

NEW INSIGHTS FROM EMERGING TYPES
OF RETAIL LOYALTY PROGRAMS

Valeria Stourm

A DISSERTATION

in

Marketing

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy

2016

Supervisor of Dissertation

Co-supervisor of Dissertation

Eric T. Bradlow, Professor of
Marketing, Statistics, and Education

Peter S. Fader, Professor of Marketing

Graduate Group Chairperson

Eric T. Bradlow, Professor of Marketing, Statistics, and Education

Dissertation Committee

Eva Ascarza, Assistant Professor of Marketing, Columbia University

David Bell, Professor of Marketing, University of Pennsylvania

Robert Meyer, Professor of Marketing, University of Pennsylvania

NEW INSIGHTS FROM EMERGING TYPES OF RETAIL LOYALTY
PROGRAMS

© COPYRIGHT

2016

Valeria Stourm (Valeria Montero Garnier)

This work is licensed under the
Creative Commons Attribution
NonCommercial-ShareAlike 3.0
License

To view a copy of this license, visit

<http://creativecommons.org/licenses/by-nc-sa/3.0/>

Dedicated to my family

Thank you to Cecilia, my son, Ludovic, Giana, and my parents Liana and Arturo for your love and support.

ACKNOWLEDGEMENT

I begin by expressing appreciation for the support and recognition that this dissertation has received. The Jay H. Baker Retailing Center awarded research funding and the Wharton Customer Analytics Initiative awarded the data used in Chapter 3. This research earned honorable mentions for the Marketing Science Institute Alden G. Clayton Doctoral Dissertation Proposal Competition and the American Marketing Association Mathew Joseph Emerging Scholar Award.

My dissertation has benefitted from the valuable expert perspectives of the members of my committee: Eva Ascarza, Bob Meyer, and David Bell. I also thank Joel Huber, Pradeep Chintagunta, and two anonymous reviewers at Journal of Marketing Research for helpful comments and suggestions to Chapter 2 of this dissertation.

To my co-advisors Eric Bradlow and Peter Fader: our research meetings have been fun and they will be missed. They have been a constant shining light that have kept me grounded in my work through both the happiest and most challenging events in my life. I extend a special thanks to Pete for encouraging me to pursue a PhD in Marketing and for sparking my research interest in loyalty programs.

Lastly, thank you to the faculty and PhD students of the Marketing Department for making our school an amazing place. It has been a pleasure.

ABSTRACT

New Insights from Emerging Types of Loyalty Programs

Valeria Stourm

Eric T. Bradlow Peter S. Fader

In a standard loyalty program, a single retailer offers rewards to customers who stockpile points up to a certain amount. While research on these archetypal loyalty programs is vast, there is an increasing trend for companies to adopt reward programs that do not explicitly incentivize customers to return in order to “cash-in” rewards. Two examples are linear and coalition reward programs. In a linear program, points can be redeemed at anytime for any amount. In a coalition program, points can be earned and redeemed across several partner retail stores.

A chapter titled “*Stockpiling Points in Linear Loyalty Programs*”, uses transaction data from a linear loyalty program in Latin America to examine why customers tend to stockpile points for long periods of time, despite economic incentives against doing so (i.e., time value of money). A mathematical model of redemption choice posits three explanations for why customers seem to be motivated to stockpile on their own, even though the retailer does not reward them for doing so: economic (value of forgone points), cognitive (nonmonetary transaction costs), and psychological. The psychological motivation is captured by allowing customers to book cash and point transactions in separate mental accounts. The results indicate substantial heterogeneity in how customers are motivated to redeem and suggest that behavior in the data is driven mostly by cognitive and psychological incentives.

A chapter titled “*Market positioning in a coalition loyalty program: the value of*

a shared reward currency” uses a model of multi-store purchase incidence to infer the market positioning among popular partners of a coalition loyalty program. The model shows how the value of a rewards currency that is shared among partner stores can explain patterns in customer-level purchases across the stores, and how these reward spillovers are driven by (1) differences in reward redemption policies among the partners, (2) product category overlap between stores and (3) geographic distance between them. By leveraging a devaluation of the program’s points that occurred in our observation period, we demonstrate how the value of coalition points influences the positioning of partner stores within the network.

Table of Contents

Acknowledgement	iv
Abstract	v
List of Tables	xi
List of Figures	xii
CHAPTER 1 : A Framework to Study the Emerging Loyalty Program Landscape	1
CHAPTER 2 : Stockpiling Points in Linear Loyalty Programs	7
2.1 Motivation	7
2.2 Program Description	13
2.3 Data Description	15
2.4 Redemption Model	19
2.5 Empirical Identification	29
2.6 Empirical Results	32
2.7 General Discussion	42
CHAPTER 3 : Market Positioning in a Coalition Loyalty Program: the Value of a Shared Reward Currency	44
3.1 Motivation	44
3.2 Data Description	48

References for Chapter 2	94
References for Chapter 3	96
References for Chapter 4	98
References for Appendix 3	99

List of Tables

1.1	Dimensions of increased reward fungibility	2
2.1	MA dimensions to evaluate points and cash	11
2.2	Heterogeneity in customer behavior	16
2.3	Summary of empirical specifications	27
2.4	Summary of $\partial z / \partial \tilde{s}$ for specifications	31
2.5	$\partial z / \partial m$ for specifications	31
2.6	Deviance Information Criterion and (-Log Marginal Density)	34
2.7	Posterior distribution of MA's population-level parameters	36
2.8	Segment means for each standardized motivation indicator	38
2.9	Financial analysis of policy changes	40
3.1	Summary statistics of purchases across stores	49
3.2	Categories of top 15 retail partners	50
3.3	Dollar value of rewards earned for \$100 spent at each store	51
3.4	Dollar value of rewards earned for \$100 spent at each tier	52
3.5	Category proximity between stores	57
3.6	Geographic distances between stores (kilometers)	58
3.7	Percent of customers who purchased at n of the top stores	60
3.8	Number of top partners visited on a purchase occasion	60
3.9	Summary of empirical models	66
3.10	Measures of overall model in-sample fit	67
3.11	Number of stores patronized by the average customer	68
3.12	Customers shared between three synergistic partners	69

3.13	Evolution of own and cross effects	70
3.14	Population-level parameters for non-reward utility and own effects . .	71
3.15	Posterior distribution of $\bar{\beta}$ (before devaluation)	72
3.16	Posterior distribution of $\bar{\beta} + \delta$ (after devaluation)	73
3.17	Support of hypotheses at the individual level	73
4.1	Parameter recovery for MA	87
4.2	Parameter recovery for Affinity Model	93

List of Figures

2.1	Example of linear vs. non-linear rewards	7
2.2	Redemptions and stockpiles over time	17
2.3	Percentage of the basket that can be reduced	18
2.4	Average monthly potential savings	18
2.5	Comparison of run-lengths	19
2.6	Cost-benefit tradeoff for MLA	25
2.7	Cost-benefit tradeoff for MA	26
2.8	Bimonthly redemptions	34
2.9	Redeeming credit card's point stockpile at Amazon checkout	36
2.10	How motivations vary by segments	38
2.11	Redemptions and stockpiles across segments	39
3.1	Timeline of policy changes to reward program	51
3.2	Heat map of spillover links	77
3.3	Vulnerability vs. Competitive Clout of Reward Spillovers	79
4.1	Net utilities when shopping needs are low versus high	92

CHAPTER 1 : A Framework to Study the Emerging Loyalty Program Landscape

In a standard loyalty program, a single retailer offers rewards to customers who stock-pile points up to a certain amount. For example, a coffee shop may reward customers with a free drink after 5 consecutive purchases. This mechanism of promised rewards explicitly incentivizes customers to return and purchase again in the future.

While research on these archetypal loyalty programs is vast, there is an increasing trend for retailers to adopt new types of reward programs that are structured in a fundamentally different way. In particular, we are observing a proliferation of reward programs in which retailers do not explicitly incentivize customers to return multiple times in order to “cash-in” rewards. As a consequence, reward currencies (or points) in these programs do not impose future switching costs on consumers. Rewards in these programs are more fungible and thus more similar to cash than those offered in classic reward schemes.

Cash is fungible because any one unit of a currency is equivalent in value to any other unit. In contrast, points in a standard loyalty program in which a single-retailer offers non-linear rewards, are not fungible. Recall our example of a coffee shop that rewards customers with a free drink after five consecutive purchases. The points earned towards the reward on a customer’s first purchase are not equivalent in value to those earned on the customer’s fifth purchase. The first points have less value because they cannot be immediately redeemed.

What are the challenges faced by managers offering these “no-strings-attached” fungible reward mechanisms? What are the characteristics of customer behavior in these

programs? Why are more retailers starting to adopt these fungible reward structures? To address these questions, this introduction chapter develops a framework to highlight key dimensions in which retailers are adapting the fungibility of rewards.

The framework shown in Table 1.1 has two dimensions in which emerging programs can increase their reward fungibility: (1) the type of reward structure and (2) the multiplicity of participating retailers. The following chapters of this dissertation discuss the managerial issues that arise in emerging fungible programs within the context of this framework.

Table 1.1: Dimensions of increased reward fungibility

	Single Retailer	Multiple Retailers
Non-Linear Rewards	Points as future rewards at a store	Points as future rewards at multiple retailers
Linear Rewards	Points as store cash	Points as cash valid at multiple retailers

The first dimension, the type of reward structure, distinguishes between programs with non-linear and linear rewards. A non-linear program typically requires customers to stockpile points up to a certain amount before they can be redeemed. Another example of a program with a non-linear redemption structure is when the redeemable value of each point increases in a staggered way, such as “when you reach 5000 points, we’ll double your points”. The fungibility of points is limited in non-linear programs because their redeemable value depends on how many points are previously earned. In contrast, a linear program offers points that are redeemable for a fixed amount, regardless of how many points the customer has previously earned. The fungibility of points in linear programs is greater because these can be redeemed at anytime for

any amount.

The second dimension is the number of retailers actively participating in the reward program. Relative to a standard single-retailer program, the fungibility of points is increased when points earned at one retailer are exchangeable for points earned at another. While this dissertation focuses on multi-retailer programs with active participation, it is important to differentiate these from programs in which multiple retailers passively participate. For example, points earned at a focal store may not be directly redeemable at another retailer, but the focal store may allow customers to redeem their points for gift cards that can be used at other retailers.

The two dimensions determine four quadrants in the 2x2 framework. Quadrant 1 encompasses single-retailer programs with non-linear rewards. While these are the most common and most widely studied programs, rewards in these are the least fungible within our framework. Our coffee shop example of the standard program fits in this quadrant because a customer with three purchases cannot immediately cash in the value of his rewards.

Quadrant 2 encompasses single-retailer programs that offer non-linear rewards. Points in these programs are analogous to cash, but their use is restricted within a single retailer. Unlike cash, one unit is not exchangeable for another unit from another retailer or from another customer.

The second chapter of this dissertation, titled “*Stockpiling Points in Linear Loyalty Programs*”, studies how customers decide to stockpile and redeem points in these types of more fungible linear programs. The chapter uses transaction data from a linear loyalty program in Latin America to examine why customers tend to stockpile points for long periods of time, despite economic incentives against doing so (i.e., the

time value of money). The chapter develops a mathematical model of redemption choice that posits three explanations for why customers seem to be motivated to stockpile on their own, even though the retailer does not reward them for doing so. These motivations are economic (the value of forgone points), cognitive (nonmonetary transaction costs), and psychological (customers value points differently than cash). The psychological motivation is captured by allowing customers to book cash and point transactions in separate mental accounts. The model is estimated on data from an international retailer from Latin America using Markov Chain Monte Carlo methods and is shown to accurately forecast redemptions during an 11-month out-of-sample period. The results indicate substantial heterogeneity in how customers are motivated to redeem and suggest that behavior in the data is driven mostly by cognitive and psychological incentives.

Quadrant 3 includes linear reward programs offered by multiple retailers. The increased point fungibility in this quadrant has implications for consumer switching costs. In single-store programs, points impose switching costs on consumers because these can only be redeemed at the stores at which they were earned (unlike cash). Thus, points effectively subsidize future purchases at the places at which they were earned, analogous to store coupons. In contrast, in multi-retailer programs, points can be earned and redeemed across several partner retail stores. Thus, points earned at one store effectively subsidize purchases at another store. In fact, in some multi-retailer programs, customers can earn points across all stores but they may not even have the option to redeem points in the store in which these were earned!

The research question of the third chapter of this dissertation focuses on programs in Quadrant 3. The chapter titled “*Market positioning in a coalition loyalty program: the value of a shared reward currency*” studies how the reward rates offered at part-

ner stores affect consumer purchase behavior across other partners in the coalition. It develops a model of multi-store purchase incidence and infers the market positioning (“landscape”) among popular partners of a coalition loyalty program. The model shows how the value of a rewards currency that is shared among partner stores can explain patterns in customer-level purchases across the stores, and how these reward spillovers are driven by (1) differences in reward redemption policies among the partners, (2) product category overlap between stores and (3) geographic distance between stores. While conventional models typically compare competitors within an industry, our model positions partners that operate in different high-end retail markets by identifying a latent affinity network between them. Markov Chain Monte Carlo is used to estimate the affinity network model on transaction data from a European coalition loyalty program. By leveraging a devaluation of the program’s points that occurred in our observation period, we demonstrate how the value of coalition points influences positioning of partner stores within the network.

Quadrant 4 includes non-linear reward programs offered by multiple retailers. Programs of this type present unique challenges that need to be addressed in future research. Multi-retailer programs with linear rewards (those in Quadrant 3) already face unique challenges of how to allocate the liabilities and costs of earned and redeemed points across retailers with different margins. These challenges are magnified when rewards are non-linear. For example, if a customer multiplies by two the value of his points stockpile through purchases at retailers A and B, but then decides to redeem points at C, how should the costs of the increased value of points be allocated across the retailers? Airline coalition programs fit in this quadrant. For example, a customer may gain gold status at the coalition by spending mostly at one carrier, but the gold status earns him special rewards across all partner carriers. The final

chapter of this dissertation overviews future research directions that are relevant to firms in this fourth quadrant. The final chapter also discusses emerging trends that continue to increase the fungibility of rewards in emerging loyalty programs.


CHAPTER 2 : Stockpiling Points in Linear Loyalty Programs

2.1. Motivation

The way a loyalty program is designed influences when customers redeem their rewards. For example, certain loyalty programs, such as those offered in the airline industry, require customers to stockpile points to earn specific rewards (i.e., “reach 30,000 points to redeem a round-trip ticket”). Similar programs offer stockpiling goals such as “when you reach 5,000 points, we’ll double your points.” In these programs, as more points are stockpiled, the redeemable value of each point increases in a staggered way, as illustrated in Panel A of Figure 2.1. We refer to the design of these reward structures as “nonlinear.”

In contrast, in “linear” reward structures (illustrated in Panel B, Figure 2.1), customers cannot increase the value of each point by stockpiling. Because each point is worth a fixed amount, typically a cash equivalent, customers can easily redeem as little as one point without the hassle of complicated rules. These retailers do not set stockpiling goals and do not require a minimum amount of points to redeem.

Figure 2.1: Example of linear vs. non-linear rewards
Panel A: Non-linear Panel B: Linear



One example of a linear program is Capital One’s Quicksilver Cash Rewards credit card. Cash rewards accumulate online and can be redeemed at any time. Amazon.com offers similar rewards in which redeemed points reduce the price that is paid. Tesco’s successful Clubcard program began by rewarding every British pound spent with one point toward store vouchers (Humby, Hunt, and Phillips 2004). What these programs have in common is that they do not explicitly reward stockpiling points.

Indeed, customers participating in linear programs often face incentives *against* stockpiling. Unredeemed points can expire, and they can lose their value if the retailer enters bankruptcy or alters the program rules. Moreover, by delaying redemption, customers also forgo the time value of money from delayed rewards. Because linear programs do not reward stockpiling and unredeemed points may lose their value, we would expect customers to redeem regularly. Yet they do not! Even in controlled laboratory experiments, people are hesitant to redeem (Kwong, Soman, and Ho 2011).

As another example, in analyzing data from the linear program studied here, we found that only 3% of all purchases have redemptions associated with them. These customers rarely redeem even though doing so could reduce their basket price by nearly 30% on average. Yet, it is not the case that customers are completely ignoring opportunities to redeem: 40% of the customers in our panel eventually redeemed at least once over a 43-month observation period. Why do customers wait so long between redemptions and stockpile cash-like points in the process? This question is the focus of the chapter.

We present a model that unites three potential motivations for persistent stockpiling in a linear program:

- The opportunity cost of forgone points,
- Nonmonetary transaction costs, and
- How points are valued differently than cash.

The first is an economic incentive: customers forgo the opportunity to earn points on purchase occasions in which they redeem. The second is a cognitive incentive: customers may find redeeming to be “costly,” even if the process is as effortless as clicking a button at checkout. The third motivation is psychological, based on mental accounting (Thaler 1985). This explanation recognizes that customers may not perceive points and cash equally. For example, in interviews with customers from a linear program, one customer expressed sorrow when redeeming points: “It makes me feel sad because I don’t have any points left on my card” (Smith and Sparks 2009).

Drèze and Nunes (2004) propose a framework in which customers keep two mental accounts: one for cash and another for points. Customers may experience disutility when paying for a purchase (Prelec and Loewenstein 1998; Zellermayer 1996), and this “pain of paying” can vary by payment method (e.g., cash vs. check vs. card) (Soman 2003). Similarly, the pain of paying may also vary by the type of currency used to pay for a purchase (cash vs. points). Drèze and Nunes’ objective is to enable firms to set prices using combinations of cash and points that will minimize customers’ disutility of paying. This objective is not necessary in linear programs, because customers are always free to pay with countless combinations of cash and points using a fixed points-to-cash exchange rate. Nevertheless, their two-account framework provides an excellent foundation to study redemption behavior in linear programs.

We introduce several changes to Drèze and Nunes’ (2004) framework to make it better suited for both the context (i.e., linear vs. nonlinear program) and the decision

that we focus on here (i.e., when to redeem given the available stockpile vs. which currencies to use). First, we explicitly incorporate nonmonetary transaction costs of redeeming. Drèze and Nunes (2004, p. 71) recognize the existence of such costs but do not incorporate them into their utility model. Second, we add the consideration of forgone points by allowing customers to value the gains from the points earned in the program, not just the losses from spending cash and points. With these two changes, the model captures the first two motivations to stockpile. Third, we also allow customers to subjectively value points less than cash at a fixed conversion rate. Fourth, we do not restrict our attention to cases in which customers already have stockpiles large enough to cover the entire price (Drèze and Nunes 2004, pp. 62, 69). This restriction is reasonable for analyzing redemption behavior in nonlinear programs because, by design, customers are encouraged and sometimes restricted to wait until they have a large stockpile. This restriction is not needed to study linear programs because customers can easily redeem when stockpiles are small. Instead, we examine how redemption choice changes with the size of the available stockpile relative to total prices.

These changes lead to a multiple-accounts (MA) model that unites economic, cognitive, and psychological explanations for why customers of linear programs can be motivated to stockpile on their own. The MA model has two dimensions: (1) whether a customer evaluates gains and losses of cash in a separate mental account from those of points (i.e., multiple accounts vs. single account) and (2) whether these valuations within an account are made with value functions that are S-shaped (concave over gains and convex over losses) or instead linear over gains and losses. Table 2.1 shows that while the MA model allows for separate S-shaped mental accounts, it nests cases in which customers have a single S-shaped account (SA), multiple linear accounts

(MLA), or a single linear account (SLA).

Table 2.1: MA dimensions to evaluate points and cash

Value functions	Multiple Accounts	Single Account
S-shaped	MA	SA
Linear	MLA	SLA

The MA model predicts that customers stockpile up to a latent threshold, which is not set by the retailer, but is instead determined by the three motivations to stockpile: forgone points, transaction costs, and how each customer values points relative to cash. By allowing for two mental accounts, the relative shapes of the two value functions can motivate a customer to stockpile. The MA model can be considered structural (Chintagunta et al. 2006) in the sense that the estimated parameters directly determine the fundamental cost–benefit trade-off that governs redemption behavior. While structural models typically rely on standard expected utility theory (Von Neumann and Morgenstern 1944), the MA model is grounded in prospect theory (Kahneman and Tversky 1979; Tversky and Kahneman 1992), which governs the value function shapes, as well as mental accounting, which considers one versus two accounts.

This chapter examines the MA model in two fundamental ways. First, we analytically examine the model’s implications for stockpiling. Afterward, we build an empirical version of the model, with a hierarchical Bayes structure, to reflect heterogeneity in how people perceive gains and losses of cash relative to gains and losses of points. We estimate the MA model along with the nested specifications given in Table 2.1 on observational data from a linear loyalty program of an international retailer and show how the three motivations to stockpile differ across customers. This analysis is useful to identify how customer segments may respond to alternative strategies for

encouraging redemptions.

Alternative explanations to those encompassed by the MA model can be drawn from the literature on nonlinear programs. In particular, while research on nonlinear programs is vast (Bijmolt, Dorotic, and Verhoef 2010), the findings often cannot be easily translated to linear programs. For example, a psychological phenomenon called the “medium effect” (Hsee et al. 2003) exists when myopic consumers would make different stockpiling choices if rewards were denominated directly in dollars instead of in points. The effect is expected when points alter the perceived monetary return of stockpiling. One driver of the medium effect is a nonlinear exchange rate between points and cash. Nevertheless, “the mere presence of a medium (points) is not sufficient to produce a medium effect” (Hsee et al. 2003 p. 3). No medium effect is expected in linear loyalty programs that offer a fixed points-to-dollar exchange rate and reward customers with a fixed number of points for every dollar spent, as studied here. As a second example, the “goal-gradient hypothesis” (Kivetz, Urminsky, and Zheng 2006) finds that consumers tend to exert more effort (i.e., purchase faster) as they advance toward a redemption goal that is explicitly set by the retailer. Linear programs do not have such goals. Furthermore, the goal-gradient effect sheds light on purchasing rather than on the redemption behavior that we are interested in; it assumes that redemption occurs when customers have stockpiled enough points to do so.

In summary, we use a model of mental accounting to empirically examine the potential motivations to persistently stockpile in linear programs and how these vary across customers. Our findings and documentation of stockpiling behavior in a linear program respond directly to a call in a recent article for research on how customers redeem: “Though redemptions are critical elements of loyalty programs, we do not

really know why loyalty program members redeem, or why they do not” (Bijmolt, Dorotic, and Verhoef 2010). Our results can be used to improve communication strategies to encourage redemptions and may also provide insight into why even in nonlinear programs, some customers persistently stockpile above and beyond explicit requirements.

The remainder of this chapter proceeds as follows: We first describe the data set to motivate both the theoretical and the empirical parts of the research. Then, we mathematically lay out MA and explain how it captures each motivation. We specify an empirical version of the utility model and demonstrate its performance, along with the other specifications of the 2 x 2 framework. We then apply MA to characterize the heterogeneity in customer motivations and assess policies that may potentially increase redemptions by addressing each of the three motivations. We conclude this chapter with a summary and directions for future research.

2.2. Program Description

Our empirical setting is a linear loyalty program that has operated for more than 20 years in a prominent supercenter chain in Latin America (the chain has asked to remain anonymous). It is the market leader in several countries for a range of high-end product categories similar to those offered by Target, Bed Bath & Beyond, and Home Depot. For example, the product categories it offers include flat screen televisions, beds, hardware items, toys, kitchen appliances, home decor items, gardening tools, and camping equipment.

The data track the behavior of a cohort of customers who first signed up for the retailer’s loyalty program in January 2008. For each visit of each individual, over a 43-month period, we observe the basket price, how many points were earned or

redeemed, and the date. From these data, we are able to infer the available stockpile of points each person had at the time of each purchase.

Consistent with our introduction and motivating example, the loyalty program offers a linear reward policy. In our setting, customers cannot get the satisfaction of paying without any cash (i.e., they cannot pay 100% with points). Stockpiled points can be redeemed anytime to reduce up to 50% of the basket price (not including items on sale) at a constant and easy-to-remember points-to-cash exchange rate. We observe that this cap potentially affects 40% of redemptions—that is, those in which a customer had more points than 50% of the basket price. When the cap is nonbinding, customers, on average, redeemed points equal to 22% of their basket price.

Furthermore, redemption is a low-effort activity. Customers who want to redeem points simply show their loyalty card and tell the cashier they want to redeem. While redemption does not require customers to keep track of their stockpiles, customers can easily check their balance at the cash register, online, or by phone. Point expiration is not an important motivation to redeem in this program because any purchase delays the expiration of a customer’s entire stockpile by one year. Earning points is also simple: a customer presents his or her loyalty card to the cashier to earn 1% of the total purchase price in points.

In this program, points are not earned on shopping trips when a customer redeems. Thus, the redeeming customer sacrifices an opportunity to earn points on that trip. This opportunity cost is an economic motivation to delay redemption. Although this program feature is incorporated in our model, it is only one of three motivations to redeem. While some programs share this feature (i.e., as commonly occurs when using points to partially pay for a plane ticket), others reward customers based on

the portion of the price paid with cash. Our model can easily be modified to capture these smaller forgone points by considering unearned points on only the redeemed amount instead of the full price.

Several features the program lacks suggest that some potential drivers (alternative explanations) of persistent stockpiling are unlikely to play a large role in this program. The first “missing” feature is that customers cannot increase the value of their points by purchasing larger baskets or by stockpiling many points. Recall that in a linear program, the economic value of points is constant over time, so it does not vary with the purchase price or with how many points are redeemed at once. Second, the program does not offer customers any special-tier status or any nonmonetary benefits based on their stockpiles. Thus, stockpiling in this program is not a way for customers to signal their status to others (Drèze and Nunes 2009). Third, when a customer forgets to bring his or her loyalty card for a given purchase, that purchase is not recorded in the data set, nor can the customer earn or redeem points without the card.

Consequently, for every purchase in our data set, the customer presented the loyalty program card, so our model estimation is not affected by purchases in which customers forget to bring their cards. This suggests that lack of interest in the program or forgetting about the program are possible but unlikely explanations for the observed stockpiling in this setting.

2.3. Data Description

The data set contains 346 customers who signed up for the program in January 2008. We tracked their 10,219 purchase occasions from January 2008 through July 2011. We use each customer’s first purchase occasion in January 2008 to calculate initial

stockpile levels but then exclude this from our analysis because no redemptions could have taken place during the first purchase because customers did not have any points to redeem yet.


Redemptions are relatively rare in this program: customers only redeemed in 3% of purchases. Over the entire observation period, 60% of customers never redeemed. We call these customers “nonredeemers.” It does not seem to be the case that the high number of nonredeemers are customers who simply stopped purchasing early during the observation period, because of the 244 customers who purchased at least once in 2011 (i.e., had more than three years of purchasing experience with the program), a majority (51.6%) were nonredeemers. Of the 137 redeemers, 69 redeemed once, 54 redeemed two to four times, 12 redeemed five to seven times, and 2 redeemed more than seven times. Customers who redeemed are also quite valuable. Relative to nonredeemers, customers who redeemed at least once over the 43 months had shopping baskets worth 25% more ($p = .008$), visited the retailer 141% more frequently ($p < .001$), and their recency since the last purchase observed in the data was 164 days shorter ($p < .001$). Table 2.2 further details how customer purchase behavior varies. It shows the distribution across customers of total redemptions, purchase frequency, total spend, and the maximum points stockpiled over the observation period. The large dispersions of these distributions suggest that people may differ in how they value cash relative to points.

Table 2.2: Heterogeneity in customer behavior

Quantiles	0%	25%	50%	75%	100%
Total redemptions occasions	0	0	0	1	15
Total purchase occasions	1	11	22	41	253
Total spend (currency units)	17	2,106	4,519	9,331	211,743
Maximum stockpile (currency units)	0	26	49	91	1,285

Figure 2.2 shows how the main variables we examine in our analysis (redemptions, available stockpiles, and basket prices) evolve over time. It shows that while mean monthly basket prices are highly variable, their level is steady over time. In contrast, stockpiles exhibit a positive trend. The low levels of redemption rates in the cohort's early months reflect that stockpiling behavior may be different during customers' early experiences with the program while they are building up their points.

Figure 2.2: Redemptions and stockpiles over time



In general, stockpiles are large enough to generate nearly 30% in cash savings. Figure 2.3 shows the distribution of the percentage of the basket price that can be reduced by redeeming stockpiled points. As mentioned, the distribution is capped because this program allows customers to cash in up to 50% of the basket during any single purchase occasion. Figure 2.4 plots the monthly averages of these savings. These tend to increase over time, with exceptions during the months of November and December.

Figure 2.3: Percentage of the basket that can be reduced






Figure 2.4: Average monthly potential savings



Next, we illustrate in an exploratory way how the extent of stockpiling differs from what we would expect if customers were randomly deciding whether to redeem (i.e., flipping a coin). For every customer who redeemed at least once, we generate 1,000 simulations of their redemption choices from a sequence of Bernoulli trials with a probability of redemption equal to that of their observed redemption rate. We then compare the average run-lengths (i.e., average number of consecutive purchases with no redemption) of the observed and simulated sequences of redemption choices. Figure 2.5 shows a histogram of the differences between each customer's observed and

expected mean run-length. The majority of customers stockpile longer (i.e., tend to have longer run-lengths) than expected by a Bernoulli sequence with the same observed redemption rate. A one-sided bootstrap Kolmogorov-Smirnov test rejects the null hypothesis ($p < 0.001$) that customers' expected and observed mean run-lengths come from the same distribution, suggesting persistent stockpiling behavior.

Figure 2.5: Comparison of run-lengths



2.4. Redemption Model

MA is presented by first developing the two dimensions (shape and number of accounts), and then by describing how persistent stockpiling behavior is captured through economic, cognitive, and psychological motivations.

Dimension 1: the shape of the value function. A value function determines customer i 's perceived value of gains and losses. It is defined over deviations from a reference point (i.e., current wealth). A function that is linear in both gains and losses is the simplest specification. Equation 2.1 shows our linear specification of $w(x)$, which values gains linearly with a slope a_{iw} and losses with a slope that is steeper than

gains by a loss aversion parameter $\lambda_{iw} > 1$.

$$w(x) = \begin{cases} a_{iw}x & \text{if } x \geq 0 \\ -\lambda_{iw}a_{iw}(-x) & \text{if } x < 0 \end{cases} \quad (2.1)$$

Alternatively a value function can be S-shaped as proposed by prospect theory: concave over gains and convex over losses. Equation 2.2 shows an empirical specification of an S-shaped value function proposed by Tversky and Kahneman (1992). It exhibits diminishing marginal utility for gains and diminishing marginal disutility for losses. This is determined by the curvature parameter $a_{iw} \in (0, 1)$. The loss aversion parameter $\lambda_{iw} > 1$ allows losses to be steeper than gains.

$$w(x) = \begin{cases} x^{a_{iw}} & \text{if } x \geq 0 \\ -\lambda_{iw}(-x)^{a_{iw}} & \text{if } x < 0 \end{cases} \quad (2.2)$$

Dimension 2: the number of accounts. We describe customers as keeping both a “points account” and a “cash account” for evaluating possible redemptions and transactions. Consider a customer who cashes in \$3 in points on a \$10 item. He does so by paying \$3 worth out of his points account and \$7 out of his cash account. The \$3 are a loss to his “points wealth,” and the \$7 are a loss to his “cash wealth.” By redeeming, he also loses the opportunity to earn \$0.10 worth of points (i.e., 1% of the price). This \$0.10 is a “foregone gain” to his points account. Such debits and credits are evaluated separately as gains and losses in each mental account. In the MA and MLA models we let customers have two different value functions, one for the cash account, $w(x)$, and another for the points account, $v(x)$. For notational simplicity, we rescale the points to the same units as cash (i.e., 1 point = 1 dollar). To consider a customer

who instead books transactions using a single mental account, let $w(x) = v(x) \forall x$.

Formally, consider customer i at the cash register on purchase occasion j , who purchases a basket of items with a total price m_{ij} . We model y_{ij} , the customer's decision to redeem conditional on a purchase occurring. When the customer does not redeem ($y_{ij} = 0$), he pays the entire price with cash and earns points on that purchase. When he instead redeems ($y_{ij} = 1$), he pays in part with his stockpiled points s_{ij} and does not earn any points.

A utility-maximizing customer with two mental accounts chooses to either redeem as many points as possible ($y_{ij} = 1$), or not to redeem at all ($y_{ij} = 0$). This binary choice between “redeem as-many-points-as-possible or nothing” is an implication derived from our model (i.e., not an assumption) when both cost functions grow at a decreasing rate (DN 2004 p. 63). The binary consideration set is equivalent to modeling a consumer searching for the optimal redemption amount.

The utility for the payoffs are described mathematically next.

Not redeeming ($y_{ij} = 0$): When a customer does not redeem he pays the entire price m_{ij} in cash. Thus, his cash wealth is reduced by $-m_{ij}$, and so he values this loss with his cash value function: $w(-m_{ij})$. He also earns points worth $r\%$ of the price, where r represents the retailer's reward rate. The gain of $m_{ij}r$ new points is valued with the points value function $v(x)$. Additionally, a customer may subjectively value points differently than cash by a fixed subjective conversion rate $h_i > 0$, so this gain may be perceived as $m_{ij}rh_i$ dollars, with utility $v(m_{ij}rh_i)$.

Equation 1 lays out the utility of not redeeming. E_{ij} denotes the utility attributed to the basket of items. The error term $\epsilon_{ij}^{y_{ij}}$ represents a shock to the customer's preferences that is observed to the customer, but unobserved by the researcher. The

errors are independently distributed around zero, so $E[\epsilon_{ij}^{y_{ij}}] = 0$.

$$u_{ij}(y_{ij} = 0) = E_{ij} + \underbrace{w(-m_{ij})}_{\text{cash loss: full price}} + \underbrace{v(m_{ij}rh_i)}_{\text{gain: earned points}} + \epsilon_{ij}^0 \quad (2.3)$$

Redeeming ($y_{ij} = 1$): When a customer redeems, he pays with a combination of points and cash. The maximum points he can redeem is naturally capped by the available points s_{ij} and a cap κ on how much of the basket price m_{ij} can be redeemed (as described earlier, in the empirical section we study a program that caps the amount redeemed at $\kappa = 50\%$ of the price). Thus, by redeeming, the customer spends $\tilde{s}_{ij} = \min(s_{ij}, m_{ij}\kappa)$ points, and uses $m_{ij} - \tilde{s}_{ij}$ dollars in cash to pay for the remainder of the price. He perceives the $\tilde{s}_{ij}$ points spent as a loss to the points account worth $\tilde{s}_{ij}h_i$, valued with his points value function as $v(-\tilde{s}_{ij}h_i)$. Separately, the cash loss of $m_{ij} - \tilde{s}_{ij}$ dollars is valued by his cash value function: $w(-[m_{ij} - \tilde{s}_{ij}])$. Equation 2 summarizes the utility of redeeming. The customer also incurs a non-monetary transaction cost c_{ij} , reflecting perceived time and effort required to redeem. Note that this customer is forward-looking to the extent that he considers how redeeming affects points available for future use.

$$u_{ij}(y_{ij} = 1) = E_{ij} + \underbrace{w(-m_{ij} + \tilde{s}_{ij})}_{\text{cash loss: discounted price}} + \underbrace{v(-\tilde{s}_{ij}h_i)}_{\text{points loss: redeemed points}} - c_{ij} + \epsilon_{ij}^1 \quad (2.4)$$

Subtracting Equation 2.3 from Equation 2.4, the customer is expected to redeem when his net utility $z(s_{ij}, m_{ij})$ (Equation 2.5) is greater than zero. Equation 2.6 re-writes this inequality to distinguish the expected benefit of redeeming on the left-hand side, from the expected cost on the right-hand side.

$$z(s_{ij}, m_{ij}) = w(-m_{ij} + \tilde{s}_{ij}) - w(-m_{ij}) + v(-\tilde{s}_{ij}h_i) - v(m_{ij}rh_i) - c_{ij} \quad (2.5)$$

$$\underbrace{w(-m_{ij} + \tilde{s}_{ij}) - w(-m_{ij})}_{\text{cash saved}} \geq \underbrace{[v(m_{ij}rh_i)]}_{\text{foregone new points}} + \underbrace{c_{ij}}_{\text{transaction cost}} \underbrace{-v(-\tilde{s}_{ij}h_i)}_{\text{points spent}} \quad (2.6)$$

The benefit of redeeming denotes the value of the *cash saved*: the customer only pays $m_{ij} - \tilde{s}_{ij}$ dollars instead of the full price of m_{ij} dollars. These savings are perceived as a reduced loss (rather than a gain) in the cash account. The right-hand side shows three costs to redeem: foregone points, a transaction cost, and redeemed points.

2.4.1 Three types of motivations

As previously mentioned, foregone points are an *economic incentive* to persistently stockpile up to a point, while transaction costs are a *cognitive* one. Note that these two incentives remain fixed as the customer accumulates more points (i.e., fixed with respect to $\tilde{s}_{ij}$). The third *psychological* incentive is captured by differences in how a customer values points relative to cash. The model allows customers to value points differently than cash in two ways: 1) through the subjective conversion rate h_i , and 2) by having separate mental accounts for cash and point transactions.

2.4.1.1 Stockpiling driven by the conversion rate

For a customer who subjectively values points less than cash by $h_i < 1$, the incentive to redeem “inferior points” for “superior cash” grows as he stockpiles more and more points (regardless of the number of mental accounts). Consider a customer who has a single linear account: $v(x) = w(x)$. For this customer, the cash saved from redeeming is $w(-m_{ij} + \tilde{s}_{ij}) - w(-m_{ij}) = w(\tilde{s}_{ij})$ (because $w(x)$ is linear), and the costs are $-w(-\tilde{s}_{ij}h_i) + [w(m_{ij}h_i r) + c_{ij}]$, so his net expected utility is $z(\tilde{s}_{ij}, m_{ij}) = w(\tilde{s}_{ij}(1 - h_i)) - [w(m_{ij}h_i r) + c_{ij}]$. Notice that the value of redeeming $w(\tilde{s}_{ij}(1 - h_i))$ is positive and grows with $\tilde{s}_{ij}$, while the costs do not vary with $\tilde{s}_{ij}$. Thus, the customer will find himself increasingly motivated to redeem as his stockpile grows. He will

stockpile until $w(\tilde{s}_{ij}(1 - h_i))$ surpasses the value of foregone points and transaction costs.

When points are not considered inferior to cash (i.e., $h_i = 1$), this incentive to redeem disappears, but a customer with multiple-accounts can still be expected to eventually redeem. To see how the relative shapes of the two accounts *alone* can motivate stockpiling, we temporarily set the conversion rate $h_i = 1$, and illustrate the incentives to redeem when the two accounts are linear and S-shaped.

2.4.1.2 Stockpiling driven by booking transactions in multiple accounts

Multiple linear accounts (MLA): Let's compare the costs and benefits from Equation 2.6. The value functions $w(x)$ and $v(x)$ are linear in x , so as redeemable points grow, the benefits grow at a rate equal to the loss slope of the cash value function $w(x)$, and the costs grow at a rate equal to the loss slope of the points value function $v(x)$. With the linear functional form shown in Equation 2.1, the benefits are

$$w(-m_{ij} + \tilde{s}_{ij}) - w(-m_{ij}) = \lambda_{iw}a_{iw}\tilde{s}_{ij},$$

and the costs are (when $h_i = 1$)


$$-v(-\tilde{s}_{ij}) + [v(m_{ij}r) + c_{ij}] = \lambda_{iv}a_{iv}\tilde{s}_{ij} + [a_{iv}(m_{ij}r) + c_{ij}].$$

When $\tilde{s}_{ij}$ is close to 0, a customer has little incentive to redeem: $z(0, m_{ij}) = -[a_{iv}(m_{ij}r h_i) + c_{ij}]$. As he accumulates points, the benefits can grow faster than the costs, depending on the parameter values. Net utility $z(s_{ij}, m_{ij})$ increases with s_{ij} when the slope of losses is greater for cash than for points: $\lambda_{iw}a_{iw} > \lambda_{iv}a_{iv}$.

Figure 2.6 illustrates how the costs and benefits of redeeming evolve with $\tilde{s}_{ij}$ when

this condition is met. In the figure, the net utility of redeeming is positive for any stockpile value greater than s^* , the point at which a customer becomes indifferent between redeeming or not.

Figure 2.6: Cost-benefit tradeoff for MLA



Multiple S-shaped accounts (MA): Similarly, when the two-accounts are instead S-shaped, a customer will eventually redeem as he accumulates points if the slope of losses is sufficiently steeper for cash than for points. Additionally, the current basket price now influences how many redeemable points are needed to entice redemption. In other words, $\partial z / \partial \tilde{s}$ is a function of not only $\tilde{s}_{ij}$ but also of m_{ij} due to the non-linearity of the value functions. Using the S-shaped functional form shown in Equation 2.2 with $h_i = 1$, the benefits are

$$w(-m_{ij} + \tilde{s}_{ij}) - w(-m_{ij}) = -\lambda_{wi}(m_{ij} - \tilde{s}_{ij})^{a_{wi}} + \lambda_{wi}(m_{ij})^{a_{wi}},$$

while the costs are


$$-v(-\tilde{s}_{ij}) + [v(m_{ij}r) + c_{ij}] = \lambda_{vi}(\tilde{s}_{ij})^{a_{vi}} + [(m_{ij}r)^{a_{vi}} + c_{ij}],$$

and so the net expected utility of redeeming is

$$z(s_{ij}, m_{ij}) = -\lambda_{wi}(m_{ij} - \tilde{s}_{ij})^{a_{wi}} + \lambda_{wi}(m_{ij})^{a_{wi}} - \lambda_{vi}(\tilde{s}_{ij})^{a_{vi}} - (m_{ij}r)^{a_{vi}} - c_{ij}. \quad (2.7)$$

As $\tilde{s}_{ij}$ increases, the benefits grow at an increasing rate and the costs grow at a decreasing rate. A growing stockpile will lead to a positive net expected utility of redeeming if the slope of losses is sufficiently steeper for cash than for points (i.e., $a_{iw} > a_{iv}$). Figure 2.7 illustrates how the cost-benefit tradeoff evolves with $\tilde{s}_{ij}$ when $a_{iw} > a_{iv}$. In the figure, the benefits surpass the costs for all stockpile levels above s^* , the point at which a customer becomes indifferent between redeeming or not.

Figure 2.7: Cost-benefit tradeoff for MA



Since the value functions are S-shaped, it may be possible that $w(x)$ is steeper than $v(x)$ for large losses but not for small losses. This occurs when $a_{iw} > a_{iv}$ and $\lambda_{iw} < \lambda_{iv}$ (i.e., the points account has a greater degree of loss aversion than the cash account). An indifference point s^* can still be reached as long as $w(x)$ is steeper than $v(x)$ for large losses (i.e., when there are many points available to redeem). See Appendix 1 for a formal proof.

Having shown how MA explains persistent stockpiling behavior through the economic, cognitive, and psychological motivations, we complete the empirical specification. The propensity to redeem can be written in closed form by assuming that the errors $\epsilon^{y_{ij}}$ are independently and identically distributed Gumbel¹. As a result, we model the redemption choice y_{ij} from a Bernoulli distribution with a probability p_{ij} , where $\text{logit}(p_{ij}) = z(s_{ij}, m_{ij})$. We also let transaction costs c_i be fixed across the purchases of each individual. The empirical MA model and its special cases are summarized in Table 2.3. Due to the linearity of the value functions, MLA and SLA lead to the same empirical model, which we will refer to only as SLA.

Table 2.3: Summary of empirical specifications

Model	$\text{logit}(p_{ij})$
MA	$-c_i - \lambda_{wi}(m_{ij} - \tilde{s}_{ij})^{a_{wi}} + \lambda_{wi}m_{ij}^{a_{wi}} - \lambda_{vi}(\tilde{s}_{ij}h_i)^{a_{vi}} - (m_{ij}rh_i)^{a_{vi}}$
SA	$-c_i - \lambda_i(m_{ij} - \tilde{s}_{ij})^{a_i} + \lambda_i m_{ij}^{a_i} - \lambda_i(\tilde{s}_{ij}h_i)^{a_i} - (m_{ij}rh_i)^{a_i}$
SLA/MLA	$-c_i + \beta_{si}\tilde{s}_{ij} + \beta_{mi}m_{ij}$

The logit of MA's redemption propensity is equivalent to the net utility shown in Equation 2.7. The model allows a_{iw} to be either greater than, equal to, or less than a_{iv} (so it does not restrict the slope of losses to be steeper for cash than for points).

¹The error term $\epsilon^{y_{ij}}$ corresponds to ϵ_{ij}^1 and ϵ_{ij}^0 from Equations 1 and 2

Note that MA is able to capture an interaction between redeemable points $\tilde{s}_{ij}$ and basket price m_{ij} and can reflect framing effects. For example, since the loss function is concave, a \$5 savings over a \$10 purchase is valued more than a \$5 savings on a \$20 purchase (Thaler 1985, Tversky and Kahneman 1981). SA sets $w(x) = v(x)$ (i.e., $\lambda_{iw} = \lambda_{iv}$ and $a_{iw} = a_{iv}$). SLA further restricts $w(x)$ to be linear instead of S-shaped.

2.4.2 Bayesian specification

We complete the empirical specification by modeling individual differences through a hierarchical Bayes framework, which allows for heterogeneity across customers in how they value points relative to cash (Rick, Cryder, and Loewenstein 2008). A hierarchical Bayes model is ideally suited to analyze behavior at the individual level in this case because it leverages data on rare redemption occasions across customers. A prior distribution on the individual-level parameters allows the model to partially pool data across individuals. For MA, let β_i represent a vector of the transformed individual-level parameters: $\beta_i = [\log(\lambda_{wi} - 1), \log(\lambda_{vi} - 1), \text{logit}(a_{wi}), \text{logit}(a_{vi}), \text{logit}(h_i), c_i]$. These transformations restrict the loss-aversion parameters to be greater than one and the curvature parameters as well as the conversion rate h_i to be between 0 and 1. The vector β_i is assumed to follow a multivariate normal prior distribution with mean μ and a precision matrix Ω (Equation 2.8).

$$\beta_i \sim MVN(\mu, \Omega) \quad (2.8)$$

The prior mean μ is chosen to be a vector of zeros. Since this prior distribution governs the transformed parameters β_i , the prior means of the untransformed parameters are 2 for λ_{wi} and λ_{vi} and 0.5 for a_{wi} , a_{vi} , and h_i . The prior precision matrix Ω is a diagonal

identity matrix, meaning that on the untransformed scale, 99.7% (i.e., within three standard deviations) of the prior draws for λ_{wi} and λ_{vi} lie between 1.05 and 21.08, and 99.7% of the prior draws for a_{wi} , a_{vi} , and h_i lie between 0.05 and 0.95.

To complete our Bayesian model specification, as is standard (Gelman et al. 2003), let μ follow a conjugate multivariate normal with mean μ_0 , and precision Ω_0 , and let Ω^{-1} follow a conjugate Wishart distribution with ρ degrees of freedom and an inverse scale matrix R . The hyperparameters are chosen as identity matrices for Ω_0 and R^{-1} , a zero vector for μ_0 and the dimension of the covariance matrix for ρ to make it proper (i.e., 5 for MA). We use the same specifications of the prior and hyperprior distributions for SA and SLA. For SA, β_i excludes λ_v and a_v . For SLA, $\beta_i = [\beta_{si}, \beta_{mi}, c_i]$ and Ω is chosen to be a diagonal matrix times 0.5.

We estimate these specifications using a MCMC sampler with OpenBUGS software. Population-level parameters are sampled from their marginal posterior distributions. These can be directly sampled due to their conjugate hyperpriors. Individual-level transformed parameters β_i are sampled from $p(\beta_i|\mu, \Omega, y)$ (Equation 2.8) with an adaptive Metropolis block sampler.

Next we describe how the net utility for each specification evolves with observed stockpile levels s_{ij} and basket prices m_{ij} . These differences provide alternative and empirically testable stockpiling mechanisms.

2.5. Empirical Identification

Due to the novelty of MA and its highly non-linear nature, we conducted a simulation study to demonstrate parameter recovery for datasets of the size and sparsity used here (see Appendix 2). Analytically, the parameters are identified by the non-

linearity of the value functions. Empirically, the parameters are identified by the variation in both basket prices and stockpiles across purchases associated with and without redemptions. To elaborate on this, we examine how redemption propensity is influenced by an increase in both points and basket price.

The equations in Table 2.4 delineate when additional points may lead to an increase or instead a decrease in a customer's net utility to redeem (z). For simplicity the subscripts are omitted on m and $\tilde{s}$. The predictions of each specification differ by how $\frac{\partial z}{\partial s}$ varies with prices and stockpiles. For SLA, the simplest specification, an increase in redeemable points can increase or decrease net redemption utility in a constant manner, depending on the sign of β_{si} . For MA and SA, $\partial z/\partial \tilde{s}$ varies with both $\tilde{s}_{ij}$ and m_{ij} , meaning that an increase in points influences redemption propensity differently depending on stockpiled points and basket price. A single-account model, as previously explained, requires points to be perceived inferior to cash to predict redemptions.

For MA, when the relative shapes of the two accounts meet the condition that $\lambda_{wi}a_{wi} > \lambda_{vi}a_{vi}$, then additional points can allow $\partial z/\partial \tilde{s}$ to become positive when $\tilde{s}$ is sufficiently large relative to price. Consider two customers who are each about to purchase a \$10 basket of goods. Customer A has an \$8 stockpile and Customer B has a \$1 stockpile. If the company gave each customer 1 additional point, MA can expect this point to increase the propensity to redeem only for A.

Table 2.4: Summary of $\partial z/\partial \tilde{s}$ for specifications

Model	$\partial z/\partial \tilde{s}$
MA	$[\lambda_{wi}a_{wi}(m - \tilde{s})^{a_{wi}-1}] - [\lambda_{vi}a_{vi}h_i^{a_{vi}}(\tilde{s})^{a_{vi}-1}]$
SA	$\lambda_i a_i [(m - \tilde{s})^{a_i-1} - h_i^{a_i}(\tilde{s})^{a_i-1}]$
SLA	β_{si}

Now we examine how an increase in basket price influences redemption propensity. Table 2.5 shows $\partial z/\partial m$ for each model. SLA predicts that an increase in price either increases or decreases net expected utility in a constant manner, depending on the valence of β_{mi} . For MA and SA, $\partial z/\partial m$ is negative regardless of the parameter values, meaning that at the individual-level, these models expect larger prices to be associated with less redemptions conditional on $\tilde{s}$.

Table 2.5: $\partial z/\partial m$ for specifications

Model	$\partial z/\partial m$
MA	$\lambda_{wi}a_{wi}[m^{a_{wi}-1} - (m - \tilde{s})^{a_{wi}-1}] - a_{vi}(rh_i)^{a_{vi}}m^{a_{vi}-1}$
SA	$\lambda_i a_i [m^{a_i-1} - (m - \tilde{s})^{a_i-1}] - a_i(rh_i)^{a_i}m^{a_i-1}$
SLA	β_{mi}

The next section discusses the empirical results and examines how different types of individuals seem to vary by how strongly they consider the economic, cognitive and psychological motivations to persistently stockpile.

2.6. Empirical Results

We separate the data of the linear loyalty program into in-sample and out-of-sample datasets. The in-sample data (January 2008 - August 2010, 7557 purchases) is used to estimate the models. Three independent MCMC chains were run from different starting values, thinning every 50 samples. Convergence was determined using the Gelman and Rubin (1992) diagnostic of between-to-within chain variance. We ran each model for 5000 iterations and the last 3000 iterations of each chain (9000 draws in total) were used for analysis.

After estimating each model using the in-sample data, the parameter draws were used to generate predictive distributions (and posterior point estimates) for a substantial 11-month out-of-sample period (September 2010 - July 2011, 2662 purchases). Obtaining accurate out-of-sample forecasts is challenging because these were not generated using the observed out-of-sample redemption choices. Instead, we used a recursive forecasting approach to ensure that forecasts are generated from stockpiles that are consistent with previously predicted redemption choices. For each draw, the stockpiles at each out-of-sample purchase occasion are updated using a customer's *forecasted* previous redemption choice as shown in Equation 2.9. This approach is analogous to how Erdem, Imai, and Keane (2003) predict ketchup purchases by updating households' latent ketchup inventories with previously forecasted purchases.

$$s_{i(j+1)} = \begin{cases} s_{ij} + rm_{ij} & \text{if } y_{ij}^{predicted} = 0 \\ s_{ij} - \tilde{s}_{ij} & \text{if } y_{ij}^{predicted} = 1 \end{cases} \quad (2.9)$$

Sustaining accurate forecasts during the long 11-month out-of-sample time frame

is particularly challenging in our setting for two reasons. First, errors compound over time through the stockpiles. Early errors in redemption predictions are carried over through the stockpile levels. Second, the models must extrapolate customer behavior as the stockpiles of some customers grow beyond their in-sample levels. Recall Figure 2, which shows mean monthly prices and stockpiles across the in-sample and out-of-sample periods. While stockpiles tend to grow smoothly over time, basket prices are more volatile and do not exhibit a simple time trend. In contrast to most longitudinal analyses, we are more reliant on using the out-of-sample period to assess model validation since redemption behavior seems substantially different at the start of the calibration period and then gradually evolves to a steadier pattern (as per Figure 2). To assess model fit across the empirical models, we break our assessment into fit at the aggregate and individual-level. The former provides an overall assessment while the latter is done to reflect the heterogeneity that may exist in customer motivations to persistently stockpile.

2.6.1 Overall Model Fit

Overall model fit is evaluated with the deviance information criterion (DIC)² (Spiegelhalter et al. 2002) and the negative of the log marginal density (LMD). LMD is calculated using the harmonic mean of the likelihood values evaluated at the posterior draws (Newton and Raftery 1994, Rossi, Allenby, and McCulloch 2005). For both DIC and -LMD, a lower measure indicates a better model fit. In our context, DIC may be the more reliable measure because the harmonic mean of the likelihood values evaluated at the posterior draws can be heavily influenced by a few small outlying draws. Table 2.6 shows these measures for three periods in the data: 2008 (the


²We utilized the median rather than the mean to compute DIC due to a few outlying draws. This may be due to the highly non-linear nature of MA.

cohort's first year in the program), the remainder of the in-sample period (January 2009 - August 2010), and the out-of-sample period (September 2010 - July 2011). These measures are supplemented with Figure 2.8, which compares actual and expected bimonthly redemptions³. As explained in the data description, the different patterns in behavior during 2008 may be due to larger psychological transaction costs of redeeming when a customer joins the program, which may decrease as customers become more familiar with the retailer.

Table 2.6: Deviance Information Criterion and (-Log Marginal Density)

	in-sample (2008)	in-sample (2009+)	out-of-sample
MA	698.0 (352.9)	1139.5 (577.5)	945.9 (872.5)
SA	734.8 (372.3)	1187.0 (597.3)	866.8 (1209.7)
SLA	646.3 (310.6)	1065.2 (506.2)	2998.8 (2052.3)

Figure 2.8: Bimonthly redemptions



In-sample, all three specifications perform similarly in aggregate. However, during

³MA can partially capture seasonal holiday variation (as shown in Figure 8) without the use of holiday dummies because it conditions on purchase behaviors (i.e., frequency and prices) which also tend to systematically vary during the holiday season.

the 11-month out-of-sample period, MA and SA sustain predictions that closely track observed redemptions across time, while SLA is less successful.

We now examine each model’s ability to forecast the aggregate out-of-sample distribution of redemptions across customers. The sum of squared errors between the actual and expected number of customers who made 0 to 5 redemptions during the out-of-sample period (5 redemptions was the maximum observed) are 25 for MA, 101 for SA, and 271 for SLA. In summary, MA then followed by SA are better able to forecast aggregate patterns in redemption behavior relative to SLA. However, despite SLA’s poor aggregate forecasts, a simple one-account model may still be the appropriate model for some individuals, a fact we explore next.

2.6.2 Differences in motivations across customers

MA’s ability to capture economic, cognitive and psychological motivations to stockpile can provide insight to managers seeking ways to better manage their programs. Many retailers closely monitor redemptions because these not only determine a program’s costs and liabilities, but also indicate a program’s effectiveness and helps them to identify valuable customers. MasterCard advisor Bob Konsewicz suggests retailers to encourage redemptions: “encouraging and driving redemptions allows members to engage in and experience the value proposition of the program, and the sooner they do that, the better! (Konsewicz 2007, Kwong, Soman, and Ho 2011). Colloquy and Swift Exchange echoes Mr. Konsewicz after finding that more than one third of the \$48 billion rewards issued in the US each year are never redeemed (Hlavinka and Sullivan, 2011). As an example, Amazon automatically prompts Discover card customers to redeem at checkout. Should Amazon change the way it highlights the reward balance (as shown in Figure 2.9), and instead frame rewards in terms of gains

and losses, such as: “Save up to \$22.07! Apply rewards to reduce your price?” Are high-value customers who have never redeemed likely to respond more favorably to monetary incentives?

Figure 2.9: Redeeming credit card’s point stockpile at Amazon checkout

The screenshot shows the Amazon checkout process. On the left, a section titled "Use Discover Cashback Bonus:" displays "\$22.07 Cashback Bonus available". Below this, there is a text box containing "\$22.07" and an "Apply" button with a mouse cursor hovering over it. On the right, a dark grey button says "Place your order". Below that, the "Order Summary" shows "Items (1):" and a total price of "\$34.99".

Customers may respond differently to various incentives depending on what motivates them to stockpile. As a first step to evaluate how each of the three motivations (economic, cognitive and psychological) differs at the individual level, we examine the posterior distributions of MA’s parameter values (Table 2.7). In the discussion that follows, parameters without the i subscript refer to the untransformed population-level parameters μ , and those with the subscript refer to individual-level estimates.

Table 2.7: Posterior distribution of MA’s population-level parameters

Parameter	Mean	2.5% Bound	97.5% Bound
λ_w	1.21	1.07	1.45
λ_v	1.78	1.28	2.74
a_w	0.347	0.273	0.415
a_v	0.020	0.007	0.042
h	0.523	0.172	0.865
c	1.794	0.937	2.499

Recall that persistent stockpiling can arise from two ways of valuing points differently than cash. The first is a fixed conversion rate h_i . Note from Equation 2.7 that h_i only appears in the MA likelihood as $h_i^{a_{vi}}$. Across individuals, the 95% posterior interval for $h_i^{a_{vi}}$ ranges from 0.978 to 0.984, suggesting some impact in how cash is valued relative to points but one that is unlikely to be a major determinant of redemption

choice. The second is differences in accounts. MA indicates a generally steeper loss curve for cash than points, since the curvature parameters are greater for cash than for points (i.e., $a_w > a_v$) at the population level. At the individual level, the 95% posterior interval for $a_{iw} - a_{iv}$ is also positive, ranging from 0.296 to 0.376 (i.e., greater than 0 for every customer).

The following indicators use the posterior means computed from the MCMC sampler draws to compare the three motivations to stockpile. The last ratio summarizes differences between the cash and points accounts. The closer this “account differences” ratio is to 0, the more similar the accounts. We evaluate the ratio at a basket size of $x = 10$.

- Economic motivation: $\text{mean}((m_{ij}r)^{a_{vi}})$ (i.e., the mean value of foregone points for each individual)
- Cognitive motivation: c_i
- Psychological motivation: h_i and $\lambda_{wi}(x)^{a_{wi}}/\lambda_{vi}(x)^{a_{vi}} - 1$

These indicators are used to segment customers using K-means clustering. Table 2.8 shows the standardized indicator means of the segments with sizes 142, 118, and 87 respectively. Although the exchange rate h_i was used to segment customers, removing it did not impact the segmentation, which is consistent with our previous finding suggesting that this factor is unlikely to be a major determinant of redemption choice.

The three scatter plots in Figure 2.10 show how the standardized indicators vary across customers. As shown in the third panel, fixed costs and account differences are highly correlated: customers in segment 3 have both low fixed costs and accounts that

are very different from each other, while customers in segments 1 and 2 have high fixed costs and accounts that are more similar to each other. A single-account model may be adequate for members of segments 1 and 2, who have relatively similar accounts and are strongly motivated by cognitive fixed costs. Segments 1 and 2 differ in the magnitude of their foregone points. Adding a fourth segment does not add further insight (segments 1 and 2 are partitioned into three groups with high, medium, and low foregone points).

Figure 2.10: How motivations vary by segments





Table 2.8: Segment means for each standardized motivation indicator

	Economic	Cognitive	Psychological	
Segment	Foregone points	Fixed costs	Account differences	Exchange rate
1	-0.67	0.44	-0.43	-0.25
2	0.73	0.51	-0.57	0.77
3	0.12	-1.41	1.47	-0.62

Figure 2.11 compares individuals by their total redemptions and maximum accumulated points. The first panel shows that customers in segments 2 and 3 have greater stockpiles than those in segment 1. The second panel shows that, although these two groups have comparable high levels of stockpiles, every segment 3 individual redeemed at least once, while only a few redeemed in segment 2.

Thus, differences in redemption behavior seem to be driven mostly by the two cognitive and psychological motivations (fixed costs and differences between accounts), which are correlated with each other. Most redeemers are in segment 3. Recall from the scatter plots that segment 3 differs mostly from segments 1 and 2 based on cognitive and psychological motivations, and less so for the economic motivation.

Figure 2.11: Redemptions and stockpiles across segments



2.6.3 Analysis of policies that target each motivation

Having examined differences in motivations across customers, we illustrate how the MA model can help managers to economically evaluate potential policy changes to the program that could lead to higher redemption rates. We consider three hypothetical policy changes to the linear program studied here. Each aims to “lift redemptions” by addressing a particular motivation, so we refer to them as the economic, cognitive, and psychological policies.

Consider the following three hypothetical policies. The economic policy rewards points for every basket price (i.e., rm), regardless of redemption choices. Removing the opportunity costs of redeeming mitigates the economic incentive to stockpile.

The cognitive policy automates redemptions. It mitigates cognitive costs by automatically reducing a customer’s basket price if his stockpile is greater than 15 points (a 25% discount for the average basket price). Customers still retain the option to redeem when their stockpile is below 15 points, but since redeeming is automatic, customers can enjoy rewards more frequently without incurring cognitive costs. Finally, the psychological policy allows customers to redeem up to 100% of the basket price, instead of the current actual policy of capping rewards at 50%. Removing the cap increases redeemable points, and as these increase, the differences in the mental accounts of cash and points lead the benefits of redeeming to grow faster than the costs, so customers may redeem more frequently.

We apply MA to analyze how implementing each of these policies at the start of the out-of-sample period would impact redemptions as well as the firm’s finances. Relative to the current policy, all policies are expected to increase total out-of-sample redemptions over the 11-month period, as compared to the current policy. The economic, cognitive, and psychological policies increase redemptions by 0.6, 17.5, and 3.1 percentage points respectively (Table 2.9). These changes in redemption rates correspond to the previous results suggesting that behavior in this dataset is driven mostly by cognitive and psychological incentives.

Table 2.9: Financial analysis of policy changes

Policy	Rate	Pts. redeemed	Pts. outstanding	$\frac{\text{change in redeemed pts}}{\text{-change in outstanding pts}}$
Current	3.5%	4549	15966	NA
Economic	4.1%	5499 (+21%)	14892 (-7%)	0.75
Cognitive	21.0%	16382 (+260%)	2360 (-85%)	0.87
Psych.	6.6%	9643 (+112%)	10439 (-35%)	0.92

*Redeemed and outstanding points are scaled by some constant

Table 2.9 also compares the financial consequences of the policies. In particular, it compares lost revenue from redeemed points and liabilities from stockpiles outstand-

ing at the end of the out-of-sample period. The last column compares the costs of reducing liabilities by one currency unit. Specifically, it shows the ratio of the change in additional points redeemed (relative to the current policy) divided by the negative of the change in outstanding points (relative to the current policy). A smaller ratio indicates a more cost efficient policy. All else equal (e.g., purchase behaviors), policies that lead to more redemptions are intuitively more costly overall, since more points are redeemed, but also lead to larger reductions in liabilities. Interestingly however, in this dataset, policies that lead to more redemptions are not necessarily the “cheapest” way to reduce firm liabilities. The economic policy is expected to reduce liabilities at the lowest per-unit cost even though it is the least successful at increasing redemptions. Also note that even though the cognitive auto-redemption policy effectively reduces the most liabilities, its cost ratio is comparable to the psychological policy, so it seems to be a relatively cost efficient approach to limit the firm’s point liabilities.

The predictions illustrate how the three policies can influence redemptions and profitability. However, MA is a redemption model that conditions on purchase behavior, so its predictions do not account for the possibility that the policies themselves may lead to changes in how frequently or how much customers purchase. To the extent that greater redemptions increase customer satisfaction and thus purchase frequency, the forecasted redemption rates are an underestimate. Analogously, the forecasts for the psychological policy do not consider potential additional utility that customers may experience from purchasing a basket for “free” by paying entirely with points (Shampanier, Mazar, and Ariely 2007).

2.7. General Discussion

Complicated program rules and undesirable rewards are often blamed for low redemption rates. However, these explanations cannot describe why even in a linear program with simple rules, redemptions are relatively rare. More generally, current research has not successfully addressed why more than \$48 billion rewards issued in the US each year are never redeemed (Hlavinka and Sullivan 2011). This chapter provides insight about redemption behavior in linear programs. An advantage of studying linear programs is that we can isolate motivations to redeem that do not depend on the explicit incentives to stockpile present in non-linear programs.

We model how economic, cognitive, and psychological incentives can motivate customers to stockpile up to a point even though the retailer does not explicitly reward point accumulation. The model is estimated on observational data from a linear loyalty program, and is used to describe how these distinct motivations differ across customer segments. We find that behavior across individuals is mostly driven by cognitive and psychological motivations to redeem (fixed costs and separate accounts), and less so by economic incentives (foregone points). For retailers seeking to improve their strategies to manage redemptions, our findings provide insight into how customers are likely to respond to communication strategies, promotions, and policy changes (i.e., changing the maximum redeemable points). Future research can consider to what extent these three motivations can explain why even in non-linear programs some customers stockpile above and beyond the retailer’s explicit incentives to do so.

Despite the three motivations considered here, we recognize that the observed stockpiling behavior may be consistent with other alternatives. Appendix 3 investigates

a rational forward-looking alternative. The results of the MA model as well as exploratory statistical analyses seem to indicate a lack of empirical support for the alternative in this data set. Studies with more detailed data are better suited to explore yet other alternatives. For example, it is plausible to posit that the types of items purchased may influence the redemption decision. Customers may stockpile to redeem as much as possible on luxury items, since paying in points instead of cash can reduce the guilt of indulging on such items (Kivetz and Simonson 2002). Assessing empirical evidence for this rival explanation would be challenging in our setting since customers redeem on an entire basket of items rather than on any single item. Survey data, if appropriately collected and matched, could be used to learn which item(s) in a basket prompted a customer to redeem. Redemption choice may also influence the types of items that customers purchase. Finally, although we do not rule out the fact that prices can be dependent on the redemption choice, our multiple-accounts model predicts a negative relationship between redemptions and basket prices (Table 2.5), and the data seems consistent with these predictions (the estimated price coefficients of SLA are negative for most customers). While not a complete account of all possible stockpiling alternatives, we hope that this research is a good step towards a more integrated understanding of redemption behavior in linear loyalty programs.

CHAPTER 3 : Market Positioning in a Coalition Loyalty Program: the Value of a Shared Reward Currency

3.1. Motivation

A typical single-brand loyalty program encourages repeat purchases by offering points redeemable for future discounts. While these programs have proliferated, many stores may find them ineffective: if customers tend to visit infrequently, they will earn rewards too slowly to care about them. Such stores may instead benefit from participating in a *coalition loyalty program*: a scheme to provide rewards that can be earned and redeemed faster and with more flexibility.

In a coalition program, customers collect and redeem points across several partner stores. For example, a shoe store may reward customers with points that can be redeemed not only at itself, but also at a clothing store that is a partner of the same coalition program. The more that a store's customers are likely to purchase at other partner stores, the faster they will earn points and the more they will value the reward currency.

How does a coalition's reward currency influence how customers purchase across partner stores? That is the question we focus on in this chapter. While points in single-brand reward programs serve as switching costs that deter purchases at competitors, points from a coalition program instead subsidize other partners since points earned at one store can be redeemed at another. As a consequence, some stores may exert positive reward spillovers to other partners while others may exert negative ones.

Measuring reward spillovers can help coalition program managers assess the mix of

partners in the coalition and to identify partners that provide indirect value to the coalition. Interestingly, successful coalition programs in Europe tend to include hundreds of competing retailers (i.e., Nectar in the UK), while emerging coalition programs in the United States tend to include few partners that do not sell the same product categories. One recent example is Plenti, an American coalition program launched in 2015 with partners such as AT&T, Exxon Mobil, Macy's, Nationwide Insurance, Rite Aid Pharmacy, and Hulu.

While academic research on coalition programs is nascent, a recent study by Dorotic et al. (2011) failed to find spillover effects of store-specific point promotions for the top five non-competing retailers of a Norwegian coalition program. Our study will use data that overcomes two limitations present in their dataset that enable us to better detect reward spillovers. First, most promotions in their dataset were contingent on a relatively high level of spend (30 euros) and so may only have changed the behavior of a small subset of customers. Second, store sales were recorded at the weekly level, so customer heterogeneity could not be accounted for in their study.

Reward spillover effects have been found in a related setting. Pancras, Venkatesan and Li (2015) show that a particular focal restaurant in a city is affected by negative cross effects of rewards from nearby restaurants but positive cross effects from restaurants further away. However, their study only considers purchases across direct competitors (26 restaurants) which each operate their own independent single-brand loyalty program (i.e., points earned at one restaurant can only be redeemed at that particular restaurant).

This chapter develops a customer-level model of reward spillovers and estimates it on a dataset containing rich variation in reward rates across the partner stores of a

European coalition program. Partners in this program are high-end retailers which operate in a diverse set of product categories, from jewelry to electronics. Some partners directly compete by selling goods within the same product categories, while others do not. Although the partners share a common reward currency, their policies on how customers can earn and redeem points are different. First, the rate at which customers earn points differs across stores. Second, points can only be redeemed at select partners. Differences in earning and redeeming policies across partners are a common feature in other popular European coalition programs (i.e., Nectar).

The program we study devalued its reward currency by three-fold during the observation period. In addition, we observe several changes to store-level reward rates at different points in time. Each of the partner stores that we observe in a particular city belongs to a national franchise chain. A franchise can choose whether all of its stores offer either a low, medium, or high level of points for each dollar spent (we refer to the currency unit as a dollar to maintain the anonymity of the program's location). Several franchises decided to change tiers during our observation period.

The devaluation, which effectively reduced reward rates across all partner stores, can be viewed as a quasi-experiment because it was determined by the third-party coalition operator and not by local store managers. These reward rate changes can also be considered exogenous to local (city-level) reward cross effects because they were implemented at all national stores belonging to a particular franchise, and furthermore, franchise managers do not have access to the coalition's data on how their customers purchase at other partners. We provide empirical evidence consistent with our identifying assumption that the changes in reward rates of any given store were exogenous to sales at other local partners.

We estimate own effects and cross effects (i.e., spillovers) of rewards with a model of purchase incidence across the diverse retailers of the coalition program. Motivated by Trusov, Bodapati, and Bucklin (2010), a study which uses latent links to determine influential customers in a social network, our approach uses links to represent reward spillovers. Specifically, we model a network where stores are nodes and links between them represent reward cross effects sent and received across them.

Latent affinity links capture both positive and negative associations in customer-level purchase incidence across multiple stores. We denote the latent links as “affinity links” because these are parameterized with store-level covariates that are utilized to explain the strength and type of reward cross effects, including measures of how similar stores are in terms of both geographic distance and product category overlap.

While conventional store choice models typically compare competitors within an industry, the affinity links in our model characterize the market positioning “landscape” of coalition partners that operate in different high-end retail markets. We apply a heat map to the links to visualize how partner stores “compete” for customer purchases through the reward rates offered to their customers. By positioning partner stores with a common reward spillover metric (i.e., the links), even though several do not sell the same product categories, we contribute to a call for research that can “detect relationships among brands (in our case, stores) that lie outside the conventional definition of product category” (Elrod et al. 2002, p. 230). Finally, the asymmetry between spillover links received versus sent is quantified with measures of competitive clout and vulnerability (Kamakura and Russell 1989). These measures summarize the degree to which each partner influences and is influenced by the reward rates of other coalition partners.

We illustrate that the value of a shared reward currency influences how partner stores compete with each other through reward spillovers. This work on coalition programs is also relevant to firms such as theme parks and casinos which reward customers across their own umbrella of services. Our results allow both managers of partner stores as well as third-party managers of the coalition program to better understand the nature of competition within the coalition.

The rest of the chapter is organized as follows. We first describe the data, including a timeline of the reward policy changes to the coalition program, and provide testable hypotheses of how changes in rewards may have affected how customers purchase across partner stores. Afterwards, we develop a multi-store purchase incidence model. We proceed to present empirical results and conclude with a summary and provide a discussion of future research directions.

3.2. Data Description

From January 2006 to December 2012, we study customer transactions across the top 15 partner stores (those with the most transactions) located in a European city that is internationally known for luxury retail. These stores belong to a coalition program with a wide partner network with presence throughout the country.

Partners are high-end retailers selling highly priced goods in what is known as one of the luxury capital cities of the world. The mean basket price across them is \$134 and the mean inter-purchase time at a store is 10 months. Table 3.1 breaks these down by store. The range of product categories sold across stores is diverse as shown in Table 3.2. Some stores sell the same product categories while others operate in non-overlapping category segments. Each store is a branch from different national franchises. Stores D and E belong to one franchise chain, stores I, L, M and N belong

to a second chain, and the rest operate under different chains. No partner stores entered nor exited the coalition program during the observation period.

We observe 1560 customers who purchased at least once at these stores. At the start of our observation period, 44% had already joined the program. Their transaction data is formatted analogous to a credit card statement, listing each transaction's date, store, and total basket amount. As is typical with such statements, we do not observe which specific items are included in the basket nor purchases made without swiping the coalition's card. The card also tracks purchases made at the other 101 partner stores in the city. We define a customer's purchase occasion as a day in which she is observed purchasing at any partner store in the city (i.e., not necessarily at one of the top 15 partner stores).

Table 3.1: Summary statistics of purchases across stores

Store	Purchases	Customers	Mean basket price	Mean IPT (days) for repeat customers
A	350	155	87	273
B	487	307	299	429
C	933	125	100	181
D	1264	420	125	263
E	367	89	72	189
F	827	178	37	198
G	415	113	178	285
H	332	162	71	284
I	901	387	128	357
J	1611	487	123	321
K	1022	396	261	353
L	374	113	131	347
M	364	110	109	276
N	630	267	152	284
O	731	407	135	358

Table 3.2: Categories of top 15 retail partners

Store	Store Categories	Accepts vouchers
A	Clothing	No
B	Watches/jewelry	Yes
C	Health	Yes
D	Clothing, Home/furniture, Shoes	No
E	Clothing, Home/furniture, Shoes	No
F	Health, Chemist supplies	Yes
G	Clothing	No
H	Clothing	Yes
I	Sport, Shoes	Yes
J	Clothing	No
K	Home/furniture, Electronics/Multimedia	Yes
L	Sport, shoes	Yes
M	Sport, shoes	Yes
N	Sport, shoes	Yes
O	Entertainment, General food, Restaurant/bar, Clothing	Yes

3.2.1. Changes to reward policies


The value of the coalition’s rewards was affected by changes to both the value of points and to the number of points awarded per dollar spent. These decisions were implemented at the national level by both franchise-level managers and the third-party coalition operator. Table 3.3 presents the reward rates (i.e., the dollar value of points earned per \$100 dollars spent) offered by each store during each of the seven time epochs during which rates remained unchanged within an epoch across all stores. The timeline in Figure 3.1 marks with a “V” the six dates in the data during which reward rates changed at any of the stores.

Table 3.3: Dollar value of rewards earned for \$100 spent at each store

Epoch	1	2	3	4	5	6	7
Start date	01/06	10/06	01/07	02/07	09/09	11/09	04/11
A	1.5	1.5	1.5	1.5	1.5	0.5	0.5
B	3.0	3.0	3.0	3.0	3.0	1.0	1.0
C	3.0	3.0	1.5	1.5	1.5	0.5	0.5
D	3.0	3.0	3.0	0.3	0.6	0.2	0.2
E	3.0	3.0	3.0	0.3	0.6	0.2	0.2
F	3.0	3.0	1.5	1.5	1.5	0.5	0.5
G	3.0	3.0	3.0	3.0	3.0	1.0	0.2
H	1.5	3.0	3.0	3.0	3.0	0.5	0.5
I	3.0	3.0	3.0	3.0	3.0	1.0	1.0
J	3.0	3.0	3.0	3.0	3.0	1.0	0.2
K	1.5	1.5	1.5	1.5	1.5	0.5	0.5
L	3.0	3.0	3.0	3.0	3.0	1.0	1.0
M	3.0	3.0	3.0	3.0	3.0	1.0	1.0
N	3.0	3.0	3.0	3.0	3.0	1.0	1.0
O	3.0	3.0	3.0	3.0	3.0	1.0	1.0

Note: The time epochs are between following dates: 01/01/2006, 10/12/2006, 01/05/2007, 02/01/2007, 09/01/2009, 11/20/2009, 04/05/2011, 12/07/2012 and are marked in the timeline of policy changes to the reward program.

Figure 3.1: Timeline of policy changes to reward program



First we describe the changes to the amount of earned points. At the start of the observation period, each partner store offered customers either a low (0.1), medium (0.5), or high level (1) of points earned per dollar spent. Seven of the top fifteen partner stores experienced changes in reward tiers throughout different points in time. These changes were implemented nation-wide by franchise managers to all branches. In addition, in September 2009, the coalition operator increased the amount earned

from purchases at low tier stores from 0.1 points to 0.2 points per dollar spent. This tier-level change only applied to stores D and E in our data.

Points themselves were devalued to a third of their original value in November 2009 by the third-party coalition operator. One hundred points became equivalent to \$1 instead of \$3. Table 3.4 illustrates the impact of this devaluation on the dollar value of rewards earned for \$100 spent at a store from each tier.

Table 3.4: Dollar value of rewards earned for \$100 spent at each tier

Partner reward tier	Before devaluation	After devaluation
Low	0.6	0.2
Medium	1.5	0.5
High	3	1

3.2.2. Identification of spillover effects

To determine whether and to what extent reward spillovers occur across stores, our identifying assumption is that changes in the reward rates of any given store were exogenous to sales at the other fourteen local partner stores. In other words, changes in reward policies at a particular store did not occur because of an anticipated effect on the purchase incidence propensities for the other fourteen stores. The nature of this setting is somewhat similar to the context of Ozturk, Venkataraman and Chintagunta (2016), who assess competitive price reactions with the identifying assumption that car dealership closures represent an exogenous shock to market structure.

First we highlight arguments supporting the identifying assumption and then we provide empirical evidence consistent with the assumption. Recall the three types of changes to the value of rewards across stores: (1) the nation-wide point devaluation, (2) nation-wide changes to low-tier firms, and (3) nation-wide franchise-level changes

of reward tiers. The first two changes (1) and (2) were made by the coalition operator which devalued points to improve its margins from transactions across hundreds of partner stores across the country. Analogously, the operator’s decision to change the number of points awarded by low-tier partners nation-wide was made to reduce the total decrease in the minimum reward rate after the subsequent devaluation. Thus, the impact of these coalition-level changes did not occur in response to a store anticipating an effect on the purchase incidence propensities of the other fourteen local stores.

Changes of the third type, the franchise-level tier changes, also did not occur in response to anticipated changes in local spillover effects in the city studied for several reasons. First, stores do not observe reward spillovers because they do not have access to the coalition’s data on how their customers purchase at other partners. Second, managers are unlikely to have anticipated the impact of potential changes in the reward spillovers from other partners in the city because their impact on store-level profitability is unknown, especially for other partners which do not sell the same categories. Third, each of the seven stores in our study which experienced tier-level changes (which affected all branches within a franchise at the national level) belonged to large franchises with 6-43 branches across the country (Table 3.2).

We present empirical evidence consistent with our identifying assumption. If our assumption holds (i.e., if changes to the value of store-level rewards were not made in response to anticipated local spillover effects), then changes in reward rates at a focal store j should not be predicted by previous sales at other partners $j' \neq j$. For each month t in our dataset, we regress an indicator of whether the value of rewards at a store j changed at time t on the mean sales at all the other stores at time $t - 1$. The slope coefficient is not significantly different from zero (p-value = 0.22), which is

consistent with our identifying assumption.

We now present hypotheses of which mechanisms drive reward spillovers across the partner stores.

3.2.3. Mechanisms driving spillovers

In our model, we will refer to *reward spillovers* or *reward cross effects* as a measure of how much the purchase propensity for one store changes with an increase in the value of rewards offered by other partners. Note that for each pair of partners A and B, there are two spillover “links”. Analogous to a sender-receiver network (e.g., Stephen and Toubia 2010), store A can be described as receiving a spillover from B but also sending a spillover to B.

Spillover links can vary in valence (positive or negative) and shape (symmetric vs. asymmetric). For example, two partners A and B that mutually benefit from each other’s reward rates have spillovers with a symmetric valence, since both directions (A to B) and (B to A) are positive. Spillovers can also vary across customers and over time. The next section introduces how our model will capture these spillover characteristics. Before moving on to that section, however, we first lay out hypotheses on four mechanisms that may potentially drive the characteristics of reward spillovers across the coalition partners.

Mechanism 1: Reward redemption policies

Periodically, the coalition operator mails paper vouchers to customers for the dollar value of their earned points. These vouchers are valid for two years. Analogous to how partner stores offer different reward rates, redemption policies also differ across stores. Throughout the observation period, stores A, D, E, G, J did not accept vouchers from

customers and the others did.

At first it seems puzzling that some partner stores reward customers with points that cannot be redeemed at their own store, particularly because doing so effectively subsidizes purchases at competing partners that do accept them. However, this is a common feature of many large coalition programs such as Nectar in the UK.

Conversations with the coalition operator managers of the program studied here acknowledged that uncertainty of the associated costs is the main reason for why some partners decide not to accept vouchers. In this program, stores that accept vouchers are only reimbursed 90% of their value by the coalition operator. Thus, the impact of accepting vouchers on a store's total margin is uncertain because it depends on the percentage of the basket price that customers save with vouchers. For example, consider a store with a 5% margin. A customer with a \$20 voucher will thus cost the firm \$2. For a given purchase, it is profitable (in the short run) for the firm if the customer redeems the voucher on a \$60 basket but not on a \$20 basket.

Despite the costs associated with voucher redemption, partners that accept vouchers effectively have their products subsidized by the rewards earned across the coalition. Thus, we expect to observe these partners benefitting from more positive reward spillovers, since increases in reward rates provide customers with more points available to spend at their stores. Correspondingly, stores that accept vouchers should also benefit more from increasing their own reward rates.

H1: Stores that accept vouchers receive more positive spillover effects from the reward rates of other stores.

In addition to reward redemption policies, the affinity between partner stores can determine the type of spillovers between them. Consider an extreme case of a pair of

stores with zero affinity. Despite sharing a reward currency, these stores are unlikely to be impacted by each other’s reward rates if they cater to completely different sets of customers. The next two mechanisms describe how two dimensions of store affinity, category overlap and geographic proximity, are expected to determine the types of spillovers across partner stores.

Mechanism 2: Category affinity

As is common in European coalition programs, many partners sell the same product categories. Among those that do, reward rates are a competitive lever analogous to prices. Just like in stand-alone programs, reward rates effectively reduce store prices. Thus, we expect stores to receive more negative spillover effects from partners that compete in the same categories.

H2: Pairs of stores that compete in the same product categories experience more negative cross effects from each other.

The overlap in product categories sold (see Tables 3.2 and 3.5) is measured with the Jaccard similarity coefficient, a measure commonly used for machine learning (e.g., Netzer et al. 2012). Let A_c^j equal 1 if store j sells category c , and 0 otherwise. Category proximity (i.e., Jaccard similarity) between stores j and k is calculated as shown in Equation 3.1 and ranges from zero to one. Stores that do not overlap in any product category have the minimum category proximity of 0, regardless of how many categories they sell. For example, store A has a category proximity of 0 with both B and L, even though B sells one category and L sells two. Analogously, stores that compete in 100% of their categories have a category proximity equal to one.

$$Proximity_{jk}^{Categ} = \frac{\sum_c \min(A_c^j, A_c^k)}{\sum_c \max(A_c^j, A_c^k)} \quad (3.1)$$

Table 3.5: Category proximity between stores

Store	B	C	D	E	F	G	H	I	J	K	L	M	N	O
A	0.0	0.0	0.33	0.33	0.0	1.0	1.0	0.0	1.0	0.0	0.0	0.0	0.0	0.25
B		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
C			0.0	0.0	0.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
D				1.0	0.0	0.33	0.33	0.25	0.33	0.25	0.25	0.25	0.25	0.17
E					0.0	0.33	0.33	0.25	0.33	0.25	0.25	0.25	0.25	0.17
F						0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
G							1.0	0.0	1.0	0.0	0.0	0.0	0.0	0.25
H								0.0	1.0	0.0	0.0	0.0	0.0	0.25
I									0.0	0.0	1.00	1.0	1.0	0.0
J										0.0	0.0	0.0	0.0	0.25
K											0.0	0.0	0.0	0.0
L												1.0	1.0	0.0
M													1.0	0.0
N														0.0

Mechanism 3: Geographic affinity

The second measure of store affinity is the geographic proximity between the stores within the city. We expect pairs of stores that are geographically closer to benefit from increases in each other’s reward rates. This may occur because customers may incur travel costs to explore new neighborhoods in the city, so nearby stores benefit when a partner attracts customers to a neighborhood through attractive rewards.

H3(a): Stores receive more positive reward spillovers from nearby partners.

Geographic proximity is calculated as $Proximity_{jk}^{Geo} = \frac{1}{geo.dist_{jk}+1}$, where $geo.dist_{jk}$ is the geographic distance between the stores. On average, two stores are located

only 2.17 kilometers (less than 1.35 miles) apart, and the maximum distance is 7.8 kilometers (see Table 3.6). If two stores are at the same location, the proximity measure equals 1.

Table 3.6: Geographic distances between stores (kilometers)

Store	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
A		0.0	0.0	1.1	1.1	0.0	0.0	0.0	1.1	0.0	0.0	6.7	4.4	3.3	0.0
B			0.0	1.1	1.1	0.0	0.0	0.0	1.1	0.0	0.0	6.7	4.4	3.3	0.0
C				1.1	1.1	0.0	0.0	0.0	1.1	0.0	0.0	6.7	4.4	3.3	0.0
D					2.2	1.1	1.1	1.1	0.0	1.1	1.1	5.6	3.3	4.4	1.1
E						1.1	1.1	1.1	2.2	1.1	1.1	7.8	5.6	4.4	1.1
F							0.0	0.0	1.1	0.0	0.0	6.7	4.4	3.3	0.0
G								0.0	1.1	0.0	0.0	6.7	4.4	3.3	0.0
H									1.1	0.0	0.0	6.7	4.4	3.3	0.0
I										1.1	1.1	5.6	3.3	4.4	1.1
J											0.0	6.7	4.4	3.3	0.0
K												6.7	4.4	3.3	0.0
L													6.7	5.6	6.7
M														7.8	4.4
N															3.3
O															

We also expect to see a negative interaction between geographic and category affinity. While a focal store may benefit from the attractive reward rates of a nearby partner if new customers are drawn to the neighborhood, the net effect is likely to reverse if the nearby store sells similar product categories, since its attractive rewards effectively lower the prices of similar goods to the focal store's own customers.

H3(b): Stores receive more negative reward spillovers from nearby partners that also compete in the same product categories, relative to nearby partners that do not sell the same product categories.

Mechanism 4: Coalition-wide policy changes to the program

The last mechanism describes how the 2009 policy changes are expected to have affected spillovers. Recall that the coalition operator devalued the points currency by a third in November 2009. To compensate for the devaluation, at the same time the operator increased the fungibility of vouchers in three ways: (1) by mailing them more frequently (monthly instead of quarterly), (2) by increasing the denominations from only \$15 vouchers to \$5, \$10, \$20, \$50, and \$100 denominations, and (3) by allowing customers to redeem the vouchers up to 100% of the basket price, instead of previously 30%. Although these two changes to points and vouchers occurred at the same time, we can discern the effects of each because some partner stores accept vouchers while others do not. This is because the devaluation should have led to a decrease in the overall magnitude of spillovers across stores. Also, rewards offered are worth less, so their impact on other stores should be of a smaller magnitude.

H4(a): The points devaluation led to a decrease in the average magnitude of spillovers across stores.

The change to the voucher policy, in contrast, makes it easier for customers to redeem, and so although H4(a) expects the level of spillovers to decrease, we expect to observe a positive *difference* between the spillovers received by stores that accept vouchers vs. those that do not.

H4(b): The increased fungibility of vouchers increased the difference between spillovers received by stores that accept voucher redemptions relative to spillovers received by stores that do not accept voucher redemptions.

The next section presents a customer-level model of multi-store purchase incidence behavior that pools information across both customers and stores to measure reward

spillovers and detect how these changed with the devaluation to the rewards currency. The high-end nature of the coalition partners enhances the need to understand customer choice across them, as well as the methodological challenges that would arise due to low co-purchasing behavior. Recall that the mean inter-purchase time of a repeat customer at the typical store is 10 months. Most customers (44%) only purchased at one of the top 15 stores during the entire observation period (Table 3.7). Customers who did purchase at more than one store rarely did so in the same day: 2% of all 15818 purchase occasions included purchases at two or more of the top stores (Table 3.8).

Although the nature of the coalition program data (sparseness and many stores) requires a statistical model able to pool data across customers and stores in order to determine whether reward spillovers occur and to what extent, we do observe some modest aggregate patterns in the data that suggest the presence of own effects and spillover effects of rewards. Consider clothing stores D and E which devalued their reward rate from 3 to 0.3 dollars per \$100 dollars spent in February 2007. Yearly sales (i.e., the number of purchases) at these two stores decreased by 57% in the following year. At the same time, yearly sales at other top partners that also sell clothes increased by 35.7%, while sales at other top partners that did not sell clothes increased only by 16.1%.

Table 3.7: Percent of customers who purchased at n of the top stores

# Stores	1	2	3	4	5	6	7	8	9	10	11	12-15
% Customers	43.9	22.2	12.9	8.3	5.4	3.3	1.7	1.2	0.4	0.4	0.1	-

Table 3.8: Number of top partners visited on a purchase occasion

# stores	0	1	2	3	4
# occasions	5620	9804	379	14	1

3.3. A Multi-store Purchase Incidence Model

We observe customer $i \in \{1, \dots, I\}$ shopping in the city on N_i purchase occasions. Let n denote her n th purchase occasion and let $t(n)$ denote the calendar time at which the purchase occasion took place. The dependent variable y_{ijn} equals one if i purchased at store $j \in \{1, \dots, J\}$ at her n th purchase occasion, and zero otherwise.

The net utility of the customer of purchasing versus not purchasing at store j in purchase occasion n is modeled with a deterministic component V_{ijn} and an error term ϵ_{ijn} that is independently drawn from an extreme value distribution. This leads to the common logit form for the propensity to purchase given by $p_{ijn} = \frac{e^{V_{ijn}}}{1 + e^{V_{ijn}}}$.

Equation 3.2 specifies V_{ijn} as a function of (1) utility that is not related to rates at which reward points are earned, (2) own effects: utility from the value of the reward rate $R_{jt(n)}$ offered by store j at the time of purchase occasion n , and (3) cross effects: utility from the value of rewards offered by other partner stores $k \neq j$. This specification is analogous to the utility model of own and cross price effects in multi-category choices within a store by Manchanda, Ansari and Gupta (1999). Given the large number of stores considered and the sparsity present in the data, cross effects in V_{ijn} are specified to explicitly capture the source of correlation among alternatives, such that the remaining errors ϵ_{ijn} can be considered independent (Train 2003).

$$V_{ijn} = \text{Non-Reward Utility}_{ij} + \text{Own Effects}_{ijn} + \text{Cross Effects}_{ijn} \quad (3.2)$$

Non-reward utility is captured with a customer-level intercept α_i , a store level intercept λ_j , and a term capturing each customer's affinity to each store θ_{ij} . The intercept λ_j of the store with the most transactions is fixed to zero to separately identify cus-

tomers and store intercepts.

$$\text{Non-Reward Utility}_{ij} = \alpha_i + \lambda_j + \theta_{ij} \quad (3.3)$$

The effect of store j 's reward rate on its own purchase propensity is denoted by $\gamma_{j \rightarrow j, i, n}$.

$$\text{Own Effects}_{ijn} = \gamma_{j \rightarrow j, i, t(n)} \quad (3.4)$$

Own effects are modeled as linear functions of a store's own rewards which can vary across customers i and over time. In particular, recall that reward rates across the coalition changed six times, leaving seven epochs during which reward rates remained unchanged.

$$\gamma_{j \rightarrow j, i, t(n)} = \omega_{j, i, t(n)}^{own} R_{jt(n)} \quad (3.5)$$

The weights $\omega_{j, i, t(n)}^{own}$ on reward rates are allowed to be different for stores that accept vouchers vs. those that do not. The covariate $Voucher_j$ equals one if j accepts voucher redemptions and zero otherwise. The κ coefficients are allowed to differ for purchase occasions that took place before and after the program-wide devaluation (i.e., between epochs 5 and 6).

$$\omega_{j, i, t(o_i)}^{own} = \kappa_{0it(n)} + \kappa_{1it(n)} Voucher_j \quad (3.6)$$

We can re-write own effects as:

$$\gamma_{j \rightarrow j, i, t(n)} = \kappa_{0it(n)} R_{jt(n)} + \kappa_{1it(n)} R_{jt(n)} Voucher_j$$

which includes a main effect for rewards and an interaction term between rewards and vouchers.

Cross effects from the rewards of other stores $k \neq j$ are modeled by summing links of an affinity network where $\gamma_{k \rightarrow j, i, t(n)}$ denotes the spillover link sent by store k to store j at the time of purchase occasion n .

$$\text{Cross Effects}_{ijn} = \sum_{k \neq j} \gamma_{k \rightarrow j, i, t(n)} \quad (3.7)$$

A link from store k to j multiplies the reward value offered by k at time $t(n)$ with a cross effect weight $\omega_{k \rightarrow j, i, t(n)}^{cross}$ (Equation 3.8) that can vary across customers and time.

$$\gamma_{k \rightarrow j, i, n} = \omega_{k \rightarrow j, i, t(n)}^{cross} R_{kt(n)} \quad (3.8)$$

The cross effect weights $\omega_{k \rightarrow j, i, t(n)}^{cross}$ are a linear function of voucher policies as well measures of the affinity between stores. Let A_{jk}^{Categ} and A_{jk}^{Geo} denote the mean-centered values (over time and across stores) of category and geographic proximity ($Proximity_{jk}^{Categ}$ and $Proximity_{jk}^{Geo}$). The last term in Equation 3.9 captures an interaction between category and geographic affinity. The next section tests the first three hypotheses using the coefficients of these weights, and hypotheses 4(a) and 4(b) using the links in Equation 3.8.

$$\omega_{k \rightarrow j, i, t(n)}^{cross} = \psi_{0it(n)} + \psi_{1it(n)} Voucher_j + \psi_{2it(n)} A_{jk}^{Categ} + \psi_{3it(n)} A_{jk}^{Geo} + \psi_{4it(n)} A_{jk}^{Categ} A_{jk}^{Geo} \quad (3.9)$$

By summing across the affinity links we can re-write the cross effects of rewards as $W_{jt(n)}\psi_{it(n)}$, where the covariates $W_{jt(n)}$ include aggregate statistics of j 's affinity network.

$$\text{Cross Effects}_{ijn} = \sum_{k \neq j} \gamma_{k \rightarrow j, i, t(n)} \quad (3.10)$$

$$= W_{jt(n)}\psi_{it(n)} \quad (3.11)$$

$$= \psi_{0it(n)} \sum_{k \neq j} R_{kt(n)} + \psi_{1it(n)} \sum_{k \neq j} R_{kt(n)} Voucher_j$$

$$+ \psi_{2it(n)} \sum_{k \neq j} R_{kt(n)} A_{jk}^{Categ} + \psi_{3it(n)} \sum_{k \neq j} R_{kt(n)} A_{jk}^{Geo} + \psi_{4it(n)} \sum_{k \neq j} R_{kt(n)} A_{jk}^{Categ} A_{jk}^{Geo} \quad (3.12)$$

3.3.1 Policy changes

Let $\vec{\beta}_{it(n)}$ denote the vector of parameters that characterize both own effects and cross effects $\{\kappa_{0it(n)}, \kappa_{1it(n)}, \psi_{0it(n)}, \psi_{1it(n)}, \psi_{2it(n)}, \psi_{3it(n)}, \psi_{4it(n)}\}$. We allow these to reflect permanent changes after t^* (November 2009), the time at which the coalition devalued points and increased voucher fungibility. Changes after the devaluation are denoted

by the vector $\bar{\delta} = \{\delta_{k_0}, \delta_{k_1}, \delta_{\psi_0}, \delta_{\psi_1}, \delta_{\psi_2}, \delta_{\psi_3}, \delta_{\psi_4}\}$. These parameters $\bar{\delta}$ capture potential changes after the devaluation to how cross effects are driven by voucher policies and store affinity.

$$\vec{\beta}_{it(n)} = \begin{cases} \beta_i & \text{if } t(n) \leq t^* \\ \beta_i + \bar{\delta} & \text{if } t(n) > t^* \end{cases} \quad (3.13)$$

3.3.2 Prior distributions

Prior distributions are specified to complete the model. The prior distributions for the non-reward utility parameters are specified as follows. The customer-store attractiveness terms θ_{ij} are normally distributed with a mean 0 and a variance $I_I \sigma_\theta^2$. Customer intercepts α_i each have a normal prior with a zero mean and variance σ_α^2 . The store intercepts λ_j each have a normal prior with a mean μ_λ and variance σ_λ^2 . For the hyperpriors, the variances σ_α^2 , σ_λ^2 and σ_θ^2 each follow an Inverse-Gamma distribution with shape and scale parameters equal to 0.5. The prior mean μ_λ has a normal hyperprior with mean zero and a standard deviation of 0.5.

The prior and hyperprior distributions for the parameters governing own and cross reward effects are specified as follows. The coefficients β_i have a multivariate Normal prior with mean $\bar{\beta} = \{\bar{\kappa}_0, \bar{\kappa}_1, \bar{\psi}_0, \bar{\psi}_1, \bar{\psi}_2, \bar{\psi}_3, \bar{\psi}_4\}$ and variance Σ_β . Let K represent the number of parameters in $\bar{\beta}$ (seven in this case). The conjugate prior for the mean $\bar{\beta}$ is a Multivariate Normal with mean zero and variance KI_k . The conjugate prior for the variance Σ_β is an Inverse-Wishart with K degrees of freedom and a location matrix equal to KI_k . Each element in δ has normal hyperprior with mean zero and a standard deviation of 0.5.

3.3.3 Benchmarks

The fit of the affinity model is compared with two nested benchmarks called OWN and NRU which both exclude cross effects. NRU only includes non-reward utility and thus does not include any information on reward rates. OWN includes both non-reward utility and own effects, but not cross effects. Both of these benchmarks are consistent with the null hypothesis of no reward spillovers across partner stores. Because the benchmarks do not model cross effects, these do not contain information on the geographic and category affinity between stores.

Table 3.9: Summary of empirical models

	Deterministic utility V_{ijn}
Affinity network	$(\alpha_i + \lambda_j + \theta_{ij}) + \gamma_{j \rightarrow j, i, t(n)} + \sum_{k \neq j} \gamma_{k \rightarrow j, i, t(n)}$
OWN	$(\alpha_i + \lambda_j + \theta_{ij}) + \gamma_{j \rightarrow j, i, t(n)}$
NRU	$(\alpha_i + \lambda_j + \theta_{ij})$

3.4. Empirical Results

A Markov Chain Monte Carlo sampler coded in the R software was used to estimate the model with two independent chains of 80,000 iterations. The last 40,000 iterations of each chain were used for analysis after thinning every 20 iterations. Convergence was determined using the Gelman and Rubin (1992) diagnostic of between-to-within chain variance.

Parameters governing the prior distributions μ_λ , σ_α^2 , σ_λ^2 , σ_θ^2 , $\bar{\beta}$, and Σ_β are sampled from their marginal posterior distributions. These parameters can be directly sampled due to their closed-form marginal posterior distributions. Each of the remaining parameters α_i , λ_j , θ_{ij} , β_i , $\bar{\delta}$ are sequentially sampled from their conditional posterior distributions using a random-walk Metropolis sampler. The step sizes for each of these

parameters is adapted during the first 20,000 iterations to maintain acceptance rates between 30% and 40%. Appendix 4 illustrates parameter recovery for the Affinity model.

First we show the model fit relative to the benchmarks. Then we examine the spillover links and use them to test the hypotheses. Third, the links are visualized with a heat map and the asymmetry between links received vs. sent is quantified with measures of competitive clout and vulnerability (Kamakura and Russell 1989).

3.4.1. Model fit

The affinity network model improves overall in-sample fit relative to both benchmarks (Table 3.10), as measured by the mean log likelihood (LL), the deviance information criterion (DIC) and the log-marginal density (LMD) (Spiegelhalter et al. 2002, Newton and Raftery 1994, Rossi, Allenby, and McCulloch 2005). A smaller magnitude indicates a better fit for each of these measures. The OWN benchmark does not improve overall fit over NRU. The two models have a nearly identical fit with all three measures. Although OWN has a slightly lower mean log likelihood over draws, the two are not substantially different. OWN achieves a larger maximum draw of the log likelihood than NRU (-25802 for OWN versus -25812 for NRU). Despite the slightly lower mean log likelihood, OWN achieves a better DIC because it has a slightly better likelihood evaluated at the posterior mean of the population-level parameters (-23117 for OWN versus -23121 for NRU).

Table 3.10: Measures of overall model in-sample fit

Model	Mean LL	DIC	LMD
Affinity network	-25049	37351	-25367
OWN	-26099	40277	-26366
NRU	-26092	40299	-26328

The affinity network model achieves a better overall fit because the spillover effects allow it to better predict patterns in how customers purchase across stores. To illustrate this, Table 3.11 compares the posterior means of three error statistics across the models. These posterior means are calculated from the error statistics from a large number of datasets simulated from the posterior distributions of model parameters (Gelman et al. 2003).

The first statistic measures transactions, the second measures store patronage, and the third measures cross-store patronage. More specifically, the first row shows the sum of squared errors (SSE) between the actual vs. predicted number of transactions that each customer made at each store. The second row shows the hit rate of whether each customer is predicted to have patronized (i.e., made at least one purchase at) each store. The third row shows the sum of squared errors between the actual and expected number of customers who co-patronized each *pair* of stores (i.e., for stores A and B, the number of customers who purchased at least once at A and at least once at B). The Affinity model achieves the best fit in each of these measures. Relative to NRU, OWN improves predictions of cross-store patronage because the own effects allow it to improve the fit across time epochs of how customers shop at each store.

Table 3.11: Number of stores patronized by the average customer

	Affinity	OWN	NRU
SSE transactions	13,510	14,500	14,650
Hit rate store patronage	89.5	87.9	88.3
SSE cross-store patronage	14,434	20,958	25,708

Later in this section we will show that stores B, C, and F are found to almost all benefit from each other’s positive reward spillovers. To further illustrate an example of how the affinity model better fits the data, Table 3.12 compares the observed vs. expected number of shared customers between all three stores, and between each pair

of stores. Numbers with an asterisk denote that the observed number falls within 95% posterior intervals.

The Affinity model contains three of the four true values within posterior intervals. However, the observed fourth value (i.e., the number customers who purchased at both B and F) which is not contained within the interval is only 1 customer greater than the upper 95% interval (73). Relative to the benchmarks, the Affinity model provides superior mean estimates for all values except for the number of customers shared by C and F.

Table 3.12: Customers shared between three synergistic partners

	Observed	Affinity	OWN	NRU
B, C, F	21	18.1*	14.5	12.9
B, C	54	49.0*	44.2*	40.5
C, F	37	43.8*	38.4*	35.3*
B, F	74	62.4	57.2	53.2

3.4.2. Spillovers

Table 3.13 shows the average own and cross effects (across customers and stores) for each of the time epochs during which reward rates in the coalition program remained unchanged. While the magnitude of own effects tends to be greater than the magnitude of the average store-to-store cross effect, for a given store, the net impact of the sum of the cross effects from the 14 other stores is greater than own effects. The last row of the table shows the average sum of cross effects received. After the program's currency devaluation (between epochs 5 and 6), we find that the magnitude of the average own effect increased from 1.07 to 1.39 while the magnitude of the average cross effect decreased from -0.22 to -0.15. Together, these changes indicate that the devaluation of the reward currency reduced the competitive reward interactions

among the coalition’s partner stores.

Table 3.13: Evolution of own and cross effects

Epoch	1	2	3	4	5	6	7
Avg. own effect $\gamma_{j \rightarrow j}$	1.19	1.24	1.13	1.06	1.07	1.39	1.39
Avg. cross effect $\gamma_{k \rightarrow j}$	-0.26	-0.27	-0.26	-0.21	-0.22	-0.15	-0.13
Avg. $\sum_{k \neq j} \gamma_{k \rightarrow j}$	-3.87	-4.04	-3.87	-3.21	-3.31	-2.2	-1.95

3.4.2.1 Parameter estimates for non-reward utility and own effects

Table 3.14 shows the posterior distributions of the population-level parameters for non-reward utility and own effects of rewards. We now evaluate the directions and significance of the own effect parameters before vs. after the devaluation. However, note that we cannot directly compare the magnitudes of the parameters before and after the devaluation, since the reward rates across the coalition changed at the same time. Before the devaluation, we find that own-effects are positive ($\bar{\kappa}_0 > 0$). Furthermore, stores that accept vouchers enjoy greater positive own-effects than those that do not ($\bar{\kappa}_1 > 0$). The posterior intervals of these positive own effects are significant because they do not contain zero. After the devaluation, own effects for stores that accept vouchers are still positive and significantly different from zero. However, the own effects for stores that do not accept vouchers become insignificant (i.e., zero is contained within the posterior interval for $\bar{\kappa}_0 + \delta_{\kappa_0}$). These findings suggest that stores that do not accept vouchers were the most affected by the devaluation of the reward currency.

Table 3.14: Population-level parameters for non-reward utility and own effects

	Mean	5%	95%
μ_λ	-1.49	-2.71	-0.24
$\bar{\kappa}_0$	0.19	0.09	0.31
$\bar{\kappa}_1$	0.37	0.21	0.54
$\bar{\kappa}_0 + \delta_{\kappa_0}$	0.08	-0.18	0.32
$\bar{\kappa}_1 + \delta_{\kappa_1}$	2.51	2.04	3.00

3.4.2.2 Tests of hypotheses on the mechanisms that drive cross-effects

Tables 3.15 and 3.16 show the posterior distribution of the population-level parameters governing cross effects, before ($\bar{\psi}$) and after ($\bar{\psi} + \delta$) the devaluation. We now compare the valence and significance of each parameter before and after the devaluation. We cannot compare the magnitudes of $\bar{\psi}$ and $\bar{\psi} + \delta$ without also accounting for changes in reward rates. The magnitude comparison is shown in the next subsection.

Overall, the baseline-level of cross effects across partner stores is negative, both before and after the devaluation: $\bar{\psi}_0 < 0$, $\bar{\psi}_0 + \delta_{\psi_0} < 0$. Everything else equal, increases in reward rates at one partner store are found to decrease purchase incidence at other partner stores. This effect is significant because zero is not contained within the 95% posterior intervals of $\bar{\psi}_0$ and $\bar{\psi}_0 + \delta_{\psi_0}$.

The remaining cross effect parameters describe how cross effects vary for stores that accept vouchers and between stores that are closer in category and geographic affinity. We use the valence of each of these parameters to test Hypotheses 1, 2, 3(a) and 3(b).

Hypothesis 1 is supported if stores that accept vouchers are expected to receive more positive (or less negative) cross effects than stores that do not accept voucher redemptions. Consistent with Hypothesis 1, both $\bar{\psi}_1$ and $\bar{\psi}_1 + \delta_{\psi_1}$ have positive posterior means, and their posterior intervals do not contain zero.

Hypothesis 2 is supported if spillovers tend to be more negative between pairs of stores that overlap in categories sold. While the posterior means of both $\bar{\psi}_2$ and $\bar{\psi}_2 + \delta_{\psi_2}$ are negative, consistent with the hypothesis, this effect is not significant for $\bar{\psi}_2$ because zero is contained within the 95% posterior interval. More specifically, 85% of the draws for $\bar{\psi}_2$ are negative, and 99% of the draws for $\bar{\psi}_2 + \delta_{\psi_2}$ are negative.

Hypothesis 3(a) is supported if spillovers tend to be more positive (or less negative) between pairs of nearby stores. Consistent with this hypothesis, the means and 95% posterior intervals of both $\bar{\psi}_3$ and $\bar{\psi}_3 + \delta_{\psi_3}$ are positive (i.e., zero is not contained within the posterior intervals). One hundred percent of the draws for both $\bar{\psi}_3$ and $\bar{\psi}_3 + \delta_{\psi_3}$ are positive.

Finally, Hypothesis 3(b) predicts a negative interaction between nearby stores that are also competitors. We observe this negative interaction before the devaluation: 100% of the draws for $\bar{\psi}_4$ are negative. After the devaluation, however, the interaction term $\bar{\psi}_4 + \delta_{\psi_4}$ is not different from zero (i.e., zero is contained within the 95% posterior interval).

Table 3.15: Posterior distribution of $\bar{\beta}$ (before devaluation)

$\bar{\beta}$	Mean	5%	95%	Significant
ψ_0	-0.16	-0.18	-0.15	yes
ψ_1	0.09	0.07	0.11	yes
ψ_2	-0.04	-0.12	0.02	no
ψ_3	0.21	0.16	0.26	yes
ψ_4	-0.53	-0.67	-0.40	yes

Table 3.16: Posterior distribution of $\bar{\beta} + \delta$ (after devaluation)

$\beta + \delta$	Mean	5%	95%	Significant
$\psi_0 + \delta_{\psi_0}$	-0.42	-0.45	-0.38	yes
$\psi_1 + \delta_{\psi_1}$	0.28	0.23	0.33	yes
$\psi_2 + \delta_{\psi_2}$	-0.23	-0.42	-0.06	yes
$\psi_3 + \delta_{\psi_3}$	0.33	0.21	0.45	yes
$\psi_4 + \delta_{\psi_4}$	0.31	-0.08	0.72	no

Table 3.17 shows additional support for these hypotheses at the individual level. Although most consumers have parameter means consistent with most hypotheses (column 3), the posterior intervals of most individuals tend to be very wide and contain zero because we observe few purchases per individual. The fourth column shows the percent of consumers for whom both (1) the hypothesis is supported and (2) zero is not contained within the 95% posterior interval. The table is consistent with the population-level parameters. Specifically, while most hypotheses are supported, (1) the negative effects of competition are mild, and (2) a negative interaction between competing stores and nearby stores is observed before but not after the devaluation.

Table 3.17: Support of hypotheses at the individual level

$\beta + \delta$	$t(n)$	Consistent	Consistent and significant
$E[\psi_{1it(n)}] > 0$	$< t^*$	82%	9.7%
$E[\psi_{2it(n)}] < 0$	$< t^*$	62%	0.9%
$E[\psi_{3it(n)}] > 0$	$< t^*$	87%	7%
$E[\psi_{4it(n)}] < 0$	$< t^*$	98%	2.6%
$E[\psi_{1it(n)}] > 0$	$\geq t^*$	100%	74%
$E[\psi_{2it(n)}] < 0$	$\geq t^*$	84%	3.7%
$E[\psi_{3it(n)}] > 0$	$\geq t^*$	94%	19%
$E[\psi_{4it(n)}] < 0$	$\geq t^*$	9.3%	0%

Overall, the empirical support for the first three hypotheses sheds light on how a diverse store portfolio benefits from a shared reward currency. In particular, we find

evidence that geographic proximity among partner stores attenuates competitive pressures from stores that do sell similar categories. In addition, the results suggest that offering more generous reward policies (in particular, accepting voucher redemptions and offering higher reward rates) helps stores to compensate for the presence of competitors within the coalition. Finally, we find that the currency devaluation seems to have reduced the intensity of competition for nearby category competitors.

We now use the affinity links to test Hypotheses 4(a) and 4(b). These predict how spillovers were affected by the coalition’s policy changes to both the value of rewards and voucher fungibility. To evaluate both of these hypotheses, we compare the value of the affinity links using reward rates across the coalition offered right before and after the devaluation (these correspond to the rates during the 5th and 6th epochs shown in Table 3.3).

Hypothesis 4(a) is supported if the average magnitude of spillovers decreased after the devaluation. We test this hypothesis with the statistic $s_{H4(a)}$, where the expectation is an average over all customers i , focal stores j , and other stores k . We expect $s_{H4(a)} > 0$.

$$s_{H4(a)} = E[|\gamma_{k \rightarrow j, i, t \in \text{epoch}5}| - |\gamma_{k \rightarrow j, i, t \in \text{epoch}6}|] \quad (3.14)$$

To test H4(b) we calculate the *difference* in spillovers for stores that accept vs. do not accept vouchers. H4(b) is supported if this difference increased after the increase in voucher fungibility. We test this hypothesis with the statistic $s_{H4(b)}$, which we expect to be negative.

$$\begin{aligned}
s_{H4(b)} = & \left(E_{\{j:V_j=1\}}[\gamma_{k \rightarrow j,i,t \in epoch5}] - E_{\{j:V_j=0\}}[\gamma_{k \rightarrow j,i,t \in epoch5}] \right) \\
& - \left(E_{\{j:V_j=1\}}[\gamma_{k \rightarrow j,i,t \in epoch6}] - E_{\{j:V_j=0\}}[\gamma_{k \rightarrow j,i,t \in epoch6}] \right) \quad (3.15)
\end{aligned}$$

The posterior distributions of the statistics $s_{H4(a)}$ and $s_{H4(b)}$ across draws support both hypotheses. The posterior distribution of $s_{H4(a)}$ rejects the null hypothesis: the mean across draws is 0.26, and the 95% posterior interval does not contain zero (0.20, 0.33). The mean posterior distribution of $s_{H4(b)}$ equals -0.005, consistent with Hypothesis 4(b), but the effect is small and zero is contained within the 95% posterior interval (-0.12, 0.13). Thus, we do not find strong evidence that the coalition’s decision to increase voucher fungibility significantly impacted reward spillovers across the partners.

3.4.2.3 Market structure based on reward cross effects

We now visualize the competitive “landscape” across the partner stores by visualizing the reward spillovers to describe the coalition’s market structure. We use a heat map to visualize the affinity links that characterize reward spillover effects. Recall that an affinity link $\gamma_{k \rightarrow j,i,t(n)}$ represents the reward spillovers sent by store k to store j . The value of the link varies across customers and across seven epochs during which reward rates across the stores remained constant.


The following illustration (Figure 3.2) visualizes the affinity links across stores (averaged across customers) during epoch 5 (right before the devaluation). Red entries denote negative spillovers and blue entries denote positive ones. The scales are asymmetric: the red-to-white scale ranges from -0.9 to zero, and white-to-blue scale ranges

from 0 to 0.09. Each non-diagonal entry represents the affinity link sent by the column store to the row store.

The heat map reveals interesting patterns of competition within the coalition. First we note the few net positive spillovers observed. Stores B, C, and F (one jewelry store and two health stores) form a “love triangle”: almost all give and receive positive reward spillovers from each other (except $C \rightarrow F$). In contrast, some pairs of stores have spillovers that are asymmetric in valence. For example, store K only receives negative spillovers from other partners, but it benefits stores B, C, and N with positive spillovers.

It is also evident that a trio of “sport and shoe” stores L, M and N form a cluster sending negative spillovers to other partners. The stores most affected by these negative spillovers are B, F, G, and O, which operate in different categories and have different voucher policies. However, each of these shares the same location, and is relatively far from the trio L, M and N. Thus, B, F, G and O seem to mark the center of a geographic hub that is hurt when L, M and N attract customers to shop further away with higher rewards.

Figure 3.2: Heat map of spillover links



3.4.2.4 Competitive clout and vulnerability of reward spillovers

The heat map shows substantial asymmetry in how links are sent and received across stores. We summarize the asymmetry of the links received and sent across stores with an approach based on Kamakura and Russell (1989), which summarizes negative cross price elasticities among brands using measures of *vulnerability* and *competitive clout*. In our setting, vulnerability captures how sensitive a store is to other’s reward rates and competitive clout captures how influential the reward rates of a store are on other partners.

For ease of illustration we suppress customer and time indices of the affinity links to define the new metrics. Competitive clout sums over the absolute value of the affinity links “sent” by a store to other’s in the network. Vulnerability sums the absolute

value of the affinity links “received” by a store from the rewards of other partners. The absolute value allows these measures to characterize how strongly firms influence each other, whether it is through negative or positive spillovers.


$$Clout_j = \sum_{k \neq j} |\gamma_{j \rightarrow k}| \quad (3.16)$$

$$Vulnerability_j = \sum_{k \neq j} |\gamma_{k \rightarrow j}| \quad (3.17)$$

Figure 3.3 uses these metrics to visualize the coalition’s market structure before and after the devaluation (i.e., using the reward rates during epochs 5 and 6). Analogous to the heat map, competitive clout and vulnerability were calculated using the links averaged across each customer’s posterior mean. Measures of net clout and net vulnerability (i.e., recalculating the measures without the absolute values on the links), which allow positive and negative spillovers to cancel out, yield similar insights since the few positive spillovers we find are small in magnitude.

Figure 3.3 illustrates that the devaluation of rewards greatly narrowed the differences in both clout and vulnerability across stores. Overall, stores C, I, N, and L can be considered the “strongest” partners, since each have relatively low vulnerability or a relatively high clout in both epochs. Similarly, D, F and O can be considered the “weakest” partners, with a relatively high vulnerability to other’s rewards and low competitive clout.

Figure 3.3: Vulnerability vs. Competitive Clout of Reward Spillovers



3.5. General Discussion

In this chapter, we have developed a multi-store purchase incidence model to measure reward spillovers sent and received across diverse partners of a coalition program. The model parameters were used to test hypotheses on which mechanisms drive the types of spillovers across the stores.

First, we found evidence that stores are more likely to receive more favorable reward spillovers if they allow customers to redeem reward vouchers. This finding has important implications for the coalition, because although stores are currently very cognizant of the costs of accepting vouchers, without access to the coalition's data across stores, they cannot observe the potential benefits received from more favorable reward spillovers.

Second, we found that the affinity between partner stores can explain the types of spillovers across them. Stores that compete which are geographically close within the city are more likely to send and receive more favorable spillovers. These findings

can be used to alleviate competitive concerns within the coalition, since a partner can offset competitive pressures by offering more generous reward policies: either by increasing the reward rate that it offers to customers or by accepting voucher redemptions.

Third, we found that the devaluation of the reward currency in 2009 decreased the magnitude of spillovers across partners. This finding highlights the value of sharing a reward currency across stores: rewards with higher value change the way stores compete for customers through reward spillovers. We also found that the coalition's measures to increase the fungibility of vouchers had a limited impact on spillovers received by stores that do accept vouchers.

The spillovers characterize the market positioning of partner retail stores of a coalition loyalty program based on how customers are observed purchasing across. A heat map is used to visually position stores with a common spillover metric although several of these do not sell the same categories and purchases across stores is sparse. Asymmetry in the spillovers are measured using competitive clout and vulnerability metrics (Kamakura and Russell 1989). The positioning heat map and asymmetry metrics help to identify stores with a potential for cross-marketing opportunities.

Our plans for future work in this area include using the model to evaluate policy counterfactuals such as imposing changes to tier-level reward rates. We also plan to incorporate demographic information to characterize differences in customer-level spillovers. Finally, while this paper focuses on a coalition program with brick-and-mortar stores, future research can explore the spillover interactions between partners with both brick-and-mortar and online outlets.

CHAPTER 4 : Concluding Reflections on Emerging Loyalty Programs

4.1. Summary of contributions

This chapter summarizes the contributions of this dissertation, discusses their implications for managers, and identifies trends that continue to increase the fungibility of rewards. The chapter concludes with recommendations for future research.

The introductory chapter developed a unifying framework to study emerging types of loyalty programs which provide increasingly fungible rewards. Chapters 2 and 3 each delved into issues pertaining to each of the two dimensions of the framework: the rate at which rewards can be redeemed (linear vs. non-linear) and the multiplicity of actively participating retailers.

Chapter 2 modeled data from a linear program to generate insights into why customers seem to persistently stockpile points. We find empirical evidence suggesting that prospect theory and mental accounting (Thaler 1985) can partially explain why customers in linear loyalty programs are motivated to stockpile points for long periods of time, despite the absence of economic incentives for doing so from the retailer. By mentally booking gains and losses of cash separately from gains and losses of points, customers can be intrinsically incentivized to save points up to a certain amount before redeeming them in a linear program. This psychological motivation is distinguished from economic and cognitive redemption costs. The data suggests that customers seem to be driven mostly by this psychological motivation as well as by cognitive fixed costs of redeeming, but less so by economic costs (i.e., the value of foregone points).

These findings have implications for understanding why even in non-linear programs,

many customers often stockpile points above and beyond the retailer’s requirements to redeem. Furthermore, they suggest that efforts to encourage redemptions should consider different ways of framing these in terms of gains and losses of cash and points.

Chapter 3 developed a multi-store purchase incidence model to measure reward spillovers across diverse partner stores in a coalition loyalty program. The model is used to test hypotheses showing that the spillovers between stores are driven by reward redemption policies and affinity between stores. Reward spillovers are visualized to position the diverse set of stores in a common map that characterizes their market structure.

Our results allow both the managers of partner stores as well as the third-party managers of the coalition program to better understand the nature of competition within the coalition. Most importantly, we illustrate that the value of a shared reward currency influences how partner stores compete with each other through reward spillovers. This work on coalition programs is also relevant to firms such as theme parks and casinos which reward customers across their own umbrella of services.

4.2. Recommendations for future research

We provide directions for future research in coalition loyalty programs by discussing three examples from the perspective of each of the three parties involved: customers, partner retailers, and the third-party coalition operator. From the perspective of customers, it would be interesting and challenging to model the evolution of the customer lifetime cycle not only at a specific retailer, but across sectors, geographic locations, and across the program. From the perspective of retailers, more work is needed to understand the impact of entry and exit to and from a coalition network (variation that was not present in our data). Finally, from the perspective of the

operator, improved store-level analytics can be provided to partners by quantifying and leveraging the value of shared data. For example, while a store may suspect attrition for a specific customer, the coalition has more specific data on whether that customer has recently shopped at nearby or competing retailers.

4.3. Emerging trends

Having summarized the contributions of the preceding chapters and avenues for future research, we conclude by highlighting trends that are continuing to increase the fungibility of rewards in emerging loyalty programs. In particular, we discuss four examples of innovations that are further increasing the fungibility of rewards which pose new challenges for the managers of loyalty programs. While this dissertation focuses on retail loyalty programs, we examine other types of firms, including airline carriers and hotel chains, to identify these trends.

First, the digitalization of rewards eases the exchange of information among parties. This trend has led some firms to lose tight control over the fungibility of their own rewards to third parties. Mileage-tracking websites are one prominent example which allow customers to keep track of various loyalty programs in one place, and even allow customers to more easily compare the value of points across competing airline carriers, hotels, and car rental companies (McCartney 2011).

Second, the existence of increasing competition pressures companies to match rewards earned at competitors. Many hotel chains and airline carriers have status matching programs which reward customers with special status if they have earned a similar status at a competitor. One example is the “Status Match ... No Catch” policy at Best Western Hotels & Resorts which matches a customer’s elite status in any other hotel loyalty program, free of charge. These programs make single-firm programs

similar to coalition programs, in which a customer can earn a special status valid across all partners even though the bulk of his spending was directed at a few firms.

Third, firms are facing pressure to adapt their programs to leverage mobile technologies. By doing so, rewards have the potential to become truly redeemable “anytime, anyplace.” Finally, competitive pressures also encourage firms to increasingly allow customers to more easily transfer points to other customers, lowering the ability of points to impose switching costs.

To conclude, this dissertation has outlined a framework to study novel challenges in emerging types of loyalty programs, provided valuable contributions with managerial implications, and highlighted promising avenues for future research to build upon our work.

Appendices

APPENDIX 1: PROOF OF EXISTENCE OF AN INDIFFERENCE POINT

This appendix presents a theorem showing that when there are two S-shaped value functions $w(x)$ and $v(x)$, an indifference point s^* can be reached as long as $w(x)$ is steeper than $v(x)$ for large losses (i.e., when there are many points available to redeem). The subscripts are dropped in the proofs for simplicity.

We begin with brief propositions that characterize how the variable benefits and costs of redeeming vary with redeemable points. These proofs stem directly from two properties of an S-shaped value function: over the loss domain ($x < 0$), a value function $u(x)$ is increasing (Property 1: $u' > 0$) and convex (Property 2: $u'' > 0$).

Proposition 1. The benefit of redeeming $b(\tilde{s}_{ij}; m_{ij}) = w(-m_{ij} + \tilde{s}_{ij}) - w(-m_{ij})$ is a strictly increasing and strictly convex function of redeemable points $\tilde{s}_{ij}$: $b' > 0$; $b'' > 0$.

Proof: The first derivative $b'(\tilde{s}) = w'(-m + \tilde{s}) > 0$ comes from Property 1 of the value function. The second derivative $b''(\tilde{s}) = w''(-m + \tilde{s}) > 0$ comes from Property 2 of the value function.

Proposition 2. Denote the variable cost of redeeming by $c(\tilde{s}_{ij}; m_{ij}) = -v(-\tilde{s}_{ij})$. The variable cost is a strictly increasing and strictly concave function of redeemable points $\tilde{s}_{ij}$: $c' > 0$; $c'' < 0$.

Proof: The first derivative $v'(-\tilde{s}) < 0$ comes from Property 1 of the value function. Thus, $c'(\tilde{s}) = -v'(-\tilde{s}) > 0$. The second derivative $c''(\tilde{s}) = v''(-\tilde{s}) < 0$ comes from

Property 2 of the value function.

Now we introduce a condition that formalizes when $w(x)$ is steeper in losses than $v(x)$ (i.e., spending cash should be “more painful” than spending points) at least for large prices:

Asymmetry condition: There exists a price m' such that for all $m > m'$, $|w(-m)| > |v(-m)|$.

Theorem 1. Given the asymmetry condition, a threshold $s^ > 0$ such that $z(s^*, m_{ij}) = 0$ is guaranteed to exist for prices m_{ij} greater than some finite level $\bar{m}_{ij}$.*

Proof: The net utility of redeeming is decomposed into a variable component and a fixed component, $z(s) = q(\tilde{s}) - f$, where $q(\tilde{s}) = b(\tilde{s}) - c(\tilde{s})$, and the fixed costs include the opportunity cost of foregone points and the transaction cost, $f = v(mr) + c$.

When $\tilde{s} = 0$, $q(0) = 0$ and thus $z(\tilde{s}) = -f < 0$. When $\tilde{s} = m$, $q(m) = v(-m) - w(-m) > 0 \forall m > m'$ by the asymmetry condition. Since $f''(m) < 0$ and $q''(m) > 0 \forall m > m'$, there exists a price $\bar{m}$ such that $q(m) > f(m) \forall m > \bar{m}$. When $q(m) > f(m)$, then $z(m) > 0$ and thus, a threshold s^* such that $z(s^*) = 0$ is guaranteed to exist by the continuity of $z(s)$.

APPENDIX 2: PARAMETER RECOVERY FOR MA

This appendix illustrates an example of parameter recovery for MA’s population-level parameters based on our observed data. We use the independent variables in the complete dataset (10219 purchase occasions from 346 individuals) to simulate a new set of redemption choices. First we generate the β_i parameters for each indi-

vidual according to Equation 2.8, where μ is set close to values estimated from our dataset: [-1.7, -0.3, -0.64, -3, 0.1, 1.8], and Ω is set to an identity matrix times 2.9. Each individual’s simulated β_i parameters, their observed prices, and their observed stockpiles are then used to generate a new sequence of redemption choices for each of their observed purchase occasions, according to the MA model (Equation 2.7). The redemption rate in the simulated dataset is 2.5%.

We ran three independent chains from different starting values, and assess their convergence using the Gelman and Rubin (1992) diagnostic of between-to-within chain variance. We ran the model for 5000 iterations and the last 3000 iterations of each chain (9000 draws in total) were used for analysis. Table 4.1 compares the actual and estimated untransformed population-level parameters. Each of the true population-level parameters is contained within 95% posterior intervals. The errors for the individual-level parameters λ_{wi} , λ_{vi} , a_{wi} , a_{vi} , h_i , and c_i are on average 0.31, -0.13, -0.05, 0.04, 0.03, and 0.04 respectively, and the mean absolute percentage errors are 25.9%, 12.5%, 20.4%, 104.9%, 14.6%, and 16.6%, indicating good model fit.

Table 4.1: Parameter recovery for MA

Parameter	Actual	Estimated	Error	2.5% Bound	97.5% Bound
λ_w	1.18	1.40	0.22	1.08	2.18
λ_v	1.74	1.51	-0.23	1.18	2.00
a_w	0.35	0.28	-0.07	0.20	0.35
a_v	0.05	0.08	0.03	0.03	0.14
h	0.52	0.56	0.03	0.20	0.89
c	1.80	1.86	0.06	1.29	2.36

APPENDIX 3: A RATIONAL EXPECTATIONS ALTERNATIVE

This appendix empirically explores the directional predictions of a Rational Expectations (RE) model, an enriched SLA model in which consumers are forward-looking

with rational expectations on the timing and size of their future shopping needs. A customer in this model views points as a means to save cash in the future, and does not inherently value points in a separate account. Behavior can be driven by foregone points and fixed non-monetary costs, but not separate mental accounts. For simplicity, we omit the i subscripts when presenting the individual-level model.

Consider a consumer who is forward looking over an infinite horizon and who values monetary incentives with a linear utility, as shown in Equation 4.1. As before, when the customer redeems, he pays $-m + \tilde{s}$ instead of $-m$, and can also incur a non-monetary cost of redeeming c .

$$u(y|s, m) = \begin{cases} -\beta_m m & \text{if } y = 0 \\ -\beta_m(m - \tilde{s}) - c & \text{if } y = 1 \end{cases} \quad (4.1)$$

As in MA, this consumer understands that redeeming affects his future stockpile s' (Equation 4.2). However, he does not value all future points equally. The value for s' depends on *how* and *when* he expects to use them. We assume rational expectations on future shopping needs m' and future inter-purchase times. Let q be the distribution of m' conditional on the current basket price m . Future payoffs are discounted by β , interpreted as the daily discount rate δ times the expected duration until the next purchase. The consumer optimizes his redemptions to maximize expected total discounted utility (Equation 4.3).

$$s' = \begin{cases} s + mr & \text{if } y = 0 \\ s - \tilde{s} & \text{if } y = 1 \end{cases} \quad (4.2)$$

$$V(s, m) = \max_{y \in \{0,1\}} u(y|s, m) + \beta E[V(s', m')|s, m, y] \quad (4.3)$$

We use Chebyshev regression (Judd 1998), a flexible non-parametric approach, to approximate the value function $V(s, m_{level})$ at each of the two price levels: *low* and *high*. For details on the implementation of this approach, see the next section of this appendix called “Value function estimation.” Discretizing m allows us to convey the intuition behind the basic tradeoffs that determine how redemptions vary with shopping needs. The estimates are used to calculate $z_{RE}(s, m_{level})$, the difference between the utility of redeeming and not redeeming in the current period (Equation 4.4). The consumer redeems when $z_{RE}(s, m) > 0$. Let s^* be the stockpile level at which $z_{RE}(s^*, m) = 0$.

$$z_{RE}(s, m) = \beta_s \tilde{s} - c + \beta [E[V(s', m')|s, m, 1] - E[V(s', m')|s, m, 0]] \quad (4.4)$$

We compare $z_{RE}(s^*, m_{low})$ and $z_{RE}(s^*, m_{high})$ to intuitively describe the model’s directional predictions under two cases: when c is 0 and when c is large. When $c = 0$, foregone points drive stockpiling behavior. Since foregone points are proportional to m , a consumer will prefer to redeem on a small price rather than a high price. In other words, s^* increases with m (i.e., $z_{RE}(s^*, m_{low}) < z_{RE}(s^*, m_{high})$). When c is substantial, both the economic and cognitive motivations are allowed to drive behavior. In this case, s^* instead decreases with m . Intuitively, he stockpiles longer to save more at each redemption. A large c may seem necessary ex-ante to explain the observed persistent stockpiling because (i) foregone points are small in magnitude (only 1% of the basket price) and (ii) individuals often neglect to consider the absence of potential

positive cash flows (Frederick et al. 2009).

To recap, RE expects redemptions to decrease with m if c is small, and to instead increase with m if c is substantial. To explore whether either of these predictions are directionally consistent with our data, we examine the individual-level estimates from SLA, a model linear in $\tilde{s}$ and m . At the individual-level, β_{si} is positive for 100% of customers (with a 95% posterior interval ranging from 0.022 to 0.250 across individuals), and β_{mi} is negative for most customers (95.7%) and not different from zero for the rest with a 95% posterior interval ranging from -0.151 to 0.003 across individuals). So, to be consistent with this data, RE implies that foregone points (the economic motivation) must be the main determinant of persistent stockpiling and *not* cognitive fixed costs. In contrast, MA can capture this empirical pattern regardless of whether a customer considers foregone points (to see why, refer back to Table 5 and set $a_{vi} = 0$).

A second exploratory analysis did not find preliminary evidence for RE in the data. We ran a logistic regression (Equation 4.5) of redemptions with the observed running means of each individual’s shopping needs $\bar{m}_{ij}$ (i.e., $\bar{m}_{ij} = \frac{1}{j} \sum_{k=1}^j m_{ik}$) and of the days between purchases $\bar{d}_{ij}$, two factors that should influence consumers with rational expectations above and beyond $\tilde{s}_{ij}$ and m_{ij} . Fixed effects α_i were added to control for differences across individuals¹.

$$\text{logit}(p_{ij}) \sim \alpha_0 + \alpha_i + \beta_s \tilde{s}_{ij} + \beta_m m_{ij} + \gamma_m \bar{m}_{ij} + \gamma_d \bar{d}_{ij} \quad (4.5)$$

¹The advantage of using running means rather than the mean across the entire individual’s observed purchases is that it allows us to control for individual-level fixed effects. This helps overcome the difficulty of disentangling model differences from heterogeneity (e.g., customers with longer durations may have systematically different discount rates, and the distributions of shopping needs may tend to vary with the value for cash across individuals).

While RE with low c expects that redemptions decrease with expected shopping needs ($\gamma_m < 0$) and increase with expected durations ($\gamma_d > 0$), both of these parameters were positive and not significantly different from zero (p-values 0.50 and 0.99 respectively). The control variables were consistent with the results from SLA: β_s is positive ($4.5 * 10^{-3}$, p-value < 0.001) and β_m is negative ($-5.2 * 10^{-4}$, p-value = 0.019) .

In summary, to be consistent with this dataset, RE implies that persistent stockpiling must be mostly driven by the economic incentive of foregone points and *not* by cognitive costs. In contrast (as discussed in the chapter), MA finds that differences in redemption behavior across individuals are mostly driven by cognitive and psychological motivations to redeem (fixed costs and separate accounts). While we do not rule out either model, the common neglect of foregone gains (Frederick et al. 2009) brings in question the plausibility of RE for this dataset.

Value function estimation


The value function is approximated using orthogonal polynomials on 5 Chebyshev nodes (i.e., the zeros of a Chebyshev polynomial) over the stockpile levels as:

$\tilde{V}(s, m_{level}) = \sum_{k=1}^n c_{k,level} T_k(s)$, where $T_k(s)$ denotes the Chebyshev basis function of order k , and $c_{k,level}$ are the coefficients. The $2n$ coefficients are estimated by numerically solving the system of $2n$ equations: $0 = \tilde{V}(s, m_{level}) - V(s, m_{level})$.

In the example shown in Figure 4.1, we used the following example parameters: $c = 0$, $\beta_m = 1$, $\beta = 0.98$, $m_{low} = 3$, $m_{high} = 7$, $q(m_{low}|m_{low}) = 0.2$, and $q(m_{high}|m_{high}) = 0.9$. The intuitive results do not change when the transition matrix q is modified so that m_{low} is more common than m_{high} . The kinks at the maximum redeemable amount ($m\kappa$) reflect that redeemable points at any purchase occasion are bounded at 50% of the basket price. The consumer persistently stockpiles up to the level s^*

at which $z_{RE}(s^*, m) = 0$. When $c = 0$, the figure shows that s^* increases with m (i.e., $z_{RE}(s^*, m_{low}) < z_{RE}(s^*, m_{high})$). The same intuition holds when the transition matrix q is modified so that m_{low} is more common than m_{high} . When cognitive costs c are substantial, the consumer stockpiles longer and redeems on basket prices that are high instead of low. To see this graphically, refer again to Figure 4.1. As c increases, both curves move down. When c is large enough so that the zero crossing point occurs once $z_{RE}(s, m_{high})$ has surpassed $z_{RE}(s, m_{low})$, the customer's threshold s^* decreases with m .

Figure 4.1: Net utilities when shopping needs are low versus high



APPENDIX 4: PARAMETER RECOVERY FOR AFFINITY MODEL

This appendix illustrates an example of parameter recovery for the parameters of the Affinity model. We used the observed covariates from the 15 stores to simulate a new set of purchase choices for 50 individuals. The true population-level parameters were chosen to be close to those estimated from the observed dataset. We ran a single chain for 10000 iterations (after thinning every 20 iterations) and used the last 8000

to compare the posterior distributions to the actual values. Table 4.2 shows that each of the true population-level parameters, as well as the true mean of θ_{ij} , is contained within 95% posterior intervals.

Table 4.2: Parameter recovery for Affinity Model

Parameter	Actual	Estimated	2.5% Bound	97.5% Bound
μ_λ	-2.00	-1.76	-2.19	-1.29
$\bar{\kappa}_0$	0.20	-0.01	-0.40	0.31
$\bar{\kappa}_1$	0.40	0.72	0.36	1.15
ψ_0	-0.20	-0.30	-0.43	-0.18
ψ_1	0.10	0.15	0.03	0.27
ψ_2	0.00	0.03	-0.10	0.16
ψ_3	0.20	0.20	0.07	0.33
ψ_4	-0.60	-0.63	-0.79	-0.46
δ_{κ_0}	-0.10	-0.05	-0.17	0.07
δ_{κ_1}	2.00	1.90	1.72	2.09
δ_{ψ_0}	-0.20	-0.18	-0.20	-0.17
δ_{ψ_1}	0.10	0.10	0.07	0.12
δ_{ψ_2}	-0.20	-0.20	-0.26	-0.15
δ_{ψ_3}	0.10	0.10	0.06	0.13
δ_{ψ_4}	0.80	0.69	0.57	0.81
$\text{mean}(\theta_{ij})$	$-5e^{-4}$	$-5e^{-5}$	$-7.8e^{-3}$	$7.7e^{-3}$

BIBLIOGRAPHY

References for Chapter 1

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

References for Chapter 2

1. Chintagunta, Pradeep, Tulin Erdem, Peter E. Rossi, and Michel Wedel (2006), “Structural Modeling in Marketing: Review and Assessment,” *Marketing Science*, 26(6), 604–616.
2. Dorotic, Matilda, Tammo H. A. Bijmolt, and Peter C. Verhoef (2011), “Loyalty Programmes: Current Knowledge and Research Directions,” *International Journal of Management Reviews*, 14(3), 217–237.
3. Dreze, Xavier and Joseph C. Nunes (2004), “Using Combined-Currency Prices to Lower Consumers’ Perceived Cost,” *Journal of Marketing Research*, 41(1), 59–72.
4. Dreze, Xavier and Joseph C. Nunes (2009), “Feeling Superior: The Impact of Loyalty Program Structure on Consumers’ Perceptions of Status,” *Journal of Consumer Research*, 35(6), 890–905.
5. Erdem, Tulin, Susumu Imai, and Michael P. Keane (2003), “Brand and Quantity Choice Dynamics Under Price Uncertainty,” *Quantitative Marketing and Economics*, 1(1), 5–64.
6. Gelman, Andrew and Donald B. Rubin (1992), “Inference from Iterative Simulation Using Multiple Sequences,” *Statistical Science*, 7(4), 457–511.
7. Gelman, Andrew, John B. Carlin, Hal S. Stern, and Donald B. Rubin. *Bayesian data analysis*. CRC Press, 2 edition, 2003.

8. Hlavinka, Kelly and Jim Sullivan (2011), "The 2011 Colloquy Loyalty Census," technical report, April 2011.
9. Hsee, Christopher K., Fang Yu, Jiao Zang, and Yan Zhang (2003), "Medium Maximization," *Journal of Consumer Research*, 30(1), 1–14,.
10. Humby, Clive, Terry Hunt, and Tim Phillips (2004). *Scoring Points: How Tesco Is Winning Customer Loyalty*. Kogan Page Publishers.
11. Kahneman, Daniel and Amos Tversky (1979), "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47(2), 263–291.
12. Kivetz, Ran and Itamar Simonson (2002), "Earning the Right to Indulge: Effort as a Determinant of Customer Preferences Toward Frequency Program Rewards," *Journal of Marketing Research*, 39(2), 155–170.
13. Kivetz, Ran, Oleg Urminsky, and Yuhuang Zheng (2006), "The Goal-Gradient Hypothesis Resurrected: Purchase Acceleration, Illusionary Goal Progress, and Customer Retention," *Journal of Marketing Research*, 43(1), 39–58.
14. Konsewicz, Bob (2007), "Does your Point Burner Campaign Really Reduce Your Loyalty Program's Point Liability?," technical report, MasterCard Advisors.
15. Kwong, Jessica Y. Y., Dilip Soman, and Candy K. Y. Ho (2011), "The Role of Computational Ease on the Decision to Spend Loyalty Program Points," *Journal of Consumer Psychology*, 21(2), 146–156.
16. Newton, Michael A. and Adrian E. Raftery (1994), "Approximate Bayesian Inference with the Weighted Likelihood Bootstrap," *Journal of the Royal Statistical Society, Series B (Methodological)*, 56(1), 3–48.
17. Rick, Scott I. , Cynthia E. Cryder, and George Loewenstein (2008), "Tightwads and Spendthrifts," *Journal of Consumer Research*, 34(6), 767–782.
18. Rossi, Peter E. , Greg M. Allenby, and Rob McCulloch (2005). *Bayesian Statistics and Marketing*. Wiley, 1 edition.
19. Shampanier, Kristina, Nina Mazar, and Dan Ariely (2007), "Zero as a Special Price: The True Value of Free Products," *Marketing Science*, 26(6), 742–757.

20. Smith, Andrew and Leigh Sparks (2009), “It’s Nice to Get a Wee Treat if You’ve Had a Bad Week: Consumer Motivations in Retail Loyalty Scheme Points Redemption,” *Journal of Business Research*, 62(5), 542–547.
21. Soman, Dilip (2003), “The Effect of Payment Transparency on Consumption: Quasi Experiments from the Field,” *Marketing Letters*, 14(3), 173–183.
22. Spiegelhalter, David J., Nicola G. Best, Bradley P. Carlin, and Angelika van der Linde (2002), “Bayesian Measures of Model Complexity and Fit,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4), 583–639.
23. Prelec, Drazen and George Loewenstein (1998), “The Red and the Black: Mental Accounting of Savings and Debt,” *Marketing Science*, 17(1), 4–28.
24. Thaler, Richard (1985), “Mental Accounting and Consumer Choice,” *Marketing Science*, 4(3), 199–214.
25. Tversky, Amos and Daniel Kahneman (1992), “Advances in prospect theory: Cumulative representation of uncertainty,” *Journal of Risk and Uncertainty*, 5(4), 297–323.
26. Tversky, Amos and Daniel Kahneman (1981), “The Framing of Decisions and the Psychology of Choice,” *Science*, 211(4481), 453–458.
27. Von Neumann, John and Oskar Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.
28. Zellermayer, Ofer (1996), *The Pain of Paying*. PhD Dissertation, Department of Social and Decision Sciences, Carnegie Mellon University.

References for Chapter 3

1. Bell, David R. (2014), “Location is (Still) Everything: The Surprising Influence of the Real World on How We Search, Shop, and Sell in the Virtual One,” New Harvest.
2. Danaher, Peter J. and Michael S. Smith (2011), “Modeling Multivariate Distri-

- butions Using Copulas: Applications in Marketing,” *Marketing Science*, 30 (1), 4-21.
3. Dorotic, Matilda, Dennis Fok, Peter C. Verhoef, Tammo H. A. Bijmolt (2011), “Do vendors benefit from promotions in a multi-vendor loyalty program?,” *Marketing Letters*, 22 (4), 341-356.
 4. Elrod, Terry, Gary J. Russell, Allan D. Shocker, Rick L. Andrews, Lynd Bacon, Josef a. Mazanec, Vithala R. Rao and Venkatesh Shankar (2002), “Inferring Marketing Structure from Customer Response to Competing and Complementary Products,” *Marketing Letters*, 13 (3), 221-232.
 5. Gelman, Andrew, John B. Carlin, Hal S. Stern, and Donald B. Rubin (2003), *Bayesian Data Analysis*, 2nd ed. Boca Raton, FL: CRC Press.
 6. Gelman, Andrew and Donald B. Rubin (1992), “Inference from Iterative Simulation Using Multiple Sequences,” *Statistical Science*, 7 (4), 457-511.
 7. Goldfarb, Avi and Catherine Tucker (2014), “Conducting Research with Quasi-Experiments: A Guide for Marketers,” working paper.
 8. Kamakura, Wagner A. and Gary Russell (1989), “A Probabilistic Choice Model for Market Segmentation and Elasticity Structure,” *Journal of Marketing Research*, 26 (November), 379-390.
 9. Manchanda, Puneet, Asim Ansari, Sunil Gupta (1999), “The ‘Shopping Basket’: A Model for Multicategory Purchase Incidence Decisions,” *Marketing Science*, 18 (2), 95-114.
 10. Netzer, Oded, Ronen Feldman, Jacob Goldenberg, and Moshe Fresko (2012), “Mine Your Own Business: Market-Structure Surveillance Through Text Mining,” *Marketing Science*, 31 (3), 521-543.
 11. Newton, Michael A. and Adrian E. Raftery (1994), “Approximate Bayesian Inference with the Weighted Likelihood Bootstrap,” *Journal of the Royal Statistical Society, Series B (Methodological)*, 56 (1), 3-48.
 12. Ozturk, O. Cem, Sriram Venkataraman, and Pradeep K. Chintagunta (2016),

- “Price Reactions to Rival’s Local Channel Exits,” *Marketing Science*, Articles in Advance, 16 Feb 2016.
13. Pancras, Joseph, Rajkumar Venkatesan, and Bin Li (2015), “Investigating the Value of Competitive Mobile Loyalty Program Platforms for Intermediaries and Retailers,” *Marketing Science Institute Working Paper Series 2015*. Report No. 15-107.
 14. Rossi, Peter E., Greg M. Allenby, and Rob McCulloch (2005), *Bayesian Statistics and Marketing*, 1st ed. New York: John Wiley & Sons.
 15. Shugan, Steven M. (2014), “Market Structure Research.” *The History of Marketing Science*. Ed. Russell S. Winer, Ed. Scott A. Neslin. Hanover, MA. World Scientific Publishing Company and Now Publishers Inc, 129-159.
 16. Stephen, Andrew T. and Olivier Toubia (2010), “Deriving Value from Social Commerce Networks.” *Journal of Marketing Research*, 47 (2), 215-228.
 17. Spiegelhalter, David J. Nicola G. Best, Bradley P. Carlin, and Angelika van der Linde (2002), “Bayesian Measures of Model Complexity and Fit,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64 (4), 583-639.
 18. Trusov, Michael, Anand V. Bodapati, and Randolph E. Bucklin (2010), “Determining Influential Users in Internet Social Networks,” *Journal of Marketing Research*, 47 (4), 643-658.

Reference for Chapter 4

1. McCartney, Scott (2011), “Two Airlines Battle Mileage-Tracking Websites – Ruckus over Online Companies That Display Frequent-Flier Miles, Hotel and Rental Car Points in One Place,” *The Wall Street Journal*, October 13, 2011.
2. Thaler, Richard. Mental Accounting and Consumer Choice. *Marketing Science*, 4(3):199– 214, 1985.

References for Appendix 3

1. Frederick, Shane, Nathan Novemsky, Jing Wang, Ravi Dhar, Stephen Nowlis (2009), "Opportunity cost neglect," *Journal of Consumer Research*, 36 (4), 553-561.
2. Judd, Kenneth L., "Numerical methods in economics," MIT Press, 1998.