed to maximize the utility derived from the good and from the reward. The dynamic aspect is important because loyalty programs shift consumers from optimizing 'myopically' each purchase occasion separately, to optimizing over longer sequences of purchases (Lewis, 2004) - i.e. taking into account the future consequences of their current actions. The different parameters of the model are typically estimated such that the likelihood of the observed sequence of choices is maximized.

Lewis (2004) uses data from an Internet grocer that rewards customers with frequent-flyer miles based on the value of the cumulative purchases they realize with one year (rewards are awarded for crossing each of the 3 thresholds set by the retailer). Given the finite-horizon in which the RP is applicable, the author uses backward recursion to solve the dynamic problem. The results show that at relatively high levels of cumulative spending, the probability of purchase increases as the expiration time (52 weeks) approaches. This is in line with the points pressure mechanism explained in section 2.2.2. However, for a low cumulative spending, the probability of spending stays almost flat for the first 35 weeks and then it decreases continuously, reflecting the lost hope of winning the reward. The policy simulations suggest that the loyalty program as effective, as its removal would drop the purchase incidence from 19.7% to 19.2% per week.

Hartman and Viard (2004) analyze a loyalty program offered by a golf course, that offers in-kind benefits to members: every 10 credits entitle members to a free round. Credits can be earned by either purchasing or using an earned round. Their primary focus is on estimating switching costs, as explained in section 2.2.1, using a dynamic structural model. In line with their expectation, they find that less frequent customers face higher switching costs when they get close to the reward, compared to customers with higher frequency purchases. However, it is not often the case that low-frequency consumers have earned enough credits to be in a high-switching cost state. Three factors contribute to this result: when holding a reward, frequent customers are more likely to use it quicker; credit expiration deadlines are less likely to affect frequent customers; and the discount effect is smaller for frequent customers, as the time between accruing the first and last credit needed for a reward is shorter.

Kopalle et al. (2007) and Kopalle et al. (2012) propose models applicable to RPs that offer two different types of benefits: free in-kind rewards (flights or hotel stays) for which the consumers decide the redemption time and status upgrades that are granted automatically to participants who accumulate enough points within a period of time. This structure creates a trade-off between redeeming rewards and amassing points in order to gain status. The structural dynamic models that they propose capture the conflicting goals that the programs create. They emphasize both the rewarded behavior and the points pressure mechanisms and conclude that both components of the programs are effective in generating incremental sales.

Study	Experimental	Transaction data	Forward-looking	Tests specific mechanisms	Endogenous redemption	Source of identification
Kivetz et al. (2006)	✓	✓	✓†	✓		random assignment; structural
Drèze and Nunes (2011)		✓	✓†	✓		structural; regression discontinuity
Taylor and Neslin (2005)		✓	✓†	✓		pre-test \post-test contast; switching regression
Sharp and Sharp (1997)				✓	NA	secondary source contrasts
Bolton et al. (2000)		✓			NA	post-test measurement \non equivalent groups
Verhoef (2003)		NA			NA	post-test measurement \non equivalent groups
Mägi (2003)					NA	post-test measurement \non equivalent groups
Meyer-Waarden (2008)		✓			NA	post-test measurement \non equivalent groups
Nako (1992)		✓		✓	NA	post-test measurement \non equivalent groups
Liu (2007)		✓			NA	pre-test \post-test contrast‡
Meyer-Waarden and Benavent (2009)		✓			NA	pre-test \post-test contrast with control group
Lal and Bell (2003)		✓				difference in differences
Drèze and Hoch (1998)		✓				difference in differences
Leenheer et al. (2007)		✓		✓	NA	instrumental variables
Hartman and Viard (2004)		✓	✓	✓		dynamics in distance from reward
Lewis (2004)		✓	✓			dynamics in distance from reward
Kopalle et al. (2007)		✓	✓	✓	✓*	dynamics in distance from reward
Kopalle et al. (2012)		✓	✓	✓	✓*	dynamics in distance from reward
Dorotic et al. (2014)		✓			✓	reduced form model of redemption and purchase

† Implicitly

‡ Considering the first month into the program as pre-test

* For programs with lumpy collection (airlines, hotels)

Table 2.2: Summary: Loyalty programs - empirical work

2.4 Discussion

Table 2.2 summarizes the extant empirical research on the effectiveness of RPs, organizing it along important dimensions. Since random assignment is credited as facilitating causal inferences (Shadish et al., 2002, p.248), one such important dimension is whether the study follows an experimental design or not. Unsurprisingly, given how businesses implement loyalty programs, there is only one experimental random assignment study (Kivetz et al., 2006). Another important feature is whether the study relies on transaction data, obtained from the records of the company implementing the RP, which is less prone to errors than other data sources, such as interviews or consumer diaries. As loyalty programs offer rewards at the end of *sequences* of purchases, they shift the consumer's horizon from single to multiple periods. Therefore, acknowledging that consumers are forward looking, is likely to better capture and explain the mechanisms of RPs, such as points pressure. The specific mechanisms through which RPs affect behavior are important because they provide more structure and allow researchers to make more tight predictions. When one just compares the molar outcomes in an RP condition to those in a control condition, there may be potential threats to internal consistency, as the RP treatment itself, or other components of the treatment packages may be responsible for the effect. However, when the molar treatment is unpackaged into specific mechanisms, more precise, and thus more falsifiable hypotheses can be advanced. Unless refuted, such hypotheses build better theories than less precise hypotheses (Chalmers, 1999).

Finally, another important feature of the research on RPs is whether it investigates programs where the reward is granted automatically (such as status upgrades), or programs where the decision to cash-in is deliberate. With two notable exceptions (Kopalle et al., 2007, 2012), the structural models reviewed here are focused exclusively on the first type of programs. However, most of the loyalty programs implemented by companies (frequent flyer, retailing or financial services schemes) allow participants to accumulate points and exchange those points for a reward at the moment they choose to. Moreover, it was observed that consumers are not very active at redeeming (Abu-Shalback Zid, 2004; Davis, 2004; Andrew, 2006; InsideFlyer, 2006) and have a tendency not to drain-out their points accounts after making a redemption. While in the case of airlines, hotels and cruise lines redemption may be limited by availability or lack of sufficient points, in the retail sector the rewards are always available and they comprise wide ranges, including more accessible rewards. Why consumers keep hoarding points in such contexts where redemption should be relatively easy and where the value of the rewards doesn't increase convexly in the number of points required is a question that current research has not addressed yet.

The models that have so far been advanced in the literature cannot account for consumers' preference for maintaining high balances of unredeemed points, again with the exception of (Kopalle et al., 2007, 2012) who focus on specific frequent flyer or hotel programs, where participants have conflicting goals - gaining status vs. cashing-in a reward. In the case of flights or hotel stays, consumers collect points in large, relatively infrequent lumps. In order to accumulate sufficient points for status tier upgrade, consumers may prefer to hold on to their points, which explains the accumulation of large balances. For example, in the frequent flyer program studied by Kopalle et al. (2007), the average customer earns 6500 miles during an average trip and they need to accumulate 25000 miles within an year for the silver tier, 50000 for the gold and 75000 for the platinum; at 50000 accumulated miles, consumers can redeem a free flight. A consumer who already has 50000 accumulated miles might not wish to cash-in a free flight, but rather pay for it and thus accumulate in one shot $\frac{1}{4}$ of the additional miles they need for the platinum status.

However, this tendency of clinging to the loyalty points and miles has also been observed in programs where redeeming a reward is much less likely to interfere with one's goals of gaining status. In retailing and financial services loyalty programs, consumers collect small amounts of points almost every week and also the choice of rewards tends to include low value options; by redeeming such a low valued option, consumers don't really threaten the objective of gaining sufficient points for a status upgrade. At the loyalty program that I am studying, consumers can redeem 95 miles for a $10 grocery gift certificate. In the same time, by refraining to use up the loyalty points and paying with cash in order to obtain the necessary 1000 points per year that would grant them gold status, they would barely earn about .5 miles. In this case substituting cash for loyalty points in order to improve one's chances of acceding to a superior status is simply not a sensible plan. Given that in US only 31% of the loyalty program memberships are in travel and hospitality, while 39% of the memberships are in retail, 21% in financial services and 9% in other industries(Berry, 2013), we need a more general model that accounts for miles accumulation in the absence of competing goals (free rewards vs. status). In the next two chapters I outline two such models.

One of the key takeaways from this chapter is that while some studies focus on the overall effectiveness of RPs, for example comparing the outcomes of different groups in before/ after quasi experiments, others provide more depth, by measuring one or all of the three main specific effects that were discussed: points pressure, switching costs and the rewarded behavior effect. The second approach is more robust as a method of investigation, as it allows researchers to point out specific mechanisms, which are more informative and more actionable than general mechanisms.

However, all these three specific mechanisms are in one way or another related to the redemption moment. Switching costs are effectively created when consumers have already accumulated a sufficient number of points that make defection to a competitor more costly. Similarly, the points pressure becomes apparent when the reward is ‘in sight’, while the rewarded behavior effect appears only after a redemption. Redemption seems thus to be a key component, that sets in motion the mechanisms which actually make RPs effective. As in most RPs consumers keep accumulating points and are free to decide when and how many points to use for a redemption, designing models that don’t impose a certain redemption, is crucial for marketers in this area.

Next I develop two models that account for the preference for high balances and make the redemption timing and size endogenous (Chapter 3) or keep the redemption timing exogenous but enodogenize the size of the reward (Chapter 4). These models are applicable to sectors where there is no trade-off between status and rewards, therefore they represent a useful and robust tool for assessing how effective the most common types of RPs are.

Chapter 3

A collection-redemption cycle: the 'cash' model

As noted in the previous chapter, research so far, especially in structural models, has little to say about redemption timing and size. The previous structural models focused mostly on rewards that are granted automatically to RP members, such as upgrades to superior status tires or other types of rewards that are transferred automatically to participants as soon as they reach certain thresholds of accumulated points in given time intervals. This lack of interest may be justified if the problem was uninteresting. This would be the case if either consumers were completely disengaged and redeemed whenever they happened to remember or, in in programs with rewards that are linear in the number of points, they redeemed as soon as they had enough points for the minimal reward. However, the data does not support either of these two conjectures. Firstly, I see a large share of consumers hoarding their points, even if rewards are not convex in the amount of LPs and for someone who discounts the future, there is no value in waiting to receive a reward later when the reward can be obtained immediately. Secondly, I observe consumers who redeem a part of their available LPs, so obviously they haven't forgotten about them.

One may say that this partial redemption can be attributed to ceiling effects (consumers simply not needing/ wanting a reward that they can afford but choose not to have). This is indeed a possibility when the rewards are large in size, and the model outlined in the next chapter makes provisions for such situations. In this chapter I focus on the situation when the rewards are relatively small and consumers have many options of spending them and therefore ceiling effects are not likely. It is difficult to imagine that someone who used their points for a \$10 grocery discount in a week simply didn't need more groceries in that week. So in this chapter I propose a model for, and analyze programs that offer rewards as in-store discounts redeemable at consumers' chosen time. I use the terms 'reward', 'gift certificate' and 'coupon' interchangeably.

This chapter is organized as following: first I lay out a theoretical model which describes how participants in the program behave. The two main aspects of behavior are points collection (specifically how points pressure and the rewarded behavior effect can influence points collection) and points redemption, which is governed by the utility of different redemption sizes. Since collectors are not myopic, the actions that they choose to take (in terms of collection and redemption) are driven not only by the immediate utility, but by 'today's' utility *and* 'tomorrow's' discounted utility. The model captures this dynamic by using a Bellman equation.

Secondly, I present the data that I am using to estimate the parameters of the model outlined in Section 3.1. This data comprises AIR MILES participants who joined the AIR MILES Cash program. Section 3.2 presents details on the specifics of this program, as well as descriptive statistics of the sample that I use. The Empirical Strategy Section presents details on the estimation method for the parameters of the model described in Section 3.1. I am using indirect inference, a method that belongs to the larger class of general method of moments estimators. The key aspect of indirect inference is the fact that it relies on simulations and on auxiliary moments to pin down the parameters of interest. This section also provides arguments to support the claim that the model is well identified, i.e. that just one set of parameters can 'explain' the patterns observed in the data.

The Results Section shows how different specifications of the model fit the data, the estimated parameters, and, most importantly counter factual analyses that assess how profitable the program is and how different changes in the structure of the program will affect its profitability. The Chapter concludes with Section 3.5, which recaps the main features of the model presented in this chapter and the main findings.

3.1 Theoretical model

Participants in an RP are not always engaged with the program (Ashley et al., 2011), especially in low-involvement categories. In the model I assume that in some periods consumers are disengaged; in these periods they are oblivious to any marketing efforts meant to attract their attention; they don't collect and don't redeem any point.

When they are not disengaged, I assume that participants are engaged at one of the two levels. The first level of engagement is that where consumers collect only 'regular' points - that is the type of points that are handed out by the retailer who offers the program according to their standard policy (for example 1 point for every \$10 spent). At the second level of engagement, program participants are also responsive to the special offers or promotions that allow them to accumulate bonus points; in these bonus periods, participants collect both regular and bonus points. For example retailers may offer bonus points to customers who spend above a certain threshold, or who buy particular products that are on a limited-time 'bonus points' promotion.

I consider that the probability of being disengaged, denoted by q_d is exogenous. When engaged, consumers can collect bonus (or promotional) points with an exogenous probability q_b. I assume these probabilities to be exogenous because they are not easily changed by the company's activities and they are more a matter of personal characteristics. For example, consumers have either high or low propensities to take advantage of promotions and they will do so across retailers or categories. I allow q_d and q_b to be different among different segments in the population reflecting the idea that people have different propensities to be disengaged with the program or to take up 'bonus points' offers. When they are not either disengaged or collecting bonus points, consumers collect only regular, non-bonus points and the probability of this state is given by $1 - q_d - q_b$. Equation 3.1 captures the distribution of points m[1] for individual i in period t in any of the three period types: disengaged, only non-bonus (or regular)

[1]Throughout the dissertation I use 'points' and 'miles' interchangeably, hence the notation m for points.

points, or both regular and bonus points.

$$m_{it} \begin{cases} = 0 \text{ with probability } q_{d_i} \\ \sim lnN(\Phi_1 + \Phi_2\rho_{it} + \Phi_3 Bonus_{it} + I[HasRedeemed_i = 1](\Phi_4 \cdot TimeSinceRedemption_{it}^{-\Phi_5}), \sigma_2) \\ \text{ with probability } q_{b_i} \\ \sim lnN(\Phi_1 + \Phi_2\rho_{it} + I[HasRedeemed_i = 1](\Phi_4 \cdot TimeSinceRedemption_{it}^{-\Phi_5}), \sigma_1) \\ \text{ with probability } 1 - q_{d_i} - q_{b_i} \end{cases} \tag{3.1}$$

When engaged (rows two and three of equation 3.1) participants have a 'baseline collection', captured by parameter Φ_1, which is equivalent to points they receive for the purchases that they would make if there was no points pressure and no rewarded behavior effect. Note that this baseline collection, as captured by Φ_1 may include purchases that are still generated by the zero-cost or effortless switching mechanism - i.e. still due to the RP. However, since in the absence of data before the enrollment into the RP there is no variation in this mechanism, I cannot identify it. RPs can be effective thorough any of these three mechanisms (points pressure, effortless switching and rewarded behavior effect). But since I am only able to identify the points pressure and the rewarded behavior effect, my estimates of the impact of the program will be lower bounds of the true effectiveness.

As they keep accumulating points, consumers find it optimal to exert more effort in acquiring LPs (Kivetz et al., 2006; Blattberg and Neslin, 2008). This is due to the fact that the cost of effort is fully offset by the increasingly large rewards that become available, which are valued non-linearly. In contrast, when they have few points in their accounts, the rewards, and especially the larger rewards, are distant into the future and heavily discounted, such that they can't compensate for the cost of effort. This is the points pressure mechanism and is captured in equation 3.1 by $\Phi_2\rho_{it}$, where Φ_2 is a parameter to be estimated and ρ_{it} is the optimal (unobserved) level of effort chosen by the consumer i in period t. I assume that effort is discrete, reflecting different actions that consumers can take to increase their collection - for example buying a small item just to make sure they spend enough to obtain a loyalty point on a specific visit, buying something that is on 'points offer' (or 'miles offer'), despite that product not being one's first choice, giving up cherry-picking among retailers, or driving longer to shop from a retailer that offers the RP instead of a competitor in the consumer's proximity. Furthermore, consumers can have a boost in collection if they come across a promotion, this effect being captured by Φ_3.

The baseline collection, the additional collection due to effort, and the effect of a bonus promotion, are all anticipated components. However, post-redemption, consumers can also exhibit unanticipated boosts in collection due to the rewarded behavior effect, because the association between the reward (the unconditional stimulus) and the RP (the conditional stimulus) is strongest at the time of redemption and fades away over time (Rothschild and Gaidis, 1981; Blattberg and Neslin, 1990, 2008). Therefore I use two parameters to capture this effect - Φ_4 and Φ_5 in equation 3.1, where $I[\cdot]$ is an indicator variable equal to 1 if the statement within the brackets is true and 0 otherwise ($HasRedeemed_{it} = 1$ means the consumer i already redeemed something by time t). Φ_5 is expected to be positive, showing that the effect of a previous redemption decays as the time since the last redemption increases. I model this effect as unanticipated because it has been theorized as a very basic phenomena that goes below consumers' conscious radar. Moreover, the distinction between anticipated and unanticipated effects is important because my model assumes forward-looking consumers with rational expectations.

I use R to denote the minimal number of points that can be redeemed. All redemptions are multiples of R; these multiples are denoted by k and referred to as 'number of redeemed certificates' or 'coupons'. However, consumers cannot redeem more points than the points they have. Denoting with S the stock of points they have at the beginning of an observation period and with m the number of points acquired in that period, the largest number of coupons that they can redeem is $\lfloor \frac{S+m}{R} \rfloor$.

Consumers enjoy redeeming LPs and even if the reward is linear (i.e., P points correspond to αP dollars in in-store discounts). I allow consumers' utility to be non-linear in the sense that redeeming 2R might feel better than redeeming R on two separate occasions. The main argument for this specification is that retail grocery discounts are relatively small and redeeming only one discount certificate might not be enough to entirely cover the price of most items available in a store. Since consumers are more willing to spend LPs on treats on which they don't like to spend money, they may have a preference for redeeming more certificates at once, such that the points can entirely cover the price of the treat being purchased. Moreover, by computing the utility not in absolute numbers of certificates redeemed, but in logs of the number of certificates redeemed, I put a countervailing pressure on utility increasing convexly in the number of certificates, which becomes apparent especially at hight levels of redemptions. Without logs, any utility function that would give a certain difference in utility between redeeming say 1 or 2 certificates, (that would match the choices observed in the data) would tend to imply an unrealistically high difference in utility between redeeming say 7 or 8 certificates, which are less frequent choices in the data, but which would hamper the ability to conduct thought experiments.

The non-linear utility function that I described above is meant to capture average, systematic effects. Depending on the parameters, it may imply for example that redeeming 2 certificates brings 3 times more utility than redeeming one certificate. However, consumers who have enough points to redeem 3 certificates may still redeem only 2, due to idiosyncratic, unobserved influences - for example they really wanted a reward that required only 2 certificates. I use ϵ_{itk} to denote the unobserved taste for a redemption of size k in period t by consumer i and I assume that ϵ_{itk} are distributed Extreme Value.

Exerting effort to acquire additional points is costly. I assume that the cost of effort increases exponentially, such that when $\rho = 0$ (i.e. there is no effort), there is zero cost as well. The exponential cost of effort is justified by the idea that people first take the low hanging fruit (i.e. they first take those actions meant to increase their collection of LPs that are the least costly). In retail RPs, an example of low cost action would be buying a small item to make sure one's bill is large enough to receive one LP, while a high cost action could be purchasing 100% of one's grocery needs from the retailer that offers the RP despite it having higher prices and being more inconveniently located than a competitor. Equation 3.2 shows the instantaneous utility that consumers have from redeeming certificates and the their disutility of effort. In this expression a_1 to a_3 are model parameters and e^{ρ} is the exponentiated effort. In practice I allow participants to have different utilities from redemption, by allowing a_1 to be different among different population segments.

Equation 3.2 specifies the utility derived by consumer i in period t who redeemed k certificates. It comprises the non-linear systematic effects, the disutility produced by effort and the idiosyncratic components discussed above. ε is a positive constant added so that the log of the number of certificates redeemed can be taken even at 0.

$$U_{itk}(\rho) = a_1 log(k_{it} + \varepsilon) + a_2 log(k_{it} + \varepsilon) log(k_{it} + \varepsilon) + a_3(e^{\rho_{it}} - 1) + \epsilon_{itk} \tag{3.2}$$

I assume that different redemption sizes are realized proportional with the utility they bring. So what matters is the relative utility of each redemption size as compared to other redemption sizes. I use the utility of a non-redemption ($\varkappa = 0$) as the baseline; the utilities for other redemption sizes are calculated relative to $\varkappa = 0$. Equation 3.3 below gives the probability that $\varkappa$ numbers of certificates will be redeemed, based on the relative utility of different redemption sizes.

$$Pr(K = \varkappa) = \frac{exp\Big(a_1[log(\varkappa + \varepsilon) - log(\varepsilon)] + a_2[log(\varkappa + \varepsilon)^2 - log(\varepsilon)^2]\Big)}{1 + \sum_{k=1}^{maxK} exp\Big(a_1[log(k + \varepsilon) - log(\varepsilon)] + a_2[log(k + \varepsilon)^2 - log(\varepsilon)^2]\Big)} \tag{3.3}$$

One thing to note in equation 3.3 is that the probability for each redemption size is given only by the differences in the flow utility, i.e. utility associated only with the redemption itself and not with the position in which the participant is left for future redemptions. It is conceivable that sometimes consumers prefer to make a smaller redemption now in order to be able to make larger redemptions into the future; but this specification does not allow for such effects. The reason is quite technical; to explain it here I will foreshadow some of the features of the model, that I will explain in more detail later. In a nutshell, I develop a dynamic model which allows program participants to put in effort 'today' even if the redemption only takes place 'tomorrow'. This infinite-horizon dynamic model is of the following form:

$$\begin{aligned} V(\text{state space today}) = max_{\text{actions\_I\_can\_take}}\ & \text{Flow\_utility\_today} + \\ & \beta \Sigma Prob(\text{tomorrow I end up in state } \int) V(\text{state space } \int) \end{aligned} \tag{3.4}$$

In other words the value of the world in which I 'wake up' today is given by the best action(s) that I can take today in order to maximize my utility today and the discounted utility for the set-up in which I will wake up tomorrow (weighted by the probability that I end up in that set-up). β is the discount factor. The fact that β is sub unitary is key in solving this type of model, because this makes the general equation above a contraction mapping, which can be solved iteratively by starting from any V^0(state space$\int$), plugging that guess into the right hand side of the equation 3.4 to obtain a new V^1(state space $\int$), until the difference between V^n and

V^{n+1} is very small. By introducing $V(\cdot)$ into the Flow_utility_today, equation 3.4 is no longer a contraction mapping and I can no longer solve the model through value function iteration.

To recap, equation 3.3 shows the probability with which each rewards size is chosen. These probabilities are influenced by the current-period utility only.

As mentioned above, the main feature of RPs is that they transform a spot market into a multi-period market where consumers take into account not only the present but also the future (Beggs and Klemperer, 1992; Dowling and Uncles, 1997). So the forward-looking consumers choose their actions not in order to maximize utility at time t (equation 3.2), but that utility plus their prospects at time $t + 1$. For example, consumers may choose to exert effort (and incur the cost of effort) in the current period even though they know that they are not likely to make a redemption in this period; however exerting effort 'today' may allow them to make a larger redemption 'tomorrow', so for a forward-looking consumer, putting in effort today is still worth, even though the benefits are delayed.

The actions consumers can choose are the following: 1) how many certificates to redeem (k), and 2) how much effort to put into collection (ρ). The first choice is limited, such that participants can't redeem more points than what they have - i.e. stock of points at the beginning of the period (S_{it}) plus whatever they accumulate during the period (m_{it}).

In each period consumes consumers know the period type (PT_{it}) - i.e. whether they are disengaged ($PT = 0$, or period type is 'disengaged'), collecting non-bonus points only ($PT = 1$ or period type is 'base points only') or collecting both non-bonus and bonus points ($PT = 2$ or period type is 'both regular and bonus points'), as well as the stock of points they have in their account (S_{it}). The stock of miles is an observed state variable, but the period type is not observed - for example if a consumer collects zero miles in a period it may be because they were disengaged ($PT_{it} = 0$) or they were engaged ($PT_{it} = 1$), but still did not collect any mile. There are another two state variables which are unobserved by the econometrician - ξ_{it} - to which I refer to as the 'luck of the draw', i.e. the realized number of points corresponding to any combination of state and choice variables for participant i in time t, and ϵ_{itk} - which is the taste for different redemption sizes, or consumers' impulses of spending.

As shown in Figure 3.1, in my specification, at the beginning of each time period participants know their stock of miles (S), the period type (PT) and their average taste for different

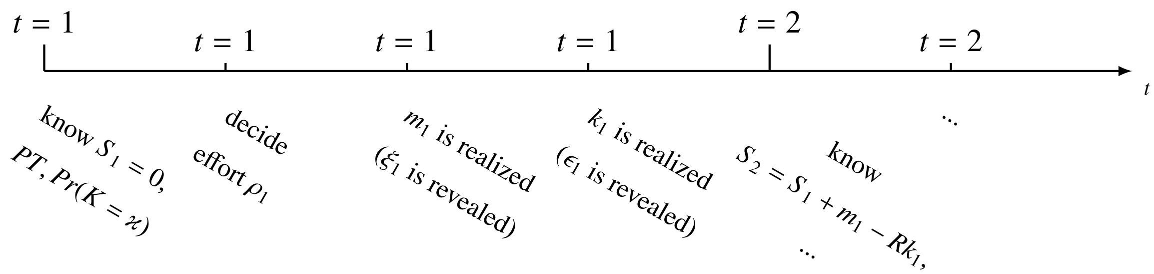

Figure 3.1: Sequence of actions

redemption sizes ($Pr(K = \varkappa)$). Knowing this, they first decide on the optimal level of effort (ρ). This corresponds to a scenario where consumers decide in advance how much effort they will put into collection (e.g. which retailer to visit), but they don't know for sure how many points this decision will bring them, as this outcome also depends on random factors such as specific offers and prices of the retailer (this randomness being captured by ξ), which the consumer doesn't know in advance. Depending on their choice of effort and on the how lucky they are in collecting miles , the number of miles collected in that period (m) is realized. After miles collection took place, participants learn the set of ϵ's - i.e. the random components that affect their taste for each redemption size and make the redemption that maximizes their utility in that period (if no redemption brings the highest utility, then no redemption is made). Note that the decision to spend is done at the end of the period, allowing consumers to spend the newly acquired points as well. In the next period (denoted by $t = 2$ in the Figure), this cycle is played again.

When they join the program, consumers' stock of miles is zero - hence $S_1 = 0$ in Figure 3.1. Equation 3.5 below shows the law of motion for the stock of miles: the next period's stock (S') is given by the previous period stock (S), plus the number of miles acquired in the current period (m) minus the number of certificates redeemed in the current period (K) multiplied by the number of miles required for one certificate (R).

$$S' = S + m - Rk \tag{3.5}$$

I assume that there is no dependency between period types over time. As shown in the transition matrix below, the probability of each period type is the same, irrespective of the period type that preceded it.

$$PT' = \begin{array}{cc} & \begin{array}{ccc} 0 & 1 & 2 \end{array} \\ PT = \begin{array}{c} 0 \\ 1 \\ 2 \end{array} & \begin{Vmatrix} q_0 & 1 - q_0 - q_b & q_b \\ q_0 & 1 - q_0 - q_b & q_b \\ q_0 & 1 - q_0 - q_b & q_b \end{Vmatrix} \end{array} \tag{3.6}$$

As alluded above, the two choices that consumers make are ρ (the effort which influences the number of miles acquired and k - the size of the redemption. The number of certificates to be redeemed is an integer between 0 and the maximum number of certificates that participants can afford given their starting stock of miles in that period (S) and the number of miles acquired in that period (m):

$$k \in \left\{0, 1, ..., \lfloor \frac{S + m}{R} \rfloor\right\} \tag{3.7}$$

Effort is constrained arbitrarily at 6 distinct levels as shown in equation 3.8. I will talk more about this arbitrary constraint in the 'Identification' subsection.

$$\rho \in \{0, .2, .4, .6, .8, 1\} \tag{3.8}$$

However, effort is costly. The cost function is given by the following expression:

$$cost = -a_3(e^{\rho} - 1) \tag{3.9}$$

where a_3 (a parameter) is expected to be negative. When $\rho = 0$, the cost is also 0. Higher levels of effort increase the cost exponentially.

Despite effort being costly, program participants are motivated to exercise effort, because it can bring them more miles and having more miles gives them the possibility to make higher redemptions, which bring positive utility. However, I assume that participants only learn their taste for different redemption sizes right before the end of the period, so when they decide the optimal effort, they only know the probability with which redemption sizes k is selected (the probabilities being proportional to the utilities) and the utility brought by each k. As explained for Equation 3.2, the utility of redemption is quadratic in the logs of the number of certificates redeemed and ε is a constant added so that I can take logarithm when $k = 0$. So the objective function is given by the expected value of the utility of the redemption sizes that are within the consumer's reach:

$$rewards = \sum_{\varkappa=0}^{max\_\varkappa(m(\rho))} Pr(k = \varkappa)[a_1 log(\varkappa + \varepsilon) + a_2 log(\varkappa + \varepsilon)^2] \tag{3.10}$$

In the equation above the summation is from 0 (consumers can always choose to redeem nothing) to $max\_\varkappa = \lfloor \frac{S+m}{R} \rfloor$. As shown in the equation the maximum number of certificates that can be redeemed ($max\_\varkappa$) depends on the number of miles acquired in the current period (m), which in turn depends on the effort ρ.

Equation 3.11 shows the Bellman equation that describes consumers' maximization problem (the i and t subscripts have been omitted to simplify exposition). In this equation $E_{PT'}$ denotes the expectation over the future period type (disengagement, collection of regular points only or collection of both non-bonus and bonus points). In principle participants can differ with respect to their propensities of being disengaged or to collect bonus points and $E_{PT'}$ are taken to accurately reflect those different propensities. $E_{m^a}(\rho)[\cdot]$ is the expectation over the number of points to be realized in the current period, but this expectation takes into account only the anticipated points (see discussion above) - in this notation subscript m refers to points and superscript a refers to 'anticipated'. The choice of effort (ρ) influences the number of acquired points, which in turn influences the maximum number of gift certificates that can be redeemed. δ is the per-period discount (including both the time discount and any risks consumers may associated with the continuation of the RP into the future).

When consumers are disengaged ($PT = 0$), they don't make any decision and the stock of points for the next period stays unchanged, so in this situation the value function is given by the discounted value of the next period's value function ($\delta V(S, PT')$). However, if consumers are engaged ($PT \neq 0$), they choose the optimal effort ρ, based on their expectations of the realized number of points $E_{m^a}(\rho)[\cdot]$ and their expectation of redemption. The bottom line in equation 3.11 represents the expected utility from redemption, given that consumers don't know the taste shocks (ϵ_{itk}) at the time when they make the choice of effort. The utility from each possible redemption is weighted by the probability that that redemption is optimal.

This is a stationary infinite-horizon dynamic problem, the working assumption here being that consumers expect to remain enrolled into the LP program over indeterminate time. In order to solve it I use both value function and policy function iterations for given sets of structural parameters. The solution of this problem is a policy function which, for each S and PT specifies the optimal effort ρ that consumers should put into collection. Note that only effort (ρ) is a solution of the dynamic problem, not k.

$$V(S,PT) = E_{PT'}\Bigg(I[PT=0]\delta V(S,PT') + I[PT \neq 0]\Bigg(\max_{\rho} E_{m^a}(\rho)$$
$$\Bigg\{\sum_{\varkappa=0}^{max\_\varkappa(m)} Pr(k=\varkappa)\big[a_1 log(\varkappa+\varepsilon) + a_2 log(\varkappa+\varepsilon)^2 + a_3(e^{\rho}-1) + \delta V(S+m-R\varkappa, PT')\big]\Bigg\}\Bigg)\Bigg) \tag{3.11}$$

Table 3.1 and Table 3.2 summarize the state and choice variables that characterize consumers' problem, respectively the structural parameters that characterize the model.

Table 3.1: Summary: state and choice variables

	State variables		*Choice variables*
Symbol	Interpretation	Symbol	Interpretation
S	balance of points at the beginning of the period	ρ	effort put into collection
PT	period type	k	number of redeemed coupons
	(disengaged, regular only, regular+bonus)		
ϵ	unobserved taste shock for different redemption sizes		
ξ	unobserved points shock		

3.2 The data

The data is provided by the AIR MILES Rewards Program, the largest customer loyalty coalition program in Canada. In a coalition program participants (shoppers) can collect points at any of the sponsors of the program - i.e. retailers that have joined the program. In the AIR MILES program the points are referred to as miles. As it is typical in such programs, the sponsors are

Table 3.2: Summary: structural parameters

Symbol	Interpretation
Φ_1	baseline points collection
Φ_2	rate of transformation of effort into points
Φ_3	effect of promotional activity on points collection
Φ_4/Φ_5	effect of previous redemption on points collection as function of time since redemption
$\sigma_1/\ \sigma_2$	scale parameters of the log-normal distribution in regular/ regular+bonus periods
a_1, a_2	quantify the utility derived from redeeming coupons
	(allowing for non-linear utility in the number of coupons)
a_3	the dis utility produced by the cost of effort

not direct competitors. For example, among gas stations, Shell is a member of the AIR MILES coalition program in Canada, but Petro Canada is not. AIR MILES sponsors span a wide range of sectors, from courier and logistics to apparel and grocery. Sponsor offer LPs on regular purchases, ranging from 1 LP for \$5 spent to 1 LP for \$500 spent, but sometimes they also have time-limited promotional offers where more points can be earned for the same amount or, in retailing, certain items are sold with bonus points. By the company's estimates, two thirds of the Canadian households are members of the program. In the word of Brian Pearson, the CEO of LoyaltyOne - the company which runs the AIR MILES Rewards Program - 'in Canada, our brand has higher penetration than any credit card, higher usage than any laundry soap, and higher retention than any wireless service' (Pearson, 2012, p. 3).

In January 2012 AIR MILES introduced a program called 'AIR MILES Cash' which allowed participants to select a share of the miles they collect which should be directed towards what they called 'the cash account'. Miles in the cash account can only be used to redeem discounts (or gift certificates/ vouchers) at partner stores in a wide range of sectors such as groceries, gas stations, pharmacy, house renovation, coffee shops, toys, movies, etc. 95 AIR MILES are worth 10 Canadian dollars in instant discounts or vouchers which the program participant can also offer to someone else should they wish.

For this chapter I use transaction and redemption data of AIR MILES program participants who have collected at least 100 miles in 2008 and who chose to direct all of their collected miles into the cash account at any point between January 2012, when the option was introduced, until 2013 when the observation period ends, without reverting that decision while observed. The first restriction reduces the likelihood that the sample contains inactive participants - i.e. shoppers who stopped using their card altogether.

The company provided data on 31,020 participants who do not have an AIR MILES branded credit card linked to their AIR MILES account and who live in one of Canada's Eastern provinces. I focus on this subset of participants because those with credit cards have many more opportunities to amass AIR MILES, so their collection patterns look different. Participants are observed from the moment they switched their collection to 'cash only' (i.e. chose that all the miles they collect are directed towards the cash account that can only be used for discounts) up to the end of the observation period - December 2013. For each account, I aggregated the total miles collected in each 2-week period, as well as the miles redeemed within the same interval. I define an observation as an unique account-period combination. I consider the two weeks period to be the optimal aggregation level that captures trends in consumers' collection and redemption patterns, without picking up much noise (as it is the risk with higher frequency aggregations) and without missing much information (as it is the case with lower frequency aggregations).

It is important to know that the distribution of miles collected by participants in each period is markedly right skewed, with a mean of 14.9 and and median of 6. In the top .1% of observations, the number of collected LPs is larger than 342. In order to acquire this number of miles within two weeks, participants should have spent about $1,700 at Holiday Inn (one of the most generous sponsors offering 1 mile per $5 spent), or around $6,840 at Metro - a grocery store. I consider collections above this threshold as outliers, reflecting one-time unusual expenses, so I remove from the data set all the 944 accounts (3%) that had at least one instance where they collected more than 342 miles within a period. Another 190 participants are eliminated for inaccuracies in their records of regular and bonus miles collections. This results in a usable sample of 29,886 participants.

Of the 29,886 participants, during the observation period 4.2% have reached a balance of miles equal to or greater than 792 LPs, which is the top 1% of the observed balances in the data. These accounts are labeled as 'censored' and they are followed only up to the moment when they reach this threshold - giving a data set of 1,202,417 observations. I use the balance cap because the balance of miles is one of the state variables and keeping the size of the state space under control helps speed up the estimation. Calculating what the best effort policy is for someone who has for example 1,000 miles in their account takes time and chances are that very few program participants are ever in this situation. Moreover, with a balance of 1,000 miles a program participant can redeem anywhere between 0 and 10 coupons; calculating the chances for each of these eleven redemption sizes is again computationally costly. However the cap doesn't imply completely ignoring these accounts. Both in the data and in the estimation, program participants who cross this limit are followed up to the moment they have crossed it.

The collection of miles shows a significant spike around the redemption moment. In Figure 3.2 I plot the average number of miles collected as a function of time *to* the moment when a redemption takes place (the left panel) and the average number of miles collected as a function of time *from* the moment when a redemption takes place (the right panel), emphasizing separately the regular and the bonus miles. Since counting always starts with 1 from a redemption moment, both series have the same value at 1 (in the redemption periods). In both cases, the collection decreases significantly with time.

This represents model-free evidence for the points pressure and for the rewarded behavior effect, especially since in both cases the decrease is progressive. It appears that the increase in collection is not associated only with the redemption period, case in which it could be that a large random collection triggers redemption, not that participants actively increase their collection in order to make a redemption. Moreover, both effects are apparent both for the baseline and for the bonus miles.

A valid concern regarding the above interpretation of Figure 3.2 is that collectors with high propensities to collect bonus miles and low propensities to be disengaged can accumulate miles more easily and thus they will always be a short distance away either *to* or *from* a redemption. Thus, the increased collection around redemption times might be the outcome of a heterogeneous sample rather than genuine points pressure and rewarded behavior effects. Separating the points pressure from sample heterogeneity is very important in order to accurately assess the effectiveness of the program. Therefore, a key aspect of the model is to take into account heterogeneous propensities to be disengaged or to collect bonus miles.

Moreover, for the left panel (accumulation of points leading to an observed redemption) another possible interpretation is that there is some selection in the timing of the reward: that participants may be more likely to choose to make a redemption after a streak of lucky collections - i.e. increases in their collection that are not due to systematic effort exertion, but to luck. So the redemption moment is *selected* based on the random draws that characterize miles collection. The model that I estimate helps to quantify these 3 different effects: true loyalty program effect, heterogeneity and selection. By teasing out the last two effects, I can be more confident that the estimated RP effect is not upward biased.

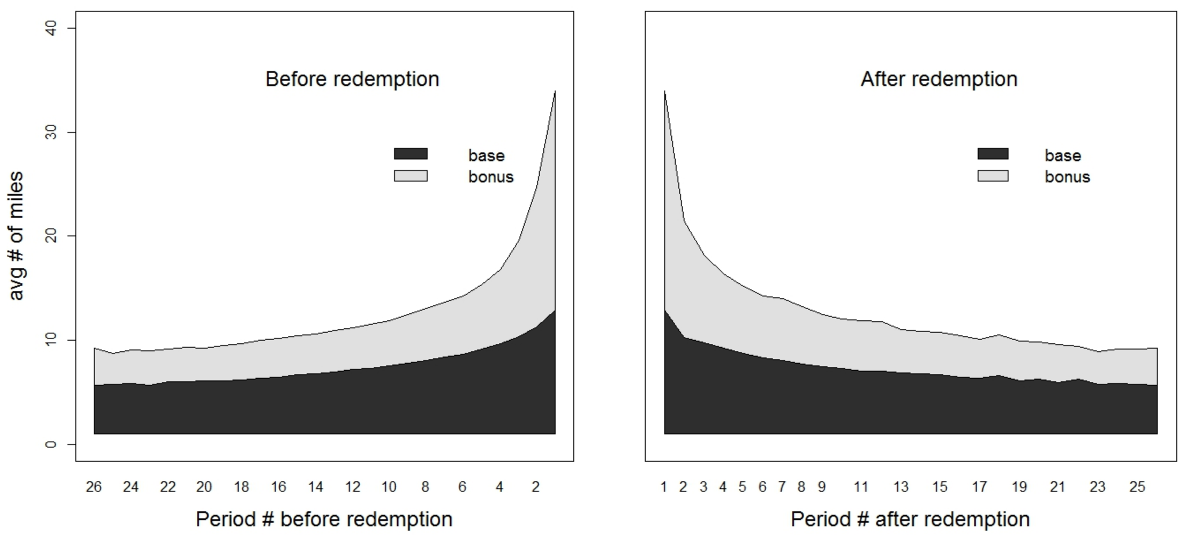

Figure 3.2: Average collection as a function of time to/ after a redemption

The upper panel of Figure 3.3 presents the distribution of miles balances at the end of a redemption period by redemption size - e.g. the leftmost box plot is the distribution of miles that participants who redeemed 1 coupon (95 miles) have in their account at the end of a redemption period. What is most striking is the fact between one quarter and half (depending on the redemption size) of the balances are larger than 95 miles which is the minimal redemption. Given that after cleaning the data the mean number of miles collected per period is 13.2 it is hard to believe that all this extra stock of miles was built after the redemption. Moreover, given that 95 miles offer a $10 discount which can be used at several retailers it is also hard to believe that participants simply didn't need to make more expenses within that period. This implies that keeping a relatively high balance is an active strategy, not the outcome of circumstantial constraints.

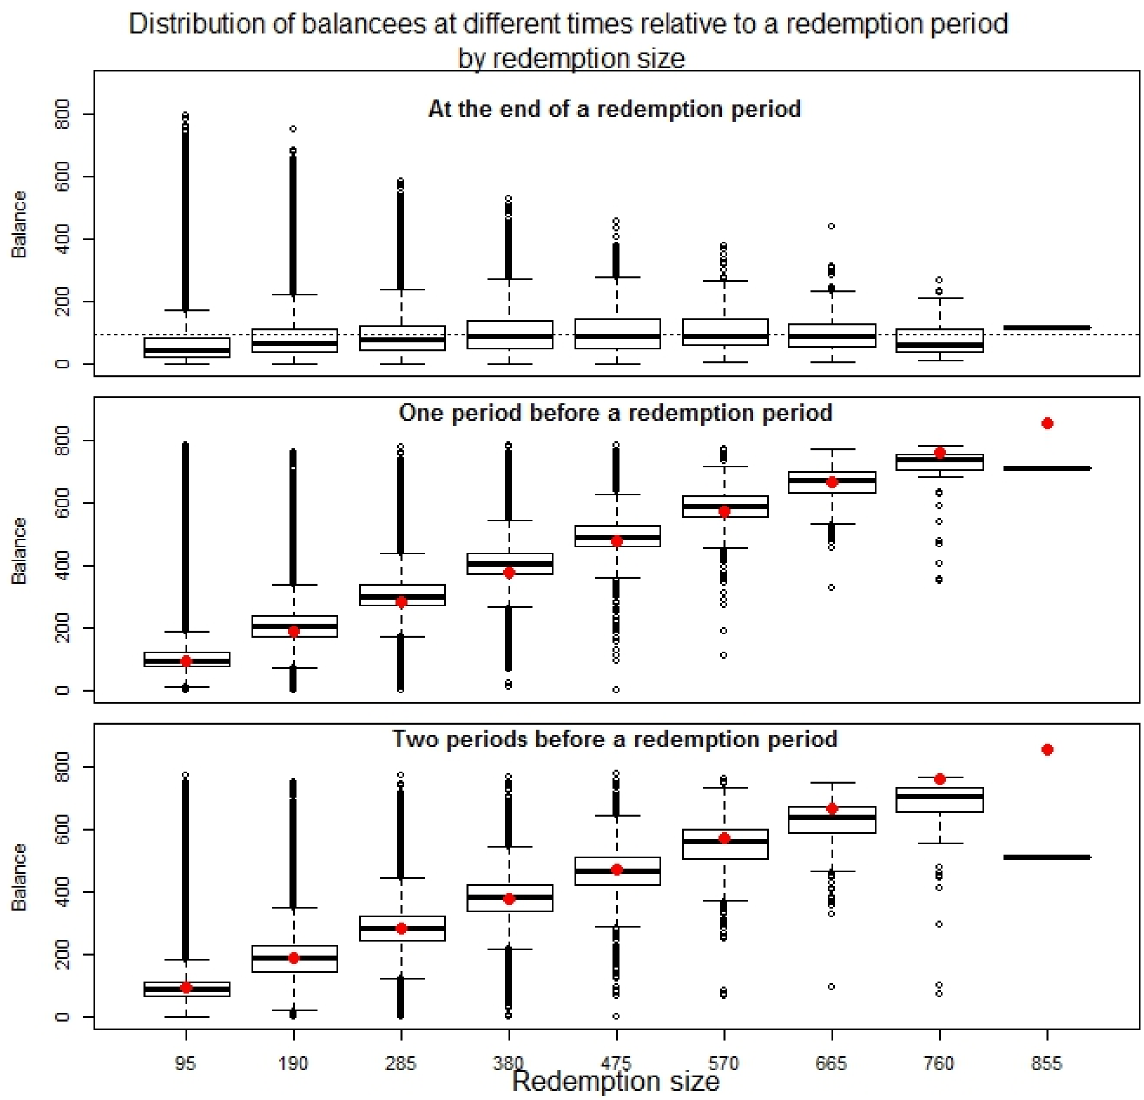

Figure 3.3: The distribution of miles balances after a redemption (top panel), one period before a redemption (middle panel) and two periods before a redemption (bottom panel) by the size of the redemption

Furthermore, the middle and bottom panels of Figure 3.3 show the distributions of the balances one and two periods before the redemption. These plots show that about half of the redemptions could have taken place 2 periods (4 weeks) before the moment they actually took place. Two mechanisms can explain this patterns: 1) participants are disengaged, or 2) they strive to accumulate high balances which allow for larger redemptions that provide larger utility, but in the same time they have impulses (or random utility shocks - ϵ's), that make them redeem 'now', even if that redemption size has a systematic utility which is lower than a larger redemption size which could be attained in the future - call this 'an option value'. The model that I propose accommodates both mechanisms: disengagement through q_d and the random utility shocks through the specification of the utility function (equation 3.2).

3.3 Empirical strategy

3.3.1 Indirect inference

The outlined model poses a challenge for estimation, as its complexity impinges on one's ability to derive closed form expressions that can be estimated through maximum likelihood. Therefore, I use indirect inference (Gourieroux et al., 1993; Smith, 2008), a method that relies on certain key features (moments) - usually denoted by $\hat{\theta}$ and simulations, in order to derive the structural parameters of interest. The moments are collectively called the auxiliary model.

Indirect inference has been previously used in economics and business, especially finance, for similarly complex problems such as showing the impact of social interactions on unemployment (Topa, 2001), of social security on saving and retirement decisions (van der Klaauw and Wolpin, 2008), estimating companies' cost of financing (Hennessy and Whited, 2007) or estimating CEO entrenchment (Taylor, 2010).

Since in the sample there are 29,886 participants, I simulate an equal number of individuals over a number of periods equal to that observed in the data (since the participants in my sample joined the AIR MILES cash program at different times, they are observed over different lengths of time). In the simulation each individual collects points and makes redemptions until they either reach a balance of 792 miles (the ceiling level mentioned above) or until their number of observation periods ends (again, the number of observations are exactly those observed for individuals in the real data). For a given set of structural parameters ($\boldsymbol{\beta}$) the choices of the simulated individuals are based on the optimal policy (equation 3.11) and their specific random errors. The optimal policy depends on the structural parameters and prescribes the best actions (in terms of effort and number of certificates to be redeemed) for each situation (state) in which the simulated individuals may be. There are three sets of errors and they remain fixed for different $\boldsymbol{\beta}$'s: 1) the period type (disengaged; only regular miles; regular and bonus miles), denoted by e; 2) the specific draw of miles from the normal distribution (ξ); 3) the specific taste shocks for different redemption sizes (ϵ).

Once the optimal policy has been calculated (for a given $\boldsymbol{\beta}$), a collection-redemption pattern is generated for each simulated individual. Now the task is to see how close the simulated data is to the real data and find that $\boldsymbol{\beta}$ which generates simulated data that is as close as possible to the real data. In order to measure this degree of closeness I use key descriptive statistics of both data sets, or moments, such as means or regression coefficients (the next section describes them in detail). The moments of the true data are denoted by $\hat{\theta}$ and those of the simulated data are denoted by $\tilde{\theta}(\boldsymbol{\beta})$. In mathematical terms:

$$\hat{\boldsymbol{\beta}} = arg \min_{\boldsymbol{\beta}} (\hat{\theta} - \tilde{\theta}(\boldsymbol{\beta}))' W (\hat{\theta} - \tilde{\theta}(\boldsymbol{\beta})) \tag{3.12}$$

where W is the inverse of the variance-covariance matrix of $\hat{\theta}$ obtained by bootstrapping, where the off-diagonal elements have been replaced by 0s.

3.3.2 Identification

The model outlined above features a disutility produced by the cost of effort (a_3) and two parameters associated with the non-linear benefits (a_1 and a_2) of redemption. Since both costs and benefits are constructs that exist only in theory (the observed data doesn't contain any measure of them), I fixed a_3 to 1. Moreover, the levels of effort are set to 6 discrete values, between 0 and 1 in .2 increments. The discretization of effort is in line with the observation that consumers increase their collection above baseline by taking certain individual actions like switching retailers, products or accepting higher prices.

The parameters that define the non-linear utility of the redemption size (a_1 and a_2) are identified mainly from the average redemption ($R \cdot E(k_{it}|k_{it} > 0)$ - where R is the number of points required to redeem one certificate and k is the number of redeemed certificates) and from the average balance left into one's account at the end of a redemption period ($E(S_{it}|k_{it} > 0)$). A more convex shape of the utility function should mean that participants make larger redemptions and also leave fewer miles into their accounts when they can afford to redeem higher numbers of certificates.

Since I allow a_1 to be different for different segments in the population - call them 'redeemers' and 'non-redeemers', I will estimate a_1^r, a_1^{nr}, as well as p_r - the percentage of participants who belong to the 'redeemers segment'. To identify these parameters I use two additional moments related to the pattern of redemption: one is the percentage of observed participants who don not make any redemption while observed, $E(I[NonRedeemer = 1])$ and the second is the average observed inter-redemption time $E(\tau|cens = 0)$ - i.e. including only the completely observed, non-censored spells.

I emphasize that 'redeemer' and 'non-redeemer' are only labels. I expect $a_1^{nr} < a_1^r$, such that participants classified as non-redeemers (NR) derive less utility from any positive redemption than the participants who are classified as redeemers (R). This decreased utility from redemption will make participants in the non-redeemer segment redeem less often, but they can still make redemptions. I use the label 'non-redeemers' for brevity, but the proper label should be 'less inclined to redeem'.

In order to identify the baseline collection (Φ_1) and the boost generated by the bonus miles (Φ_2), I look separately at the mean collection in periods when no-bonus miles are collected ($E(m_{it}|PT_{it} < 2)$) and the mean collection in periods when both regular and bonus miles are collected ($E(m_{it}|PT_{it} = 2)$). To pin down the two scale parameters of the log-normal distributions (σ_1 and σ_2), I rely on the variance in collections: $Var(m_{it}|PT_{it} < 2)$ and $Var(m_{it}|PT_{it} = 2)$.

However, the observed average collection is also influenced by the effort which consumers put into it, which, in line with the theory, is expected to increase as participants approach a redemption moment. Thus, in order to disentangle the effect of effort (Φ_3) from the baseline collection (Φ_1) and bonus (Φ_2) I use a regression (denoted R1) which includes only the subset of observations that contain collection streaks that eventually lead to an observed redemption. For example if a participant is observed for 40 periods and makes only one redemption in period 20, I use only the data up to and including period 20. In this regression the LHS is the number of collected miles and the explanatory variable is the logarithm of the number of the period starting with 1 from a redemption period and going backwards up the the last redemption or to the moment when participants switched their collection to the 'cash account' ($log(Count\_back)$). As effort is expected to increase as $log(Count\_back)$ decreases (i.e. as the redemption moment is getting closer), the coefficients of this regression are key in identifying the the points pressure mechanism. Furthermore, the rewarded behavior effect (Φ_4 and Φ_5) are identified from a regression of the number of miles collected on the logarithm of time since the most recent redemption ($log(Count\_fwd)$), using the subset of observations that follow a redemption (regression R2 in Table 3.3).

Without heterogeneity in the propensities to be disengaged or to collect bonus points, in each period each participant to the program has a probability equal to q_d of being disengaged ($PT = 0$), a probability equal to q_b of collecting bonus points ($PT = 2$) and a probability equal to $1 - q_d - q_b$ of collecting only regular miles ($PT = 1$). By introducing heterogeneity, I allow q_d and q_b to be specific for each segment. The shares of the different segments are captured by parameters denoted by Q^I and Q^{II}; the share of the third segment is simply $1 - Q^I - Q^{II}$. In order to identify the share of these different segments, as well as the specific characteristics of each type (e.g. the probability that an individual of type I is disengaged - q_d^I, or the probability that an individual of type II collects bonus miles - q_b^{II}), I rely on several other moments. Two of these moments are simple averages of the instances when either zero miles or bonus miles are accumulated - using $I[\cdot]$ as an indicator variable, these moments are simply $E(I[m_{it} = 0])$ and $E(I[PT_{it} = 2])$. The other moments require input obtained at the individual level - namely the percentage of times an individual collects zero miles $\mu_i^0 = E_i(I[m_t = 0])$, the percentage of times an individual collects bonus miles $\mu_i^b = E_i(I[PT_t = 2])$, the average number of miles collected in a bonus period $\mu_i^{mb} = E_i(m_t|PT_t = 2)$ and the average number of miles collected in a non-bonus period, conditional on a positive number of miles being acquired $\mu_i^{mr} = E_i(m_t|PT_t = 1 \text{ and } m_t > 0)$. The difference between the average miles collected by each individual in bonus periods and the average miles collected in non-bonus, non-zero-collection periods is denoted by Δ_i, so $\Delta_i = \mu_i^{mb} - \mu_i^{mr}$.

Using these individual measures, I can calculate eight additional moments: $Var(\mu_i^0)$, $Var(\mu_i^b)$, $Cov(\mu_i^0, \mu_i^b)$, $\frac{Q_{25}(\mu_i^0)}{Q_{75}(\mu_i^0)}$ and $\frac{Q_{20}(\mu_i^b)}{Q_{80}(\mu_i^b)}$, $E(\Delta_i)$, $Q_{25}(\Delta_i)$, $Q_{75}(\Delta_i)$, where Q_x stands for the x^{th} quantile of the distribution. Since these moments are used in estimation, I need them to be responsive to even small changes in the sample and quantiles typically do not change easily with small changes in the sample (e.g. removing one individual). For this reason I use a ratios of quantiles. Also, for the zero miles individual indicator, μ_i^0, the 20^{th} quantile observed in the data is zero, which renders any information provided by the 80^{th} quantile useless - therefore I used the 25^{th} and 75^{th} quantiles. The eight moments discussed in this paragraph, together with the intercept of regression R1 allow identifying up to 3 distinct segments (each segment being characterized by the probability of being disengaged (q_d) and the probability of collecting bonus miles (q_b) - while the probability of collecting only regular miles is the balance up to 1). The last three moments are especially important for distinguishing between variance in the collection of miles in the regular-only and bonus periods (σ_1 and σ_2) and a more systematic source of variance driven by membership into specific subsegments.

To recap, the auxiliary model that I use for estimation allows me to identify up to 6 distinct latent segments: 2 in terms of redemption propensity x 3 in terms of the propensity to be disengaged, collect regular miles only or collect both regular and bonus miles. These segments are orthogonal - there is no correlation structure between the two factors (redemption propensity and collection type propensity). In other words a redeemer is as likely as a non-redeemer to be in any of the 3 'engagement' level segments.

In theory, the number of latent segments can be expanded by including more moments into the auxiliary model. However, in doing so I encounter two problems. The first (and easier to solve) is computational power, since calculating the optimal effort and redemption for is done separately for each segment. The second and more difficult problem is identification. With more latent segments it becomes difficult to separately identify each segment's parameters and size, because different latent segments become more and more similar to each other. To differentiate among them, even more auxiliary moments are needed.

A case in point is the inclusion of the last three moments ($E(\Delta_i)$, $Q_{25}(\Delta_i)$, $Q_{75}(\Delta_i)$). The model with 6 latent segments has 19 parameters, so in theory 19 informative moments would have been sufficient to identify the parameters in a just-identified model. However, when I conducted an exercise trying to recover known parameters and confirm that the model is identified (details in Appendix A) using only the first 19 moments listed in Table 3.3, I noticed that these were not sufficient. Relying only on them failed to properly identify the 'natural' variance in collection (σ_1 and σ_2) from the shares of different latent segments (Q^I and Q^{II}). $E(\Delta_i)$, $Q_{25}(\Delta_i)$, $Q_{75}(\Delta_i)$ were specifically chosen to help identify these parameters and, as shown in Appendix A, their inclusion renders the model with 6 latent segments well identified.

3.4 Results

Table 3.3 shows the data moments ($\hat{\theta}$), their weights, as well as the simulated moments for 3 different specifications. In all the three specifications there are two distinct segments that differ in their utilities for redemption a_1^r and a_1^{nr}. The difference between the 3 specifications lies in the number of assumed latent segments with different propensities to be disengaged or collect bonus points. The bottom line shows the value of the objective function (equation 4.3.2), the number of parameters used in each specification and the moment and model selection criteria (MMSC) suggested by Andrews and Lu (2001) to select among competing models. The MMSC-BIC criterion is analogous to the BIC criterion used for maximum likelihood estimation; it is computed as

$$MMSC - BIC = J_N - (|c| - |b|)ln(N) \tag{3.13}$$

where J_N is the J-test statistic, (the minimized value of the objective function), $|c|$ is the size of moments vector (22 in this case), $|b|$ is the size of the parameter vector (13, 16 or 19 - depending on the specification), and N is the number of observations. The second term of the equation above is called 'the bonus term' by Andrews and Lu (2001). The idea is that models with less parameters fit the data worse (J_N is larger) so a higher 'bonus term' is subtracted in the case of these models to compensate for the fact that they use fewer parameters. The best model is the one with the smallest MMSC-BIC, in this case the model with 3 distinct segments that differ in their propensities to collect bonus miles or to be disengaged.

Several aspects stand out in Table 3.3. First there is wide variance in terms of the weights attached to the moments. Failing to match closely those moments with higher weights is penalized more heavily in the objective function than failing to match moments with lower weights. The different weights reflect partly different orders of magnitude and partly different signal to noise ratios in the data: those moments that vary little among different bootstrapped samples

Table 3.3: The data moments, their weights and the simulated moments under the assumption of 1, 2 or 3 latent segments in terms of propensity to be disengaged or to collect bonus miles

	Data			Simulation		
	Moment		Weight	Number of latent segments		
				1	2	3
1	$E(I[m_{it}=0])$	0.1813	5,213,973.36	0.0350	0.1846	0.1778
2	$E(I[PT_{it}=2])$	0.3403	3,566,749.65	0.3533	0.3542	0.3452
3	$Var(\mu_i^0)$	0.0474	25,941,859.74	0.0009	0.0365	0.0480
4	$Var(\mu_i^b)$	0.0508	29,394,953.33	0.0060	0.0416	0.0494
5	$E(m_{it}\|PT_{it}<2)$	5.1735	6,658.79	5.6197	5.6104	5.6787
6	$Var(m_{it}\|PT_{it}<2)$	73.9440	1.13	38.5401	81.8024	78.9505
7	$E(m_{it}\|PT_{it}=2)$	28.6410	1,258.38	27.4059	27.9105	27.7448
8	$Var(m_{it}\|PT_{it}=2)$	970.0255	0.04	1002.0170	1093.1159	908.9943
9	R1 intercept	29.5357	5,913.00	28.5268	29.2824	29.4459
10	R1 $log(Count\_back)$	-7.2306	17,984.40	-7.6953	-7.4223	-7.3330
11	$RE(k_{it}\|k_{it}>0)$	138.8113	41.52	141.7780	140.6499	138.3464
12	$E(S_{it}\|k_{it}>0)$	74.1151	11.70	48.0090	69.0921	66.5987
13	R2 intercept	29.4471	1,749.28	30.5328	29.4467	29.6445
14	R2 $log(Count\_fwd)$	-7.7196	3,181.87	-8.7309	-7.8834	-7.9078
15	$Cov(\mu_i^0, \mu_i^b)$	-0.0319	141,159,332.23	-0.0003	-0.0358	-0.0329
16	$\frac{Q_{25}(\mu_i^0)}{Q_{75}(\mu_i^0)}$	0.0759	1,750,820.58	0.0692	0.0702	0.0692
17	$\frac{Q_{20}(\mu_i^b)}{Q_{80}(\mu_i^b)}$	0.2456	479,566.74	0.6922	0.2089	0.2176
18	$E_i(I[NonRedeemer=1])$	0.2642	273,715.34	0.2320	0.3173	0.2747
19	$E(\tau_{it}\|cens=0)$	9.8707	4,509.34	10.4069	10.0782	10.2739
20	$E(\Delta_i)$	19.4621	538.89	21.3075	20.5216	20.5311
21	$Q_{25}(\Delta_i)$	9.8461	453.36	13.5885	13.3534	13.1296
22	$Q_{75}(\Delta_i)$	24.6995	147.63	26.6769	25.2383	25.6036
	objective function			500,314.63	20,476.53	11,096.58
	number of parameters			13	16	19
	MMSC-BIC			500,188.63	20,392.53	11,054.58

have higher weight than those which vary a lot. Secondly, the model which does not include heterogeneity produces moments that are noticeably further from the true moments in the data, than the two models that include heterogeneity.

The models (especially the one that uses 3 latent segments) produces moments that resemble well the true moments observed in the data and the fit is comparable to other models that use a similar approach (e.g. Goettler and Gordon (2011)). The most dramatic difference is for $Q_{25}(\Delta_i)$, the 25^{th} quantile of the individual differences between miles collected in bonus periods and miles collected in non-bonus periods with positive collection; in the real data this value is 9.85, while in the simulation it is 13.13. The reason is that in the simulation, the collection in bonus periods is systematically and consistently higher than in non-bonus periods, while in the real data, there is less consistency to this pattern. That said, the 3 segments model produces moments that are reasonably close to those observed in the data.

The parameters in Table 3.3 were chosen purposefully to match these moments. Another test is to assess how well the model does on other moments, that were not used in the estimation. Table 3.4 shows a sample of 8 such moments. The first one is the average balance across all periods; it is lower in the data than in the simulation (127.8 vs 140.8) and the whole distribution of balances seems shifted to the right in the simulation. A possible explanation for this discrepancy is the fact that in the model I assumed that being a non-redeemer type is orthogonal to the disengaged/ core/ bonus types. Likely, in reality collectors who collect many bonus miles are less likely to be non-redeemers, so those who are in a position to accumulate large balances are also less likely not to spend them. However, when looking at the percentage of times different redemption sizes are redeemed, the simulated data appears in closer agreement to the true data.

Table 3.5 shows the estimated parameters with their standard errors. In order to calculate the standard errors, I first calculated numerically the matrix of derivatives of each moment with respect to each parameter at the value of the parameters that minimize the objective function (the Jacobian, denoted with J). The standard errors are the squared root of the diagonal of the following matrix: $(\Lambda \cdot W \cdot \Lambda')^{-1}$, where W is the weighting matrix used in the estimation. With the exception of the last parameter (segment 3 propensity to be disengaged) all the other parameters have narrow confidence intervals and are more than two standard deviations away

Table 3.4: True and simulated moments not used in the estimation

Moment	In the data	In the simulation
average balance across all periods	127. 8	140.8
Q25 balance across all periods	43	49
Q50 balance across all periods	86	93
Q75 balance across all periods	160	190
% redemptions equal to minimum redemption (95)	72.32	71.31
% redemptions equal to 190	17.24	19.07
% redemptions equal to 285	5.91	5.86
% redemptions larger than 285	4.53	3.76

from 0.

An interesting observation is that the estimated rate of transformation of effort into points (Φ_2) decreases with the number of assumed latent segments, suggesting lower estimates for the points pressure effect in the model with 3 segments. Also, the parameters that quantify rewarded behavior effect (Φ_4 and Φ_5) seem to point to a lower effect in the model with 3 segments, compared to the other two models: the initial effect is smaller and the decay rate of the effect is markedly higher in this model. These findings support the hint expressed when analyzing the descriptive statistics (Section 3.2), that heterogeneity among program participants can be mistaken for positive effects of the program. In Section 3.4.1, I will come back to this aspect and assess whether the implied effectiveness of the program is higher when not considering heterogeneity (1 segment model) or considering it in a more limited fashion (2 segments model).

Looking at the last column of Table 3.5 it appears that almost 47% of the observed participants are 'redeemers' ($p_r = 0.4672$), while the remaining 53% are'non-redeemers'. Orthogonal on this classification, there are 3 other distinct groups in terms of propensities to be disengaged or collect bonus miles. I call the first group, comprising about 15% of the population, the 'disengaged', because almost 64% of the times they are disengaged: do not collect any mile and do not make any redemption. The second group, almost 38% of the population

Table 3.5: Estimated structural parameters

		Number of latent segments		
Parameter	Description	1	2	3
Φ_1	baseline collection	1.3777	1.5771	1.5989
		(0.0011)	(0.0004)	(0.0040)
Φ_2	rate of transformation of effort into miles	1.1420	1.0984	0.8823
		(0.0056)	(0.0106)	(0.0029)
Φ_3	effect of bonus periods on collection	1.3077	1.2368	1.2156
		(0.0027)	(0.0029)	(0.0074)
σ_1	scale parameter; distribution of miles in non-bonus periods	0.4820	0.8533	0.8308
		(0.0004)	(0.0033)	(0.0075)
σ_2	scale parameter; distribution of miles in bonus periods	0.4676	0.6582	0.5768
		(0.0024)	(0.0023)	(0.0025)
a_1^r	utility redeemers	-4.4569	-3.6787	-2.6242
		(0.0235)	(0.0270)	(0.0298)
a_1^{nr}	utility non-redeemers	-8.7220	-8.8437	-6.6646
		(0.0635)	(0.0734)	(0.0330)
p_r	share of the 'redeemers' segment	0.6714	0.6595	0.4672
		(0.0036)	(0.0024)	(0.0033)
a_2	utility non-linear term	3.9030	2.5622	2.4932
		(0.0321)	(0.0171)	(0.0198)
Φ_4	rewarded behavior	1.1425	0.6181	0.5553
		(0.0060)	(0.0046)	(0.0096)
Φ_5	decay in rewarded behavior	1.1295	1.5308	2.4999
		(0.0039)	(0.0206)	(0.0442)
q_b^I	bonus probability for 'the disengaged'	0.3523	0.1161	0.1466
		(0.0003)	(0.0008)	(0.0022)
q_d^I	disengagement probability for 'the disengaged'	0.0238	0.3914	0.6356
		(0.0005)	(0.0012)	(0.0020)
Q^I	share of 'the disengaged' segment		0.3917	0.1520
			(0.0017)	(0.0043)
q_b^{II}	bonus probability for 'the core'		0.5077	0.1476
			(0.0004)	(0.0028)
q_d^{II}	disengagement probability for 'the core'		0.0199	0.1643
			(0.0005)	(0.0018)
Q^{II}	share of 'the core' segment			0.3781
				(0.0097)
q_b^{III}	bonus probability for 'the bonus' segment			0.5706
				(0.0021)
q_d^{III}	disengagement probability for 'the bonus' segment			0.0010
				(0.0009)

is characterized by less extreme behavior - they collect bonus miles in about 15% of the period times and are disengaged in about 16% of the period times; I call them the 'core'. The remaining 47% are referred to as the 'bonus' collectors, because 57% of the times they collect bonus points. The combination between redemption and collection propensities gives rise to 6 different segments that will be analyzed in more detail in section 3.4.1.

In Figure 3.4 I further analyze the distinct effects of heterogeneity, choice of redemption timing and true program effects. All the lines in this figure are based on *simulated data*, generated under different assumptions, as I will explain below. The lines show the expected number of miles collected as a function of time *to* or *from* redemption, obtained by regressing the collected number of miles on the logarithm of the period number to or from redemption. The different lines reflect the different effects of timing selection, heterogeneity and true effects of the loyalty program (points pressure and rewarded behavior). Looking at the left panel of the figure, the 6 thin lines show the collection pattern of each of the 6 latent groups (redeemers/ non-redeemers x bonus/core/disengaged) *when there is no points pressure and no rewarded behavior effect*. All the 6 groups show increased collection towards redemption (the redeemers more than non-redeemers and the bonus redeemers showing the most pronounced effect) which is only induced by timing selection: participants are more likely to make a redemption after a streak of higher than usual collections. This is the opposite of the points pressure effect which states that participants collect more points *in order to make a redemption*.

Furthermore, the black thick solid line of the left panel represents the expected collection for the entire sample - simply aggregating over the 6 latent segments. The slope of this line captures both the time selection and the heterogeneity effects. Because 'bonus' collectors collect more miles due to their propensity to acquire bonus miles, they are also more likely to be closer to a redemption moment at any given point in time. So the increased collection in periods that are closed to the redemption period is also partially given by the fact that high collectors are more likely to always be close to a redemption.

The thick broken line adds the true points pressure effect to the timing selection and unserved heterogeneity effects and the thick dotted line also adds the rewarded behavior effect. The fact that adding the rewarded behavior effect increases the slope of the collection *towards* redemption suggests that serial collectors (those who make frequent redemptions) are benefiting from both effects.

The right panel of Figure 3.4 tells a similar story. It plots the expected number of miles collected as a function of the period number after a redemption, again using a regression of the numbers of miles collected on the logarithm of the period number from redemption. Since post-redemption there is no selection effect, it analyzes only the effect of unobserved heterogeneity as opposed to true program effects. The solid line shows that when there are different types of collectors, those high types are more likely to be ‘right after a redemption’, creating average collections that are higher in periods that are closer to a previous redemption. The broken line shows the expected collection when both the unobserved heterogeneity and the rewarded behavior effect are at work, but there is no points pressure. Finally the dotted line allows for the true rewarded behavior effect to be manifest along heterogeneity and points pressure. This is the steepest line and it shows that the points pressure effect plays a significant role even *after* redemptions: as I will discuss later, about 65% of the points that are collected due to the program are actually collected due to points pressure, as opposed to the rewarded behavior effect. The fact that different latent segments have different propensities to exert points pressure reinforces the effects of the unobserved heterogeneity on the pattern of post-redemption collections.

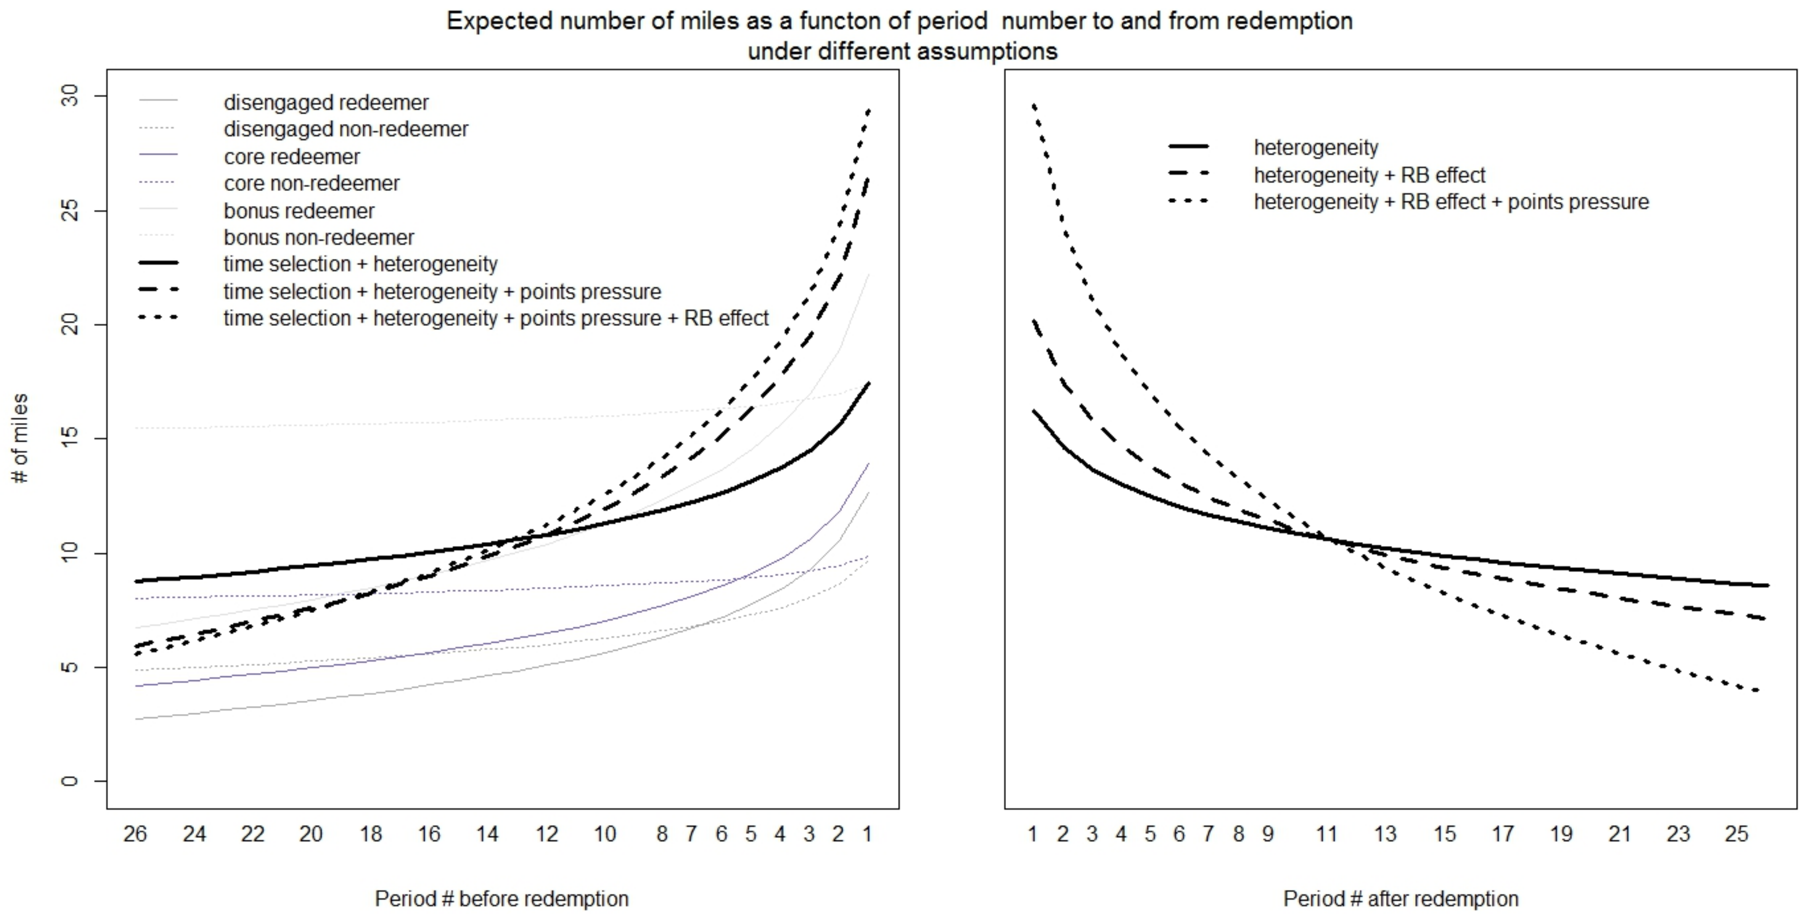

Figure 3.4: Disentangling the effects of redemption timing selection, heterogeneity, points pressure and rewarded behavior effect.

Another interesting implication of the results is the shape of the static utility functions (Figure 3.5), for both redeemers and non-redeemers which I calculated using the point estimates of parameters a_1^r, a_1^{nr} and a_2 taking $\varepsilon = .9$[2]. It appears that for redeemers the average utility derived from redeeming 1 certificate is less than that of redeeming nothing; at redemptions larger or equal to 1 certificate the utility starts to increase linearly in the redemption size. For non-redeemers, even the utility from 10 certificates appears lower than the utility of redeeming nothing.

These are the shapes which best mimic the pattern of redemptions in the data. They suggest that consumers have a fixed cost of making redemptions, which encourages them to keep accumulating miles. As seen in Figure 3.3, participants do not rush to exhaust their available balance when making a redemption and they wait, sometimes a few periods, to make a redemption even though they already have the necessary number of miles. While on average consumers would have higher utilities from redeeming larger amounts, they ask many times the question 'Should I redeem or not ?' before getting to build those high balances needed for larger redemptions; in each period when asking this question, consumers also have idiosyncratic taste shocks for the different redemption sizes which they can afford, so chances are that in one of these instances they will have a relatively large shock for one of the redemption sizes with a low mean utility, so they make this redemption, because in that specific period it offered the highest utility.[3]

Figure 3.6 below shows how the estimated utility functions translated into different probabilities of redemption both for redeemers and non-redeemers. When not enough miles are available, there is no redemption. However, when participants can afford to redeem one certificate, only in 35% of the cases they actually make the redemption and in 65% of the cases they choose to redeem nothing. However, if, for example, a consumer needs 4 periods to build his balance from 95 to 189 miles, and on each period he asks the question 'Should I redeem or not?, there is only an 18% ($.65^4$) chance that he doesn't make a redemption by the time he

[2]A different value of ε would have likely lead to slightly different estimated parameters of the utility function: a_1^r, a_1^{nr} and a_2. Neither ε, nor these parameters are important in themselves - they only serve to produce redemption probabilities that are observed in the data.

[3]This logic is similar to Benartzi and Thaler (1995)'s explanation for the equity premium puzzle: people end up making poor decisions because they make the evaluation to frequently.

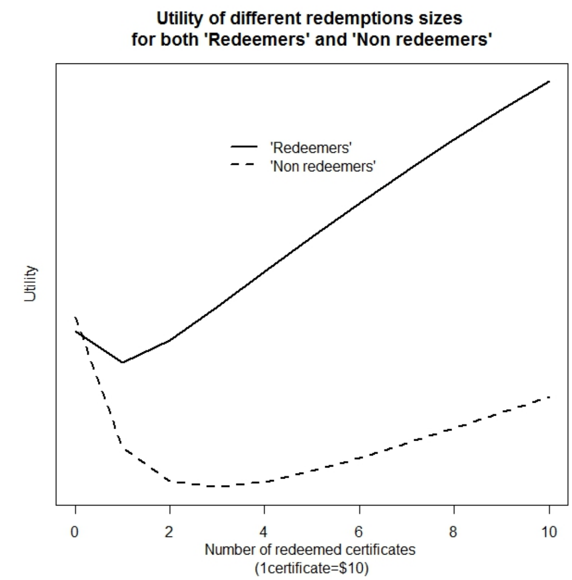

Figure 3.5: The mean utility of each redemption size from 0 to 10 for both 'Redeemers' and 'Non redeemers'.

accumulated 189 miles in his account. These patterns of redemption are very similar among the bonus/core and disengaged types, so here I only plotted the patterns for the core segment.

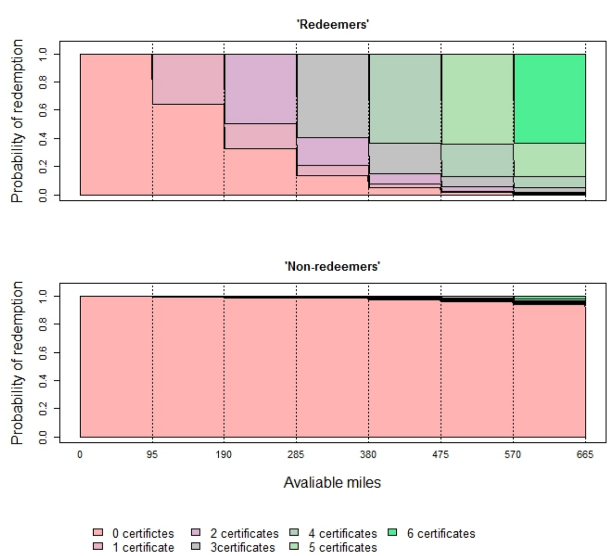

Figure 3.6: Probability that different rewards are redeemed by the number of available miles.

Table 3.6: Fixed parameters

discount factor (δ)	0.9784
effort cost (a_3)	1
effort choice (ρ)	0 to 1 in .2 increments

The outcome of the model is that it specifies the optimal effort to be put into collection as a function of the state variables: balance of miles at the beginning of the period and period type - i.e. periods when only regular miles are available for collection and periods when bonus miles are also available for collection for each of the 6 distinct segments. Unsurprisingly, one finding is that it pays more to put in effort during the bonus periods, than during the non-bonus periods. However, before reaching a stock of 95 miles, it is not optimal to put in any effort - which was again expected, since redeeming 1 certificate brings the least average utility (Figure 3.5). For balances higher than 190 miles, when they barely have enough miles for one redemption size, participants do not put in any effort in the non-bonus periods, but as they approach a new threshold, they are motivated to put in effort, since in this region the utility increases linearly with the redemption size.

It is worth noting that, even if the effort level is constrained between 0 and 1 in .2 increments, it is rarely optimal to put in efforts that are either not the minimum or the maximum. I think this is the case because the stock of miles changes fast relatively to the redeemable amounts. As I will show in the next chapter, when participants need many periods to accumulate miles for a redemption, the optimal policy is a gradual increase in effort.

Table 3.6 shows the parameters that were not estimated. The discount factor (δ) was set to 0.9784 for a two weeks period, a low value compared to the interest rates or discounts used previously in the literature. However, this choice reflects the fact that the discount factor in this problem also captures any risks that consumers associate with the reward program, such as it changing its redemptions policy to make redemption more difficult, or being canceled altogether.

3.4.1 Policy experiments

As explained in Chapter 2, the main interest is in disentangling the miles collected without effort from those collected due to the effort and those collected due the good will generated after a redemption, in order to assess how profitable the program is. The other important goal is to assess how changes in the program that are within merchant's power to implement can affect the profitability of the program.The estimated parameters can be used to achieve these goals. I present the results below.

Program evaluation

By setting to 0 the 3 parameters (Φ_2, Φ_4 and Φ_5) that capture the two effects of the program (points pressure and rewarded behavior effect), I create a counterfactual where the only way in which the loyalty program can influence consumers is through effortless switching. Throughout this section I assume that the effortless switching is zero and equate this counterfactual to the 'no loyalty program' situation. By contrasting this counterfactual with the observed reality, I can gauge the effect of the program. Table 3.7 shows the result of this comparison for each of the 6 distinct segments, where 'R' stands for redeemer and 'NR' stands for non-redeemer. The first two rows show the average number of extra miles that the typical consumer in each segment collects *because of the program* - i.e. for which they spend money they would have not otherwise spent. The last two rows show the percentage of these extra miles that is due to points pressure, as opposed to rewarded behavior effect, obtained by building a counterfactual where participants are only allowed to put in effort to collect more points (points pressure), but there is no rewarded behavior effect.

Several things stand out in Table 3.7. First, as expected, redeemers changed their behavior more than non-redeemers, who are only affected by the rewarded behavior effect. As seen in Figure 3.5, the estimated utility function for non-redeemers is very low for each positive redemption, so for the participants in this segment the rewards are not a strong enough incentive to elicit effort over the observed range of available stock of miles. There is nothing in the model that constrains these participants to abstain from effort. Second, participants increase their collection more in bonus periods than in periods when only regular miles are available, reflecting the fact that putting in effort in a bonus period pays off more than putting effort in a non-bonus period. Thirdly, it appears that for redeemers, who are affected by both forces, points pressure is stronger than the rewarded behavior effect. Taking into account each segment's size and the customers in each segment's different propensities to find themselves into a bonus period, it turns out that 65% of the extra miles acquired due to the loyalty program are actually acquired due to the points pressure effect.

Table 3.7: The effect of the loyalty program on each of the 6 distinct segments

		Disengaged		Core		Bonus	
		R	NR	R	NR	R	NR
Segment share		7.11%	8.11%	17.65%	20.14%	21.95%	25.04%
% of bonus periods		14.67%	14.68%	14.76%	14.77%	57.06%	57.04%
Extra miles	non-bonus period	0.100	0.006	0.490	0.069	1.577	0.174
due to program	bonus period	6.384	0.069	8.314	0.252	14.004	0.521
% due to	non-bonus period	45%	0%	21%	0%	29%	0%
points pressure	bonus period	87%	0%	77%	0%	67%	0%

The results above are expressed in miles, but the real interest is in estimating a monetary value of the program. To do that, I need additional assumptions, because what I observe in the data is miles collected, not money spent by the consumers. There are three assumptions that I need to make: 1) the amount of money that program participants need to spend to acquire a regular, non-bonus mile; 2) the amount of money that program participants need to spend to acquire a bonus mile; 3) the share of regular miles accumulated in a bonus period.

The first assumption is the most straightforward. Since most of the sponsors of the program offer 1 mile for either \$20 or \$30 spent, I make the assumption that a mile collected in a 'regular mile only' type of period is associated with \$25 of sales. Establishing the amount of sales associated with a mile in a 'regular and bonus miles' type of period is more complicated, because the sponsors offer bonus miles for amounts of expenditure that cannot always be easily quantified. For example, Metro (a grocery stores chain) offers 95 bonus miles when spending \$95 or more on Thursdays and Saturdays, while Goodyear offers 100 bonus miles when you spend \$50 or more. There are also examples where the amount spent to obtain a bonus mile may be more clear: Metro offered 5 bonus miles to anyone who buys the Bertolli olive oil in a specified period. The price of this product was \$8.99, which implies that the consumers needed to spend around \$1.80 for a bonus mile. However, it is difficult to tell whether Metro's bonus miles promotion caused consumers to buy oil in the first place (case in which indeed consumers spend \$1.8 to obtain a bonus mile) or it merely determines them to shift to a more expensive product, say \$1 more expensive, (case in which consumers only need to spend \$.2 - the \$1 divided by 5 miles - to obtain a bonus mile). Given the difficulty of establishing the amount of money that consumers need to spend to acquire a bonus mile, I use the following range of values $\{\$.2, \$.5, \$1, \$1.5\}$ and I evaluate the program under all the 4 different assumptions.

Finally, since I do not model separately the bonus and regular miles, I need to make an assumption on the mix of regular and bonus miles that participants collect in a bonus period. The model shows that in an average, non-bonus period, participants collect 23% of the miles they collect in an average period when they also collect bonus miles. However, this result dilutes the value of miles obtained in non-bonus periods, as it includes zero collections. By considering only non-bonus periods with positive collection of miles, it appears that the collection in these periods is almost 30% of the collection in bonus periods. Furthermore, looking at data, where I have precise recordings of each type of miles collected in each period, it appears that 50% of the miles collected in a bonus period are regular miles. The data seems to suggest that bonus miles, instead of crowding out regular miles, actually make participants collect more regular miles.[4] Since I have no strong argument for any particular mix of regular and bonus miles in the bonus periods, again I use a range form 20% to 50% regular miles in bonus periods, in 10% increments.

Table 3.8 shows specific outcomes, for each segment for the particular assumption that a bonus mile is associated with \$.2 of sales and that 30% of the miles collected in a bonus period are regular miles. The additional revenue is calculated using these assumptions and by taking into account the miles acquired by each segment *due* to the program and their different

[4]For each participant I calculated the average number of regular miles collected both in bonus and in non-bonus periods; it appears that on average they collect 3 regular miles more in bonus periods than in non-bonus periods.

propensities to be in a bonus period. The last column, aggregates over the 6 distinct segments taking into account their shares, giving thus the result for the average participant in the program. The second row, the cost, is the cost of providing rewards. The assumption I used here is that each redeemed mile cost 12 cents (9.5 cents for the reward and 2.5 cents for administrative costs of running the program). Also, this row is built using the estimated redemptions (which are in line with the observed redemptions); but the observed redemptions are only about 60% of the observed collection, which gives a breakage rate of 40%. Using the observed redemptions, it appears that over the whole sample, the cost of the program is 4.3% of the sales it generates. A 20% breakage (or non-redemption) rate implies that the cost of the program is 5.7% of the sales it generates; a breakage rate of 10% implies hat the cost of the program is 6.4% of the sales it generates, while 0 breakage implies a cost of 7.1% of sales.

The cost as percentage of extra sales generated is in effect a break-even margin. If the retailer's margin is less than this value (4.3% in this case), then all the additional revenue is not enough to cover the cost of the program, so the retailer's contribution is smaller in the scenario where there is a loyalty program in place then in the scenario where there is no loyalty program. Alternatively, when the margin is higher than the break-even value, the additional revenue is more than enough to cover the cost of the program, so it it makes a positive contribution to the bottom line.

The bottom part of Table 3.8 shows explicitly the revenue in the two scenarios, for each segment. Unsurprisingly, those segments that require the lowest margin for the program to be profitable are those who are most responsive to the program, i.e. whose share of the total purchases that is attributable to the program is the greatest: the bonus redeemers, followed by the core redeemers and then followed by the disengaged redeemers. Overall 13.6% of the current sales are generated by the loyalty program through either points pressure or rewarded behavior effect. This should be a lower bound of the true program effect, as it does not include any switching - cases when consumers effortlessly switch to the retailer who offers the loyalty program.

The share of the revenue generated by the loyalty program appears to be in between the values estimated for two other loyalty programs. Based on the results provided by Lewis (2004), I calculated that that the program he studies (offered by an online grocery), generates

2% of the sales, while the results shown by Kopalle et al. (2012) who study the loyalty program offered by a hotel chain, imply that 24% of the sales are generated by the loyalty program.

To further check the face validity of these results, I assessed the implied effectiveness of the program also based on the results of the models where instead of 6, there were only 2 and respectively 4 latent segments (2 types of redeemers x 1 and respectively 2 types of collectors). The reason for which I did this exercise is that, according to Figure 3.4, sample heterogeneity gives rise to patterns that resemble the points pressure and rewarded behavior effects patterns. This implies that the model(s) where all program participants have the same propensity to collect bonus miles or to be disengaged should give upward biased estimates of the true effect of the program. To see if this is the case, I conducted policy simulations (with the current program vs. in its absence) for the two simpler models, where either everyone has the same probability of collecting bonus miles and being disengaged, either there are only two segments with different propensities to collect bonus miles and be disengaged. As expected, the model which assumed homogeneity implied that the program generates \$44.33 in additional sales (as opposed to \$22.68 for the model with 3 segments), giving a break-even margin of only 2.4%. However, the second model, which assumed 2 segments, was very close to the model with 3 latent segments, implying even a slightly lower additional revenue generated by the program (\$21.9).

The results presented so far use the assumption that a bonus mile is associated with \$.2 of sales and that 30% of the miles collected in a bonus period are regular miles. Tables 3.9, 3.10, show how the break-even margin and respectively the share of revenue generated by the program vary with the assumptions. The break-even margin ranges from 5.64% in the least favorable case, to 2.84% in the most favorable, while the share of revenue generated by the program (Table 3.10) ranges from 12.2% in the least favorable case to 15.65% in the most favorable case. It is worth noting that when it is assumed that regular miles make up a large share of the miles acquired in bonus periods, the results are less sensitive to the specific revenue associated with a bonus mile; in other words, there is less variation in the last column than in the first column of each of the two tables.

One of the advantages of using a structural model is that it allows obtaining such crisp measures of the minimum required margin to sustain a profitable program and the share of

Table 3.8: Detailed results assuming a bonus mile is associated with $.2 of sales and that 30% of the miles collected in a bonus period are regular miles

	Disengaged		Core		Bonus		Average
	R	NR	R	NR	R	NR	
Segment share	7.11%	8.11%	17.65%	20.14%	21.95%	25.04%	
Additional revenue	9.28	0.20	19.82	1.75	77.98	4.14	22.68
Cost	0.45	0.05	0.94	0.19	2.62	0.67	0.98
Cost as % of extra sales	4.9%	23.0%	4.7%	10.9%	3.4%	16.1%	**4.3%**
Revenue with program	71.83	62.73	164.77	146.57	247.92	173.90	166.77
Revenue w/out program	62.55	62.52	144.95	144.82	169.95	169.77	144.09
%revenue due to program	12.9%	0.3%	12.0%	1.2%	31.5%	2.4%	**13.6%**

revenue attributable to the program for those who participate. Moreover it can be used to see how these measures vary with the assumptions (as shown in Tables 3.9 and 3.10). Whether businesses should continue with the program/ adopt such a program depends on their margins. A survey conduced by PricewaterhouseCoopers, a consultancy, and the Retail Council of Canada revealed that in 2013 the median gross margin for a grocery retailer was 30% and the median gross margin for retailers who were neither in groceries nor selling apparel was 41.4% (PricewaterhouseCoopers, 2013). These benchmarks suggest that the loyalty program that I have studied requires a relatively low margin (5.64% using the least favorable assumptions) to be profitable. Therefore, most likely the ARI MILES sponsors that participate improve their profitability considerably as a result.

Policy changes

In this subsection I analyze how changes that retailers could implement affect the efficiency of the program. It is important to note that the previous analysis of outcomes under different assumptions was not a policy change analysis. For example the results above showed that the retailers are better off when a bonus mile is associated with $1.5 of sales than when it is associated with $1 of sales. However, we cannot conclude from here that the best recommendation

Table 3.9: Break-even margin as a function of revenue associated with a bonus mile and percentage of non-bonus miles collected in bonus periods

	Percentage of non-bonus miles collected in a bonus period			
Revenue associated with a bonus mile ($)	20%	30%	40%	50%
0.2	5.64%	4.33%	3.52%	2.96%
0.5	5.48%	4.25%	3.47%	2.93%
1	5.23%	4.12%	3.39%	2.89%
1.5	5.01%	3.99%	3.32%	2.84%

Table 3.10: Share of revenue generated by the program as a function of revenue associated with a bonus mile and percentage of non-bonus miles collected in bonus periods

	Percentage of non-bonus miles collected in a bonus period			
Revenue associated with a bonus mile ($)	20%	30%	40%	50%
0.2	12.20%	13.60%	14.65%	15.47%
0.5	12.35%	13.70%	14.72%	15.51%
1	12.60%	13.86%	14.82%	15.58%
1.5	12.84%	14.02%	14.93%	15.65%

for the retailer is to give out less bonus points in order to increase the value of sales associated with 1 bonus point (e.g. offering 1 bonus point instead of 5 for the Bertolli oil), because this change will likely affect the pattern of collection and implicitly the effectiveness of the program.

There are two types of policy changes that I analyze. The first one changes the multiple of miles that consumers can use to make a redemption and adjusts the size of the reward to maintain the same 'rewarded dollars'/mile ratio. The current program allows consumers to make redemptions in multiples of 95 miles and offers \$10 for each multiple of 95. This set up can be changed to offer either \$5 for each multiple of 48 miles, \$15 for reach multiple of 145 miles or \$20 for each multiple of 190 miles (I adjusted slightly the multiples, so that the multiples are round, easy to remember numbers). The second change is related to the 'rewarded dollars'/mile ratio; it keeps fixed the 95 miles multiple but assumes that this gives consumers either \$9, \$11 or \$12 in rewards.

Allowing consumers to make redemptions at every multiple of 48 miles, effectively discourages points pressure. With the current program, the points pressure starts to kick in at a stock of 92 miles in bonus periods at a stock of 171 miles in non-bonus periods. By giving the consumers the opportunity to redeem sooner, many of them will do so, therefore they will no longer be able to build the balances at which it becomes optimal to put in effort, this being especially true for redeemers, who put in effort in the first place. The non-redeemers are slightly advantaged, as they too redeem more often and thus benefit more from the rewarded behavior effect, but this increase is not enough to compensate for redeemers putting in less effort. As shown in Table 3.11, this policy generates \$164.54 in revenue for the average participant on a two weeks period, while generating costs of \$1.06 per participant, per period. In this set up, the revenue is smaller than the revenue with the current policy (\$166.77), and the cost is larger than the cost of the current policy (\$0.98), so there is no reason for which it should be adopted.

One may think that I should allow the consumer to redeem only at higher multiples. I conducted the exercise by allowing consumers to redeem every multiple of 145 miles for \$15 dollars. The results showed that this policy decreases revenues even further. The reason is that, according to the estimated utility function (Figure 3.5), the average utility of redeeming \$15 is still lower than the average utility of making no redemption, so program participants are not motivated to put in effort in a bonus period until they have build a stock of 137 (as opposed to 92 with the current program) and 222 in the no-bonus period (as opposed to 171 currently). Furthermore, the fact that the number of miles required to make a redemption increases implies that there are less redemptions per unit of time, therefore less rewarded behavior effect. This policy is estimated to bring a revenue of $158.67 per participant, per period at a cost of $.80 per participant per period. Compared to the current policy, it appears that this alternative policy would be better at margin rates less than 2.2%. But at this level of the margin, no program is actually the best alternative.

Increasing the multiple even more (from 145 miles to 190 miles) with a corresponding increase in the reward from $15 to $20 improves the outcomes, but not by much. The reason is that, since now it takes longer to build the stock for a minimal redemption, redeemers only start to put in effort at a stock of 140 miles in the bonus periods and 284 miles in regular miles only periods. More importantly, since on the one hand redeeming 190 miles brings more utility than redeeming nothing, and on the other hand waiting to accumulate 380 points is not appealing, in this scenario the chances of redeeming the minimum allowable quantity increase significantly. Thus participants have less chances to put in effort. The sales per the average program participant in this set-up are $161.69 and the costs $.82. Compared to the current policy, it appears that this alternative policy would be better at margin rates less than 3.15%. But at this level of the margin, no program is actually the best alternative.

From the analysis of the 3 policy changes above it appears that AIR MILES has chosen the best redemption multiple. However another change that could be implemented is in the amount of money offered for a 95 miles redemption. I evaluated changes to 9, 11 and 12 dollars. As expected, offering $9 instead of $10 reduces the effort that participants put into collection, thus the average revenue per participant decreases from $166.77 to $156.25. The redemption rates under this alternative program change only slightly form the redemption rates with the current program; however, the cost is reduced from 12 cents per redeemed mile to 10.8 cents. This program would be better than the current program at margins lower than 2.17%, but at this level no program at all is the best option.

Increasing the value of the reward to either $11 or $12 increases both the revenue and the cost, as participants put in more effort and the cost, as the cost of a redeemed mile increases.

These programs are better than the current one only at higher margin rates - 8.25% and respectively 7.72%. These results imply that for margins less than 7.72% retailers should prefer the current version of the program, while for margins higher than this threshold, they should implement a program which rewards consumers with $12 for each multiple of 95 miles. It is interesting to note that the break-even margin between a policy with $11 in rewards and a policy with $12 in rewards (not shown in Table 3.11) is 7.3%. When choosing only among these two policies, at rates below 7.3%, rewards of $11 produce higher contribution. However, when $10 is also an option, at margins of 7.3% this is the policy which maximizes contribution.

Table 3.11: Policy changes - summary of findings

Policy change		Revenue ($)	Cost ($)	Break-even margin vs. no program	Break-even margin vs. current program
	No program	144.09			
	Current program	166.77	0.98	4.33%	
Redemption	48	164.54	1.06	5.16%	-
multiples (miles)	145	158.67	0.80	5.51%	2.20%
	190	161.69	0.82	4.67%	3.15%
Value	9	156.25	0.75	6.20%	2.17%
of	11	168.20	1.10	4.57%	8.25%
reward ($)	12	169.97	1.23	4.75%	7.72%

In the analysis of the policy changes I assumed that the rewarded behavior does not change with changes in the amount of reward per redeemed mile and that in their utility consumers actually care about the amount of money they receive for each redemption, not the number of certificates redeemed - e.g. consumers derive the same utility from redeeming one 95 miles certificate which rewards them with $10 as from redeeming two 48 miles certificates which reward them with $5 each. I also assumed that changes in the structure of the rewards do not determine any participants of the 'redeemer' type to become non-redeemers or vice versa. Finally, I assumed that there are no competitive reactions.

3.5 Discussion

In this chapter I developed a model for loyalty programs where consumers accumulate frequent small amounts of points, which can be used for relatively small in store discounts at many participant stores. A main feature of this context is that ceiling or idiosyncratic effects (not needing the reward in a given period) can be safely excluded. Thus I think that the most plausible explanation for the observation that sometimes, even when engaged with the RP, consumers fail to make redemptions or make only partial redemptions is that the loyalty points are kept in a separate mental account - i.e. consumer's regular currency and the loyalty points are not fungible (Shefrin and Thaler, 1988). Otherwise consumers should be spending the LPs as soon as they have the opportunity and save up the regular money which is in fact an universal means of exchange and on which they can earn interest. Moreover, consumers have a propensity to spend spend loyalty points on indulgences (Kivetz and Simonson, 2002). The separate mental accounts and the preference to use miles in order to purchase hedonistic items for which consumers have different taste shocks in each period provide a parsimonious explanation for the observed patterns in the data, where sometimes consumers postpone their redemptions or make redemptions smaller than their available amount. I translate this set of assumptions in my model by allowing the utility of each redemption size in each period to be influenced both by a systematic component, as well as a random utility shock (ϵ). In periods with low shocks for redemption consumers won't redeem, while in periods with a large shock for a low redemption size they will make the small redemption even though they may have enough miles for a larger redemption.

The model comprises both the collection and redemption aspects. As discussed above, the redemption of rewards is driven by both a systematic and a random component. The collection of miles is driven by a baseline factor (miles corresponding to purchases that would have been realized even in the absence of the program or when only effortless switching is at work) and by the two effects of interest: points pressure and the rewarded behavior effect. These components introduce dynamics into the model, as the two effects are are stronger around the redemption time.

To estimate the model I used data provided by AIR MILES comprising only participants who opted to collect all their miles into the'cash' account that can only be used for in-store discounts. Given its complexity, I used indirect inference, an estimation method based on simulation, in order to pin down the structural parameters. The model achieved very good results in terms of fit, similar to other models that use a similar estimation method - e.g. Topa (2001); Hennessy and Whited (2007); van der Klaauw and Wolpin (2008); Goettler and Gordon (2011).

The results of the estimation showed that consumers exhibit different collection and redemption patterns. In terms of collection, post-hoc I labeled the two groups 'redeemers' and 'non-redeemers', while in terms of collection (more precisely propensity to collect bonus miles or be disengaged), I labeled the three groups 'bonus', 'core' and 'disengaged'. As shown in Figure 3.4, heterogeneity can be confounded with true program effects, so it needs to be teased apart, in order to be able to evaluate the program accurately. The other effect that I account for for in order to obtain accurate evaluations of the program is the fact that even in the absence of points pressure, consumers are more likely to make a redemption in a period with a large collection of points that is due only to chance.

I use the estimated parameters to quantify he changes in sales induced by the program. However, translating the changes in miles collected to changes in money spent requires further assumptions. For this reason Tables 3.9 and 3.10 present ranges of possible outcomes, rather than a single outcome in terms of the minimal contribution margin needed to sustain a profitable RP in its current design and respectively the share of sales that are attributable exclusivley to the loyalty program.

Based on the counterfactual analysis (and using the assumption of a bonus mile being associated with $.2 of sales and 30% of the miles collected in a bonus period being regular), it appears that the current program is more profitable than a situation with no loyalty program at all at margin rates higher than 4.33%. At margin rates higher than 7.72% offering $12 instead of $10 for the same amount of redeemed miles should further boost profitability. Furthermore, it appears that for the average program participant 13.6% of the sales are generated *because of the RP*. Without knowing what percentage of customers are members of the RP, it is difficult to estimate the role that the program plays in retailer's overall sales.

In Chapter 5 I further elaborate on the contributions of this model, its limitations and future avenues of research in this area, but before that, in the next chapter I turn to analyzing the 'dream rewards'.

Chapter 4

Dream rewards

The previous chapter focused on the particular case when rewards are homogeneous in-store discounts or gift certificates. However, many loyalty programs offer consumers rewards that are much more varied in terms of size and also more substantial. This chapter presents a model that applies to this type of programs. The main conceptual difference between the two setups is the way participants think about redemptions. In the first case consumers manage an on-going process where they balance the impulse to use their miles for a redemption in the current period and the tendency to keep accumulating the miles so that they can obtain a larger, more meaningful redemption in the future. In the other case consumers do not decide in every period whether to use their points for a redemption or not and how large should the redemption be, but they form redemption goals (Kivetz et al., 2006) and try to achieve those goals - for example they might decide that they'll use their points to buy airline tickets to fly home for Christmas. The size and nature of the rewards dictates how often consumers make the redemption decision. It is likely that for small treats, consumers decide in every period whether they will have the treat or not - i.e. whether they will spend their points to obtain the treat or not, but for large rewards, consumers make decisions less frequently and stick to those decisions over longer periods. In the case of programs that offer in-store discounts, consumers are likely to ponder on every shopping trip on whether they should use their available points to buy, for example, ice cream or not. With programs that offer larger rewards this is not likely the case. Building on the previous example, if the consumer has decided that he or she will use the points to fly home for Christmas, it is unlikely that every other week or so, the consumer in my hypothetical example will ponder on whether to use them for a gift, to buy furniture or for a cruise.

So the key aspect of the model proposed in this chapter is that consumers form specific redemption goals in terms of *reward size* and *timing*. Participants look forward to the moment when they make that redemption; beyond the redemption time, they know that they will form new redemption goals, so they still care about the points that they have in their account (as these will be used for those future goals), but they are not aware of the specifics of any future redemption beyond the closest one. Again using the example above, if the airline tickets costs 10,000 points the participant would be happier to have 7,000 still left into his or her account than to have nothing left; once the tickets have been purchased, the consumer will start thinking about a new redemption and having 7,000 points offers better prospects than having no point.

The estimable version of the model that I present in this chapter adds the amount of points left at the end of a redemption period to the flow utility of the program participants, emphasizing that consumers treat loyalty points as if they have intrinsic value (Hsee et al., 2003; van Osselaer et al., 2004). In other words, at the end of a redemption period having points left into one's account increases utility because these points allow participants to have a head start in their collection for the next goal. Moreover, having available points into one's account is valuable because they allow participants to entertain the thought of future redemptions and thus provide anticipatory utility for the reward to be redeemed (Loewenstein, 1987).

Another important assumption of this model is that, whenever consumers fail to attain their goal (they do not have enough points to cover the cost of their desired reward), they will use all the available points they have and make up for the difference using cash. This is in agreement with the observation that most loyalty programs allow consumers to pay for the rewards with a combination of points and cash (Drèze and Nunes, 2004).

To summarize, the main difference between the new model described in this chapter and the model presented in Chapter 3, is the way consumers think about the rewards, given their different scales. This difference, which might appear minor, actually triggers substantive differences in the way the model plays out. Given that the rewards in this set-up are larger and more hedonistic, I call this model the 'dream rewards' model to emphasize the difference from the previous model where the rewards were meant to be spent pretty much like cash.

In this chapter I start by laying out a theoretical model of points collection and redemption.

The points collection is influenced by the same type of factors that were outlined in the model for 'cash' collectors: a baseline rate of collection plus the influence of the RP through points pressure and the rewarded behavior effect. The points pressure introduces dynamics to the model, as effort 'today' becomes more and more valuable as the collector gets closer to the redemption moment. One of the important features of this model is the fact that the inter-redemption times are given and consumers choose only the size of the reward that they aim to redeem. At the end of a redemption period consumers care about the balance of points that they have left into their accounts since this balance allows them to make other redemptions into the future. The theoretical model that I outlay in Section 4.1 is likely to be difficult to estimate due to a combination of computational intensity, data requirements and identification issues. Therefore, in Subsection 4.1.1 I discuss three ways in which the model is simplified in order to be able to estimate it.

Section 4.2 describes the AIR MILES data used in the estimation. Section 4.3 is divided into two subsections. In Subsection 4.3.1 I discuss the identification strategy and bring arguments that the model is well identified - i.e. the data used in the estimation is rich enough such that only one set of parameters can best fit it. In Subsection 4.3.2 I revisit indirect inference, the estimator that I used for the model in Chapter 3.

The results are presented in Section 4.4. There are two sets of results that I examine. The first set was derived by applying the model in Subsection 4.1.1 on the whole sample, while the second set was derived by applying the same model to 3 different sub-samples, that were obtained by a priori partitioning the data using k means analysis. Contrary to my expectations, the two sets of result imply the same level of profitability of the RP. I further used the estimated parameters to assess how much the efficiency of the program would improve if the program provider managed to persuade participants to decrease their inter-collection times. Section 4.5 concludes with a discussion of the main features of the estimated model and of the results.

4.1 Theoretical model

In this model I consider τ the timing between two consecutive redemptions (the inter-redemption time, or spell) which is driven by observable individual characteristics Z_i (for example whether the participant agreed to receive promotional emails or not). t denotes the amount of time that is left out of spell τ until redemption, so t takes values between 0 and $\tau - 1$. For example if there are 30 periods between two consecutive redemptions, then $\tau = 30$; in the redemption period $t = 0$, one period before the redemption period $t = 1$ and so forth.

Equation 4.1 shows the law of motion for t, where t' denotes the value that t will take in the next period. It shows that in every non-redemption period, the next period's t is decreased by 1. After a redemption period at the end of spell τ (i.e. when $t = 0$), there is a new redemption spell τ', so the next period's t becomes $\tau' - 1$.

$$t' = \begin{cases} t - 1 \text{ for } t \in [1, \tau - 1] \\ \tau' - 1 \text{ for } t = 0 \end{cases} \tag{4.1}$$

The first step is to use data to estimate a survival model for the inter-redemption spells. Depending on how frequently participants make redemptions, long observation periods may be needed in order to obtain reliable estimates for the distribution of τ. Should this data be unavailable, I describe a simplified version of the model (see section 4.1.1) that circumvents this step. But first I will lay out the complete model.

While the inter-redemption spells are exogenous to the structural model (hence the need to estimate their distribution from the data), the redemption goals themselves (denoted by G^*_τ and expressed in number of points) are a choice variable. I assume that the goals are distributed $\Gamma(k, \theta)$. The important aspect of this choice is that it is only made when joining the program or after a redemption. Subscript τ indicates that the choice of the goal G is made at the beginning of the spell τ and superscript $*$ indicates that this choice is optimal. In all the other periods G^*_τ is fixed, so this is why in equations 4.7 and 4.9, as I will discuss later, it appears as a state variable.

As described above, G^*_τ is optimal from the point of view of a consumer for whom only the value of the reward matters, not the specific content of the reward - i.e. all it matters is that the reward is worth 10,000, not that it consists of flight tickets or a cruise. However, in reality participants in reward programs are likely to have strong preferences for the particular type of reward that is being redeemed and the offer of rewards is finite, so they will end up making trade-offs between, on the one hand, goals that are a good deal from the point of view

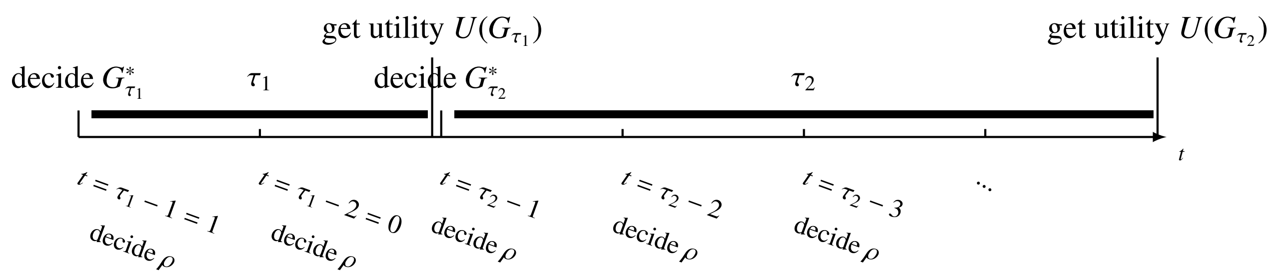

Figure 4.1: Sequence of spells. The goal (G*) is chosen at the beginning of each spell. Effort is chosen in each period. The utility of redemption is realized at the end of the spell.

of redemption spreading over time and matching of goals to existing stock of points and on the other hand goals that are appealing to them for idiosyncratic reasons. I model this compromise by stating that the chosen goal is distributed according to a normal distribution around the optimal goal G^*_τ. The dispersion around G^*_τ is captured by σ_ξ - a parameter to be estimated.

Effort put into collection is the other choice variable. It is contingent on the time left from the current spell and is decided in every period. By putting in more effort, the participants can collect more points and increase their chances of being able to attain their goals and if the goals are achieved, they increase the stock of points which will enable them make new, perhaps larger redemptions into the future. But putting in effort is assumed to be costly, because it means that participants would change their shopping routine, the routine being the course of action which was precisely the least costly for them. I denote the cost of effort by $c(\rho)$. Equation 4.2 below shows the particular functional form for this cost: when $\rho = 0$, the cost is also 0, but as effort increases, the cost increases exponentially.

$$c(\rho) = e^\rho - 1 \tag{4.2}$$

Figure 4.1 shows a hypothetical example of two redemption spells (τ_1 and τ_2). The key aspects of this figure are the fact that the goals are chosen only at the beginning of each spell, while the effort to be put into collection (ρ) is chosen in every period. The benefits of redeeming a reward are accrued only at the end of an inter-redemption spell.

I assume that the number of points collected in period t (m_t) is distributed log-normal with scale parameter σ and location parameter μ.

$$m_t \sim Log - normal(\mu_t, \sigma) \tag{4.3}$$

There are three components that affect μ_t. One is the baseline collection (denoted by ϕ_0), which captures the points collected for amounts of money that would have been spent with the focal retailer (or coalition of retailers) even in the absence of the loyalty program. In practice ϕ_0 can vary with specific consumer characteristics. In the specific data that I am using, an important factor which affects the points collection is whether the program participant has a credit card *linked to the program* through which they can collect points. Since consumers can make a large share of their purchases using credit cards, those who collect points with their credit cards will have overall higher collection patterns. Credit cards linked to the RP program allow participants to 'double-dip': they collect points both for the purchases they make and for the fact that they use the credit card linked to the program to pay for those purchases. I denoted the baseline collection for credit card holders with ϕ_0^{cc} and the baseline collection for non credit card holders with $\phi_0^{\overline{cc}}$.

The second component which affects μ_t is ρ_t - the effort that participants put in to change their behavior and collect more points due to the points pressure mechanism (Kivetz et al., 2006; Blattberg and Neslin, 2008). This effort is transformed into points at a rate ϕ_1 - a parameter to be estimated. These two components form the anticipated component of μ_t. Besides these components, there is also an unanticipated component, that stems from the rewarded behavior effect (Rothschild and Gaidis, 1981; Blattberg and Neslin, 1990, 2008): $\phi_2 e^{\phi_3 u_t}$, where ϕ_2 and ϕ_3 are parameters to be estimated (ϕ_2 expected to be positive and ϕ_3 expected to be negative) and u_t is the amount of time that has passed since consumer's last redemption at time t. This parametrization captures the idea of the rewarded behavior effect: right after a redemption consumers are expected to have a boost in collection; this boost decays gradually as time passes from redemption.

$$\mu_t = \underbrace{\phi_0 + \phi_1 \rho_t}_{\text{anticipated}} + \underbrace{\phi_2 e^{\phi_3 u_t}}_{\text{unanticipated}} \tag{4.4}$$

Redeeming points for rewards generates utility in the redemption period. Equation 4.5 shows a simple specification for the utility function. S_0^i denotes the stock of points that is available at $t = 0$, i.e. in the redemption period, before the redemption actually takes place. In general, I use the subscript t on the stock of points, (in this case 0), to describe the stock of points at the end of period t. Here I added the superscript i to suggest an 'intermediary' stock, i.e. the stock after all the points have been acquired in period 0, but before the reward was redeemed. In other words $S_0^i = S_1 + m_0$.

According to equation 4.5, if the participant has enough points to achieve his goal, the utility is given by the value of the goal multiplied by α_1 - a parameter to be estimated. However, if the stock of points is lower than the goal, the participant is assumed to use all the available stock to acquire a reward and supplement the difference with cash; the utility in this case is given by $\alpha_2 S_0$. The key insight here is that α_2 should be less than α_1, showing the fact that, even at similar amounts of points spent, being able to achieve one's goal by only using loyalty points is better than being able to cover only partially the cost of the reward.

$$U(G_\tau) = \begin{cases} \alpha_1 G_\tau, & \text{if } S_0^i \geq G_\tau \\ \alpha_2 S_0^i, & \text{if } S_0^i < G_\tau \end{cases} \tag{4.5}$$

Equation 4.6 below shows the law of motion for the stock of points. S_τ is the stock of points available right before the spell that leads to a redemption starts. Right before joining the program, participants don't have any points, so this quantity is 0 for their first spell. However, if a previous reward has been redeemed, the quantity of points available right before the beginning of a new spell is the quantity that was available at the end of the previous spell (S_0). The second branch shows that the stock at the end of a redemption period, S_0, is 0 if the participant could not achieve their goal by just using points, or whatever is left after the redemption, if the goal was achieved. The last branch shows that in all the periods that follow a non-redemption period, the new stock of points is simply the sum between the the stock of points at the end of the previous period and the points acquired in the current period.

$$\begin{cases} S_\tau = \begin{cases} 0, \text{ if } \tau \text{ is the first spell since joining the program} \\ S_0, \text{ otherwise} \end{cases} \\ S_0 = max(0, S_1 + m_0 - G_\tau) \\ S_t = S_{t+1} + m_t \text{ for } t \in [1, \tau - 1] \end{cases} \tag{4.6}$$

Equation 4.7 shows the utility of the consumer in the redemption period. The 3 state variables that characterize the utility in every period are 1) the number of points at the beginning of the period; 2) time left until redemption; and 3) the redemption goal. At the beginning of a redemption period ($t = 0$), the consumer has S_1 points available into their account, i.e. the same stock that was available at the end of the previous period. I make the assumption that the rewards are redeemed at the end of a redemption period, so the consumer can use towards redemption all the points acquired in the redemption period.

The number of points acquired in the current period, right before redemption, ($m(\rho_0)$) depends on the 3 effects outlined in equation 4.4, but consumers can only affect one of the components, namely the effort. This choice affects in expectation the number of points that will be acquired ($E_{m^a(\rho_0)}$), which in turn affects the utility. Moreover, consumers only count on the anticipated points (see equation 4.4), hence the a superscript. So by choosing an effort level ρ_0, consumers incur the cost $c(\rho_0)$, where $c(\cdot)$ is increasing in the level of effort. But higher levels of effort are also likely to increase the utility $U(G_\tau)$ and also the the value of the prospects after a redemption as captured by V.

$U(G_\tau)$ is the utility from redemption described in equation 4.5. $E_{\tau'}$ is the expectation over the distribution of the inter-redemption spells and E_ξ is the expectation over the distribution of the deviations from the optimal goal. As discussed above, the idea here is that participants make a compromise between selecting on the one hand goals that are attainable at reasonable levels of effort and that provide steady inflows of utility and on the other hand goals that are appealing to them personally. For example for someone who already has 45,000 points it would make sense to pursue the goal of redeeming for a cruise that costs 50,000 in a one year's time, but they may prefer to use their points to buy a jewelery that only costs 20,000 points. In this particular case ξ, the idiosyncratic shock for the size of next period's reward, would be negative.

$$V(S_1, 0, G_\tau) = \max_{\rho} -c(\rho_0) + E_{m^a(\rho_0)}\Big\{U(G_\tau) + \delta E_{\tau'} E_\xi V(S_0, \tau' - 1, max(G^*_{\tau'} + \xi, G_{min}))\Big\} \quad (4.7)$$

In the equation above the utility at the end of an inter-redemption spell is given both by the utility of the redemption being made ($U(G_\tau)$) and by the prospects that the participant has, as captured by the value function $E_{\tau'} E_\xi V(S_0, \tau', max(G^*_{\tau'} + \xi, G_{min}))$, discounted by the per-period discount factor δ. A larger stock of points at the beginning of the new inter-redemption interval (S_0) leaves the participants better positioned to pursue higher, more rewarding goals, so V is increasing in S. However V is not monotonic in τ'. The reason is that while indeed lower τ's mean shorter waiting times before the next redemption, it may also mean less time available to work towards one's goal and thus decreased chances of achieving it, which, as shown by equation 4.5, comes with a penalty. The optimal goal is chosen at the very beginning of a new inter-redemption spell, as shown in equation 4.8. However, since the idiosyncratic shock for the size of next period's reward (ξ) can be negative, the next period's actual goal will be the maximum value between $G^*_{\tau'} + \xi$ and the minimal redeemable amount G_{min}.

$$G^*_{\tau'} = \arg\max_G V(S'_\tau, \tau', G) \quad (4.8)$$

In all the other periods, when no redemption takes place, i.e. for $t \in [1, \tau - 1]$, consumers do not derive any direct utility. They still need to decide the level of effort that they put into collection (ρ_t). Higher levels of effort are costlier, but the collector may be better 'tomorrow' or in the period when he will eventually make the redemption that brings utility (equation 4.9).

$$V(S_{t+1}, t, G_\tau) = \max_{\rho} -c(\rho_t) + \delta E_{m^a(\rho_t)}\{V(S_{t+1} + m_t, t - 1, G_\tau)\} \quad (4.9)$$

It is worth noting that the redemption goal does not change after a non-redemption period (G_τ in equation 4.9), but only after a redemption ($G^*_{\tau'} + \xi$ in equation 4.7). The other two state variables, (time to redemption and the stock of points) also have different laws of motion depending on whether a redemption takes or does not take place in the current period, as shown by the branches of equations 4.1 and 4.6.

For any given set of parameters, V is solved for through a combination of value function iteration and backward recursion. Start with a guess for V, call it V_0, determine $G^*_{\tau'}$ and plug it into the right hand side of equation 4.7, to obtain a new value for V (V_1) at $t = 0$, i.e. at the redemption time. Then, use equation 4.9 recursively to calculate V_1 at $t = 1$, then at $t = 2$ and so forth until $t = \tau_{max} - 1$, until V_1 is calculated on the whole domain. Use V_1 the same way V_0 was used to obtain V_2. Stop when the distance between V_k and V_{k+1} is very small.

4.1.1 An estimable model for dream reward collectors

In this subsection I describe a simpler version of the model above, which I can estimate with the data that I have. There are 3 aspects which I simplify: 1) the need to estimate a distribution for the inter-redemption times; 2) the need to estimate σ_ξ - the variance of redemptions due to idiosyncratic preferences for certain types of rewards and 3) the assumption that participants are always optimizing their collection. I discuss them in turn.

As stated before, in order to obtain a reliable estimation of the distribution of the inter-redemption times, one needs a data set that stretches over a long period of time. With rewards being expensive, the program participants need time to accumulate the required number of points, so the number of observed redemptions may not be enough to accurately estimate the distribution of the inter-redemption spells. Moreover, some of the observed inter-redemption times will be either interval or right censored. Interval censored observations are those where a redemption is being observed after t_0 periods since the beginning of the observation period. Knowing for how long the consumer has been enrolled in the loyalty program at the moment of the redemption (T), one can only say that the inter-redemption spell up to the redemption time is between t_0 and T - hence the interval censoring. After the observed redemption, if no other redemption is observed until the end of the observation window (t_w), the next inter-redemption spell is also censored; all one can say is that this second spell is larger than $t_w - t_0$.

In the simplified version of the model, consumers no longer form expectations on the length of their next inter-redemption spell. Instead, each new spell represents a new ‘game’ and the spells used as input into the model are the spells actually observed in the data. A key feature of the more complex model above was that in a redemption period utility was derived both from the reward being redeemed and from the future prospects that the participant had after the redemption, as captured by the value function V. This key feature is still retained in the simple version, but instead of capturing the value of the prospects with V, I capture it by changing the utility function as shown in equation 4.10. Since V was increasing in S, in the new specification, the flow utility was amended to include a second term that captures all the future benefits that the remaining stock of points offer. α_2 is a parameter to be estimated, expected to be positive, since more points left after a redemption leave consumers in a better position to achieve future goals. I take logarithms on both branches of equation 4.10 because additional remaining points after making a redemption are expected to provide decreasing benefits. In other words, consumers want to have *something* left for future redemptions; the difference between medium and very large stocks of points available at the end of a redemption period does less to improve one's prospects than the difference between no mile left into one's account and a few hundred points. Also, when consumers fail to achieve their goal (the lower branch of equation 4.10) any additional mile spend towards the goal that still fails to cover the cost entirely does little to improve one's utility. ε is a positive constant that allows taking logarithms at 0 stocks of points.

$$\mathbf{U}(G_\tau) = \begin{cases} \alpha_1 G_\tau + \alpha_2 log(S_0^i - G_\tau + \varepsilon), & \text{if } S_0^i \geq G_\tau \\ \alpha_2 log(S_0^i + \varepsilon), & \text{if } S_0^i < G_\tau \end{cases} \tag{4.10}$$

I modify the value function in the redemption period accordingly, as shown in equation 4.11. Equation 4.9 remains the same, since the problem remains the same in the non-redemption periods. As the flow utility now includes the residual component that was before associated with V, V no longer shows up in the right hand side of equation 4.11. In this new specification the value of the future beyond the redemption time is concentrated into a parameter, α_2, and an observable quantity, $log(S_0^i - G_\tau + \varepsilon)$. Therefore the infinite horizon problem described earlier is transformed into a finite-horizon problem - each redemption is a separate 'game' now. This problem is easier to solve, as it only requires backwards recursion, starting from $t = 0$ up to $t = \tau - 1$.

$$V(S_1, 0, G_\tau) = \max_{\rho} -c(\rho_t) + E_{m(\rho_t)}\{\mathbf{U}(G_\tau)\} \tag{4.11}$$

To recap, the dynamic problem has now been recast as a finite-horizon problem, from $t = \tau - 1$ to $t = 0$, as shown in Figure 4.1. The result of this problem is a policy function which specifies the optimal action - in this case effort ρ. The action of choosing the reward is also subject to some optimization, but this is done outside the frame of the dynamic problem. The reward from redemption is realized at the end of the inter-redemption spell and the particular functional form is presented in Equation 4.10. The costs associated with effort are incurred in each period and the functional form is shown in equation 4.2.

The second aspect which I simplify is the choice of the reward. In the previous specification it was assumed that the chosen reward is a compromise between the best reward in terms of size and best reward in terms of 'fit', where fit is given by unobserved, idiosyncratic preferences for one reward or another. This required the estimation of the dispersion of the chosen reward around the optimal reward (σ_ξ). In the simplified version, I replace this parameter with an assumption. The reason is that σ_ξ might be difficult to identify in the absence of information on how consumers actually choose the rewards they redeem. Typically marketers have data on the number of points acquired and number of points redeemed in each period by each participant, with little insight into the driving forces behind these actions.

One idea is to regress the number of points redeemed on the available stock of points at the beginning of the inter-redemption spell, the duration of the spell and the square of the duration of the spell and use the mean square error (MSE) of this regression to identify σ_ξ. The logic is that if there is very little room for idiosyncratic preferences to play a role in choosing the reward, then participants who have a similar stock of points and a similar spell of time before the redemption should choose rewards that are quite similar in terms of size; this gives a small MSE. On the other hand, if individual preferences for particular types of rewards play a significant role, then we should expect to see a high dispersion of rewards (high MSE) for participants with similar prospects at the time when the goal is adopted.

Another idea, which I use, is to assume that goals have a chance of being selected proportional to their value function. A consumer who has S_τ points available and is $\tau - 1$ periods away from a redemption can select among N_G available goals. The probability of selecting goal G_i

is given by

$$p(G_\tau = G_i) = \frac{e^{V(S_\tau,\tau-1,G_i)}}{\sum_{k=1}^{N_G} e^{V(S_\tau,\tau-1,G_k)}} \tag{4.12}$$

Finally, in the full specification outlined in the introduction of this chapter, I assumed that consumers always optimize their points collection with an eye towards redemption. However, given that the inter-redemption spells can be quite long (six or seven years for example), it is difficult to imagine that during this whole time consumers are preoccupied with optimizing for a rather distant event. Instead, I make the simplifying assumption that the optimization does not start earlier than t_{max} periods before redemption. When more than t_{max} periods are left before a redemption, the consumers do not actively change their behavior to accelerate the collection; in other words, the choice of effort ρ is always zero in all the time periods that are far away from a redemption. The program may still have an influence on consumers through the rewarded behavior effect, but there will be no points pressure. While this assumption apparently complicates the model, it actually reduces the state space and thus helps speeding up the estimation.

Table 4.1 and Table 4.2 summarize the state and choice variables that characterize consumers' problem, respectively the structural parameters that characterize the model.

Table 4.1: Summary of the finite-horizon problem: state and choice variables

	State variables		*Choice variables*
Symbol	Interpretation	Symbol	Interpretation
t	time left to a redemption	ρ	effort put into collection
S_t	stock of points at the end of the period t		
G	the size of the target redemption		
	(the realized redemption can be smaller than G)		

Once the problem has been formulated, the next step is to solve it. In typical reward programs that offer consumers substantial rewards, both the stock of accumulated points and the observed redemptions reach the levels of tens of thousands. Moreover even by limiting the planning horizon (t_{max}) to one year and considering a time period to comprise two weeks, the state space still reaches a size of the magnitude of billions. The solution I used was to solve

Table 4.2: Summary: structural parameters

Symbol	Interpretation
$\phi_0^{cc}/\ \phi_0^{\overline{cc}}$	baseline collection for credit card holders and non holders
ϕ_1	rate of transformation of effort into points
ϕ_2 , ϕ_3	parameters capturing rewarded behavior effect
σ	the scale parameter of the log-normal distribution of points collected
k, θ	parameters of the distribution of targeted rewards
α_1	utility associated with a point redeemed when a goal is achieved
α_2	utility associated with a point redeemed when a goal is not achieved
	or with a point that will be used for a future redemption
c	cost of effort
δ	per period discount factor

the problem only for a subset of points in the target reward domain (call the set of these points G° and the stock of points domain (S°) and use interpolation for the points where the value and the policy functions hadn't been calculated. Specifically, I used interpolation on 2 levels: 1) interpolation on V for those $S \notin S^{\circ}$ when calculating V backwards; 2) interpolation on the policy function for values of S and G that are not in S° and respectively not in G°. I present the specific details of interpolation in Appendix B.

4.2 The data

The data for this model is also provided by the AIR MILES Reward Program and includes daily collection and redemption data for 25,898 individuals from Eastern Canada, observed between 2009 and 2013. I aggregated both the collection and redemption data at two weeks intervals and calculated each individual's balance at the end of each period. 594 appeared to have a negative balance (redemptions higher than the available stock of points, or 'miles' as they are called in this program) at at least one point during the observation period, so they were removed; 60 did not have data on the stock of miles at the beginning of the observation period - i.e. December 31st 2008 - so they were also removed.

Finally, since I want to analyze and predict the behavior of the average program participant, I removed individuals with unusual, perhaps idiosyncratic patterns of collection and redemption. The highest observed balance at the beginning of the observation period was 143,000, that is about 14,000 dollars waiting to be spent! I removed the 253 participants that had the top 1% of balances, i.e. all those that had in their account more than 13,891 miles as of December 31^{st} 2008 .

I also removed those who showed unusual large collections in at least one period. The highest observed collection in a two weeks interval was 35,000 miles. Given that participants usually need to spend at least \$20 to obtain a mile, such an influx of miles is thus associated with an expenditure of at least \$700,000 within two weeks. The 99.9^{th} percentile in the distribution of accumulated miles was 800, so I chose this threshold to eliminate all the collectors who gained more than 800 miles in at least one interval. Such large amounts of collected miles might be associated with one of a kind situation - for example buying or selling a \$400,000 house through Century 21, a real estate agency which is a sponsor of the program, allows collectors to earn 800 miles. The final number of collectors used in the estimation is 23,602 and Table 4.3 provides a summary of the number of individuals eliminated for each of the reasons discussed above.

Table 4.3: Cleaning the Data

Number of accounts	Condition
594	Had a negative balance at at least one point
60	No data on the balance at the end of 2008
253	Starting balance larger than 13,891 (99^{th} percentile)
1,389	Collected more than 800 miles (99.9^{th} percentile) in a two weeks interval
23,602	Remaining accounts included in the estimation

While the available data spans over 5 years. I use the first and last year only to assess whether each observed participant makes any redemptions in these years and focus fully on the 3 years in the middle. Knowing whether someone has recently made a redemption or is about to make one is important for the model outlined above because recent redemptions create the rewarded behavior effect and future redemptions create points pressure. So the descriptive statistics that I present below are focused on 3 years of observation of 23,602 individuals. As I split each year into 26 two-week periods, I rely on 1,840,956 observations.

Over the focal 3 years I observe 14,462 redemptions. Of the 23,602 observed participants, 15,440 (65%) do not make any redemption, 4,880 (21%) make one redemption, 1810 (8%) make two redemptions and the remaining 1472 (6%) make more than two redemptions. By considering the whole observation interval of 5 years, 55% of the sample participants do not make any redemption, 22% make one redemption, 10% make two redemptions and the remaining 13% make more than two redemptions.

The size of the rewards redeemed vary widely, from a minimum of 25 miles to a maximum of 19,120 miles. The distribution is shown in Figure 4.2 below. The mean redemption is 1,460 miles and the standard deviation is 1640. The types of rewards being redeemed also varies considerably. The company uses 5 types of categories to classify rewards: 'Flight', 'Gift certificate', 'Leisure and Entertainment' (concert tickets), 'Merchandise' and 'Travel' (car rentals or hotel stays).

Figure 4.3 shows the observed redemptions as a percentage of the available stock of miles before redemption, focusing on those redemptions smaller than 10,000 (99.5% of the redemptions being less than 10,000 and 77% being less than 2,000). This figure shows that, at least for small and medium rewards, there is wide variation with respect to the miles redeemed as percentage of the available stock. This suggests that the idiosyncratic preference for certain types of rewards is likely to play an important role. If this was not the case, participants with similar stocks of miles would choose similar goals, would work equally hard to achieve them and as a result the miles redeemed as a percentage of the available miles would be less spread out.

The observed redemptions seem to display both the points pressure and the rewarded behavior, as shown in Figure 4.4. The left panel of this figure shows two types of points pressure: towards the redemption moment (with red) and towards achieving the number of miles which will eventually be redeemed (with blue). When collectors redeem the whole stock of miles that they have available (or at least redeem a reward for which they did not have sufficient miles one period before redemptions), the acceleration towards redemption coincides with the acceleration towards the goal. This is the case in only 560 out of the total 14,462 observed redemptions.

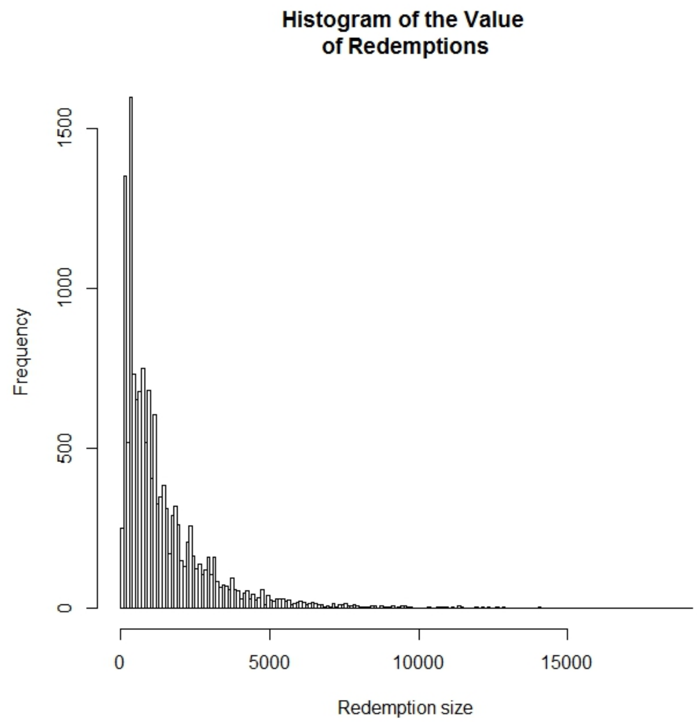

Figure 4.2: Distribution of the observed redemptions

In the remaining 13,902 cases, the number of miles which is redeemed was achieved before the redemption period. Furthermore, out of these 13,902 instances, in 5,469 cases the goal was achieved either during a previous spell or outside the observation period, so these cases are not accounted for in Figure 4.4. As shown in Figure 4.5, the number of observations used to determine the average collection as a function of 'time to goal' is much less than the number of observations used to determine the average collection as a function of 'time to redemption' precisely because often times the participants kept collecting for long periods before making a redemption, or made small redemptions relative to their balance. This should explain the less smooth averages across periods for the 'towards goal' averages. However, the large spike observed at the time of attaining the goal is also partly explained by selection: on average, in 27% of the periods participants do not collect any miles; however, the period when one attains the goal is necessarily a period of non-zero collection. And I say only 'partly explained' because the average collection in a period with positive collection is about 24 miles, but the average collection at the moment when a goal is achieved is almost 82 miles.

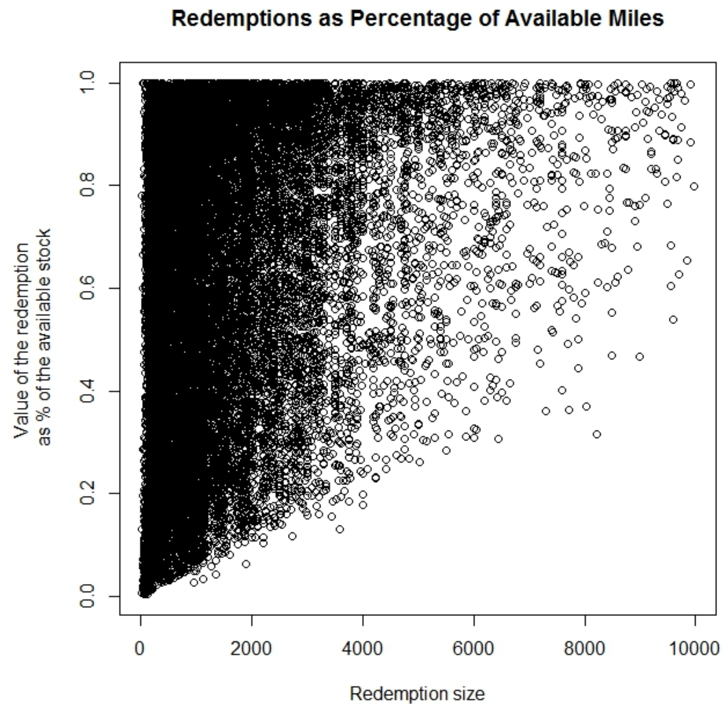

Figure 4.3: Observed redemptions as percentage of the stock of miles available before redemptions (focusing on redemptions smaller than 10,000 miles)

The right panel of Figure 4.4 shows the average number of point collected as a function of period number after a redemption. The gradually declining slope after redemption is consistent with the rewarded behavior hypothesis.

4.3 Empirical strategy

4.3.1 Identification

Similar to the Cash model in Chapter 3, I rely on indirect inference to estimate the structural parameters of interest, collectively denoted by $\boldsymbol{\beta}$. The moments are chosen to serve to the identification of the structural parameters in Table 4.2. The parameter associated with the utility of achieving one's goal (α_1) is identified mainly from the coefficients of a regression of the number of miles collected as a function of the logarithm of the number of periods left until

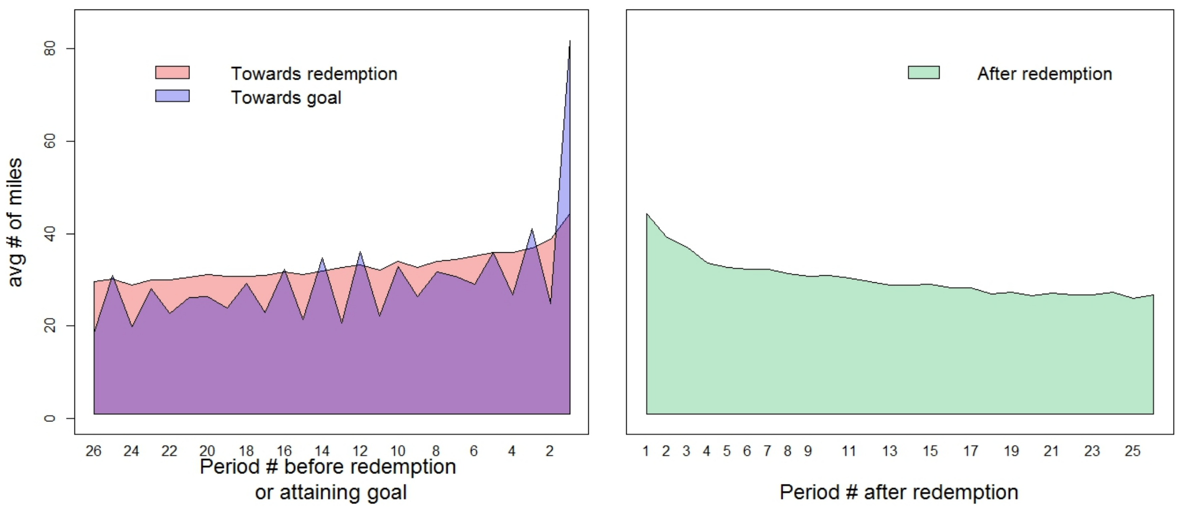

Figure 4.4: Average number of miles collected as a function of the number of periods up to a redemption or up to achieving the number of miles that will eventually be redeemed (left panel) and after a redemption (right panel)

redemption - basically a parametric version of the red surface in the left panel of in Figure 4.4. Since credit card holders have a higher overall collection and a higher peak as they approach redemption, I run this regression separately for credit card holders and non credit card holders. The parameter which characterizes the value of a mile kept for future redemptions or spend to cover a goal for which not enough miles were accumulated (α_2) is identified from the average number of miles that program participants have in their account at the end of a redemption period and also from the coefficients of a regression of the number of miles collected as a function of the logarithm of the number of periods left until achieving the goal - a parametric version of the blue surface in the left panel of in Figure 4.4. Higher α's should elicit steeper slopes in collection leading to redemption/ achieving goals, while a higher α_2 should elicit lower redemptions and leave participants with more miles in their account at the end of a redemption period.

Steeper slopes (increased collection both towards redemption and towards achieving one's goal) can also be triggered by a lower cost of effort. In order to be able to identify α_1 and α_2, I

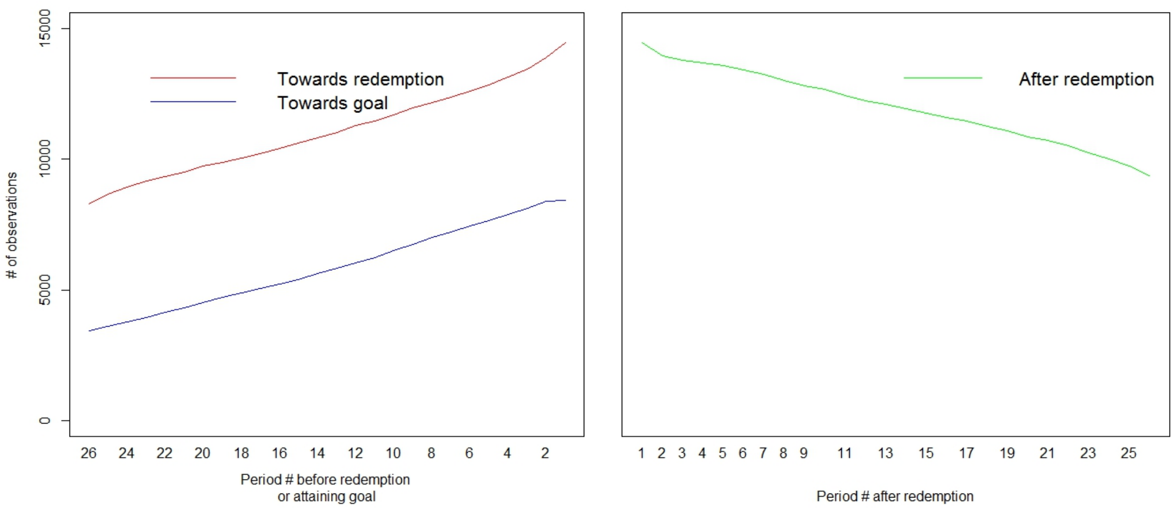

Figure 4.5: The number of observations used to compute the averages shown in Figure 4.4

fixed the cost of effort. Effort can take 11 distinct values (ρ) on an arbitrary scale form 0 to 2 in .2 increments[1] and the cost of effort increases exponentially as shown in equation 4.2:

By assumption, participants who are further than 1 year (26 periods) from redemption do not put in effort to speed up collection. This assumption, together with a moment that captures the average collection of miles per period, separately for credit card holders and non credit card holder, pin down parameters ϕ_0^{cc} and $\phi_0^{\overline{cc}}$ - the baseline collection rates, or miles corresponding to purchases that are not generated by the loyalty program. The variance in the numbers of miles collected per period identifies σ - the scale parameter of the log-normal distribution that the number of collected miles per period is assumed to follow.

[1]The reason for which in this model I use an effort scale from 0 to 2, while in the cash model presented in Chapter 3 I used a scale from 0 to 1, has to do with the discreteness of the choice of effort and the need to smooth out this discreteness in order to be able to calculate the standard errors of the parameters. Discrete choices may not be sensitive to small changes in the parameters. However, in order to calculate the standard errors of the estimated parameters I need to take the numerical derivatives of the moments in the auxiliary model with respect to the parameters and these derivatives need to be non-negative. In earlier versions of the model I had difficulties with such zero derivatives. By choosing effort from a larger set, the chances increase that even small changes in parameters change the optimal effort for some regions of the state space. The change in optimal effort further triggers changes in the auxiliary moments and renders non-negative derivatives.

Finally, the decay in the collection rate after a redemption (a parametric version of the right panel of Figure 4.4) identifies ϕ_2 and ϕ_3. k and θ (the parameters of the Gamma distribution that characterize the targeted or planned redemptions) are identified from the observed mean and variance of the redeemed rewards. δ, the per-period discount factor is set to 0.9784.

To make sure that these moments properly identify the parameters of interest, I conducted an identification exercise similar to the identification exercise for the cash rewards model. More precisely, I generated data according to known parameters and then tried to recover the known parameters. The results, reported in Appendix C, support the claim that just one set of parameters best fit the auxiliary model.

4.3.2 Estimation

As in Chapter 3, I use indirect inference to estimate the parameters of interest. This method relies on simulating the behavior of the program participants under different sets of structural parameters, recording the moments described in section 4.3.1 for each set of parameters (these simulated moments are denoted by $\tilde{\theta}$), and checking how close they are to the true moments observed in the data (denoted by $\hat{\theta}$). Those parameters that produce a simulated data set which is closest to the original data set are kept and used for policy analysis. Equation 4.13 below is the same as the equation used in the previous chapter. Again, W is the inverse of the variance-covariance matrix of $\hat{\theta}$ obtained by bootstrapping, where the off-diagonal elements have been replaced by 0s.

$$\hat{\boldsymbol{\beta}} = arg\,\min_{\boldsymbol{\beta}}(\hat{\theta} - \tilde{\theta}(\boldsymbol{\beta}))' W(\hat{\theta} - \tilde{\theta}(\boldsymbol{\beta})) \tag{4.13}$$

Since I observe 23,602 participants over 78 periods (3 years), I simulate an equal number of individuals over an equal number of periods, for each one using the observed redemption timings and the observed starting balances as inputs into the model. For a given $\boldsymbol{\beta}$ the choices of the simulated individuals are based on the optimal decision (or policy), as given by eqs. (4.9), (4.11) and (4.12), and the specific random errors of each simulated individual. For example, the optimal policy for a given constellation of the state variables may be to put in maximum effort into acquiring miles, but the realized number of miles is given by both this systematic influence ('maximum effort') and a random error, which varies across simulated individuals who find themselves at the same constellation of the state variables. The two types of random errors that I used are related to miles collection, as explained above, and to choosing goals, since, as shown in equation 4.12, the goals are chosen probabilistically, according to their value. The two sets of random errors are kept unchanged throughout different sets of $\boldsymbol{\beta}$.

4.4 Results

Table 4.4 shows the value of both the true moments and simulated moments at the minimum value of equation 4.13. The parameters that produce these results, together with their standard errors are shown in Table 4.6. Simulated moments that have a higher weight (i.e. those that vary less when bootstrapping the data) are in general closer to the true moments, than those moments with high high variance in the data. For example, the fifth moment, the slope of acceleration towards a redemption for collectors who do not have a credit card associated with the program ($CC = 0$) is the one with the highest weight: it is -2.33 in the data and -2.70 in the simulation - a relatively small difference. However, the twelfth moment, the average stock of miles at the end of a redemption period, with a weight of 0.002, is about 1224 in the data and 1022 in the simulation.

Table 4.5 shows additional moments, not used in the estimation. In this table 'Q' stands for quantile so for example 'Q25' is the 25^{th} quantile. In terms of collection it looks as the model does not emulate particularly well the distribution of the collected miles. In the data, in 27% of the instances participants did not collect any mile, while in the simulated data, this happens only 8% of the time. Also, the simulated data is more concentrated: 50% is in the interval [2,10] while 50% of the true data is in the interval [0,16]. The mean balance across all periods is matched quite closely, but the distribution is more concentrated in the simulation that in the real data. Finally, the distribution of the number of miles redeemed is shifted to the right in the simulation compared to the real data.

The discrepancies outlined above point to possible avenues in which the model could be improved. To capture the spike at zero in terms of miles collected, the model can be augmented with engagement/ disengagement periods similar to the model in Chapter 3. Furthermore, in

Table 4.4: Matched moments, i.e. the moments used in estimation (both the values observed in the data ($\hat{\theta}$) and the simulated values ($\hat{\theta}$) together with the weight attached to each moment)

		Moment	True moments	Simualted moments	Weight
1		$E(m_{it}\|cc = 0)$	10.65	10.74	267.006
2		$E(m_{it}\|cc = 1)$	29.50	29.31	27.050
3		$Var(m_{it})$	1718.66	1820.67	0.004
4	$m_{it}\|toredemption, CC = 0$	Intercept	23.60	22.53	42.931
5		log(Count_back)	-2.33	-2.70	537.910
6	$m_{it}\|toredemption, CC = 1$	Intercept	63.77	63.15	9.696
7		log(Count_back)	-6.53	-6.74	113.264
8	$m_{it}\|togoal$	Intercept	49.81	48.91	19.045
9		log(Count_back)	-8.05	-8.34	209.009
10		$E(G_{it})$	1471.35	1523.47	0.010
11		$Var(G_{it})$	1641.97	1406.66	0.004
12		$E(S_0)$	1223.75	1021.61	0.002
13	$m_{it}\|fromredemption$	Intercept	43.64	43.85	40.911
14		log(Count_fwd)	-5.75	-5.57	424.539
	objective function				534.639

order to better match the distribution of redemptions and balances available after a redemption, the utility function (equation 4.10) can be augmented with an additional parameter to allow for non-linearities in the utility of goals. However, this added flexibility comes at the cost of increased number of moments used in the estimation.

Table 4.5: Moments that were not used in the estimation

	True Moments	Simulated moments
% periods 0 miles collected	26.73	8.32
Q25 miles collected	0	2
Q50 miles collected	5	5
Q75 miles collected	16	10
mean balance across all periods	2049	2114
Q25 balance after redemption	96.25	277.00
Q50 balance after redemption	345.00	505.00
Q75 balance after redemption	1348.00	937.50
Q25 miles redeemed	395.00	450.00
Q50 miles redeemed	950	1094
Q75 miles redeemed	1900	2215

While the point estimates (Table 4.6) are not very meaningful in themselves, they can be used to gain meaningful insights. For example, as shown in equation 4.10, participants are assumed to balance high targeted redemptions against the risk of not attaining them, since not being able to attain one's goal comes at a penalty. Parameters α_1 and α_2 control how much utility participants derive from a mile used to attain one's goal as opposed to a mile which is spent for an unattained goal or kept for future redemptions. Using the estimates in Table 4.6, Figure 4.6 below shows the utility for three different goals (395 - the black line, 950 - the red line and 1900 miles - the blue line) at different levels of available miles. As expected, being able to achieve one's goal is associated with marked jumps in utility. But an interesting aspect is that for small goals, not being able to achieve one's goal brings higher utility than barely attaining the goal (the dip in the black line at 395 miles available). This seems counter-

Table 4.6: Point estimates and standard errors in parenthesis below

Parameter	Description	Entire sample
$\phi_0^{\overline{cc}}$	baseline collection for non credit card holders	1.5982
		(0.0102)
ϕ_0^{cc}	baseline collection for credit card holders	2.3980
		(0.0095)
ϕ_1	rate of transformation of effort into miles	0.7232
		(0.0022)
σ	scale parameter; distribution of miles collected	1.3453
		(0.0054)
α_1	utility point when goal achieved	0.3190
		(0.0032)
α_2	utility log(point) when goal not achieved	61.1255
		(3.2822)
k	target rewards, shape parameter	0.5298
		(0.0039)
θ	target rewards, rate parameter	0.0004
		(0.0000)
ϕ_2	rewarded behavior	0.6261
		(0.0200)
ϕ_3	rewarded behavior decay	-0.0441
		(0.0058)

intuitive, but my interpretation is that the implicit goal that program participants have is to make a redemption *and* have something left into their account for future redemptions. Failing to attain this second, implicit part of the goal causes reductions in utility. At high values of the goals, however, participants are more willing to forgo this additional objective of a non-zero balance after redemption (no dip in the blue line). In other words, someone who barely has enough miles to buy movie tickets may derive less utility from their redemption thinking that they have already spend all their available miles for a rather low-value reward than someone who only covers the half of the cost of their movie tickets using miles, thus feeling that their out-of-pocket cost was reduced by half.

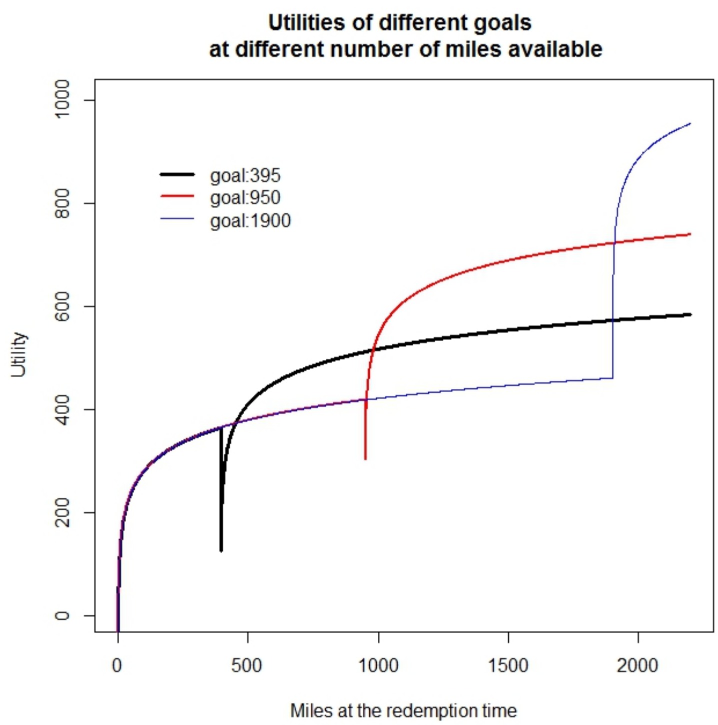

Figure 4.6: Utilities for three different goals (395, 950 and 1900 miles) at different levels of miles available at the redemption time

Furthermore, the utility of redemption dictates how much effort is optimal to put into collection. Figure 4.7 shows an example of the optimal effort for a credit card holder who has the goal of redeeming 2000 miles, in any state of the world (in terms of distance to the redemption time and balance of miles) she may find herself. The figure shows the stock of miles only in multiples of 100, but as explained in Appendix B, I use interpolation to determine the optimal level of effort for stocks of miles that are in between these multiples. Looking at the bottom right corner, if the consumer has more than 2300 miles 25 periods before the redemption period, it is never optimal to put in effort, because she already has enough miles for her goal and since the utility of remaining stock of miles is taken in logs, increasing that stock is not worth the effort. If the consumer has 1500 miles 25 periods before redemption, the optimal policy is to put in the minimum effort, which is .2. In the next period (that is 24 periods before redemption or the second row from bottom to top), the consumer will have let's say 1550 miles. The optimal policy is to put in .4 effort at a stock of 1500 miles and .2 at a stock of 1600 miles, so by linear interpolation the consumer will put in a level of effort equal to .3. When the consumer has a balance lower than 600 miles 25 periods before redemption, the optimal policy is to put in the maximum amount of effort. It is interesting to also look at the periods right before redemption (the most upward row of Figure 4.7). In general it pays off to put in effort in the periods that immediately precede a redemption. One exception is when one already has more than 2200 points (i.e. enough for the goal plus something extra); the other exception is around balances of 800-1200. The reason for the second exception is that when one has only one or two periods left to redemption no amount of effort will be enough to accumulate the miles needed to attain one's goal; also, since the utility of miles redeemed as part of an unfulfilled goal enter the utility function in logarithmic form, putting in effort to increase that amount is still not worth it.

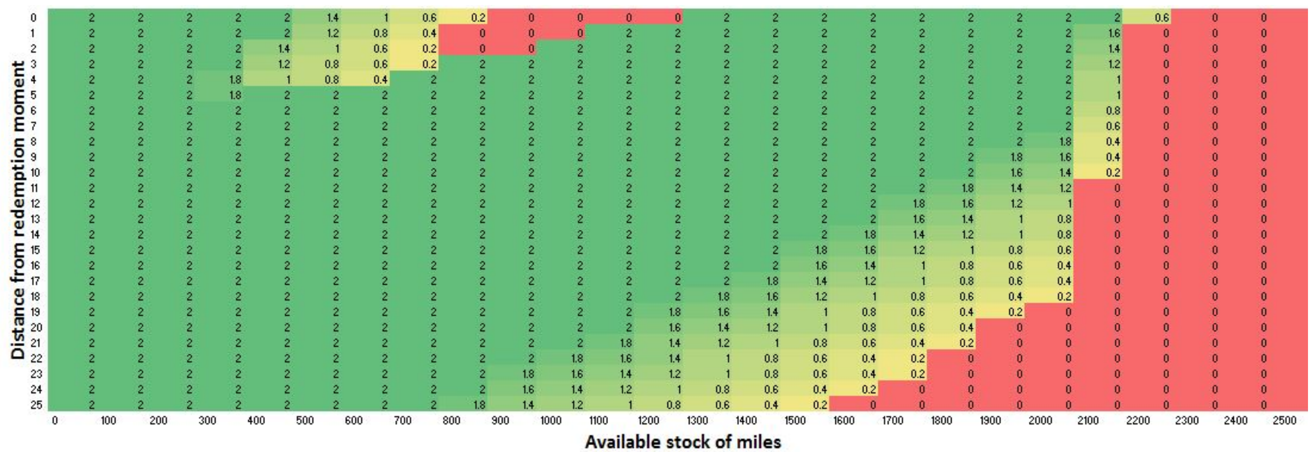

Figure 4.7: Optimal effort for a credit card holder who has the goal of redeeming 2000 miles, as a function of the available stock of miles and closeness to the redemption moment

The previous chapter showed that is important to distinguish between heterogeneity and true program effects. The results outlined so far distinguish between credit card holders and non credit card holders, but they are oblivious to differences in propensities to collect bonus miles or to be disengaged. A first step that I take in addressing this other source of heterogeneity is to estimate the model separately for a priori defined segments. The ideal approach is to use latent segments estimated flexibly within the model, as I did in Chapter 3, rather than determine the segments exogenously, outside the model. However, in doing so the model becomes more complex, it requires an extended auxiliary model to match the increased number of parameters and requires more computational resources to estimate the parameters. Therefore, here I adopt an approach that is easier to implement to see if addressing the heterogeneity in collection patterns changes the results obtained so far.

In the next subsection I run the model separately for each of the 3 identified segments. I then use both sets of results to see if taking into account participants' distinct propensities to collect bonus miles or to be disengaged changes the insights derived from a model which only distinguishes between credit card holders and non credit card holders.

4.4.1 Results for a priori defined segments

For each participant I calculated the percentage of times they collected bonus miles, the percentage of times they did not collect any mile - the remaining periods being those when they collected positive sums of regular miles. I used this data to split participants into segments, by using k-means analysis, which revealed that 3 groups strike the right balance between parsimony and minimizing within group variation. Table 4.4.1 below shows the characteristics of each of the 3 identified groups and labeled post hoc as 'Core', 'Bonus' and 'Disengaged'.

Segment	% in the population	% of times collecting bonus miles	% of times collecting zero miles	% of times collecting only non-bonus miles
'Core'	43.71	23.44	18.74	57.82
'Bonus'	26.70	57.57	5.13	37.31
'Disengaged'	29.59	9.42	57.16	33.42

Table 4.7 is the analog of Table 4.4 - it shows the true moments used in estimation, for each identified segment, together with the simulated moments and their weights. As was the case before, the moments for each sub-segment are fit reasonably well. In this table R1 stands for ‘regression 1’, where the y-variable is the number of miles acquired and the x-variable is the log of the distance to the redemption moment for participants that don’t have credit cards. R2 uses the same variables but it includes only credit card holders. R3 regresses the miles accumulated on the logarithm of the period number until the goal is attained and R4 regresses the miles accumulated on the log of the period number from redemption.

Table 4.7: Matched moments, i.e. the moments used in estimation (both the values observed in the data ($\hat{\theta}$) and the simulated values ($\hat{\theta}$) together with the weight attached to each moment), separately for the 'Bonus', 'Core' and 'Disengaged' segments

	Bonus			Core			Disengaged		
	True	Simulated	Weight	True	Simulated	Weight	True	Simulated	Weight
$E(m_{it}\|cc = 0)$	22.13	22.27	44.8636	8.48	8.14	378.3239	3.65	3.64	440.0349
$E(m_{it}\|cc = 1)$	52.47	51.07	8.5987	26.79	25.59	28.9609	12.17	11.69	25.9418
$Var(m_{it})$	3169.70	3365.13	0.0006	1328.55	1461.08	0.0034	636.79	566.79	0.0032
R1: Intercept	35.82	34.98	13.1333	13.48	13.68	75.6558	8.98	9.51	33.3019
R1: log(Count_back)	-3.13	-3.22	125.0462	-1.12	-1.16	774.2070	-1.21	-0.95	396.0506
R2: Intercept	58.57	59.29	12.3929	30.09	30.73	28.3857	17.08	18.87	18.5832
R2: log(Count_back)	-5.89	-5.58	91.9505	-3.25	-2.97	233.0212	-2.73	-2.18	227.0568
R3: Intercept	63.45	63.09	4.5216	38.74	37.60	9.5920	32.51	29.79	3.4814
R1: log(Count_back)	-9.39	-8.84	39.4269	-6.16	-5.80	79.6300	-6.76	-5.87	31.2370
$E(G_{it})$	1717.26	1779.24	0.0025	1301.08	1342.56	0.0047	1093.26	1150.67	0.0022
$Var(G_{it})$	1802.50	1431.07	0.0011	1501.13	1250.67	0.0012	1254.70	1015.17	0.0006
$E(S_0)$	1649.25	1359.81	0.0005	906.69	841.28	0.0015	647.12	653.23	0.0007
R4: Intercept	83.53	83.71	4.1473	48.19	47.82	6.2607	25.35	29.57	2.1691
R4: log(Count_fwd)	-7.96	-8.25	34.5896	-3.98	-3.92	60.0684	-1.94	-2.02	24.7277

As I did for the aggregated model, I present below other moments that were not used in the estimation. As was the case before, the model produces fewer periods with no miles collected than the data and a narrower distribution of the miles collected; in terms of redemptions, the model does not capture very well the extreme right tail in the balance after redemptions, while the distribution of the miles redeemed is captured somewhat more closely.

Table 4.8: Moments that are not used in the estimation for the 'Bonus', 'Core' and 'Disengaged' segments

Moments	Bonus		Core		Disengaged	
	True	Simulated	True	Simulated	True	Simulated
% periods 0 miles collected	5.27	0.82	18.74	14.19	57.33	37.75
Q25 miles collected	7	5	1	1	0	0
Q50 miles collected	15	11	5	3	0	1
Q75 miles collected	33	22	13	8	4	3
mean balance across all periods	2957	2614	1922	1927	1421	1654
Q25 balance after redemption	155	395	79	158	57	75
Q50 balance after redemption	636	847	248	334	165.5	189
Q75 balance after redemption	2041	1402	1050	739	677.2	532
Q25 miles redeemed	550	655	350	380	320	410
Q50 miles redeemed	1125	1346	800	943	640	837
Q75 miles redeemed	2268	2614	1650	1962	1360	1655

Table 4.9 shows the point estimates and the standard errors for each of the 3 distinct segments. As was the case for the model applied on the whole sample, the asymptotic standard errors are small in general, rendering parameters that are significantly different from zero. One exception is the α_2 parameter for the 'Disengaged' segment - estimated to be around 19.47 with a standard error of 13.97. Given the smaller sample size, it is not surprising that one of the parameters became statistically insignificant.

Table 4.9: Point estimates and standard errors in parenthesis below, separate for the 'Bonus', 'Core' and 'Disengaged' segments

Parameter	Description	Bonus	Core	Disengaged
$\phi_0^{\overline{cc}}$	baseline collection for non credit card holders	2.4090	1.2915	0.4155
		(0.0431)	(0.0096)	(0.0174)
ϕ_0^{cc}	baseline collection for credit card holders	3.0340	2.2367	1.4106
		(0.0277)	(0.0116)	(0.0276)
ϕ_1	rate of transformation of effort into miles	0.7156	0.7251	0.9439
		(0.0500)	(0.0132)	(0.0275)
σ	scale parameter; distribution of miles collected	1.0171	1.4798	1.7918
		(0.0800)	(0.0316)	(0.0355)
α_1	utility when goal achieved	0.0994	0.3694	0.6995
		(0.0072)	(0.0273)	(0.0190)
α_2	utility log(point) when goal not achieved	45.6986	32.9135	19.4665
		(2.8826)	(8.8897)	(13.9728)
k	target rewards, shape parameter	1.2144	0.6520	0.8812
		(0.0210)	(0.0039)	(0.0186)
θ	target rewards, scale parameter	0.0007	0.0005	0.0007
		(0.0000)	(0.0000)	(0.0000)
ϕ_2	rewarded behavior	0.4806	0.4173	0.7063
		(0.0544)	(0.0540)	(0.0745)
ϕ_3	rewarded behavior decay	-0.0112	-0.0201	-0.0604
		(0.0044)	(0.0027)	(0.0134)

Program evaluation

Table 4.10 shows the number of miles that are attributable to the program and uses this number to calculate the break-even margin for the program an the share of sales attributable to the program. The first part of this table presents results in terms of miles which are actually observed in the data. It shows that the average non credit card collector acquires 1.56 extra miles per period due to the program, while the average credit card collector accumulates 8.25 extra miles per period due to the program. The share of this 'extra' collection that is attributable to points pressure, (as opposed to rewarded behavior effect) is 49.1% and 32.6% respectively, markedly less than 65% which was the case for 'cash' collectors. This result seems to suggest that 'dream rewards' elicit more goodwill after the redemption than points pressure before redemption, as opposed to the 'cash rewards' for which the points pressure plays the larger role in changing behavior.

In order to assess how effective the program is, again I needed to make assumptions on the monetary value associated with each collected mile. Since in this model I don't distinguish between bonus and non bonus periods, I assume that each collected mile has a 30% probability of being collected in bonus period and 70% probability of being collected in a non-bonus period, since these are the fractions observed in the data. In line with the assumptions that I used in Chapter 3, I assume that each mile collected in a bonus period is associated with \$7 of sales (70% bonus miles each associated with \$.2 sales and 30% regular miles each associated with \$25 sales imply \$7.64 sales per mile in a bonus period; I round this number down to \$7). For the regular miles, I tried to further distinguish between miles acquired due to purchases with the credit card vs. those acquired in other ways, because credit cards collectors need to spend between \$10 and \$20 to acquire a mile. As in the previous chapter, I assumed that regular miles for non credit card holders are associated with \$25 of sales. For credit card holders I assumed that the difference between their average and non credit card collectors' average is attributable to credit card collection which rewards each multiple of \$10 with 1 mile, while the non credit card collectors' average was assumed to be associated with \$25 of sales.

Using these assumptions, I calculated that the average program participant spends \$58 more during each two weeks period compared to the situation where there is no loyalty program. This large amount is driven mainly by credit card collectors. The break-even margin is 2.3% and the share of sales attributable to the program is around 27%.

Table 4.10: Program evaluation. Applying the model to the whole sample

	No credit Card	Credit Card	Total
Share in population	64.1%	35.9%	
Extra miles due to program	1.56	8.25	3.96
Average redemption			11.09
% due to points pressure	49.1%	32.6%	
Additional revenue	30.63	106.80	58.00
Cost			1.33
Break-even margin			2.29%
Revenue with program	211.2	380.8	272.1
Revenue w/out program	180.6	274.0	214.2
% of sales due to the program	17.0%	39.0%	27.1%

These results seem to imply that the 'Dream rewards' program is much more efficient than the 'Cash rewards' program. In order to further test the robustness of these results, I also did analyzed the counter-factual for each of the 3 distinct segments and aggregated over them, in the idea that perhaps the results above miss-attribute some of the heterogeneity in the propensities to acquire bonus miles or to be disengaged to the true program effects. Table 4.11 presents the results. If anything, it shows that by running the model separately on each of the 3 segments, the implied efficiency of the program is slightly increased: the additional revenue is now \$64.1 instead of \$58, while the cost remains approximately the same. In Section 4.5 I further elaborate on the reasons which might have lead to this result.

Policy Change

Intuitively, encouraging program participants to make redemptions more often should boost the loyalty program's effectiveness, since more frequent redemptions are likely to trigger more often both the points pressure and the rewarded behavior effects. However, this strategy is not

Table 4.11: Program evaluation. Applying the model to each of the 3 pre-defined segments and aggregating the results

	Bonus	Core	Disengaged	Total
Additional revenue	146.9	45.6	16.6	64.1
Cost	2.7	1.1	0.4	1.3
Share in the population	26.7%	43.7%	29.6%	
Break-even margin				2.08%

without risk, as more frequent redemptions may discourage consumers from aiming to substantial goals that elicit effort or may leave less time for consumers to put in the effort. Also, persuading consumers to shorten their inter-redemption times may be difficult or costly. In this exercise I try to see how the cost of the program and the sales it generates would change if retailers managed to convince consumers to shorten their inter-redemption cycles. An important question is whether retailers should aim to convince many participants to slightly decrease their inter-collection times or whether they should try to convince a few participants to redeem much more frequently. Should the results suggest that the program's efficiency is significantly improved by shortening the inter-redemption timings with any of these two strategies, retailers can start looking into tactics that they can use to persuade consumers to redeem more frequently. By comparing the cost of those tactics with the extra benefits (if any) that are gained, retailers can decide whether encouraging redemption is a worthwhile strategy or not.

In the structural model presented above the inter-redemption times (τ's) are taken to be exogenous. So in order to asses how shorter spells between redemptions would affect the program outcomes, I simply modify the input to the model. There are two types of changes that I consider. Firstly, each inter-redemption time that is larger than 1 (i.e. redeeming in two consecutive periods) has a 50% chance of being reduced by 20%. I label this 'Policy I'. With this policy, for example an inter-redemption spell of 55 periods, if selected to be changed (with 50% probability), becomes a spell of 44 periods; if it is not selected to be changed, it stays 55. Secondly, each inter-redemption time that is larger than 1 (i.e. redeeming in two consecutive periods) has a 20% chance of being reduced by 50%. This is 'Policy II'. With this policy, for example an inter-redemption spell of 55 periods, if selected to be changed (with 20% probability), becomes a spell of 27 periods; if it is not selected to be changed, it stays 55. Importantly, in both specifications, the censored inter-redemption times are left unchanged.

Table 4.12 shows the outcomes for the current set up and the two changes described above. It appears that trying to change the behavior of many participants in a small way (Policy I) increases the revenue generated by the program from \$57.99 under the current set-up to \$60.43, but requires a slightly higher margin to be profitable (2.34% vs. 2.29%). Trying to change the behavior of a few participants in a more substantial way (Policy II) also increases sales generated by the program to \$59.00 but does not require a higher margin in order to be profitable. By comparing Policy I and Policy II, it appears that at margins higher than 7.64% Policy I should be preferred. However, whether any of the two policies should be implemented depends on the costs of the marketing campaigns needed to achieve these changes. I can only use the model to obtain insights on how the reduction in inter-redemption times is likely to affect sales and profits, but managers have better information on the costs of any campaign that aims to change participants' behavior.

Table 4.12: Program outcomes under the current set up and two alternative set ups where the inter-redemption times are reduced

	Current program	Policy I	Policy II
Extra miles due to program	3.96	4.14	4.03
Average redemption	11.09	11.81	11.25
Additional revenue	57.99	60.43	59.00
Cost	1.33	1.42	1.35
Break-even margin	2.29%	2.34%	2.29%
Revenue with program	272.15	274.34	273.10
Revenue w/out program	214.16	214.16	214.16
% of revenue due to the program	27.08%	28.25%	27.56%

4.5 Discussion

In this chapter I presented a model applicable to loyalty programs that offer substantial rewards, for which consumers need time to accumulate the points and for which, presumably, the consumers decide upon some time in advance. Following the exposition of the full, complex model, I presented a simpler model that is less demanding in terms of data and computation power and that can be estimated more easily. Specifically, the simplified model no longer requires estimating the distribution of the inter-redemption spells, but uses the spells observed in the data. It assumes that different goals have a chance of being picked that is proportional to their value at the time when the choice is made. A final simplifying assumption is that the points pressure mechanism does not come into play before the collector has more than one year before redemption.

As was the case for the 'cash' collectors model presented in the previous chapter, this model too addressed both the collection of miles and the redemption of rewards. In terms of redemption, the inter-redemption spells were assumed to be exogenous and I used the spells observed in the data as inputs for the model. However, I endogenized the size of the redemption: in my model participants are more likely to pick goals that have a higher value function associated with them at the beginning of the inter-redemption spell. The simplified model assumes that each redemption is a separate 'game'. In the periods leading up to the redemption consumers are motivated to put in effort to increase their chances of attaining their goal. The closer the redemption time, the less heavily discounted the reward and the higher chances that it pays off to put in additional effort to collect more miles. After a redemption consumers are under the rewarded behavior effect, so they are still likely to spend more than they would have spent in the absence of the RP.

I estimated the parameters for this model using data from AIR MILES' 'Dream rewards' program. In one specification I applied the model on the whole available sample. In the second specification, I first split the data into 3 distinct segments, based on participants' propensities to collect bonus miles or not to collect miles at all; then I applied the model to each of three a priori identified segments and aggregated over the segments, to obtain results comparable to the ones obtained in the first specification. In both specifications, in terms of fit, it seems that the model could be improved by allowing for an engaged/not engaged switch (similar to $PT = 0$, or period of type ‘disengaged’ that I had in the model presented in Chapter 3), which would help to better match the observed spike at 0 in the collection pattern. The assumption that different redemption goals are selected proportional to their value at the moment of the redemption timing also seems too rigid; by relaxing this assumption and introducing additional parameters to better capture the value of different goals, the match of the distribution of the redeemed rewards could also potentially be improved. However, these changes (and particularly the second one) may entail difficulties in terms of estimation and identification.

The reason behind applying the model separately to the 3 segments was the insight from Chapter 3, which showed that sample heterogeneity can be confounded with true program effects, so it needs to be teased apart in order to obtain unbiased estimates of the effect of the program. Therefore, my expectation was that this specification should imply a higher break-even margin for the program to be profitable and a lower share of sales that is attributable to the program. Contrary to this expectation, the results revealed that in the second specification, if anything, the program appears to perform even better.

This unexpected result can be explained in several ways. First, latent heterogeneity may play a smaller role among dream rewards collector than it played among ‘cash’ collectors. One idea in support of this conjecture is the fact that for the cash collectors I used the binary variable of whether they have a credit card linked to the loyalty program through which they can collect points for all the payments they make with that card, to stipulate different baseline collection rates for credit card holders and non credit card holders. This observable variable may have well captured the a large share of the variance in individual collection patterns, and since it was already accounted for, further dividing the sample into a priori latent groups did not change much the results.

The second explanation is that by a priori dividing the sample into latent segments and then applying the model to each segment was not a good technique to partial out the effects of heterogeneity. As opposed to deriving latent segments within the model (the technique which I used in Chapter 3), this method is more rigid: when applying the model, individuals have been already assigned with a probability equal to 1 to each of the sub-samples.

Finally, in this Chapter, I changed the inter-redemption times used as inputs for the model.

The results showed that shorter inter-redemption times boost retailer's revenue while also increasing the costs. However, whether launching a campaign that encourages redemptions is a good recommendation depends on the campaign's costs.

In the final chapter of the dissertation I further discuss the differences between this model and the model presented in the previous chapter, as well as the different estimates of profitability that they imply. I also discuss the limitations of this model and the ways in which it could be improved.

Chapter 5

Conclusion

5.1 Contributions

In this dissertation I reviewed the literature on reward programs, clearly differentiating between the general concept of customer loyalty and the loyalty fostered through reward programs. I addressed both the theoretical mechanisms that can drive increased sales for merchants who offer RPs (mainly points pressure and the rewarded behavior effect), as well the empirical studies that aimed to quantify the effectiveness of the RPs on driving sales.

I proposed two structural models for analyzing loyalty programs which do not automatically grant rewards to qualifying members and which do not put participants in a situation where they need to make trade-offs between using their points to gain status upgrades and using their points to redeem rewards. There are several advantages to using structural models. Using a combination of data and assumptions, they are like 'conceptual X-rays' that allow researchers and practitioners to understand how different mechanisms concur to create the patterns in the observed data. This deeper understanding implies that the results of these models are in a better position to be used to make cogent managerial recommendations.

The particular structural model estimated in Chapter 3 for 'cash' collectors showed how participants selecting the time of the redemption and the sample heterogeneity can be confounded with the effect of the loyalty program (Figure 3.4). Thus, in order to evaluate the impact of the program on sales, first I parsed out these two effects. The results of the model help managers decide whether the RP is profitable for them and assess how changes in the structure of the program will affect profitability. I used thus the results of the structural model like an experimental laboratory to answer questions such as what is the optimal redemption multiple or what is the optimal reward per redeemed mile? The results suggested that that the current multiple of 95 mile is optimal but that, depending on the contribution margins, increasing the amount of money offered per redeemed mile (e.g. offering $12 instead of $10 to consumers who redeem 95 miles) may increase profitability. The structural model estimated in Chapter 4 for dream rewards collectors was again used to assess the profitability of the program and the increase in profitability that could be obtained by persuading participants to shorten their inter-redemption times.

The two models described in this dissertation are the first structural models in the literature that are applicable to the typical programs in the retail and financial sectors, where most of the recent expansion of RPs has taken place (Berry, 2013). The main two features of the programs in these industries are that consumers are able to decide the size and timing of their rewards and that the customer tier component is less important and not likely to play a role in the redemption decision. Lewis (2004) and Hartman and Viard (2004) develop structural models for loyalty programs with a relatively simple structure, where the reward is granted automatically to the qualifying program participants. Kopalle et al. (2007) and Kopalle et al. (2012) develop structural models for loyalty programs that are specific to the airline and hospitality industries: in these set-ups consumers face the trade-off between using the accumulated points to make a redemption and hoarding them in order to receive a tier upgrade.

However, unlike frequent flyers or business travelers who accumulate airline (and hotel respectively) points in lumpy, infrequent amounts, participants in RPs offered by retailing or financial services companies accumulate points in a more continuous way and make redemptions that do not crowd out their opportunities to collect points. For example a frequent flyer who redeems for a free flight forgos collecting a substantial amount of points that could have brought him a status upgrade, for example from silver to gold. Unlike the frequent flyer, a participant into the AIR MILES program who uses her points to redeem for a free flight may at most forgo collecting a few hundred points if she has a credit card linked to her AIR MILES account, so for her there is no significant trade-off between redemption and status upgrade motives. These fundamental differences in the structures of the programs imply that the models developed by Kopalle and his co-authors cannot be simply tweaked and applied to industries other than airlines and hotels. The main contribution of this dissertation is therefore to develop and estimate structural models to evaluate the impact of RPs that can be applied to the typical programs in the growing sectors of retail and financial services.

The conceptual difference between the two models that I have presented and estimated stems from the way consumers think about the two types of rewards that they offer: with cash-like rewards they evaluate on an on-going basis whether to make a redemption or not (conditional on having enough points for the minimal redemption), while with dream-like rewards consumers are more likely to form a stable redemption goal and accelerate their collection to achieve their goal and improve their prospects after the redemption. In estimation, one of the major differences between the two models was the way heterogeneity was incorporated. In the case of 'cash' collectors (Chapter 3), I used the model to estimated 6 latent segments that differ with respect to their propensities to redeem and their propensities to collect bonus or only regular miles or to be disengaged. In the model for dream rewards collectors (Chapter 4) I divided the sample a priori into segments that I labeled 'Core', 'Bonus' and 'Disengaged', using k means analysis and then I applied the model separately to each of these 3 segments. While the approach that I took in Chapter 3 is has significant advantages, I regard the second approach as a useful exercise to gain more insights into how heterogeneity plays out with the true RP effects. Using latent segments in the dream rewards model would have required more computational power and a more complex auxiliary model (additional moments) which could be prone to difficulties with the estimation and identification issues.

The results imply that the 'Dream rewards' program is more effective than the 'Cash rewards' program. For 'Cash rewards', depending on the assumptions, the break-even margin ranged from 2.84% to 5.64%, while for the 'Dream rewards', under the assumptions listed in Chapter 4, this rate is 2.29%. The share of revenue attributable to the program ranged between 12.20% and 15.65% (again depending on the assumptions) in the case of 'Cash', while for the 'Dream', under the current assumptions it is estimated to be 27.1%.

For both models, in the model development and analysis, I considered the general framework of one retailer that runs the RP. While this is a general and flexible framework, the data that I used to estimate the parameters of the model is provided by AIR MILES - a *coalition* RP.

In this coalition, AIR MILES issues the points and sells it to its sponsors; the sponsors give away the points to their patrons as they see fit. The cycle is closed by the program participants who redeem their rewards from AIR MILES in exchange for the points that they have accumulated. In the coalition set-up, the residual claimant, or the party that bears the risk of low breakage or high redemption, is the party which administers the program - AIR MILES, in this case. The individual sponsors pay for all the miles that they issue and are unaffected by the redemption rates. In my results, I adopted AIR MILES' point of view and calculated the cost of the redemptions based on the observed redemptions (which are about 60% of the observed collections). However, the retailers who participate in the program pay for all the miles that they issue, not only for those that are redeemed. As I showed in Chapter 3, the redemption rates have a significant impact on the minimum margin rate that make the program profitable: at the current level of redemptions this rate is 4.3% but when 100% of the issued points are redeemed, the rate is 7.1%. In other words, the average program sponsor who pays upfront for the cost of all the miles they issue, needs to have a contribution margin which is at least 7.1% to find the program attractive.

In sum, this dissertation presents two structural models that can be used to evaluate the typical loyalty programs offered in retail and financial services. The unique features of RP's in these industries are the endogenous nature of redemption and the fact that the prospect of upgrades to superior customer tier does not act as a counter-balance to the tendency of making redemptions. These features require models that different substantively from the ones developed so far in the literature. My models disentangle between true program effects and two important confounds (selection of the redemption timing and sample heterogeneity), in order to provide unbiased estimates of the effectiveness of the program and insights into the changes that can further bolster their profitability. Therefore they are a valuable managerial tool that can be used to evaluate RPs and assess different RP designs.

5.2 Limitations

A significant challenge with any model that tries to assess the causal impact of a loyalty program on sales is is the risk of confounding the true program effects with a spurious relationship between the increase in miles collection around the redemption time and the redemption act. In Chapter 3 I showed that even if the RP does nothing to change the behavior of consumers there is a positive correlation between collection amounts and the redemption moments, driven by sample heterogeneity and by the fact that the redemption moment is more likely to be selected after a streak of high-collection periods that were produced by chance, in the absence of any systematic effect.

However, there may be factors that encourage consumers to both accumulate and redeem more miles. Around any special occasions, like anniversaries or holidays, consumers are more likely to both increase their collection of miles *and* make a redemption. If this case, it is the special occasion that drives the positive correlation observed between collection and redemption, and not the loyalty program. So models that don't take into account this special occasion effect miss attribute the observed correlation to the program. This situation can lead to more or less bias in the estimated effectiveness of the program, depending on how large the increase in collection around holidays and other personal important events is and on how strongly these moments are correlated with an increased tendency to redeem.

The solution of simply controlling for special occasions is not easily implementable, since different consumers have different special occasions that are not observable by the researcher. A first step could be accounting for any 'December effect' when people tend to spend more money. However, if program participants genuinely prefer to make redemptions at these special occasions and are thus more likely to put in more effort into collection precisely in these periods, then the 'December effect' or 'special occasion' control swipes away completely the RP effect and leads to the opposite kind of bias, where the effectiveness of the program is under-estimated.

I believe that such confounds are not likely to play a major role in the cash rewards model, but they may be a more serious limitation of the dream rewards model, especially for those collectors that have a credit card linked to their RP account. The credit card is a powerful tool for collecting points, that allows participants to have a wide variation in their collection, i.e., allows them to increase substantially their collection during those special occasion periods when they spend more.

To illustrate how credit card holders are more prone to such confounding effects, a scenario could unwind as following: a program participant uses her credit card to pay for flights for the family holiday, accumulates an increased amount of points as a result, and then uses the points to make hotel bookings for the same holiday. In this scenario, it is the holiday that drives both the accumulation of an increased amount of points and the redemption, so the correlation between increased collection of points and redemption is spurious. This type of effect may explain why the model presented in Chapter 4 implies such a high share of revenue attributable to credit card holders. Unfortunately, controlling for it requires richer, survey data where consumers provide details on the context in which they made their redemptions - such as whether the redemption was part of a suite of expenditures generated by an exogenous event - holiday, anniversary or any other special event.

Another limitation that has to do with the data used to estimate the model for dream rewards collectors in Chapter 4 is the fact that I only use 3 years of data for the actual estimation and keep the 1 year before and 1 year after to collect information on whether participants make redemptions within these periods. During the 3-year observation window, only 35% of the participants make at least one redemption. This short observation window allows for the possibility that those participants who have higher baseline collections are more likely to also be redeemers. If this is the case, the model that I estimated erroneously attributes the increased collection to the loyalty program, when in fact there may be a selection effect. A solution to alleviate this concern is to observe participants over longer periods of time. In this way, the group of redeemers becomes more heterogeneous (collectors with both higher and lower propensities to collect baseline miles) and they will be less likely to be confounded with collectors with high baseline collection. Thus it is possible that by using data that stretches over longer periods of time, the estimated profitability of the dream rewards program to be revised downwards.

5.3 Future research

An immediate area for further research is to increase the complexity of the models that I have developed and try to achieve a better fit to the data, which would lend more credibility to the obtained results. Specifically, for the cash collectors model such attempts could focus on modeling a correlational structure among the 6 identified segments, increasing the state space to account for the seasonal effects ('December') and check how the results obtained with that specification compare with the results reported in Chapter 3, or expanding the sample to include the participants who have a credit card linked to their RP account.

Another avenue for future research is testing the bounds of the assumption that the propensity to acquire bonus miles or be disengaged from the program is unchanging over time and independent from the stock of miles that one has already accumulated. To achieve this, the auxiliary model should be extended to include two additional regressions: in the first additional regression the dependent variable is whether an instance of zero collection is observed and in the second the dependent variable is whether bonus miles are acquired or not - while the explanatory variable is the distance to the redemption moment. Two additional parameters should be included in the model. Specifically, instead of capturing the propensity to be disengaged by q_0, it can now be a function of the stock of miles accumulated: $q_{01}e^{-q_{02}S}$, where S is the stock of miles. Similarly, instead of capturing the propensity to be disengaged by q_b, it can also become a function of the stock of miles. The additional moments in the auxiliary model will help identify the additional parameters that confer extra flexibility to the model.

For the dream rewards model there is even more opportunity for further research. This can range form incremental improvements to the model (introducing an engaged/ not engaged switch; allowing for a more flexible choice of the redeemed rewards; estimating latent segments within the model; accounting for seasonal effects), to more radical improvements, that rely on augmenting the observed data data with survey data, which collects information on the specific context in which redemptions took place, for example for how long it has been planned or whether it coincided with any special event that would normally determine the program participant to accumulate more points.

The results suggest that the 'Dream rewards' program is more effective in increasing sales than the 'Cash' rewards program. More research is needed to explain this difference and investigate whether the dream reward program elicits more engagement or simply attracts program participants who are more responsive to the incentives of the program, while those who are not as responsive to the incentives of the program opt for the 'Cash' set-up.

I estimated the models using data from AIR MILES, Canada's largest coalition loyalty program. However, I did not exploit at all the 'coalition' feature of the data. On the one hand this makes the models more general, applicable to programs run by individual retailers. On the other hand, this leaves room for asking further questions. The models I used assumed that consumers have a cost of changing their behavior and collect more points in order to generate points pressure. It may be the case that this cost is different among different sponsors, so some may benefit more from the points pressure effect than others. It may also be the case that the rewarded behavior effect is stronger for the retailer where the consumer made the redemption than for other sponsors of the program. In this case retailers where consumers are more like to cash in their miles are advantaged. These sponsor-specific effects, if they exist, could have major implications for how the coalition loyalty program is run: it may turn out that some sponsors are benefiting more from the program, being subsidized by the others. If this is the case, AIR MILES can asses the possibility of differentiated tariffs for its sponsors. So a major avenue for future research is to assess whether such differentiated effects exist among sponsors and designing a differentiated pricing scheme that could increase AIR MILES' profitability.

In my analysis I assumed that the retailer's objective is always profitability, not sales volume. It may be the case that sometimes retailers use loyalty programs for other short-term objectives, like increasing sales or helping with inventory management, so retailers may sometimes sacrifice contribution in the short term. Loyalty programs are then a tool which allows achieving this flexibility, so they may have additional value beyond that of increasing contribution. Exploring how retailers actually use RPs and quantifying the benefits of the flexibility that they provide can also be an area ripe for future research.

In a more general sense, it would be interesting and managerially relevant to see how the RPs' design affect their efficiency. In this dissertation I studied RPs where the redemption in endogenous - i.e. barring expiration, it is at the participants' latitude when and how much to redeem, while much of the previous structural models in this area focused on RPs where the reward is granted automatically. Comparing the performance of these different designs would be valuable for a retailer deciding which type of program to implement. In theory both RP structures have advantages and disadvantages: the endogenous redemption allows participants to choose the rewards that they best like and are thus likely to be motivated by them to increase their collection effort, while the automatic redemption avoids the problem of participants being indecisive and may therefore keep them better engaged.

Furthermore, other details of the RP design may be important for its success. For example, in the case of programs with automatic redemption a key question is what should the amount of points required for a reward be? Very large required amounts may discourage most of the participants, while small amounts may end up rewarding participants who do not change their behavior at all in response to the program. In the case of programs with endogenous redemption it is important to know what are the rewards that best motivate participants - e.g. do vacation-related rewards elicit more points pressure and/ or a higher rewarded behavior effect than merchandise rewards? Another relevant question is what is the right balance between flexibility and complexity? Offering large sets of available rewards increases flexibility and the chances that each program participant finds a reward that is motivating them to increase their points collection. But on the other hand such large sets of rewards might increase the burden of choice and leave the consumers incapable of selecting and committing to a goal. In the absence of a specific goal consumers may not be very keen to increase their collection of points.

Finally, even though in this dissertation I only studied points or miles as a currency that consumers accumulate and spend, I hope that the insights gained here, especially those that referred to how consumers seem to save up points for occasions where they might enjoy them more, can spark further research in areas of broader interest, for example saving behavior.

Appendix A

Identification for the 'Cash' model in Chapter 3

To ensure that the model developed in Chapter 3 is identified - i.e. that the same set of moments cannot be generated by different sets of parameters, I conducted the exercise described below.

I generated data (and moments) using known parameters. The parameters that I picked are shown in Table A.1. In order to obtain a confidence interval around these parameters, one needs the Jacobian matrix (the numeric first order partial derivatives of each moment with respect to each parameter) and a weighting matrix (W). In the spirit of the generalized methods of moments literature, the weighting matrix can be any positive-definite matrix. So I used W obtained as the inverse of the variance-covariance matrix of the moments in the real data, with the off-diagonal elements replaced by 0. The last column of Table A.1 shows the standard errors obtained in this fashion. The more time-consuming alternative would be to do bootstrapping on the generated data to obtain W.

Then I picked 4 sets of random starting points within $+\backslash-$ two standard errors from the true parameters (shown in Table A.2) and different sets of idiosyncratic errors (e, ξ and ϵ) and tried to recover the true parameters. Table A.3 shows the obtained estimates for each of the four starting points and Table A.4 shows how far (in the metrics of standard errors) the obtained parameters are from the true parameters.

The results of the exercise show that in all the 4 tests, with few exceptions, the recovered parameters are within $+\backslash-$ two standard errors from the true parameters, lending support to

Table A.1: Parameters used to generate the data

	true parameters	std errors
1 Φ_1	1.5717	0.0067
2 Φ_2	1.1422	0.0097
3 Φ_3	1.4214	0.0052
4 σ_1	0.7195	0.0063
5 σ_2	0.4115	0.0046
6 a_1^r	-3.6977	0.0408
7 a_2	2.8101	0.0467
8 Φ_4	0.5566	0.0064
9 Φ_5	3.2486	0.1024
10 q_b^I	0.1303	0.0044
11 q_d^I	0.5676	0.0113
12 Q^I	0.1945	0.0065
13 q_b^{II}	0.2723	0.0035
14 q_d^{II}	0.1346	0.0052
15 Q^{II}	0.4546	0.0110
16 q_b^{III}	0.6359	0.0017
17 q_d^{III}	0.0031	0.0011
18 a_1^{nr}	-7.7545	0.1417
19 p_r	0.7464	0.0053

the claim that the model is well identified.

Table A.2: Starting points - within + \ − two standard errors from the true parameters

	Starting points			
Φ_1	1.5855	1.5699	1.5603	1.5720
Φ_2	1.1482	1.1450	1.1430	1.1535
Φ_3	1.4160	1.4171	1.4134	1.4187
σ_1	0.7114	0.7175	0.7242	0.7206
σ_2	0.4129	0.4059	0.4114	0.4092
a_1^r	-3.6933	-3.6903	-3.7037	-3.7267
a_2	2.7253	2.7410	2.8019	2.8631
Φ_4	0.5550	0.5494	0.5534	0.5605
Φ_5	3.1513	3.1116	3.3578	3.2568
q_b^I	0.1322	0.1306	0.1277	0.1218
q_d^I	0.5704	0.5690	0.5799	0.5771
Q^I	0.1969	0.1989	0.1867	0.1985
q_b^{II}	0.2753	0.2717	0.2776	0.2743
q_d^{II}	0.1449	0.1391	0.1289	0.1329
Q^{II}	0.4389	0.4572	0.4616	0.4673
q_b^{III}	0.6378	0.6349	0.6315	0.6366
q_d^{III}	0.0033	0.0039	0.0030	0.0029
a_1^{nr}	-7.9105	-7.6978	-7.7719	-7.9618
p_r	0.7419	0.7560	0.7396	0.7502
Objective function	4192.5200	9405.8400	16574.6300	784.3500

Table A.3: Estimated parameters

	Ending points			
Φ_1	1.5816	1.5800	1.5622	1.5690
Φ_2	1.1500	1.1592	1.1350	1.1482
Φ_3	1.4133	1.4130	1.4303	1.4241
σ_1	0.6995	0.7051	0.7204	0.7192
σ_2	0.4069	0.4054	0.4146	0.4118
a_1^r	-3.6995	-3.7010	-3.6982	-3.7097
a_2	2.8199	2.7626	2.8519	2.8116
Φ_4	0.5552	0.5651	0.5606	0.5571
Φ_5	3.1603	3.1411	3.3455	3.2630
q_b^I	0.1277	0.1281	0.1220	0.1224
q_d^I	0.5638	0.5651	0.5755	0.5701
Q^I	0.1977	0.1981	0.1879	0.1911
q_b^{II}	0.2741	0.2820	0.2791	0.2787
q_d^{II}	0.1351	0.1291	0.1325	0.1340
Q^{II}	0.4476	0.4724	0.4744	0.4672
q_b^{III}	0.6320	0.6470	0.6412	0.6385
q_d^{III}	0.0045	0.0014	0.0011	0.0019
a_1^{nr}	-7.7719	-7.8288	-7.6680	-7.8309
p_r	0.7461	0.7541	0.7385	0.7484
Objective function	12.7500	12.2600	21.2900	17.1400

Table A.4: Distance (in standard errors) of the estimated parameters form the true parameters

Φ_1	1.4831	1.2453	-1.4265	-0.4070
Φ_2	0.7994	1.7500	-0.7411	0.6092
Φ_3	-1.5608	-1.6067	1.6923	0.5123
σ_1	-3.1827	-2.2908	0.1417	-0.0548
σ_2	-1.0001	-1.3328	0.6745	0.0593
a_1^r	-0.0428	-0.0807	-0.0118	-0.2930
a_2	0.2115	-1.0176	0.8966	0.0330
Φ_4	-0.2227	1.3289	0.6243	0.0765
Φ_5	-0.8617	-1.0496	0.9462	0.1404
q_b^I	-0.5948	-0.5117	-1.8758	-1.7902
q_d^I	-0.3407	-0.2262	0.6984	0.2216
Q^I	0.4920	0.5624	-1.0160	-0.5224
q_b^{II}	0.5122	2.7871	1.9494	1.8350
q_d^{II}	0.1101	-1.0495	-0.3881	-0.1150
Q^{II}	-0.6336	1.6137	1.8006	1.1420
q_b^{III}	-2.2694	6.4489	3.1058	1.5082
q_d^{III}	1.2271	-1.5335	-1.8123	-1.0355
a_1^{nr}	-0.1228	-0.5242	0.6103	-0.5391
p_r	-0.0456	1.4417	-1.4796	0.3731

Appendix B

Interpolation for the 'Dream' model in Chapter 4

Target rewards (G's) are modeled as a distributed $\Gamma(k, \theta)$; for each set of k and θ that is tried out in the process of finding the best fitting parameters, I calculate every 5^{th} percentile of the distribution and use that value as an interpolation point - relying thus on 21 points for G. Thus, G° is the set of percentiles.

In order to calculate the set of interpolation points for the stock of miles, S°, I follow several steps. First, I calculate the maximum value of the balance that a simulated individual could conceivably attain during the simulation window (t_w). This situation happens when the simulated individual starts with a high balance, denoted $maxS_0$, has just made a redemption right before being observed, so their collection is boosted by the rewarded behavior effect; will make a new redemption right at the end of the observation period, putting in all the possible effort ($\rho = \rho_{max}$) and is always very lucky receiving the highest number of miles, given the location and scale of the distribution from which the miles are drawn. $maxS_0$ is observed in the data. The total rewarded behavior effect (**P**) is calculated over the entire t_w period is calculated as following:

$$\mathbf{P} = \int_1^{t_w} \phi_2 e^{\phi_3 u} du = \frac{\phi_2}{\phi_3}(e^{\rho_3 t_w} - e^{\phi_3}) \tag{B.1}$$

So the maximum location parameter that could hypothetically be achieved over the observation window is

$$\mu_{max} = \frac{1}{t_w}\Big[\phi_0 t_w + \mathbf{P} + \phi_1 \rho_{max} t_{max}\Big]. \tag{B.2}$$

Thus, the maximum stock of miles that can be attained during the observation period is $maxS$, given by the sum of the maximum stock at the beginning period and the maximum number of miles that can be attained per period, multiplied by the number of observed periods:

$$maxS = maxS_0 + t_w \cdot F^{-1}(.9999) \tag{B.3}$$

In the equation above $F(\cdot)$ is the cumulative distribution function of the log-normal distribution with location parameter μ_{max} and scale parameter σ.

Given that $maxS$ captures the very extreme case when all the random factors collude to create a large observed balance, I use a combination between an exponential scale and a linear scale to determine the interpolation points for the stock of miles between 0 and $maxS$. I use in total 51 points of interpolation for S (so the size of S° is 51). The first 21 points are linear - starting from 0 up to 2000 in 100 increments. The remaining 30 are at an increasing distance form each other. I divide the interval $[log(2000), log(maxS)]$ into 30 equally spaced intervals and use as an interpolation point for S, all the values that correspond to the ends of these intervals. This combination of linear and exponential grid has the advantage that has a higher density for the lower values of S, where most observations lie, but in the same time it covers all the interval of values that the stock of miles could imaginably take.

In solving the finite horizon problem, I first solve for the utilities in the redemption period (at $t = 0$) for all the $21 \cdot 51$ interpolation points. Using these values, I solve for the value and policy functions at all $t \in [1, t_{max}]$ and at all G° in the grid, by interpolating only for S. For example, in order to calculate $V[S_2, 1, G_\tau]$, I need to know all $V[\varsigma, 0, G_{tau}]$, where ς represents all the possible levels of the balance at $t = 1$, i.e.the current balance S plus all the possible number of miles that can be accumulated in a period. Since not all ς's are in the interpolation set S°, I use the set $V[\cdot, 0, G_{tau}]$ at the interpolation points.

Though, in the current version of the model I use only rewards that are in the interpolation set, for rewards that are outside this set the optimal level of effort for can be computed by interpolating the optimal effort at the points within the interpolation grid.

Appendix C

Identification for the 'Dream' model in Chapter 4

To ensure that the model developed in Chapter 4 is identified, I conducted an exercise similar to the one described in Appendix A.

I generated data (and moments) using known parameters. The parameters that I picked are shown in Table C.1. In order to obtain a confidence interval around these parameters, again I used the Jacobian matrix and a weighting matrix (W), obtained as the inverse of the variance-covariance matrix of the moments in the real data, with the off-diagonal elements replaced by 0. The last column of Table C.1 shows the standard errors obtained in this fashion.

Then I picked 4 sets of random starting points within $+ \setminus -$ two standard errors from the true parameters (shown in Table C.2) and different sets of idiosyncratic errors (the two sets of idiosyncratic errors are related to the shocks in miles collection and the shocks in the size of the redeemed rewards) and tried to recover the true parameters. Table C.3 shows the obtained estimates for each of the four starting points and Table C.4 shows how far (in the metrics of standard errors) the obtained parameters are from the true parameters.

The results of the exercise show that the recovered parameters are in general within $+ \setminus -$ two standard errors from the true parameters, lending support to the claim that the model is well identified. The major exception is the shape parameter of the distribution of rewards in the first test, which is 15 standard errors larger than the true parameter. However, in the other 3 tests this parameter is within reasonable distance from the its true value: 3.17, -.70 and -.70.

Table C.1: Parameters used to generate the data

	True parameters	std errors
$\phi_0^{\overline{cc}}$	2.318269	0.008061
ϕ_1	0.739684	0.010297
σ	1.1455	0.006801
α_2	37.8436	2.063653
k	1.265785	0.011629
θ	0.000547	1.86E-05
ϕ_2	0.57684	0.095859
ϕ_3	-0.04074	0.008019
α_1	0.113131	0.001896
ϕ_0^{cc}	2.835605	0.022307

Table C.2: Starting points - within + \ − two standard errors from the true parameters

Parameter				
$\phi_0^{\overline{cc}}$	2.3127	2.3164	2.3173	2.3242
ϕ_1	0.7400	0.7410	0.7574	0.7462
σ	1.1424	1.1490	1.1378	1.1362
α_2	36.9270	39.2058	35.0475	38.1293
k	1.2692	1.2784	1.2431	1.2537
θ	0.0005	0.0005	0.0006	0.0005
ϕ_2	0.6260	0.5820	0.6137	0.5970
ϕ_3	-0.0409	-0.0406	-0.0227	-0.0325
α_1	0.1118	0.1120	0.1116	0.1165
ϕ_0^{cc}	2.8261	2.8291	2.8371	2.8455
Objective function	58.48	6.86	538.02	531.53

Table C.3: Estimated parameters

$\phi_0^{\overline{cc}}$	2.3094	2.3160	2.3086	2.2941
ϕ_1	0.7403	0.7409	0.7599	0.7661
σ	1.1407	1.1491	1.1030	1.1606
α_2	39.6398	39.2153	39.3289	36.6905
k	1.3349	1.2791	1.2678	1.3700
θ	0.0006	0.0005	0.0005	0.0006
ϕ_2	0.6001	0.5825	0.6004	0.6095
ϕ_3	-0.0370	-0.0408	-0.0282	-0.0348
α_1	0.1135	0.1122	0.1050	0.1042
ϕ_0^{cc}	2.8175	2.8286	2.8339	2.7960
Objective function	2.45	1.57	14.92	5.33

Table C.4: Distance (in standard errors) of the estimated parameters form the true parameters

$\phi_0^{\overline{cc}}$	-1.104657	-0.275719	-1.199176	-3.003085
ϕ_1	0.057480	0.116210	1.959899	2.565022
σ	-0.713010	0.525466	-6.253172	2.223130
α_2	0.870380	0.664704	0.719736	-0.558768
k	5.946741	1.141570	0.169079	8.965637
θ	1.606681	0.140721	-1.856587	2.068021
ϕ_2	0.242404	0.058794	0.245298	0.341053
ϕ_3	0.463631	-0.004990	1.558901	0.741222
α_1	0.213874	-0.487856	-4.272067	-4.718497
ϕ_0^{cc}	-0.813118	-0.314112	-0.077197	-1.774889

Bibliography

Aaker, D. (1991). *Managing Brand Equity: Capitalizing on the Value of a Brand Name*. Free Press.

Abu-Shalback Zid, L. (2004). Loyalty doesn't always pay. *Marketing Management*, 13(3):4.

Allaway, A. W., Berkowitz, D., and DSouza, G. (2003). Spatial diffusion of a new loyalty program through a retail market. *Journal of Retailing*, 79(3):137–151.

Andrew, P. (2006). Consumers choose rewards credit cards but don't always take rewards.

Andrews, D. and Lu, B. (2001). Consistent moment and model selection procedures for GMM estimation with application to dynamic panel data models. *Journal of Econometrics*, 101(1):123–164.

Ashley, C., Noble, S. M., Donthu, N., and Lemon, K. N. . (2011). Why customers won't relate: Obstacles to relationship marketing engagement. *Journal of Business Research*, 64(7):749–756.

Beggs, A. and Klemperer, P. (1992). Multi-period competition with switching costs. *Econometrica*, 60(3):651–666.

Benartzi, S. and Thaler, R. (1995). Myopic loss aversion and the equity premium puzzle. *The Quarterly Journal of Economics*, 110(1):73–92.

Bendle, N., Farris, P. W., Pfeifer, P. E., and J., R. D. (2016). *Marketing Metrics*. Pearson Education.

Berry, J. (2013). The 2013 colloquy loyalty census.

Bettman, J. R., Luce, M., and Payne, J. (1981). Constructive consumer choice processes. *Journal of Consumer Research*, 25(3):187–217.

Blattberg, R. and Neslin, S. (1990). *Sales Promotion - Concepts, Methods, and Strategies*, chapter 2 - Theories and Concepts of Consumer Behavior Applied to Sales Promotion, pages 19–64. Prentice Hall.

Blattberg, R. and Neslin, S. (2008). *Database Marketing: Analyzing and Managing Customers*, chapter 22 - Frequency Reward Programs, pages 549–578. Springer.

Bolton, R., Kannan, P., and Bramlett, M. (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.

Chalmers, A. (1999). *What Is This Thing Called Science?* Open University Press.

Chaudhuri, A. and Holbrook, M. (2001). The chain of effects from brand trust and brand affect to brand performance: The role of brand loyalty. *Journal of Marketing*, 65(2):81–93.

Che, H. and Seetharaman, P. (2009). “Speed of replacement”: Modeling brand loyalty using last-move data. *Journal of Marketing*, 77(5):21–36.

Chintagunta, P. T., Erdem, T., Rossi, P., and Wedel, M. (2006). Structural modeling in marketing: Review and assessment. *Marketing Science*, 25(6):604–616.

Davis, G. (2004). In store marketing: Loyalty - is the end as we know it? *Marketing Magazine*.

Davis, O. (1959). The economics of trading stamps. *Journal of Business*, 32(2):141–150.

Day, G. (1969). A two-dimensional concept of brand loyalty. *Journal of Advertising Research*, 9(3):29–35.

Deighton, J. (2000). *Handbook of Services Marketing & Management*, chapter 24 - Frequency Programs in Service Industries, pages 401–408. Sage Publications.

Dick, A. and Basu, K. (1994). Customer loyalty: Toward an integrated conceptual framework. *Journal of the Academy of Marketing Science*, 22(2):99–113.

Dorotic, M., Verhoef, P., Fok, D., and Bijmolt, T. (2014). Reward redemption effects in a loyalty program when consumers choose how much and when to redeem. *International Journal of Research in Marketing*, 31(4):339–355.

Dowling, G. and Uncles, M. (1997). Do customer loyalty programs really work? *Sloan Management Review*, 38(4):71–82.

Drèze, X. and Hoch, S. J. (1998). Exploiting the installed base using cross-merchandising and category destination programs. *International Journal of Research in Marketing*, 15(5):459–471.

Drèze, X. and Nunes, J. (2004). Using combined-currency prices to lower consumers' perceived cost. *Journal of Marketing Research*, 41(1):59–72.

Drèze, X. and Nunes, J. (2009). Feeling superior: The impact of loyalty program structure on consumers perceptions of status. *Journal of Consumer Research*, 34(6):890–905.

Drèze, X. and Nunes, J. (2011). Recurring goals and learning: The impact of successful reward attainment on purchase behavior. *Journal of Marketing Research*, 48(2):268–281.

Ehrenberg, A., Uncles, M., and Goodhardt, G. (2004). Understanding brand performance measures: Using Dirichlet benchmarks. *International Journal of Research in Marketing*, 57(12):1307–1325.

Financial Times (2006). Loyalty programs.

Goettler, R. and Gordon, B. (2011). Does AMD spur Intel to innovate more? *Journal of Political Economy*, 119(6):1141–1200.

Gourieroux, C., Monfort, A., and Renault, E. (1993). Indirect inference. *Journal of Applied Econometrics*, 8(S1):S85–S118.

Hansen, L. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054.

Hartman, W. and Viard, V. (2004). Frequency reward programs create switching costs? A dynamic structural analysis of demand in a reward program. *Quantitative Marketing and Economics*, 6(2):109–137.

Heckman, J., LaLonde, R., and Smith, J. (1999). *Handbook of Labor Economics*, chapter 31 - The Economics and Econometrics of Active Labor Market Programs, pages 1865–2097. North Holland.

Hennessy, A. and Whited, T. (2007). How costly is external financing? Evidence from a structural estimation. *The Journal of Finance*, 62(4):1705–1745.

Hsee, C., Yu, F., Zhang, J., and Zhang, Y. (2003). Medium maximization. *Journal of Consumer Research*, 30(1):1–14.

Imbens, G. and Wooldridge, J. (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, 47(1):5–86.

InsideFlyer (2006). What about miles: Inside the mind of the mileage junkie.

Jacoby, J. and Chestnut, R. (1978). *Brand Loyalty Measurement and Management*. Wiley.

Keane, M. (2010). Structural vs. atheoretic approaches to econometrics. *Journal of Econometrics*, 156(1):3–20.

Kim, B., Shi, M., and Srinivasan, K. (2001). Reward programs and tacit collusion. *Marketing Science*, 20(2):99–120.

Kivetz, R. and Simonson, I. (2002). Self-control for the righteous: Toward a theory of precommitment to indulgence. *Journal of Consumer Research*, 29(2):199–217.

Kivetz, R., Urminsky, O., and Zheng, Y. (2006). The goal-gradient hypothesis resurrected: Purchase acceleration, illusionary goal progress, and customer retention. *Journal of Marketing Research*, 43(1):39–58.

Klemperer, P. (1987). Markets with consumer switching costs. *The Quarterly Journal of Economics*, 102(2):375–394.

Kohli, A. K. and Jaworski, B. J. (1990). Market orientation: The construct, research propositions, and managerial implications. *Journal of Marketing,*, 54(2):1–18.

Konsewicz, B. (2007). Does your point burner campaign really reduce your loyalty program's point liability?

Kopalle, P. and Neslin, S. (2003). The economic viability of frequency reward programs in a strategic competitive environment. *Review of Marketing Science*, 1.

Kopalle, P. K., Sun, Y., Neslin, S., Sun, B., and Swaminathan, V. (2007). A dynamic structural model of the impact of loyalty programs on customer behavior.

Kopalle, P. K., Sun, Y., Neslin, S., Sun, B., and Swaminathan, V. (2012). The joint sales impct of frequency reward and customer tier components of loyalty programs. *Marketing Science*, 31(2):216–235.

Kumar, V. and Shah, D. (2004). Building and sustaining profitable customer loyalty for the 21st century. *Journal of retailing*, 80(4):317–329.

Lal, R. and Bell, D. (2003). The impact of frequent shopper programs in grocery retailing. *Quantitative Marketing and Economics*, 1(2):179–202.

Leenheer, J., Van Heerde, H. J., Bijmolt, T. H., and Smidts, A. (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.

Lemon, K., Rust, R., and Zeithaml, V. (2001). What drives customer equity? *Marketing Management*, 10(1):20–25.

Lewis, M. (2004). The influence of loyalty programs and short-term promotions on customer retention. *Journal of Marketing Research*, 41(3):281–292.

Liberman, Y. (1999). Membership clubs as a tool for enhancing buyers patronage. *Journal of Business Research*, 45(3):291–297.

Liu, Y. (2007). The long-term impact of loyalty programs on consumer purchase behavior and loyalty. *Journal of Marketing*, 71(4):19–35.

Liu-Thompkins, Y. and Tam, F. (2013). Not all repeat customers are the same: Designing effective cross-selling promotion on the basis of attitudinal loyalty and habit. *Journal of Marketing*, 77(5):21–36.

Loewenstein (1987). Anticipation and the value of delayed consumption. *The Economic Journal*, 97(378):666–684.

Mägi, A. (2003). Share of wallet in retailing: the effects of customer satisfaction, loyalty cards and shopper characteristics. *Journal of Retailing*, 79(2):97–106.

Meyer-Waarden, L. (2008). The influence of loyalty program membership on customer purchase behaviour. *European Journal of Marketing*, 42(1/2):87–114.

Meyer-Waarden, L. and Benavent, C. (2009). Grocery retail loyalty program effects: Self-selection or purchase behavior change? *Journal of the Academy of Marketing Science*, 37(3):345–358.

Nako, S. (1992). Frequent flyer programs and business travelers: An empirical investigation. *Logistics and Transportation Review*, 28(4):395–414.

Nielsen (2013). Nielsen survey: 84 percent of global respondents more likely to visit retailers that offer a loyalty program.

O'Brien, L. and Jones, C. (1995). Do rewards really create loyalty? *Harvard Business Review*, 73(3):75–82.

Oliver, R. (1999). Whence consumer loyalty? *Journal of Marketing*, 63:33–44.

Passingham, J. (1999). Grocery retailing and the loyalty card. *Journal of the Market Research Society*, 40(1):55–63.

Pearson, B. (2012). *The Loyalty Leap*. Penguin.

PricewaterhouseCoopers (2013). Measuring up: Retail benchmarking survey.

Read, D., Loewenstein, G., and Rabin, M. (1999). Choice bracketing. *Journal of Risk and Uncertainty*, 19(1):171–197.

Reichheld, F. (1996). *The Loyalty Effect*. Harvard Business School Press.

Rothschild, M. L. and Gaidis, W. (1981). Behavioral learning theory: Its relevance to marketing and promotions. *Journal of Marketing*, 45(2):70–78.

Rust, R., Lemon, K., and Zeithaml, V. (2004). Return on marketing: Using customer equity to focus marketing strategy. *Journal of Marketing*, 68(1):109–127.

Shadish, W., Cook, T., and Campbell, D. (2002). *Experimental and Quasi-experimental Designs for Generalized Causal Inference*. Houghton Mifflin.

Sharp, B. and Sharp, A. (1997). Loyalty programs and their impact on repeat purchase loyalty patterns. *International Journal of Research in Marketing*, 14(5):473–486.

Shefrin, H. and Thaler, R. (1988). The behavioral life-cycle hypothesis. *Economic inquiry*, 26(4):609–643.

Shugan, S. M. (2005). Brand loyalty programs: Are they shams? *Marketing Science*, 24(2):185–193.

Sirdeshmukh, D., Jagdip, S., and Sabol, B. (2002). Consumer trust, value, and loyalty in relationale xchanges. *Journal of Marketing*, 66(1):15–37.

Smith, A. (2008). *The New Palgrave Dictionary of Economics*, chapter Indirect Inference. Basingstoke: Palgrave Macmillan.

Sopanen, S. (1996). Customer loyalty schemes: The bottom line. *European Retail Digest*, pages 12–19.

Taylor, G. A. and Neslin, S. (2005). The current and future sales impact of a retail frequency reward program. *Journal of Retailing*, 81(4):293–305.

Taylor, L. (2010). Why are CEOs rarely fired? Evidence from a structural estimation. *The Journal of Finance*, 65(6):2051–2087.

The Wall Street Journal (2010). Have miles but can't get a seat? How to snare one.

Topa, G. (2001). Social interactions, local spillover and unemployment. *Review of Economic Studies*, 68(2):261–295.

van der Klaauw, W. and Wolpin, K. (2008). Social security and the retirement and savings behavior of low-income households. *Journal of Econometrics*, 145(1):21–42.

van Osselaer, S., Alba, J., and Manchanda, P. (2004). Irrelevant information and mediated intertemporal choice. *Journal of Consumer Psychology*, 14(2):257–270.

Verhoef, P. (2003). Understanding the effect of customer relationship management efforts on customer retention and customer share development. *Journal of Marketing*, 67(4):30–45.

Zhang, Z., Krishna, A., and Dhar, S. (2000). The optimal choice of promotional vehicles: Front-loaded or rear-loaded incentives? *Management Science*, 46(3):348–362.

Curriculum Vitae

Name:	Mihaela Alina Nastasoiu
Post-Secondary Education and Degrees:	Academy of Economic Studies Bucharest, Romania 2004 - 2008 B.A.
	National School of Political Science and Public Administration Bucharest, Romania 2004 - 2008 B.A.
	Central European University Budapest, Hungary 2008 - 2010 M.A.
	Western University London, ON Canada 2011 - 2016 Ph.D.
Honours and Awards:	AMA-Sheth Consortium Fellow (2015) Wharton Customer Analytics Initiative - data recipient (2012) Outstanding M.A. Thesis Award, Economics, Central European University (2010)

Publications:

Bendle, Neil and Mihaela Alina Nastasoiu (2014) *Primary Elections and US Political Marketing* in Political Marketing in the US, Ed. Lees-Marshment, J., B. Conley and K. Cosgrove, Routledge