Distribution Agreement

Understanding Customer Dynamics in a Data-rich Environment

By

Nuo Xu
B.S. Fudan University, China, 2008

Business

Jagdish Sheth, Ph.D.
Advisor

Douglas Bowman, Ph.D.
Committee Member

Ramnath Chellappa, Ph.D.
Committee Member

Yi Zhao, Ph.D.
Committee Member

Accepted:

Lisa A. Tedesco, Ph.D. Dean of the James T. Laney School of Graduate Studies

Date

Understanding Customer Dynamics in a Data-rich Environment

By

Nuo Xu
B.S. Fudan University, China, 2008

Advisor: Jagdish Sheth, Ph.D.

An abstract of
A dissertation submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
in Business
2016

Abstract

Understanding Customer Dynamics in a Data-rich Environment

By

Nuo Xu

Consumers' behaviors are dynamic. The understanding of these dynamic processes is critical for managerial decisions. In a data rich environment, an effective use of the data can help us uncover these dynamics and deliver actionable intelligence. These three essays all focus on providing solutions to important managerial questions in a data rich environment.

The first essay looks into the banking industry where customers' life cycle plays a critical role in their financial activities. The goal of the first essay is to provide a solution based on Cusum control chart to detect a life change of interest using observed customer activities. The recovery of this life change information can help managers better customize direct marketing efforts based on customers' life events.

The second essay extends the application of Cusum control chart to detect changes in the market's responses to marketing stimulus. This essay presents another fine property of the Cusum control chart: the test statistic can help managers to trace the time point when changes occur. This time point provides a road map for managers to conceptualize and understand the underlying cause of the change. The plan is to use a dataset provides by Dominic supermarket chain to illustrate the effectiveness of the proposed solution in detecting changes in the pattern of market responses when new brands or new products enter the market.

The third essay examines customers' shopping behaviors using a coalition loyalty program in Europe. The goal of the essay is to model the synergies among the merchants participating the coalition loyalty program in order to understand the impact of cross buying on customers' purchase behavior at the focal merchant store overtime.

In sum, these three essays provide solutions for managers in a data-rich environment to transfer massive data into customer knowledge and actionable intelligence.

Understanding Customer Dynamics in a Data-rich Environment

By

Nuo Xu
B.S. Fudan University, China, 2008

Advisor: Jagdish Sheth, PhD

A dissertation submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
in Business
2016

Acknowledgement

PhD is a long journey. It would be an impossible mission without the support and encouragement from the many wonderful people around me: my mentors, my family and my friends. Because of them, I never felt alone during this journey.

I would like to give my deepest gratitude to my advisor, Professor Jagdish Sheth. I went to Professor Sheth for advices at the most confusing time of my PhD life, a period when I was struggling to figure out my dissertation and an advisor to work with. Professor Sheth rendered his generous supports to me. He is wise, experienced, visionary and has a keen sense for the big picture and managerial implications. He guided me to see the research problems that I was working on from a higher-level perspective, and he constantly encourage me to get more exposure to industrial practice. During the process of my dissertation work, Professor Sheth is also kind, understanding and patient. He knows that a PhD student grows through trials and errors. He encouraged me to take time to explore and work on the questions that interest me, and he gave me the time I needed, listened to me explaining the ideas, and helped me polish them. Beyond the guidance on the academic side, Professor Sheth also guided me to be a happy and confident person. He is positive, cheerful and knows how to motivate people in their most difficult times. When there were obstacles in the way, he told me to see the big picture and remind me the value of my work. Most importantly, Professor Sheth encouraged me to take time out with the and especially with my daughter. This greatly helped me to balance my life and work, and I became more and more energetic. Through all our interactions, Professor Sheth has demonstrated the qualities of a good mentor. He is and will always be the role model and the inspiration for me.

I am also very thankful to Professor Yi Zhao. I met Professor Zhao in the Georgia Research Symposium in Spring 2012 and I later audited his class on Bayesian method. Professor Zhao and I share the same research interests in statistics and customer analytics, and we have started to work together. Through his class and our work together, I gained a deeper understanding of statistical models and learnt to interpret results from a more comprehensive view. During my job market preparation, Professor Zhao read various versions of my job market paper and sat in various mock presentations, and he gave me his honest opinions. I have benefited a lot from his comment.

I would also like to thank my committee members, Professor Douglas Bowman and Professor Ramnath Chellappa for the many invaluable suggestions. Professor Bowman took me as an apprentice in teaching. He guided me through the process of preparing a class, helped me revise the slide deck and sat in the classroom to observe me teaching. This is a valuable experience for me that prepares me for teaching in the future. I also appreciate the supports and suggestions from Professor Ryan Hamilton, Professor Michael Andrews, Professor Sandy Jap, Professor TI Kim, Professor Manish Tripathi, Professor David Schweidel and Professor Mike Lewis for their helpful feedbacks.

I am also happy to have made many friends during my graduate studies. Many thanks to Xin Zhang, Mingxuan Li, Yu Wang, Shuting Liang, Beth Fossen, Tony Koschmann, Mike

Palazzolo, Hyeyoung Hah, Deidre Popovich, Harry Antonio, Yanwen Wang, Zhe Han, Ning Zhong, Jinran Zhao, Pranay Reddy, Hulya Karaman, Xin Zheng and Weishi Jia. Thank you for all your cheers and hugs. I feel so blessed to have your companies throughout the journey.

I could not have completed this dissertation work without the great support and sacrifice of my parents and my in-laws. They respected my choice of this difficult path, and they always believe in me. During this time, they travelled across the ocean to take care of us and our baby daughter. This has been a long journey for them as well. Thank you very much for your unconditional love.

Finally, I would like to thank my husband, Yun Zhu. We met in Atlanta, and together, we build a family. I feel really lucky to have Yun in my life because he not only takes good care of me, but also helps me become a better and stronger person.

This dissertation is dedicated to the sunshine of my life, my daughter, Blair Zhu. You are going to have an amazing journey. Enjoy!

Table of Contents

List of Tables

List of Figures

Chapter 1

Overview

Customer analytics is now growing into a prominent aspect of business management. With its emphasis on the analysis of rich secondary data, customer analytics offers rich insights by describing past buying behaviors, accessing effectiveness of marketing tools and predicting future behaviors. Companies are increasing their spending on customer analytics. A survey of 1139 of CMOs and CIOs by IDG Enterprise showed that 80% of the enterprise companies and 63% of the small and medium businesses have or plan to implement data-driven projects in 2015. In average, the spending on these data driven initiatives is expected to reach $13.8 million for an enterprise company in 2016 and $1.6 million for a small or medium size business (Columbus, 2015). Similarly, the CMO survey conducted by Deloitte, the American Marketing Association and Duke University in February 2016 showed that senior management are transforming their organizations to be more data-driven and analytics. Their survey of 289 senior level marketing managers showed that the spending on marketing is estimated to increase by 66% in the next three years and more and more companies are using advanced analytics to assist decisions on customer acquisition, customer retention, promotion, branding and cross-channel marketing. Advanced analytics of data is becoming a new drive of growth and innovations in organizations.

One central topic in customer analytics is dynamics in customers' behaviors. Existing studies on customer behaviors have documented various ways that customers' behaviors can evolve overtime and many theoretical models have been proposed to suggest how these patterns change as a result of individual and environmental factors such as customers' learning and experience (Erdem and Keane, 1996; Heilman, Bowman and Wright, 2000), life cycle, relationship with company (Netzer, Lattin and Srinivasan, 2008), companies' marketing activities (Li, Sun, and Montgomery, 2011; Luo and Kumar, 2013) and competitors' actions (Moon, Kamakura and Ledolter, 2007).

In these studies, researchers have applied various methods to account for the dynamic nature of behaviors. For example, one most commonly applied approach for the managers is to use time since last purchase, frequency of past purchase and average amount of spending as key indicators of customers' future purchase behaviors. Other widely applied methods in the academic field includes hazard model, random-coefficient model, state space model, hidden Markov model and dynamic programming. In spite of various methods applied in these studies, the essential idea of these analyses is to consider behaviors as a result of a certain underlying state, which can be inferred using past behavioral data. Each method is unique in its own way to capture the connection between past and future behaviors.

The development in the understanding of dynamics of customers' behavior has great implications for managerial practice. Customer base is a crucial asset for a company, and a more comprehensive understanding of customers' behaviors is linked to improved evaluation of the value of a company's customer base via the measure of customer lifetime value. Advances in the customer analytical tools also provide guidance for marketing campaigns. Insights from customer data can improve the effectiveness of marketing efforts by targeting the right customer at the right time and right place using the right message.

Today, the advances in the information technologies and low cost in data storage allow us to access rich data on customer behaviors, bring us new opportunities to gain deeper insight in dynamics in customer behaviors as well as new challenges in analytics. The richness of the data is more than its sheer volume. It is informative and complex in two ways. First, companies are able to track various behaviors of an individual customer and link the behavior data to the customer's account. Companies observe not only customers' transaction records, but also customers' wish lists, inquiries and complaints at various encounters with the company through in-store, call-in and online channels. The data allows companies to gain a more comprehensive understanding of an individual customer. Second, companies are capturing the data about customers and market place continuously, and the data contains up-to-date information about the customers and the market. To obtain this up-to-date information, an analytics tool needs to be able to incorporate new data as they come in, and the computation needs to be simple in order to deliver the results quickly. These new features of data bring new challenges in data analysis, and managers today need advanced analytical tools to leverage the rich data available to them.

This dissertation addresses the increasing need for advanced analytical tools and proposed effective and generalizable solutions to several important managerial questions related to dynamics in customer behaviors. This dissertation is composed of three essays, each focusing on one decision-making scenario that managers frequently encounter in the area of customer management.

The first essay is titled "Detection of Customers' Life Change Using Control Chart Approach". The first essay looks into a scenario where customers' major life events play a critical role in their financial activities. As customers move through life milestones, their needs and aspirations also change, which ultimately shapes what they value in a product or service. It also affects customers' income level and budget constraints when they make purchasing decisions. However, it is difficult for managers to customize direct marketing based on major life events because they lack information on customers' life events. This is a prevalent phenomenon in practice. Luckily, big data make it feasible to infer life changes from observed information such as customers' activities and individual characteristics. The goal of this paper is to develop an effective solution to detect customers' life changes using the rich information collected through customer management system. Drawing from the literature from the field of statistical process control, I develop a sequential hypothesis test of life changes based on intuitions from Cusum control chart. The proposed solution made dramatic change in the original design of the Cusum control chart in order to accommodate the complexity in customers' behaviors and marketing context. The recovery of this life change information can help managers better customize direct marketing efforts based on customers' life events.

The ability of the proposed solution to detect a major life change of interest is tested using data sponsored by a bank in the United States. Managers of the bank want to detect a specific type of major life change, a career change, using information about customers' financial activities, their communication with the bank and their individual characteristics such as their experience in their last job. We compare the performance of our method with the hidden Markov model using both empirical data and simulation. In both settings, we are able to show better performance than that of the hidden Markov framework.

The second essay is titled “Quick Detection of Changes in Market Conditions”. It extends the application of Cusum control chart to detect changes in the market’s responses to marketing stimulus. In today’s turbulent and dynamic market, adaptability is considered as a new competitive advantage, and a key component of which is to read and act on signals. Although existing literature has identified many factors that cause the change and proposed indexes that help managers to monitor the marketplace, they all suffer from two common limitations. First, none of these studies provide clear guidance on what qualifies as a significant change. They all require managers’ own judgment to filter the noises to identify the signal of change. Second, these studies generally focus on one aspect of the product such as brand premium. They lack the consideration for the change in the whole pattern of market response. The goal of this paper is to propose an effective solution to detect changes in market response as quick as possible. Another fine property of the Cusum control chart is employed in the second essay: the test statistic can help managers to trace the time point when changes occur. This time point provides a road map for managers to conceptualize and understand the underlying cause of the change. This nice property allows managers to trace back to the time point when abnormal behaviors first emerge. This information about time helps managers to narrow down possible causes of the change to a few events happening around the identified time point. Thus, the knowledge of the time serves as an initial point for insight generation where brainstorming and investigation start.

To demonstrate the applicability of our solution, we apply it to the tooth paste category in the The Dominick's database. During the observation period of the data set, several new products and brands enter the market. Previous studies based on the data have shown significant impacts of these new product entries to the incumbents in the market. We propose to illustrate the effectiveness of our method by examining how the proposed method helps to identify trends that are difficult to observe.

The third essay is titled "Synergies among Partners in a Coalition Loyalty Program". It examines customers' shopping behaviors using a coalition loyalty program in Europe. A Coalition loyalty program is a loyalty card platform or system that allows customers to earn rewards from two or more merchants. This mechanism provides the convenience to the customers in redeeming their rewards and thus has quickly gained popularity among the customers. However, debates exist on the benefits of these coalition loyalty programs. The concerns of the coalition loyalty program arise from the comparison with the traditional one-vendor loyalty program, in which rewards are used to encourage repetitive purchase and cultivate customer loyalty. It is believed that this link between reward and patronage is diluted by the presence of other merchants in the program and the rewards largely foster the loyalty to the reward program, instead of the merchants. On the other hand, people in support of the loyalty program argue that coalition loyalty help to reduce the cost of managing the loyalty program and allow merchants to share customer information. More importantly, coalition loyalty program makes use of the network effect among the merchants and attract customers to shop within the network. Thus, the objective of the study is to understand whether synergies exist among merchants participating in the coalition program. The research question is, do customers who cross-buy at multiple merchants yield higher value for the company? The knowledge on the synergies among partners help merchants in the network to identify potentially valuable customers through their cross-buying behaviors. Being selective in choosing the customers to serve is a critical issue. When a merchant has access to a whole pool of customer information in the coalition loyalty programs, it is a common and tempting mistake to market to all customers. However, untargeted cross selling is not beneficial. Previous studies had shown that in general, one in five customers who cross buy are unprofitable and they account for the 70% of customer loss. Thus, it is critical to identify the valuable partners within the network of the coalition loyalty program and conduct effective marketing contacts base on underlying associations among merchants in their customer resources.

In sum, the goal of this dissertation is to provide managers with effective and practical tools that transform customer data into insights on customer dynamics and improve the effectiveness of marketing efforts.

Chapter 2

Essay 1

Detection of Customers' Life Change Using Control Chart Approach

2.1 Introduction

As customers move through life's milestones—including graduation, career change, marriage, parenthood, home ownership, and chronic disease—their needs and aspirations also change. Major life changes ultimately reshape what customers value in a product or service. Such changes also affect the income levels and budget constraints that influence their purchasing decisions. Yet, although life events have important implications for managerial decisions, it is difficult for managers to apply these insights because they lack information on individual customers' life changes. In practice, customers do not update their life events in a firm's database, at least not in a simultaneous and immediate fashion. The result is that managers sit on a pool of outdated customer profiles that were usually collected on the first day of registration. Managers need a way to detect new changes in customers' lives.

Big data makes it feasible to infer life changes from observed information such as customers' activities and individual characteristics. Customer management systems today capture rich information about each customer. Banks, for example, have records on customers' financial portfolios and their contacts with the banks through walk-ins, phone calls, mail, e-mail, and visits to websites. Banks also acquire information about customers' financial activities at competitor banks through third-party data services, such as IXI. These pieces of information, when combined, can provide us with critical inferences on customers' life changes. Consider the case of a family preparing to purchase a home. The family is likely to increase their savings for the down payment, make more frequent visits to financial advisors, and browse information on mortgages and home insurance. In this way, a life change causes a systematic shift in a customer's behavioral pattern. By detecting these behavioral signals in the data, we can infer a family's intention to buy a home. The question is, how do we quickly detect a shift in a customer's behavioral pattern that signals a life change of interest?

This task is challenging, for three reasons. The first challenge is that customers typically show large variances in their behavior over time. The systematic shift that can be attributed to the life change of interest, however, is small. It is difficult to detect the signal of a small shift in an individual's behavioral pattern amid the noise of that individual's other behavioral variances. The second challenge is the selection of the optimal window for change detection. Because the shift is small, optimizing the time window for change detection is crucial. In an ideal case, it is most efficient to test for the shift starting from the actual change point. However, one cannot know in advance whether and when a customer's life change will occur. To ensure the efficiency of the test, one must identify the most likely change point based on the data available. The third challenge is that, in real-life application, data arrives every new period. To obtain timely intelligence on a customer's life change, an algorithm needs to be able to incorporate new data as they arrive and produce the most up-to-date results.

The objective of this study is to develop a solution to detect a specific life change of interest while addressing these three challenges. (Note, however, that our solution is not specific to a certain type of life change. This framework can be applied to detect any type of life change that can cause systematic changes in customer behavior.) Drawing from the literature of the field of statistical process control, we develop a sequential test of a life change of interest based on the framework of the CUSUM control chart. In the proposed solution, we construct the problem as one of hypothesis testing, the goal of which is to test for the shift in a customer's behavior that signals the life change of interest. The test statistic accumulates deviation in the direction of interests over time as evidence of life change. The efficiency of the test is enhanced by selecting the optimal window of observations for the testing and by modifying this optimal window dynamically as new data arrive. We thereby render the test statistic sensitive to a shift in behavior pattern.

To the best of our knowledge, our research is the first to introduce the CUSUM control chart for change detection into a general marketing context. The CUSUM algorithm is designed for real-time analysis. It offers a recursive equation describing relationships between the statistic at time t and t-1, which simplifies the computation as new data arrive. Our proposed solution extends the original design of the CUSUM control chart to accommodate the complexity of the customer management context.

It is a straightforward task to measure the parts and calculate their deviations from the designed norm in a typical quality control setting. It is more difficult to gauge deviations in customer behavior, however. In this paper, we present a detailed solution for constructing the likelihood function in order to extract typical pattern before and after life change and to use the likelihood to evaluate deviation from the typical pattern before change. We further extend the test statistic to adjust for individual and circumstantial differences in the probability of changes in customer behavior. Despite the addition of these complexities to the model, we are still able to maintain the simplicity of the CUSUM method by deriving a recursive formula for the test statistics between time t and time t-1. Success in deriving this recursive formula is crucial for applying the method to the big data scenario. In this way, this study also contributes to the literature of statistical process control by allowing the CUSUM control chart to accommodate more variety in the data.

We demonstrate the applicability of our method using data sponsored by a Fortune 500 financial services company. Managers of the bank wish to detect a specific type of major life change—career change—using information about customers' financial activities, their communications with the bank, and individual characteristics such as their lengths of time in the job. We evaluate the ability of our solution to detect life changes using both empirical data and simulation. In both settings, we are able to show better performance than that of the benchmark model.

In the next section, we review the literature on life changes and models for regime change. We then discuss the empirical context and the data to establish the context for the pro-posed solution. We then describe the details of the model and present results. Finally, we summarize our methodological and managerial contributions and discuss directions for future re-search.

2.2 Literature Review

The Importance of Major Life Changes

Major life events are valuable information for marketing managers. For example, PRIZM, a well-known system for customer segmentation, incorporates information on customers' life cycles with customers' life styles along with geographic information to effectively segment US customers. Existing studies have established the impact of major life events on a wide range of customer behavior, including consumption level (Gourinchas and Parker 2002), brand preferences (Andreasen 1984; Mathur et al. 2008) and financial behaviors, such as investing (Cocco et al. 2005), buying insurance (Wilkes 1995), and loan payment (Baek and Hong 2004). Researchers have also found that life stages classified based on major life events provide meaningful interpretations of customers' consumption patterns (Du and Kamakura 2006; Lansing and Kish 1957). The occurrence of major life events are found to be related to demographic factors, such as age, education level, family structure, employment opportunities, and economic resources (Benzies et al. 2006; Kreyenfeld 2010).

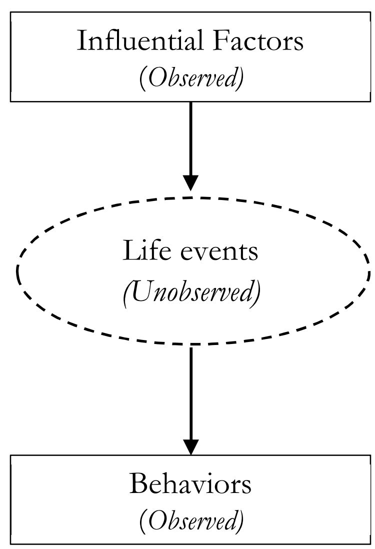

Figure 1: Conceptual Framework for Proposed Detection of Life Change

Our research takes a different path. Unlike previous research, our study does not observe life changes directly. Our goal is to develop an efficient solution to detect life changes after they occur. Based on existing knowledge about the connections between life events and customer behavior, we use changes in customer behavior as indicators of life changes. We also exploit demographic information to account for individual differences in customers' propensities to undergo life changes. Figure 1 presents a conceptual map of the problem, as well as available information. The challenge of this task is that the observed variables, when examined individually, are weak indicators of life changes. Advanced technique is needed in order to quickly detect life changes using rich customer data accumulated over time.

Evaluating and Modeling Changes

The detection of life changes can be framed as a problem of change point detection or a problem of classification (labeling observations as "no life change" or "life change"). In this broad sense, several methods in the existing literature are related to this problem; they can be categorized into four groups.

The first category contains studies using event study method and studies on structure break (Sood and Tellis 2009; Wiles et al. 2010; Perron 1989). Event study is a statistical method to evaluate the impact of an event on the value of a public firm, which is generally evaluated by the stock price. For example, studies in the past has investigated the impact of product placement in successful films (Wiles and Danielova, 2009) and marketing alliances (Thomaz & Swaminathan, 2015) on financial outcomes such as stock price and risk. A comprehensive explanation of the method can be found in the comprehensive review by McWilliams and Siegel (1997). The essence of the method is a regression analysis with dummy variables indicating the occurrence of the event. Researchers typically apply these two methods when the event is known to have occurred and the date of the occurrence is also known. In some cases, the date of the event is difficult to determine. In such a case, a common strategy is to evaluate all possible dates and choose the one that most favors the hypothesis of regime change. This remedy, however, are difficult to implement when the occurrence of the event is uncertain and the possible dates of the event cover a long period of time.

The second category of detection methods contains studies using cluster analysis, which is an algorithm for categorizing objects into groups by minimizing within-group distance and maximizing between-group distance (Liao 2005). For example, researchers have applied this method to create taxonomy of buyer-seller relationships in business market based on key factors such as information exchange, legal bond and the level of cooperation (Cannon and Perreault, 1999). The clustering algorithms is an iterative process that screens through the entire data set for the ideal partitions. Such algorithms, however, can be time consuming when applied to big data and are not suitable for real-time analysis.

The third category contains studies using logistic regression, discriminate analysis, and machine learning methods, such as decision-tree algorithm (Morrison 1969; Punj and Stewart 1983). In these methods, coefficients, or weights, are estimated for all factors in order to calculate propensity scores for group memberships. Both of these methods are designed for cross-sectional analysis and are not methods for time series data.

The fourth category contains two time-series methods: the survival model (Helsen and Schmittlein 1999) and the hidden Markov model (HMM). These methods have two merits: (a) both incorporate time-varying variables to infer the propensity of an event, and (b) both provide simple computation schemes that allow newly arrived data to be easily incorporated in the analysis. The HMM framework provides a more flexible way than the survival model to simultaneous model the different relationships on how influential factors impact the transition process and how behavior reflects changes in the underlying states (Fader et al. 2004; Schwartz et al. 2014; Schweidel et al. 2014; Wedel 2000; Schweidel 2011).

The hidden Markov model is the state-of-the-art for modeling underlying processes

(Netzer et al. 2008). It has also been widely applied in the marketing literature to model unobserved processes that guide customer behavior, such as the status of customers' relationships with firms (Netzer et al. 2008) and competitors' actions (Moon et al. 2007). However, the HMM is less sensitive to the change because it utilizes all past data to recover underlying states. The major disadvantage of using all past data is that the test statistics will take in all previous behavior variations in detecting the current life-change event. Consequently, we may end up with a low thus unconvincing probability in the life-change state because of the dilution from the previous behaviors variations. In other terms, in order for us to confirm the focal shift in the test, it will either require stronger signal to even up the prior variation or take longer time to detect the shift.

In sum, no extant method is able to dynamically select optimal observation windows for change detection while remaining simple and feasible for real-time analysis. Our solution, based on the CUSUM control chart, fills this gap.

2.3 Empirical Context

This study was conducted in the context of a Fortune 500 financial institution. (We thank the Wharton Customer Analytic Center as well as the sponsoring company for offering this data set.) Its managers are interested in detecting a specific type of major change in customers' career trajectories, a change that has great implications for managing a customer's portfolio. For reasons of confidentiality, I cannot reveal the specific career change of interest and the name of the bank in this case. Examples of this type of career change include leaving a previous job to attend graduate school; leaving a previous job to start a new company; and retiring. Such major career changes can fundamentally alter a person's financial situation, resulting in new needs for financial products. The solution is developed under the following data conditions, which are also generalizable to other customer management settings.

1. ***The type of life changes to be detected is given.*** In this setting, managers have identified the type of life change that is important for the business. I can therefore extract the typical behavior, before and after life changes, from the historic data and use these patterns as signals of life changes.

2. ***The shift in behavior due to life change is small compared to the size of variances in behavior.*** Customers show large variance in behavior, both among customers and within a single customer's data. Customers tend to conduct many of the same activities at different times for different purposes, and these purposes are not necessarily related to the focal life events. A record of a customer buying a baby play yard, for example, is not a strong indicator of parenthood because the customer can purchase the same play yard for his or her friend's baby shower. A record of a series of purchases of items such as baby formulas, diapers, and toys over a month, however, *is* a strong indicator of parenthood. Consistent and systematic changes over a wide range of a given customer's behavior effectively distinguish a major life event from a one-time event. An effective solution, therefore, should utilize holistic behavioral patterns and accumulated evidence over time.

3. ***It is not known whether and when a change will occur in a customer's life.*** In an ideal case, it is most efficient to test for changes in behavior starting from the change point. In our setting, however, the change point is not known when conducting the detection, and the temporal range when the change point might occur can span over one or two decades. Because the shift is small, it is critical to select the optimal observation window for life change detection. While a short observation window might not contain enough behavioral evidence to confirm a life change, a long observation window might include observations before a life change. Lumping behavior that precedes a life change together with behavior following a life change can dilute evidence of change. An efficient solution should dynamically select an appropriate window tailored for each individual customer.

4. ***The analysis should be able to incorporate new data as they arrive and generate actionable intelligence in real time.*** The marketing data, such as those from customer management systems and social media, are generated continuously. It is desirable for the company to obtain the most up-to-date intelligence about customers. The algorithm therefore needs to be scalable for application to real-time analysis.

5. ***The data contain different types of factors that are indicative of a consumer's propensity to undergo the life change of interest.*** In the customer management context, I observe two types of factors that are related to life changes. One type of factors is behaviors; changes in behaviors reflect changes in customers' lives. The other type of factors is the conditions that influence the propensity for life changes. Examples of this factor are individual characteristics such as age, gender, and work experience. Analysis can exploit both factors to detect life changes.

6. ***Observed behavior data contain both continuous and discrete variables.*** Customer data contain variables of different types, including continuous and categorical variables. In order to capture the holistic pattern of behavior before and after life changes, analysis needs to account for correlations among variables of different types, as well as auto-correlations of behavior over time.

7. ***Historic data are available and contain information on the actual time of the life change for purposes of validation.*** The data set is then divided into two. A sample of the historic data can be used for calibration to capture behavior pattern before and after life changes. Another sample of the historic data can be used for testing to evaluate the performance of the proposed solution.

2.4 Data

The anonymized dataset contains observations on 98,088 randomly selected customers over seventeen months from January 2012 to May 2013. Only 12,982 customers remain in the study; the rest are excluded from analysis because of missing information. The majority (81.44%) of the excluded samples lack information about customers' career changes.[1] This high percentage indicates that the company's managers have very little knowledge of customers' career changes, even though they considered this knowledge to be critical.[2]

In this data set, I observe a wide range of customer behavior on a monthly basis. These observations can be categorized into two types. One considers whether customers possess financial products at the bank; the other considers the number of customer contacts with the bank regarding financial products. I further group the financial products by their functionality: basic banking products, investments, loans, and insurances. I thereby obtain eight variables on customers' ownership and communications regarding each type of banking product. I single out the possession of auto insurance and checking accounts as two variables because these are the two most popular products and attract more than half of the bank's customers.

Table 1 presents the descriptive statistics of these variables before and after change. It shows a vivid feature of the data: the variances in customers' financial activities are large in comparison with the small shifts in behavior that result from career changes. For example,

---

[1] Other missing information includes dates when customers start their careers and dates when they first become customers of the bank.

[2] Our data sponsors have put in great efforts to gather information on the career change statuses of their customers. A reasonable guess, therefore, is that the percentage of missing data on career change status is even larger than what we observed in this data set.

the mean frequency for customers to contact the bank is 1.579 before career changes. This number drops to 1.442 after a career change, showing a 0.107 decrease in frequency. However, the variance within the group of customers before a career change is 1.465 and after a career change is 1.469. No single variable, therefore, can serve as a strong indicator of customers' career changes. The challenge is to extract information from all the weak behavior indicators and accumulate the evidence over time to create an effective indicator of customers' career changes. This requires an advanced technique.

Table 1: Descriptive Statistics of Observed Customer Activities

	Description	Before Change	After Change	Difference in Mean
Products owned by customers	Whether the customer owns any auto insurance product at the bank	0.695 *(0.460)*	0.588 *(0.492)*	0.107
	Whether the customer owns any checking account at the bank	0.588 *(0.492)*	0.517 *(0.499)*	0.071
	Whether the customer owns any basic financial product at the bank	0.652 *(0.736)*	0.597 *(0.725)*	0.055
	Whether the customer owns any investment product at the bank	0.149 *(0.453)*	0.131 *(0.436)*	0.018
	Whether the customer owns any loan product at the bank	0.159 *(0.412)*	0.136 *(0.730)*	0.023
	Whether the customer owns any insurance product at the bank	0.637 *(0.824)*	0.552 *(0.884)*	0.085
Contacts between customers and the bank	Frequency of contacts regarding basic financial service	1.579 *(1.465)*	1.442 *(1.469)*	0.137
	Frequency of contacts regarding investment	0.355 *(0.856)*	0.324 *(0.827)*	0.031
	Frequency of contacts regarding insurance	0.276 *(0.730)*	0.257 *(0.711)*	0.019
	Frequency of contacts regarding loans	0.529 *(0.884)*	0.563 *(0.884)*	-0.034

Note: the numbers in bold font represent means; the numbers in brackets represent standard deviations.

The data also provide information on dates when customers first started their original career and changed their career. Based on the data, the marginal probability of a career change at different times of their career is calculated and presented in Figure 2. As we can see in the diagram, customers at different stages of their careers have vastly different propensities for career change. In particular, the probability of leaving the original career trajectory peaks in the fourth, fifth, sixth, and eighth years. Overall, the probability of a career change decreases as a customer's time in the career increases. Customers who stay on their original career path for more than fifteen years are very likely to stay on the same path until their retirement. The length of time since first taking the career can be therefore considered to be a factor that influences career change.

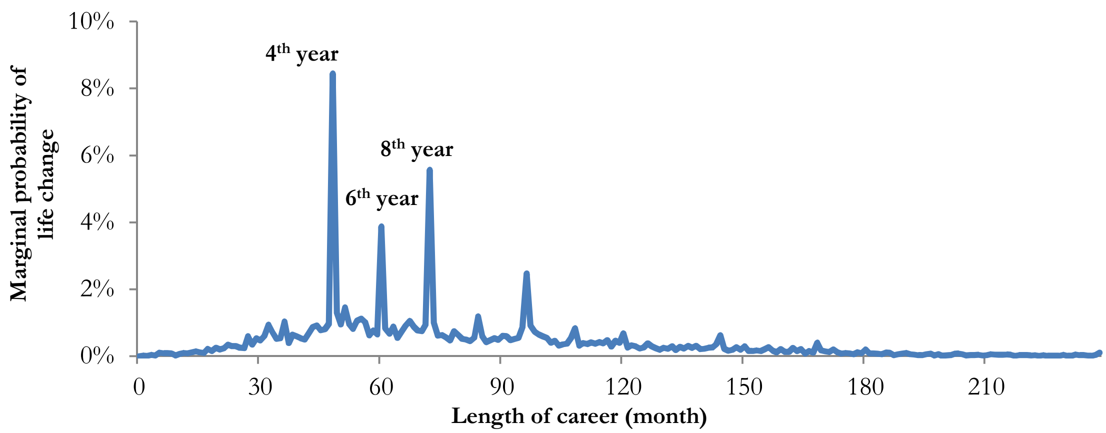

Figure 2: Marginal Probability of Career Change over the Length of the Career

2.5 Method

Engineers solve a problem in the field of quality control similar to our task of detecting small systematic shift in customers' behavior. While machines produce parts with random errors, engineers need to detect consistent, small shifts away from the design standard to avoid deterioration in quality.[3] The CUSUM control chart is considered to be one of the most efficient tools for this problem. Because of its efficiency and simplicity, the CUSUM chart is also wildly applied in computer science (Lu and Tong 2009) and public health (Chandola et al. 2013) to monitor massive data for abnormal turmoil. For example, the CUSUM control chart is used in the Real-time Outbreak and Surveillance System (RODS) in Pennsylvania and Utah for public health surveillance. Its task is to monitor data from hospitals for anomalous patterns of syndromes outbreak (Tsui et al. 2003).

The CUSUM control chart is built on the sequential probability ratio test (SPRT) (Wald 1945). Unlike traditional hypothesis testing, in which the number of observations is determined in advance, SPRT allows the test statistic to be updated as new data become available. Given the level of type I and type II error, SPRT has been shown to be an optimal test because it requires the smallest expected number of observations (Wald and Wolfowitz 1948). This makes SPRT particularly fit for real-time analysis. The CUSUM control chart, based on SPRT, further improves its sensitivity of change detection by modeling a change point in the likelihood. This change point is unobserved; it is estimated from the data. The beauty of the CUSUM test is that this complicated formula eventually reduces to a simple scheme. I describe a simple example from the context of manufacturing to provide a concrete view of the CUSUM control chart and its underlying logic.

A Univariate Example of the CUSUM Control Chart

Suppose a machine is designed to punch a hole one centimeter in diameter, but the machine produces holes with small errors. While random errors are inevitable, one-sided deviations are undesirable because they indicate a change in the machine's condition that

[3] The essence of the task is to identify deviation from the norm, which can be either deterioration in quality or improvement in quality. Any control chart for the detection of deterioration can also be used to detect improvement. In this paper, we use only quality deterioration as an example.

requires corrective attention. To detect one-sided deviations, products are constantly sampled and the holes are measured. Table 2 presents two sequences of results produced by two machines, respectively.

Both sequences have the same results until period 7. Five of the first seven observations are larger than one, indicating possibility of one-sided deviation. The procedure of testing this hypothesis by CUSUM is as follows. Let μ be the mean of the holes. The goal is to test the null hypothesis, $\mu = 1$, meaning the holes are produced as designed, against the alternative hypothesis, $\mu > 1$, meaning the holes are larger than designed.

Table 2: Examples of Two Sequences of Observations from Two Machines

Period	Machine 1	Machine 2
1	1.02	1.02
2	0.95	0.95
3	1.01	1.01
4	1.01	1.01
5	0.97	0.97
6	1.05	1.05
7	1.02	1.02
8	0.95	1.05
9	1.01	1.01
10	1.02	1.02
11	0.98	1.02
12	1.03	1.03
13	0.99	0.99
14	0.98	1.02
15	1.01	1.01
16	0.99	1.01

Let x_t be the deviation from the design of observation at time *t* (measure of the hole -1), and the test statistic in the CUSUM control chart, Q_t, is as follows.

$$Q_t = \frac{\max\limits_{1 \le k \le t} \prod_{l=1}^{k-1} L_0(x_t) \prod_{l=k}^{t} L_1(x_t)}{\prod_{l=1}^{t} L_0(x_t)}$$

Here, $L_0(\cdot)$ is the likelihood function for x_t before change and $L_1(\cdot)$ is the likelihood function after change. The change point is represented by k. Because k is estimated by choosing the time point among all past time points that maximizes the likelihood function under the alternative hypothesis. Assuming that the observation x_t follows identical independent normal distribution, the test statistic becomes:

$$Q_t = \max\{Q_{t-1} + x_t, 0\}$$

When t = 0, $Q_0 = 0$. Once Q_t is larger than a predetermined threshold, the alternative hypothesis is accepted; otherwise, monitoring of the production continues. Figure 3 presents the plot of the CUSUM control chart.

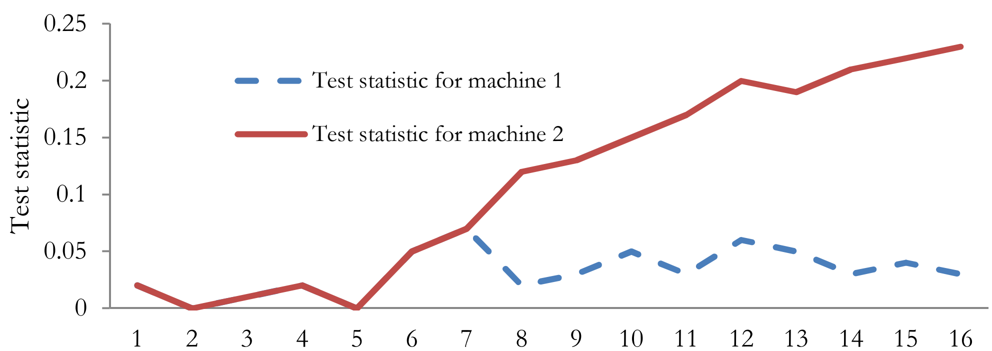

Figure 3: Control Charts for the Two Sequences of Observations

The formulation of this test statistic is in line with practical heuristics. Deviations from design are cumulated and summed over time as evidence of an upward shift in mean. In this way, deviations due to random errors in production cancel out, leaving the test statistic approximate to zero in the long run. In contrast, a consistent shift in mean will produce deviations consistently larger than zero and a test statistic larger than zero. Random errors are thereby distinguished from a consistent shift from the mean. Because cumulative sum is used as the test statistic, this method is named cumulative sum control chart.

Furthermore, the formulation of the statistic allows an automatic inference of the most likely point at which the machine's condition had changed. Because the test statistic is the maximum between $Q_{t-1} + x_t$ and zero, any evidence supporting a downward deviation is discarded. In this way, the test statistic dynamically determines the time point when the evidence should start to be accumulated. Last but not least, although the observation windows are dynamically selected, the calculation of the CUSUM statistic remains simple. When new data arrive, a test statistic can be calculated based on the new data and the test statistic from the previous period. This feature is a great fit for real-time analysis as data continuously arrive.

As illustrated in this example, the CUSUM chart was originally developed for monitoring a single feature in the manufacturing process, which typically follows an independent and identically normal distribution. Its ability to accommodate rich observations and to adjust for individual or circumstantial factors that influence changes is therefore limited. These limitations largely constrain its applicability to customer management. While deviation from design can be measured in a manufacturing setting, it is unclear how to transfer various customer activities into a measure of deviation. Furthermore, although individual differences are seldom an issue in the manufacturing field, these factors can be informative in a customer management context as illustrated in the previous conceptual map (Figure 1), because many life changes (such as career change and marriage) are related to factors such as education and age.

The Proposed Solution

In this section, I provide a detailed description of how the test statistic can be extended to accommodate the problems faced in customer management settings. I first show how the module of likelihood function can be extended to extract behavioral patterns from a mix of continuous and categorical variables, such as binary and count variables. Second, I show how factors that influence life changes can be incorporated.

(a) Likelihood of Multivariate Observations

Suppose we observe customer i's activities A_{il} at each of the time points l. Without loss of generality, suppose A_{il} contain three variables: A_{il}^1, A_{il}^2 and A_{il}^3. To capture the behavioral pattern before and after a life change, they are modeled as follows.

A_{il}^1 is a continuous variable, and is modeled using the simple linear regression.

$$A_{il}^1 = \alpha^1 + \beta^1 x_{il} + \gamma^1 A_{il-1}^1 + \varepsilon_{il}^1$$

A_{il}^2 is a binary variable, and is modeled using the probit model.

$$U_{il}^2 = \alpha^2 + \beta^2 x_{il} + \gamma^2 A_{il-1}^2 + \varepsilon_{il}^2$$

$$A_{il}^2 = \begin{cases} 0, & U_{il}^2 < 0 \\ 1, & U_{il}^2 \geq 0 \end{cases}$$

A_{il}^3 is a count variable and we can group the value of A_l^3 into *s* categories based on the distribution of A_l^3. The variable is then modeled using the ordered probit model.

$$U_{il}^3 = \alpha^3 + \beta^3 x_{il} + \gamma^3 A_{il-1}^3 + \varepsilon_{il}^3$$

$$A_{il}^3 = \begin{cases} 0, & U_{il} < 0 \\ 1, & 0 < U_{il} < \theta_1 \\ & \quad \dots \\ s, & U_{il} \geq \theta_s \end{cases}$$

We account for the auto-correlation between A_{it} and A_{it-1} by adding A_{t-1} as a covariate of A_{it}. We can also control for the heterogeneity in customers' tendencies to conduct an activity by adding personal characteristics, x_{il}, as covariates of A_{it}. To allow for correlations between different activities, we assume that ε_{il}^1, ε_{il}^2, and ε_{il}^3 follow multivariate normal distribution.

$$\begin{pmatrix} \varepsilon_{il}^{1} \\ \varepsilon_{il}^{2} \\ \varepsilon_{il}^{3} \end{pmatrix} = N(\mathbf{0}, \mathbf{\Sigma})$$

where $\mathbf{\Sigma}$ is a 3×3 variance covariance matrix.

The likelihood is:

$$L_i = f(\varepsilon_{il}^{1}) \cdot \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(\varepsilon_{il}^{2}, \varepsilon_{il}^{3} \mid \varepsilon_{il}^{1}) \, d\varepsilon_{il}^{2} d\varepsilon_{il}^{3}$$

This integral can be calculated through the GHK simulator.

The model that we describe here provides a simple and intuitive way to capture customers' behavioral patterns. It does not assume any structural constraints on customers' behavior specific to the context of financial activities. Although we illustrate this using a three-variable case, our method can be easily extended to cases with more variables. This model thereby provides a general framework for modeling a set of customer behavior of different types. For specific problems, other models might capture customer behavior more precisely. In those cases, we can simply replace the specification of the likelihood with a specification from a better model, without affecting other parts of the proposed solution.

(b) Sequential Test of Life Change given only Information of Customers' Behavior

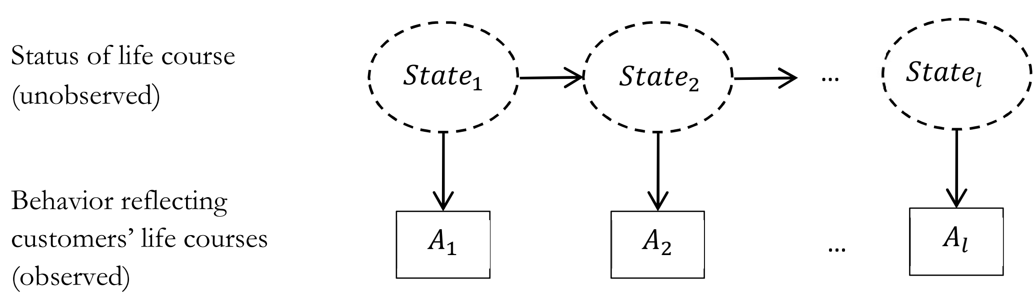

Figure 4: Data Structure given Only Information of Customer Behavior

Suppose we only observe information on customer behavior. The data structure is presented in Figure 4. The goal of the test is to determine whether a change has occurred by the current time point, t. Let C_i represent the time when a life change occurs for customer, i. We define the null hypothesis and alternative hypothesis as follows.

H_0: a life change has not occurred until current time point, t, i.e., $C_i > t$.

H_1: a life change has occurred before current time point, t, i.e., $C_i \leq t$.

For ease of representation, at time l when no change occurs, we define the probability distribution of A_{il} as $P_0(A_{il} \mid l < C_i)$. After a life change occurs to the consumer at time, C_i, the activity pattern changes and we define the new probability distribution of A_{il} as $P_1(A_{il} \mid l \geq C_i)$. Therefore, the maximum likelihood of customers' behavior under the null hypothesis is:

$$\prod_{l=1}^{t} P_0(A_{il} \mid C_i > t)$$

Under the alternative hypothesis, the life change occurs at time, k. The change point k is unknown from the data and is estimated by selecting the time point that maximizes the likelihood under the alternative hypothesis. The maximum likelihood under the alternative hypothesis is:

$$\max_{1 \leq k \leq t} \prod_{l=1}^{k-1} P_0(A_{il} \mid C_i = k) \prod_{l=k}^{t} P_1(A_{il} \mid C_i = k)$$

The ratio of the two likelihoods, Λ^{it}, is defined as:

$$\Lambda^{it} = \frac{\max_{1 \leq k \leq t} \prod_{l=1}^{k-1} P_0(A_{il} \mid C_i = k) \prod_{l=k}^{t} P_1(A_{il} \mid C_i = k)}{\prod_{l=1}^{t} P_0(A_{il} \mid C_i > t)}$$

The terms before time k cancel out. Λ^{t} becomes:

$$\Lambda^{it} = \max_{1\le k\le t}\prod_{l=k}^{t}\frac{P_1(A_{il} \mid C_i = k)}{P_0(A_{il} \mid C_i > t)}$$

In this way, the time window for the test statistic is dynamically selected: evidence of a change is accumulated from the time point at which the change most likely had occurred, and previous records are excluded from future analysis. If a change occurred at time, k, the observations can generally be better described by the post-change model $P_1(\cdot \mid C_i = k)$ than the pre-change model $P_0(\cdot \mid C_i > t)$, making Λ^{it} is greater than 1. Following the same line of reasoning, when a change did not occur before time, *t*, Λ^{it} is smaller than 1. This allows us to create a test statistic S_{it} as follows:

$$S_{it} = \max\{\ln\Lambda^{it}, 0\}$$

We are able to derive a recursive equation to describe the relationship between S_{it} and S_{it-1} that allow for agile computation when new data come. The derivation is as follows.

$$\max_{1\le k\le t}\sum_{l=k}^{t} ln\frac{P_1(A_{il}\mid C_i = k)}{P_0(A_{il}\mid C_i > t)}$$

$$= \max\left\{\max_{1\le k\le t-1}\sum_{l=k}^{t} ln\frac{P_1(A_{il}\mid C_i = k)}{P_0(A_{il}\mid C_i > t)},\ ln\frac{P_1(A_{it}\mid C_i = t)}{P_0(A_{it}\mid C_i > t)}\right\}$$

$$= \max\left\{\max_{1\le k\le t-1}\sum_{l=k}^{t-1} ln\frac{P_1(A_{il}\mid C_i = k)}{P_0(A_{il}\mid C_i > t)} + ln\frac{P_1(A_{it}\mid C_i = t)}{P_0(A_{it}\mid C_i > t)},\ ln\frac{P_1(A_{it}\mid C_i = t)}{P_0(A_{it}\mid C_i > t)}\right\}$$

$$= \max\left\{\max_{1\le k\le t-1}\sum_{l=k}^{t-1} ln\frac{P_1(A_{il}\mid C_i = k)}{P_0(A_{il}\mid C_i > t)},\ 0\right\} + ln\frac{P_1(A_{it}\mid C_i = t)}{P_0(A_{it}\mid C_i > t)}$$

$$= S_{it-1} + ln\frac{P_1(A_{it}\mid C_i = t)}{P_0(A_{it}\mid C_i > t)}$$

The resulting recursive equation of the test statistic resembles that of the CUSUM control chart. Once a negative number appears, $S_{it-1} + \ ln\frac{P_1(A_{it}|\ C_i=t)}{P_0(A_{it}|\ C_i>t)}$ is immediately replaced by zero. This means that the data before time *t* show no tendency to shift upward and are discarded for the purpose of this test. In this way, the formula dynamically changes the window of observations for the evaluation of the alternative hypothesis with new data added to the test statistic each period, ensuring agile detection of a life change once it happens. The derivation of this recursive formula is critical for the implementation of the solution in big data and real-time analysis. Without this recursive computation scheme, the computation of the statistic would be a tedious job that requires the comparison of test statistics using different time points as the change point.

(c) Sequential Test of Life Change with Additional Information on Influential Factors

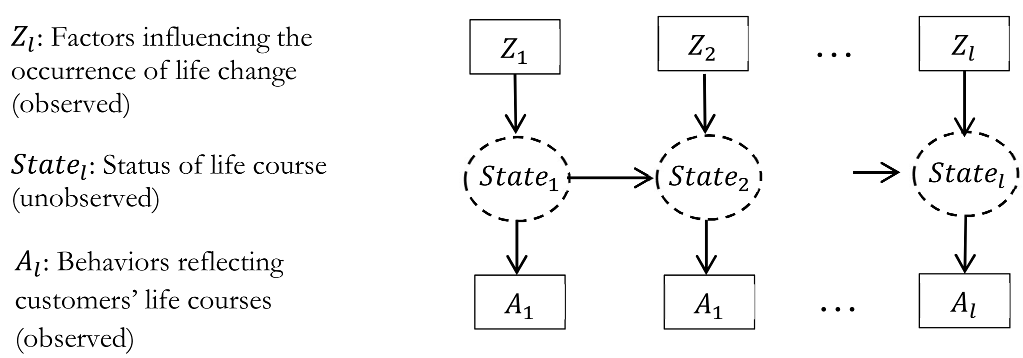

Figure 5: Data Structure given Both Influential Factors and Customers' Behaviors

In some cases, companies also observe factors that influence the probability of a consumer experiencing a life change, such as age and education. A visualization of the data structure is presented in Figure 5. The relationships between these influential factors and life changes are typically modeled using a hazard model. Let Z_{il} represent a vector of factors that influence individual i's life changes.

$$H(Z_{il}) = P(C_i = l | C_i > l - 1; Z_{il})$$

The structure of this model is different from the way we model the relationship between life changes and behavior, and it cannot be directly incorporated into the CUSUM framework. To exploit information on influential factors and behavior, we derive the test statistic as follows.

$$\Lambda^{it} = \max_{1 \le k \le t} \frac{P(A_{il}, Z_{il}; l = 1, \dots, t |\ C_i = k)}{P(A_{il}, Z_{il}; l = 1, \dots t |\ C_i > t)}$$

$$= \max_{1 \le k \le t} \prod_{l=k}^{t} \frac{P_1(A_{il} \mid C_i = k)}{P_0(A_{il} |\ C_i > t)} \times \frac{P(Z_{il}; l = 1, \dots, t | C_i = k)}{P(Z_{il}; l = 1, \dots, t | C_i > t)}$$

$$= \max_{1 \le k \le t} \prod_{l=k}^{t} \frac{P_1(A_{il} \mid C_i = k)}{P_0(A_{il} \,| C_i > t)} \times \frac{P(C_i = k | Z_{il}; l = 1, \dots, t\ ) \, P(Z_{il}; l = 1, \dots, t) / P(C_i = k)}{P(C_i > t | Z_{il}; l = 1, \dots, t\ ) \, P(Z_{il}; l = 1, \dots, t) / P(C_i > t)}$$

Because $P(C_i = k | Z_{il}, l = 1, \dots, t) = \prod_{l=1}^{t-1}(1 - H(Z_{il})) \times H(Z_{it})$, and $P(C_i > t | Z_{il}, l = 1, \dots, t) = \prod_{l=1}^{t}(1 - H(Z_{il}))$. We further derive the statistic as follows.

$$= \max_{1 \le k \le t} \prod_{l=k}^{t} \frac{P_1(A_{il} \mid C_i = k)}{P_0(A_{il} \,| C_i > t)} \times \frac{H(Z_{ik}) \,/ P(C_i = k)}{\prod_{l=k}^{t}(1 - H(Z_{il})) \,/ P(C_i > t)}$$

The test statistic therefore becomes:

$$S_{it} = max \left\{ \max_{1 \le k \le t} \ ln \prod_{l=k}^{t} \frac{P_1(A_{il} \mid C_i = k)}{P_0(A_{il} \,| C_i > t)} \cdot \frac{H(Z_{ik}) \,/ P(C_i = k)}{\prod_{l=k}^{t}(1 - H(Z_{il})) \,/ P(C_i > t)}, 0 \right\}$$

$$= max \left\{ \max_{1 \le k \le t} \left\{ \sum_{l=k}^{t} ln \frac{P_1(A_{il} \mid C_i = k)}{P_0(A_{il} \,| C_i > t)(1 - H(Z_{il}))} + \ln \frac{H(Z_{ik}) \, P(C_i > t)}{P(C_i = k)} \right\}, 0 \right\}$$

We are able to derive the recursive relationship between the test statistics S_{it} and S_{it-1} with the additional component of individual characteristics. Define S_{it}^* as $S_{it}^* = \max_{1 \le k \le t} \left\{ \sum_{l=k}^{t} ln \frac{P_1(A_{il} \mid C_i = k)}{P_0(A_{il} \,| C_i > t)(1 - H(Z_{il}))} + \ln \frac{H(Z_{ik}) \, P(C_i > t)}{P(C_i = k)} \right\}$. The recursive relationship is derived as follows.

$$S_{it}^{*} = \max_{1\le k\le t}\left\{\sum_{l=k}^{t} ln\frac{P_1(A_{il} \mid C_i = k)}{P_0(A_{il} \mid C_i > t)(1-H(Z_{il}))} + \ln\frac{H(Z_{ik})\, P(C_i > t)}{P(C_i = k)}\right\}$$

$$= \max\left\{\max_{1\le k\le t-1}\left\{\sum_{l=k}^{t} ln\frac{P_1(A_{il}\mid C_i = k)}{P_0(A_{il}|C_i > t)\ (1-H(Z_{il}))} + \ln\frac{H(Z_{ik})\, P(C_i > t)}{P(C_i = k)}\right\}, ln\frac{P_1(A_{it}|C_i = t) \quad H(Z_{it}) \qquad P(C_i > t)}{P_0(A_{it}|C_i > t)(1-H(Z_{it})\,)P(C_i = t)}\right\}$$

$$= \max\left\{\max_{1\le k\le t-1}\left\{\sum_{l=k}^{t-1} ln\frac{P_1(A_{il}\mid C_i = k)}{P_0(A_{il}\mid C_i > t)\ (1-H(Z_{il}))} + \ ln\frac{P_1(A_{it}|C_i = k)}{P_0(A_{it}|C_i > t)\ (1-H(Z_{it}))} + \ln\frac{H(Z_{ik})P(C_i > t-1)P(C_i > t)}{P(C_i = k)P(C_i > t-1)}\right\}, ln\frac{P_1(A_t|C_i = t) \quad H(Z_{it}) \quad P(C_i > t)}{P_0(A_t|C_i > t)(1-H(Z_{it})\,)P(C_i = t)}\right\}$$

$$= \max\left\{\max_{1\le k\le t-1}\left\{\sum_{l=k}^{t-1} ln\frac{P_1(A_{il}\mid C_i = k)}{P_0(A_{il}|C > t)\ (1-H(Z_{il}))} + \ln\frac{H(Z_{ik})\, P(C_i > t-1)}{P(C_i = k)}\right\} + \ln\frac{P(C_i > t)}{P(C_i > t-1)} + \ ln\frac{P_1(A_{it}|C_i \le t)}{P_0(A_{it}|C_i > t)\ (1-H(Z_{it}))}\ , ln\frac{P_1(A_{it}|C_i = t) \quad H(Z_{it}) \quad P(C_i > t)}{P_0(A_{it}|C_i > t)(1-H(Z_{it})\,)P(C_i = t)}\right\}$$

$$= \max\left\{\max_{1\le k\le t-1}\left\{\sum_{l=k}^{t-1} ln\frac{P_1(A_{il}\mid C_i = k)}{P_0(A_{il}|C_i > t)\ (1-H(Z_{il}))} + \ln\frac{H(Z_{ik})\, P(C_i > t-1)}{P(C_i = k)}\right\}, ln\frac{H(Z_{it})\, P(C_i > t-1)}{P(C_i = t)}\right\} + \ ln\frac{P_1(A_{it}|C_i \le t)}{P_0(A_{it}|C_i > t)\ (1-H(Z_{it}))} + \ln\frac{P(C_i > t)}{P(C_i > t-1)}$$

$$= \max\left\{S_{it-1}^{*}, ln\frac{H(Z_{it})\, P(C_i > t-1)}{P(C_i = t)}\right\} + \ ln\frac{P_1(A_{it}|C_i \le t)\, P(C_i > t)}{P_0(A_{it}|C_i > t)\ (1-H(Z_{it}))\ P(C_i > t-1)}$$

This recursive relationship simplifies the computation scheme, making the proposed solution fit for real-time analysis in big data.

Test Procedure

The implementation of our solution requires three steps. First, typical patterns of customers' behavior before and after life changes are extracted from a calibration sample. Second, the test statistic is calculated using the recursive formula derived above. Third, the test statistic is compared to a predetermined threshold, h. The decision rules are as follows.

If $S_{it} \geq h$, *reject the null hypothesis and report a life change;*

If $S_{it} < h$, *do not reject the null hypothesis and continue monitoring for a change.*

A given value of the threshold corresponds to a pair of true positive and false positive rates. A higher threshold h makes it easier to detect a life change, but induces a higher risk of falsely reporting a life change. Inversely, a lower threshold h reduces the risk of falsely reporting a life change while also making it more difficult to detect a life change.

We have now provided a general framework for the detection of major life changes in the context of customer management. This framework dynamically selects the optimal window for accumulating information from customers' behavior and individual characteristics for the detection of major life changes. More importantly, the framework simplifies the computation scheme so that the resulting test statistic fits within a big data scenario. In real-time analysis, where data is continuously flowing in and agile decision support is needed, the recursive formula that we derive allows us to effortlessly compute the updated statistic at time t given the test statistic from the previous period.

2.6 Empirical Analysis

Parameter Estimation

Our test statistic is composed of three parts and their parameters are estimated respectively. The first part extracts behavioral patterns before and after the career change. We jointly model the ten behavioral indicators listed in Table 1 using the framework proposed previously. Particularly, the four count variables are transformed into ordered categorical variables of three levels. The top level contains the possession of two or more products in the category because most people own no products or one product in a product category, as illustrated in Table 3.

A sample of 3,035 customers are set aside as a calibration sample. Two sets of estimates are obtained based on observations from before and after a life change, each describing behavioral patterns before or after a career change. Similar to what we observed in Table 1, the difference in observations from before and from after a career change regarding each particular activity is small. We use the remaining records of 9,947 customers as a test sample to validate the method.

Table 3: Number of Products Owned by Customers in Each Product Category

	Number of products owned in the category	Percentage
Basic products	0	51.47%
	1	33.55%
	2	14.98%
Investment	0	89.4%
	1	7.51%
	2	3.09%
Loan	0	86.3%
	1	12.25%
	2	1.45%
Insurance	0	58.88%
	1	26.05%
	2	15.07%

The second part of our test statistic captures the hazard of change at different times in a customer's career. As we observed in Figure 2, the relationship between the length of time in the career and the probability of change is quite complex and is hard to describe using simple models such as a linear or log linear relationship. Without further knowledge of the context, we estimated a non-parametric hazard model so that we do not need to impose any assumption when describing this relationship.

The third part of our test statistic is the parameters regarding the marginal probability of career change each month. In this study, we assume that customers change their career paths at the same rate, and model it as the average monthly probability of change. In cases where the change rates are different across months, the differences can be incorporated in the test statistic by modeling the marginal probability in either a parametric or non-parametric way.

Benchmark Model

As stated before, we choose the HMM as the benchmark model because it is considered the state-of-the-art framework for modeling behavior and its unobserved processes. Although the HMM has not previously been applied to infer customers' life changes, a popular approach in practice is to use logistic regression, and logistic regression can be considered a simple special case of the HMM. With additional flexibility in its structure, the HMM will perform better than the logistic regression. We therefore consider the HMM to be a good benchmark model and apply it in this study to evaluate the performance of the proposed solution. Using the HMM, life states (no change vs. change) are modeled as the unobserved states that guide customers' financial behavior. Given the life states, the calibration of the HMM reduces to the modeling of two independent components of the binary transition process and the conditional likelihood. The conditional likelihood describes customers' behavior before and after life changes. The binary transition process is modeled by the non-parametric model, which describes the relationship between the tendency to change and time in career. In this way, the HMM is applied on the same data structure and uses the same information as the proposed model does. The unobserved states are then inferred by calculating the probabilities of being in different life states (see details about the inference in Netzer et al. 2008). The inferred probability is then compared with a predetermined threshold. Customers with probabilities higher than the cut-off point are considered as having changed their careers. Similarly, when using the HMM, the determination of the threshold is also a trade-off between type I and type II errors.

The performance of the HMM and the proposed sequential test are only comparable when the transition process between life stages are modeled. The proposed solution is still able to function without modeling the transition process. In contrast, the HMM requires imposing additional assumptions in order to model the transitional process. In this case, the efficiency of the model would be harmed by wrong assumptions about the transition process.

Performance Comparison

To compare the performance of the proposed sequential test with that of the HMM, we create receiver operating characteristic (ROC) curves for both methods. ROC curves have been widely applied in the statistics field as a tool to evaluate the performance of a binary classification system when the threshold varies (Hanley 2005; Metz 1978). To create an ROC curve, the true positive rate is plotted against the false positive rate at various levels of the threshold.

Figure 6 presents the results. The line of no discrimination—the dotted diagonal line stretching from the left bottom to the top right corners—represents results from random guesses, regardless of the base rate of life changes. A curve above this line indicates better performance than a random guess, whereas a curve below the line indicates a worse performance. The results from the HMM are plotted using dashed lines, showing a substantial improvement compared to random guess. This fine performance is expected from the HMM since it is the state-of-the-art model for understanding the unobserved underlying process of customer behavior. With the efficient design of the test statistic, the performance of the proposed solution surpasses that of the HMM. The area between the curve of the proposed solution and the HMM is 0.0556, representing an 8.39% increase in performance.

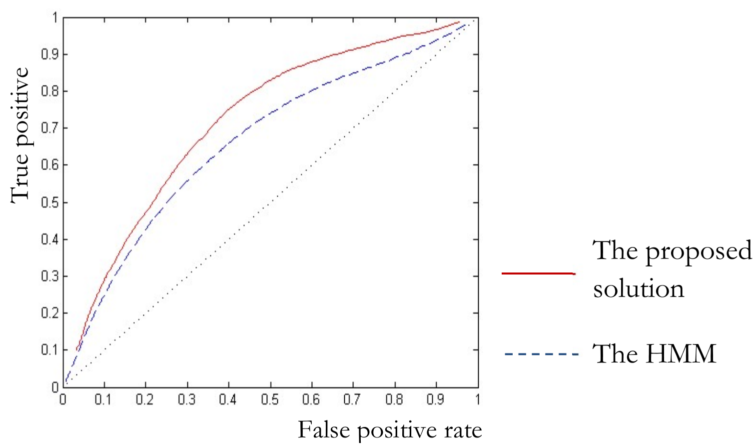

Figure 6: Comparing the Performance of the Proposed Sequential Test with the Hidden Markov Framework using Empirical Data

Simulation Study

Our data is limited in both its length and variety of observations. Consequently, we cannot assess the performance of our solution in different contexts. We therefore resort to the simulation experiments. The simulation starts when customers become members of the bank and their behavior is recorded. To conduct the simulation, we first generate customers' length of experience at work at the time they became customers of the bank. Based on these individual differences, we generate the time when customers undergo life changes. Finally, we generate customers' behavior before and after change based on the parameters from the empirical data.

Assessing the Effectiveness of the Proposed Sequential Test of Life Change

The purpose of the first simulation is to validate the results using simulation experiments and to assess the performance of the proposed solution in the long term. We simulated the behavior of 1,000 customers at the monthly level and monitored their behavior for indicators of life changes over five years.

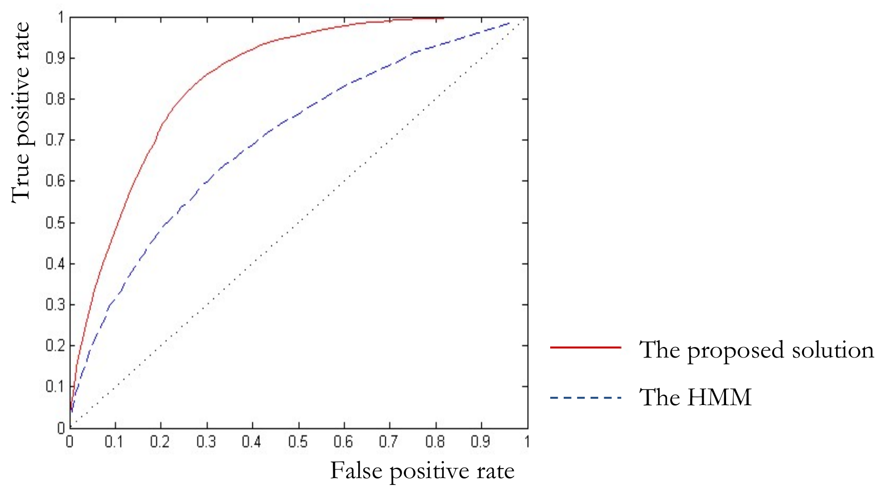

Figure 7: Comparing the Performance of the Proposed Solution with the Hidden Markov Model in Simulation Experiment

The results are presented in Figure 7. In the plot, the curve of the proposed solution is above that of the HMM, meaning that given the same level of false positive rate, the proposed model correctly detects more life changes than the HMM. The area between the curve of the proposed solution and the HMM is 0.134, representing an 18.55% increase in performance. This result validates that the proposed solution performs better than the HMM in detecting life changes. To further demonstrate the difference in the performance of the proposed model and the HMM, we provide examples of two cohorts of customers in Figure 8. The first cohort is composed of 101 customers who undergo life changes at the fortieth period of the surveillance in the simulation study (Figure 8-a). We observe that fourteen customers are reported by the proposed solution as changed before period 40.

These false detections take place only a few periods before the actual change. In contrast, thirty-seven customers are reported by the HMM as having changed and the false detections take place long before the actual change as early as the second period of the surveillance. After the occurrence of a life change, the proposed sequential test quickly detects the changes within nine periods. It takes the HMM, however, another thirty-five periods to detect all changes. A similar pattern is observed in Figure 8-b, which features customers who undergo life changes at period 70.

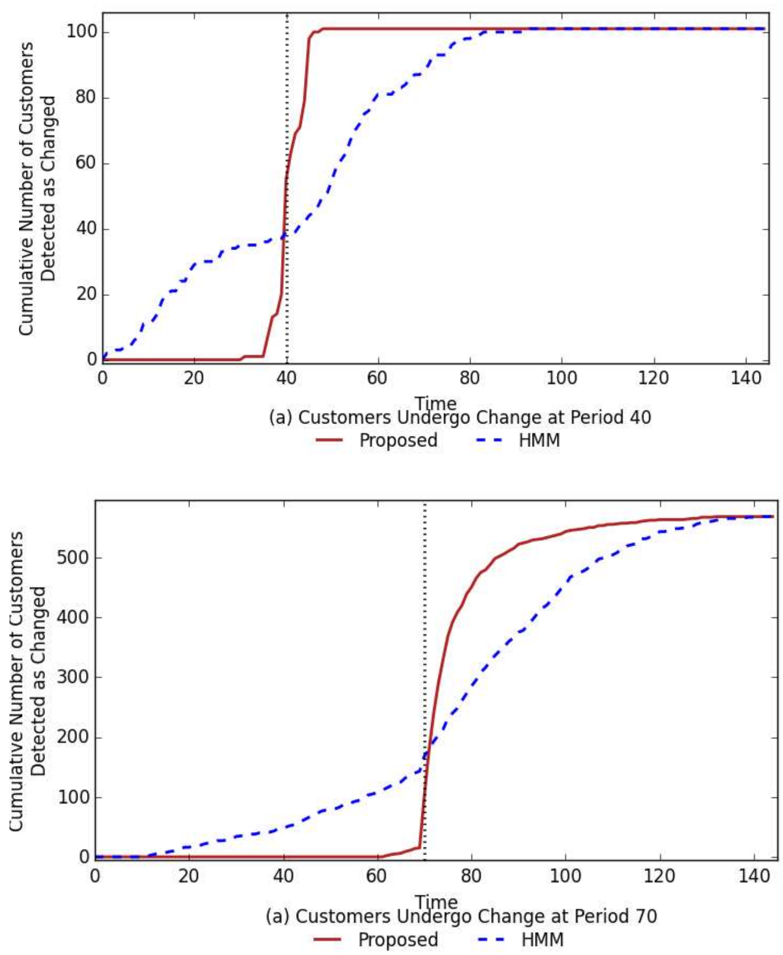

Figure 8: A Illustrative Case

The HMM performs worse than the proposed model because of fundamental differences in how the two statistics are constructed. The first difference is that while the HMM employs all past history to infer the change, the proposed solution dynamically selects the time window to test changes in behavior. Because the shift in patterns of behavior after the life change is small, when the HMM is applied, the evidence of change is vulnerable to dilution by observations before a change. The second difference is that while the proposed solution uses the hazard of change to gauge the baseline probability of change at a given time point, the HMM uses the cumulative probability of change. In this specific case, the hazard of change at the forty-eighth month of a career on average is thirty times that of any previous periods. However, in terms of cumulative probability, the overall probability of change increases from 20.29% to 28.76% at period 40, which is a much weaker signal of change compared to that of the marginal probability. These differences distinguish the two methods in their ability to detect changes based on the same information.

Assessing the Performance of the Proposed Solution when the Probability of Change is Low

We then evaluate the performance of the proposed solution at different levels of probability of change. In these simulation experiments, customers are monitored for 240 periods, and their marginal probabilities of life change per period are the same. We adjust the marginal probability of change to 1/300 and 1/2500 respectively. We present the results in Figure 9.

When the marginal probability is 1/300 (Figure 9-a), the area between curves is 0.0464. This area increases to 0.1586 when the marginal probability of change decreases to 1/2500 (Figure 9-b). The area between the two curves becomes bigger as the marginal probability of change decreases. The result shows that the performance gained from choosing the proposed solution over the HMM becomes larger when the probability of change decreases.

These simulations demonstrate the advantage of the proposed solution in a customer management setting. In general, in such a setting, the probability of life change at a given time point is low, and the possible time for life changes span over a long period of time. It is in such a setting that conducting real-time analysis and extracting information from big data becomes valuable. Otherwise, in extreme cases where the possibility of life change can be narrowed down to one or a few given time points, there is no need to constantly monitor customer behavior and dynamically identify the window for detection.

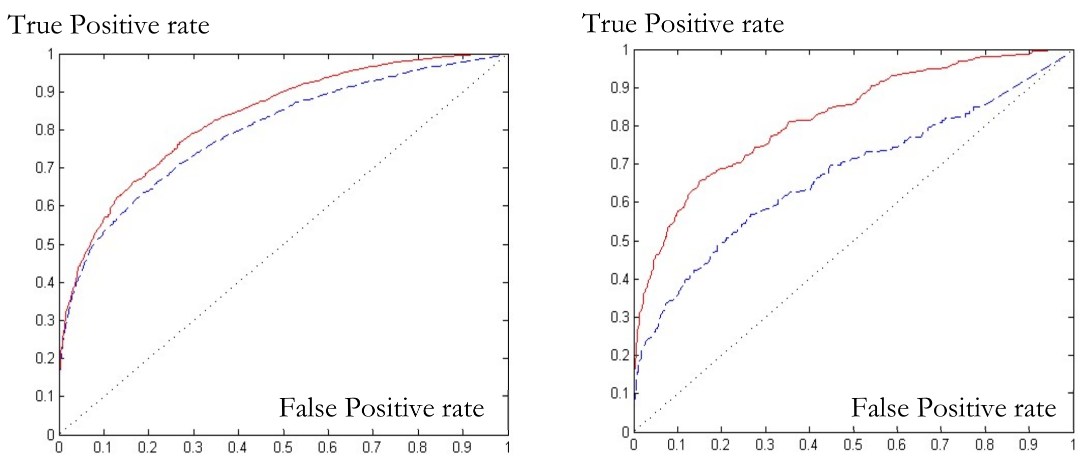

(a) Marginal Probability of Change = 1/300

(b) Marginal Probability of Change = 1/2500

Figure 9: ROC Curves When Probabilities of Change Are Different

Assessing the Performance of the Proposed Solution when given Shift in Behavior Patterns Due to Life Change is Large

We continue to evaluate the performance of the proposed solution when shifts in behavioral pattern due to life changes become large. In these experiments, we vary the shift in behavior patterns by increasing the differences in the estimates of the intercepts before and after changes, while holding other parameters the same. The difference in intercept is increased to 0.05 and 0.30 respectively. The results are shown in Figure 10.

The area between curves is 0.0913 given the original set of parameters. This area reduces to 0.0471 when the difference is increased by 0.015. The area continues to shrink to 0.0266 when the difference is increased to 0.3. As the shifts in behavioral patterns increase, the area between the two curves becomes smaller.

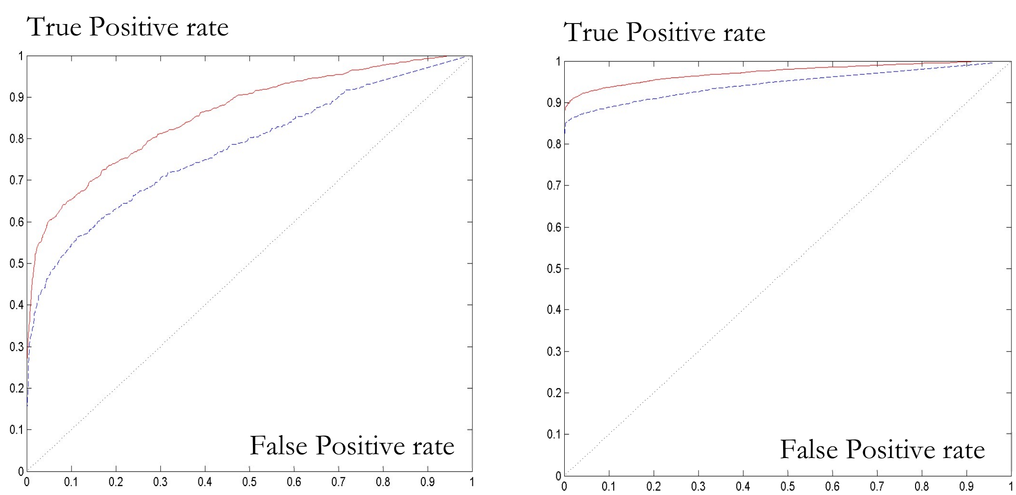

(a) +0.05 Shift (b) +0.3 Shift

Figure 8: ROC Curves When Shifts in Behavior Patterns are Different

This set of simulations demonstrates the advantage of the proposed solution when the shift in behavioral patterns is small. Small shifts in behavior are more difficult to detect when records from before and after a life change are lumped together. This becomes less of a problem when shifts in behavior become large. Strong behavioral signals improve the performance of both methods. In such cases, the performance of the HMM is already good, leaving little room for further improvement. The performances of the two methods thus converge.

2.7 Managerial Insights

We conduct an additional simulation to demonstrate the application of the proposed solution to assisting managerial decision making. The purpose is twofold. First, we evaluate the gain in profit when the proposed method is applied. Second, we demonstrate how managers can decide whether to apply the proposed solution to detect life changes and how managers can determine the optimal value of the threshold in a specific problem.

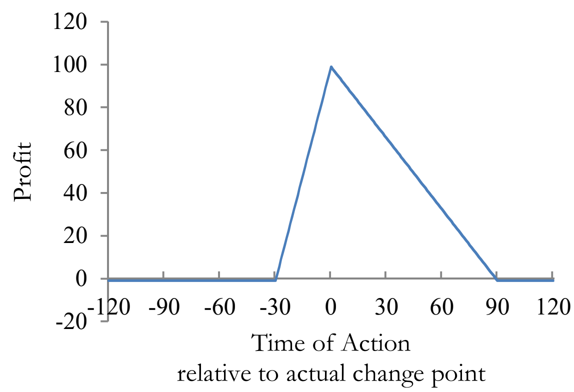

(a) Scenario 1

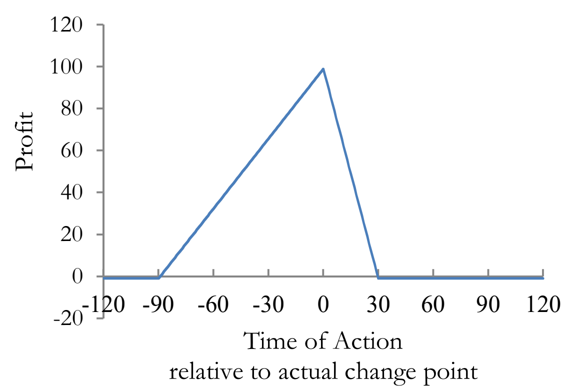

(b) Scenario 2

Figure 9: Loss Functions in Different Scenarios

In these simulations, we simulate 10,000 customers and monitor their life changes for 240 periods. The marginal probability of undergoing a life change is 1/1000 for each period. Suppose the companies take actions immediately after receiving alerts of life changes, and each action costs one dollar. We conduct two sets of simulations by manipulating the loss function in two ways.

In the first set of these simulations, we evaluate the proposed solution under different shapes of loss function. We assume the maximum revenue is one hundred dollars and can be obtained if changes are detected once they occur. As the gap between the detection time and the actual change point increases, the revenue decreases. We vary the rate that the return decreases and obtain two different scenarios, which are illustrated in Figure 11. In scenario 1 (Figure 11-a), the opportunity window after life change is long, while taking action before life changes does not earn much reward. In contrast, in scenario 2 (Figure 11-b), the opportunity window soon close after life change, but taking actions before life change can earn more reward.

To obtain the optimal value of the threshold, we first conduct simulations under different values of the threshold and then identify the optimal value through grid search. The results from the scenarios, including the maximum profits and their corresponding false positive rates, are summarized in Table 4.

The results show that in both scenarios, applying the proposed solution brings in more profit than the HMM method. Compared to the HMM method, the proposed solution brings an increase of 76.98% in profit in scenario 1 and an increase of 53.44% in scenario 2. Table 4 also shows how the value of the threshold should be adjusted based on different loss function. In scenario 1, where taking action before life change receives little reward and the opportunity window is long after life change, the threshold should be set higher to avoid type I error. In scenario 2, where taking action before a life change is rewarding and the opportunity window closes soon after life change, the threshold should be set lower to avoid type II error.

Table 4: Profit under Different Loss Function

(a) Scenario 1

	The proposed solution	The hidden Markov model
Optimal threshold	1.57	.67
False positive rate	.36	.35
True positive rate	.80	.67
Maximum profit ($)	41,872	23,659

(b) Scenario 2

	The proposed solution	The hidden Markov model
Optimal threshold	1.15	.5
False positive rate	.57	.77
True positive rate	.9	.9
Maximum profit ($)	32,132	20,941

We further illustrate the relationship between the threshold and the profit in Figure 12. Figure 12 shows how the change in the threshold influences the expected profit given the loss function in scenario 1. The relationships between the false positive rate and the expected profit show an Inverted U-shaped pattern. For the HMM, the profit reaches its peak at 23,659 when the false positive rate is set at 0.35. For the proposed sequential test, the profit reaches its peak at 41,872 when the false positive rate is set at 0.36. In sum, each value of the threshold represents a different trade-off between type I and type II error. Managers need to identify the optimal value of the threshold in their specific applications.

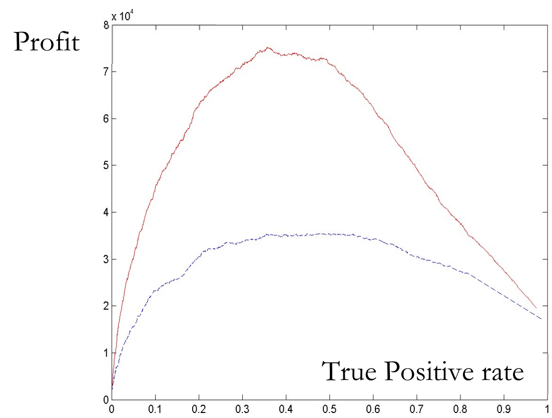

Figure 10: Profit at Different Levels of False Positive Rate in Scenario 1

The results are shown in Figure 13. When the maximum revenue is as high as twenty dollars, the total profit from the pool of 10,000 customers reaches 2,709 dollars. The profit decreases as the expected revenue from the opportunities of life changes decreases. When the maximum revenue per case further decreases to twelve dollars, the maximum profit, which is zero, is reached when the false positive rate is zero. This means that under this cost and revenue structure, the costs of false identifications is larger than the benefits of correct identifications, making it unprofitable to apply the method to detect changes.

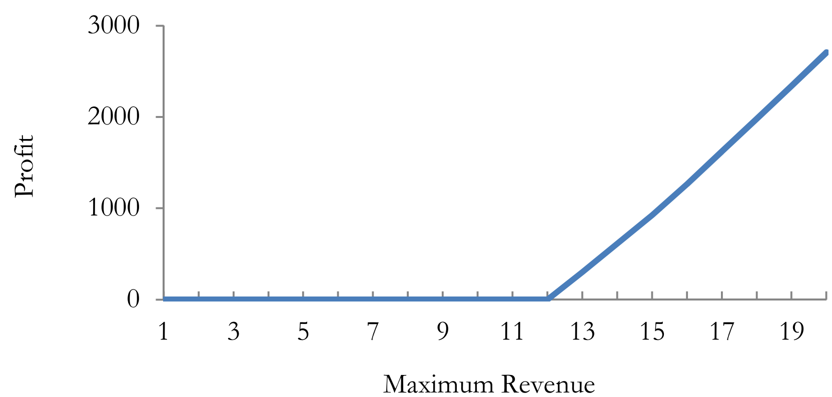

Figure 11: Total Profits from a Pool of 10,000 Customers

In sum, managers need to weigh the costs and benefits in order to determine whether to use a data-driven approach to identify potential business opportunities. Based on the revenue and cost structure, managers must then decide the optimal trade-off between false positive and true positive rates. Generally, in customer management, the return is high compared to the cost of action and the proposed solution is useful.

2.8 Discussion

Advanced big data analytics are becoming a critical driver of growth in customer value. Big data enables us to transform massive quantities of customer information into useful customer intelligence in real time to assist managerial decisions and thereby capture valuable business opportunities. Our study contributes to big data analytics by transforming copious customer data into critical information on customers' life changes. To achieve this goal, we propose a sequential test based on the framework of the CUSUM control chart and extend the test statistic in order to apply it to customer management. This solution provides a general framework for the problem of life change detection. In doing so, this paper introduces to the marketing field a new perspective in the area of change detection: a scalable approach for dynamically optimizing the window for detecting change at the individual level. We make this solution applicable to the context of customer management by accounting for the impact of individual factors in life transitions, as well as by illustrating ways of modeling multiple behavioral indicators of different types along with their correlations and autocorrelations. We thereby also contribute to the field of statistical process control by providing a way to adjust for individual and circumstantial differences in the test statistic. Despite the additional complexity, we are still able to derive a recursive formula for the test statistic, making the proposed solution particularly fit for application to big data. Our solution offers superior performance compared with the benchmark model. We also illustrate the value of its improved performance compared to the HMM under different loss functions.

Our study benefits managers by providing a practical tool for monitoring customers' major life changes in real time. This intelligence opens up precious opportunities for managers. Knowing, for example, that the customer has a newborn baby in the house, a grocery store can target the customer for diapers and formula, which are often "destination" products that drive traffic to the store. This intelligence is even more important for durable goods, such as cribs and car seats. Households purchase these products only once or twice in a lifetime and there is hardly any historical data for retailers to use to customize their marketing efforts at the individual level. The beauty of our solution is that it builds on the intuitive idea of the control chart, a notion that is easy for managers to grasp. Its easy computation scheme also makes it friendly for real-life application. Beyond detecting life changes, the proposed method can also be used to detect other types of change that are of interest to managers. Examples include detecting changes in customers' preferences in grocery shopping and changes in customers' risk of defaulting on credit card and mortgage payments.

In the specific context of banking data, the performance of our model can be further improved in two ways. One way would be to obtain more data, including a longer observation window and more detailed information on customers' activities. Our data lasts only seventeen months, which is a short window compared with customers' lifetimes at the bank. It limits our ability to incorporate other dynamics in customers' behavior when modeling behavior before and after a career change. Our data on customers' financial products is also limited to the possession of products in a given month. For example, in real life, banks observe a much richer set of information, including account balance and activities such as deposits, withdrawals, and transfers. Such additional information is also indicative of career changes. The other way is to include more existing knowledge about the financial products and specific career changes when modeling customers' behavior. It is not our aim, however, to dig deep into the data on customers' financial behavior. We keep the modeling of customers' behavior simple because our goal is to provide a general framework applicable to the type of problem that requires detection of a major life change. In a real-life application, our framework can be extended to incorporate managerial insights as well as rich findings from existing literatures to improve the performance of the test.

When applied to specific problems, the proposed solution is subject to two limitations. One limitation is in its ability to capture high dimensional data, containing thousands of variables. In such cases, a Bayesian method can be used to capture the correlations among the large number of variables. Another limitation is that we assume that customers take on a new behavioral pattern at the time of life change. In real life, customers might gradually migrate to the new pattern or present some abnormal patterns before settling down with a new pattern. Future research can account for these patterns to further improve the solution's performance.

Chapter 3

Essay 2

Quick Detection of Changes in Market Conditions

3.1 Introduction

In today's turbulent and dynamic market, adaptability is considered as a new competitive advantage, and a key component of which is to read and act on signals (Reeves and Deimler 2011). The changes can be triggered by many causes, such as the markets' intensified competition, the products' progressing life cycle and customers' evolving preferences and expectation. To keep pace with the market and get an early warning of changes in the market, companies are actively seeking, gathering and analyzing information related to their products and customers to generate actionable intelligence that allow the company to intervene with market operations in real time.

In general, two tasks are required to quickly react to changes in the market. One task is to detect changes in the market, which is to filter the information and separate signals from noises in the market. The other task is to make sense of the underlying causes of the change, which is to construct plausible theories and causes of the change in order to act to it. Adding to the complexity of the task is the explosion of information, which increases the difficulty to extract intelligence, and also the blur line of competition, which increase the difficulty to identify relevant information. The managers are in need of tools that support them through decision making in the turbulent market. This is the motivation of this study.

Existing literature has taken a theory-based approach to discover factors that causes the change. The general procedure of this approach is to construct hypothesis on factors causing changes in the market and using available data to test the hypothesis. Through these studies, researchers have uncovered many factors that influence the performance of a company or a product in the market. While these insights are great, it can be problematic for managers to apply in practice. In the turbulent market, countless events that are potentially important take place every day, and not all of them actually relate to the focal product or company. It is difficult for managers to investigate all potentially important events. Furthermore, the method adopted in these studies cannot be applied by managers to monitor the market in real time because these studies take an off-line approach to test the hypothesis. That is, these studies wait for the events and the consequences to unfold and test for the effects in retrospect. In contrast, managers need a tool that analyzes data and reports changes in real time.

A few studies propose some indexes that help managers to monitor the marketplace. Examples are the brand premium index proposed by Ailawadi et al. (2003), the brand equity index proposed by (Sriram et al. 2007), and the Google trend index proposed by Rex Yuxing et al.(2015). These indexes can be calculated for each period and is very helpful in supporting managers to make better decision in real time. However, a major inconvenience of these indexes is the lack of clear guidance in what qualifies as a significant change of these indicators. These indicators are obtained through surveys or sales data, both containing random errors. Studies have shown that human's judgment is often interfered by these noises, resulting in over- or under-reaction to the changes in the market place. Managers need tools to evaluate systematic changes in the indicators with the presence of the noises. Thus, the research question is, how to detect changes in market response as quick as possible?

In this study, our goal is to develop a solution to quickly detect changes in the pattern of the market's response to marketing stimuli. To achieve this goal, we incorporate the market response model into the Cusum control chart and develop the control chart of changes in the market. The proposed solution is composed of two steps. The first step is a sequential estimation to establish the norm of the market's response, which is to incorporate new data in the observations as they arrive. The updated parameters describe the norm of the market's responses. The next step is a sequential test to identify changes in the market place. Once the norm is established, data from the next period are compared with the norm. The residuals, which is the difference between the predicted value using the norm model and the actual observation, is used as a measure of the market's deviance from the previous pattern. These deviances are cumulated and form the test statistic. Once the test statistic passes a pre-determined threshold, managers will receive an alert of change in the market.

Our solution contributes to marketing research in two ways. First, we enrich managers' toolbox of change management by introducing a new perspective of control chart in sales management. Unlike previous proposal, our solution largely reduce managers' burden in judging whether there is a change in the marketplace through the use of Cusum test statistic. In this way, managers are more alert with the changes relevant to their brands and products. We develop this solution based on the wildly used model of market response model. A problem with many models developed in the academic field is that they receive little acceptance in practice, which is often a result of managers' unfamiliarity of the model or the complexity of the model. By leveraging existing models, we make our method more approachable to the managers. Second, the set-up of the control chart allows managers to trace back to the time point when abnormal behaviors first emerge. This information about time helps managers to narrow down possible causes of the change to a few events happening around the time point. This piece of information on time is very critical. The time serves as an initial point or the direction of insight generation where brainstorming and investigation start. The ability to provide direction and start discussion is a critical skill of leadership in the time of change.

To demonstrate the applicability of our solution, we apply it on the Dominic dataset, which contains store-level sales data. During the observation period of the data set, several new products and brands enter the market. Previous studies based on the data have shown significant impacts of these new product entries on the incumbents in the market. We propose to illustrate the effectiveness of our method by examining how the proposed method helps to identify trend that are difficult to observe.

In the next section, we briefly review the literature regarding drivers of market changes in section 2.1 and tools to monitor the market place in section 2.2. We then specify the detailed solution proposed for detecting changes in the market in section 3 and describe the data sets we will be using in section 4.

3.2 Literature Review

Two streams of studies are related to the management of dynamics in the market. One steam of studies focuses on identifying the causes of dynamics in the market. A brief review of these factors reveals that managers are unable to track and react to all potentially influential events because these events are either unobservable or too many to track. A more practical approach is to take a data-driven approach and monitor the market place for signs of changes. Once a systematic change is identified, efforts will be focused on identifying the ad hoc cause for the change.

Understanding the Causes of Dynamics of Market Responses

Studies in marketing have conducted intensive search on factors that drive changes in the market and how customers' response to marketing efforts. A major force of the changes is the intrinsic evolution of the market place as customers gradually familiarize and learn about the product or the category. In the aggregate level dada, researchers observe that for new product entering the market, the sales of the product reaches plateau after a quick rise in sales (Bass model, change-point model). At individual level, researchers found that when making a choice, customers trade between gaining information about the product category and reducing the risk of the choice. This trade-off give rises to the dynamics of the customers' decision-making. As conceptualized by Heilman, Bowman and Wright (2000), customers go through three stages in their purchasing experience: an information collection stage when purchases are focused on big-brand products, an expanded information collection stage when products from less well-known brands are sampled and a consolidation stage when products with greatest value are selected. These transitions are generally governed by an unobserved process in the market, and it is often times not obvious to the managers. It is a small shift difficult to detect. Thus, need a tool for detection.

Another major force is the marketing efforts by the company or other players in the market. Marketing campaigns can help to expand the category for the companies and cannibalize the sales of its competitors. When companies in the category all take a more aggressive approach towards marketing and increase their advertising spending, the ones that keep advertising at the original low level will suffer a decrease in their brand equity. Beyond the competition within the same category, the sales of the products are also influenced by products in other categories. As the market becomes more complex, the lines between product categories become fuzzier, and a product or can engage with competition with products in other categories. An example is the competitions for lifestyle branding. Chernev, Hamilton and David Gal found that when brands reposition themselves as means for self-expression to avoid head-on competition in the product category, they expose themselves to a broader, cross-category competition for a share of customer' identity. Other cross-product categories are observed in the competition for self-space and etc. Beyond sales, companies also compete in creative ways to engage the customers and an example is the product placement ads in movies. Researchers found that after the concept of product placement ad emerge, more and more products adopt the strategy. As a result, the effectiveness of the product placement ad experiences an inverted U-shape relationship between the number of exposures and ad response (Karniouschina, Uslay and Erenburg, 2011). Studies on marketing and finance interface also identify many marketing activities that can result in abnormal returns in stock. As we can see in this brief summary, there are many factors that can result in the dynamics in the market. The impacts of these factors are also conditioned on the circumstantial factors and the position of the company in the market. It is difficult for managers to keep track of these events and many of these events might not be relevant to the company. To complement manager's ability to identify critical events that are relevant to the product, a data-driven tool is necessary.

Tools for Managing the Dynamics of Market Responses

Efforts on monitoring the market have been focused on developing a valid measurement for brand equity. Based on the source of the data, the indexes can be grouped into four categories: (a) index based on survey, (b) index based on sales, (c) index based on financial market and (d) index based on social media and search data. The major emphasis of these studies is on the construct validity of the measure. Researchers have done great work in laying down a rigorous theoretical foundation for these measures and conducted great amount of empirical work to validate the measure by examining their correlations with other commonly available measures. As a result, we now have a set of measures that are grounded in the theory and complement each other to encompass all facets of brand equity. The validity of these measures ensures that these measures reflect the improvements and deteriorations in the performance of brands and products.

However, beyond the computation of the measures, there is little guidance in the literature on how managers should interpret these indexes. Managers need to rely on their own judgment to decide whether a systematic change have taken place in the market. Unfortunately, studies have found that human are often limited in the judgment of regime change. Researchers in the field of decision-making have conducted studies of human's capability in the detection of regime change under different settings. In these studies, a sequence of data is presented to the respondents and they are asked to judge whether there is a change in the underlying regime or the probability that the data is drawn from a certain regime. Respondents either receive the entire sequence of data at once, or in a sequential fashion. The studies also vary the environment of the decision making, such as the payoff structure and the frequency of the stimuli. The results show that while respondents generally respond in the right direction, their answers are far from optimal. Respondents show the tendencies of both under-reaction and over-reaction. Thus, it is necessary to provide the managers with a more rigorous tool to help interpreting the existing measures and detect changes in the market environment. In the following sessions, we first describe how we envision the process to monitor and react to changes in the market. Then we provide a detailed solution on how managers can detect changes in real time.

3.3 A Control Chart Approach to Monitor Market Responses

The Procedure for Monitoring Market Changes Using Control Chart

Managers can adapt the procedure of quality control to monitor changes in the market. The procedure contains five steps, which are: define, measure, analyze, improve, and recalibration. To implement a monitoring process using control chart, managers start by defining a model for the outcome variables of interests. This model includes the outcome variable of interest, such as sales volume and market share and factors that can influence these outcome variables, such as price and promotion. Managers then proceed to define the type of deviation from the original model and the level of deviation that is of interest to the management. With these clear objectives, we can then construct an efficient index to serve the purpose of quickly detecting the deviation as new data arrives. Based on the alerts from the control chart, managers can quickly take actions in response to the change in the market. After a change in the market place, managers often need to recalibrate the model to adjust for differences in the market. We provide a roadmap of this process in Figure 14.

The procedure illustrated in Figure 1 is applicable to any type of changes in the market. To effectively detect different types of changes, we need to construct the control chart index differently. Generally speaking, a large and sudden deviation can easily be noticed by managers. In the control chart paradigm, this type of change is typically monitored using Shewhart Chart. In comparison, a small but persistent change in the market is more difficult to be identified. The construction of an efficient index for the detection of small and persistent changes in the market is thus the focus of this paper.

***Define** the scope of the monitoring system clearly. The outcome of this step is a market response model that clearly defines the relationship between marketing inputs and key performance outcome and the scale of the deviation that is of interest to the management.*

↓

Measure the deviation of the market's performance from existing model using control chart index. Once the deviation is higher than the pre-determined threshold, we alert managers of a change in the market.

↓

Analyze possible causes of the change based on the direction from the control chart. The top candidates of the causes are selected for further validation.

↓

Improve the effectiveness of marketing activities based on the identified causes in order to prevent further losses in the market.

↓

Recalibrate the model after the occurrence of change.

Figure 12: The Procedure for Monitoring Market Changes Using Control Chart

Beyond the recognition of change, another important step is "analyze", which is to identify possible causes of the change. The understanding of the possible causes helps us generate actions in response to the change. To generate speculations of the causes, it may require further data analysis, intuition of the managers and extensive searches for relevant events, and it requires skills and efforts beyond the discussion of the control chart method and the scope of this paper. As the discussion and analysis unfold in the following sections, we would address how the control chart index can facilitate us in the search of possible causes of the event.

A Basic Framework of Control Chart and its Limitation in Monitoring Sales

We construct the index based on the Cusum index, which has its root in the sequential likelihood ratio test. In its very basic form, the Cusum control chart is a univariate test, and the goal is to test whether the distribution in the observation, o_t, has changed from the original distribution $\boldsymbol{f_0}(\cdot)$ to a new distribution $\boldsymbol{f_1}(\cdot)$ by the current time $\boldsymbol{t}$. We can define the null and the alternative hypothesis as follows.

H_0: a life change has not occurred until current time point, t.

H_1: a life change has occurred before current time point, t.

Under the null hypothesis, the observations follow the original distribution from time 1 to current time t, and thus the maximum likelihood of the observation under the null hypothesis is:

$$\prod_{l=1}^{t} \boldsymbol{f_0}(o_t)$$

Under the alternative hypothesis, a change occurs at time, k. Therefore, the observations follow the original distribution $\boldsymbol{f_0}(\cdot)$ before time k and a new distribution $\boldsymbol{f_1}(\cdot)$ after time k. The change point k is unknown from the data and is estimated by selecting the time point that maximizes the likelihood under the alternative hypothesis. The maximum likelihood under the alternative hypothesis is:

$$\max_{1 \le k \le t} \prod_{l=1}^{k-1} \boldsymbol{f_0}(o_t) \prod_{l=k}^{t} \boldsymbol{f_1}(o_t)$$

The control chart index is then constructed as the log of the ratio of the two likelihoods under different assumptions.

$$\boldsymbol{Q_t} = \boldsymbol{log} \frac{\max_{1 \le k \le t} \prod_{l=1}^{k-1} \boldsymbol{f_0}(o_t) \prod_{l=k}^{t} \boldsymbol{f_1}(o_t)}{\prod_{l=1}^{t} \boldsymbol{f_0}(o_t)}$$

If $\boldsymbol{f_0}(o_t)$ **and** $\boldsymbol{f_1}(o_t)$ are specified as normal distribution with the same variance, the mean of the original distribution $\boldsymbol{f_0}(\cdot)$ equals to μ_0, and the mean of the new distribution equals to μ_0 , we can derive the following equations. To monitor the upward derivation, the index becomes:

$$\boldsymbol{Q_t} = \boldsymbol{max}\{\boldsymbol{Q_{t-1}} + (o_t - \mu_0) - \frac{\mu_1 - \mu_o}{\mathbf{2}} \text{ , } \mathbf{0}\}$$

To monitor downward derivation, the index becomes:

$$\boldsymbol{Q_t} = \boldsymbol{min}\{\boldsymbol{Q_{t-1}} + (o_t - \mu_0) - \frac{\mu_1 - \mu_o}{\mathbf{2}} \text{ , } \mathbf{0}\}$$

There are several problems that prohibit us from directly apply the above formula on sales figures in order to monitor market responses. First, sales feature is a function of marketing efforts, such as features of the product, price, advertising, and sales promotion. Not accounting for the impact of these known influential factors of sales can resulted in creating more noises in the control chart index and reducing its efficiency in detecting the systematic changes in the pattern. Furthermore, many changes in the market are not only reflected in the mean level of the sales, but also the responses to the marketing stimuli. For example, when more companies adopt a more intensive promotion schedule, the effectiveness of promotion can decrease. When market response model is incorporated, the control chart index can also capture the decrease in sales level caused by decreased effectiveness of promotion. Second, in practice, the true parameters for the market response models are unknown. In order to establish the norm of the market's behaviors as the baseline for comparison, we propose to dynamically update the norm of market's responses as new data arrive. Third, the control chart requires the specification of the alternative hypothesis. For example, in the univariate case, it requires the specification of the μ_1

conduct the test. In practice, however, managers would want to identify any types of changes that take place in the market. We design a solution to simultaneously solve these problems.

A Control Chart Index for Monitoring Changes in Market Responses

Suppose at week t, the sales of a product i is S_{it}. The sales, S_{it}., is related to the features of the products such as size and flavor, and marketing efforts, such as price and promotion. The variables of product features and marketing efforts are represented by the vector, X_{it}. The relationships can be defined by the model below.

$$S_{it} = f(X_{it};\ \beta_i)$$

$$\varepsilon_{it} \sim N(0, \sigma_i)$$

In this equation, β_i is a vector of coefficients representing the market's preference of the corresponding product features and marketing mix variables. The specification of the market response model captures managers' knowledge and expectation of market responses, which is the basis of managers' decision. By calculating the difference between the actual sales data and the expected sales, we are able to gauge any deviation from the existing pattern of market responses. When analyzing the sales data, this baseline model for the existing market condition is unknown to the researchers, and we propose to uncover and approximate the existing model by dynamically updating the baseline model using past sales records since time $t-1$, which are notated using the super scribe t-1 (e.x, β_i^{t-1}), and then generate the prediction for time t.

$$\hat{y}_{it} = f(X_{it};\ \beta_i^{t-1})$$

The advantage of this approach is that managers can start monitoring sales data with small amount of historic data. When the market is stable, incorporating new data can gradually improve the estimation of the baseline model.

The deviation from the model is defined as the difference between actual sales and predicted sales using the model estimated based on data before time t-1. We normalize the measure of the deviation by the standard deviation of the error term, $\hat{\sigma}^{t-1}$.

$$D_t = \frac{y_{it} - \hat{y}_{it}}{\hat{\sigma}^{t-1}}$$

Because managers are interested in any deviation from the original model, the test statistics for monitoring upward trend based on the framework of the Cusum control chart becomes:

$$R_t = max\{R_{t-1} + D_t , 0\}$$

The test statistics for monitoring downward trend becomes:

$$R_t = min\{R_{t-1} + D_t , 0\}$$

In a most basic sense, this control chart index can be considered as the sum of the deviation from the baseline model. When there is no change in the market, the noises in the random errors will cancel out and the index would fluctuate around zero; When there is a upward or a downward trend, the systematic deviation will be accumulated by the control chart index, allowing us to observe a more clear pattern. This characteristic of the control chart index also facilitates us to identify the time when the change begins. Once we detect a considerable deviation from the original model, we can use the control chart plot to identify the time point when the index starts to raise or descend. The event that causes the change in the market is very likely to take place around this point.

3.4 Simulation

To access the effectiveness of the proposed method, we conduct a simulation study. We mainly focus on two qualities. The first quality that we looked for in the proposed model is whether the method is able to identify changes in the market response due to introduction of new products. The second quality is whether the proposed method is able to identify the time when the change first take place in the market. Ideally, the time point identified by the model should be close to the time when the new products were introduced into the market.

Simulation of the Data

We conduct simulations of the sales data under three scenarios: (a) no change happens during the observation period, (b) a decrease in the intercept of the model takes place in the middle of the observation period; (c) a decrease in the price coefficient takes place in the middle of the observation period. For all three scenarios, we simulate a sequence of data of 100 periods and changes in the underlying model take place in period 50 for scenario (b) and (c). To conduct the simulation, we first generate the price and the promotion variables. Based on the different models and parameters specified under different scenarios, we generate the sales figures. The details of the models under different scenarios are presented in Figure 15. The panel A in each of the set presents the comparison between simulated sales figure and the prediction of the sales figure. The panel B in each of the set present the index calculated using the proposed method. Because we create scenarios in which changes in parameters cause downward trend in sales, we only calculate the control chart indexes for monitoring downward trend. The same results applied to the cases when upward trend is of interest.

Comparing the sales figures and the predicted sales figures in the three scenarios, we can see little differences in the fit of the prediction using our bare eyes. This is because random errors are a part of the model and we add the errors term in addition to the deterministic part comprising of price and promotion. The small increase in the gap between the predicted sales and the actual sales caused by changes in the parameters in scenario b and c, when mixed with the random errors, become difficult to distinguish.

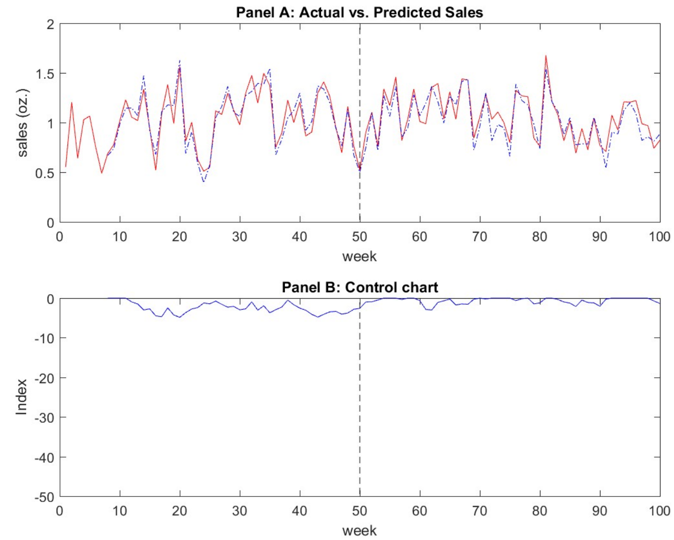

Throughout the observation period:

$$Sales = 1 - 0.5 \times Price + 0.6 \times Promotion + \varepsilon$$

(a) No Change

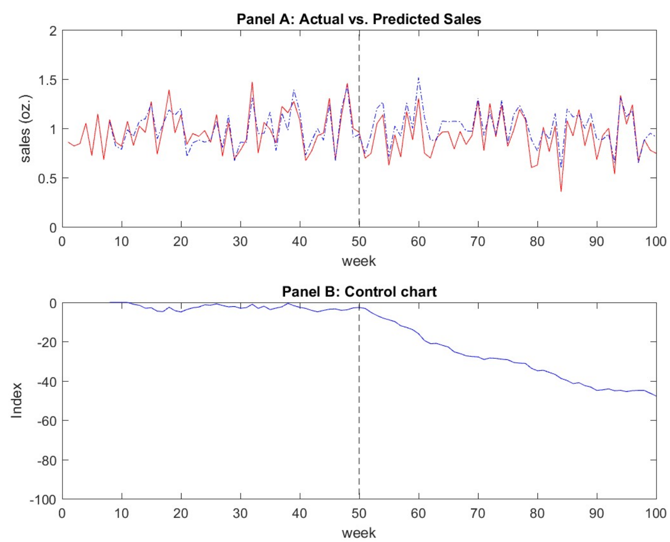

Before change: $Sales = 1 - 0.5 \times Price + 0.6 \times Promotion + \varepsilon$

After change: $Sales = 0.8 - 0.5 \times Price + 0.6 \times Promotion + \varepsilon$

(b) The Intercept Change at Period 50

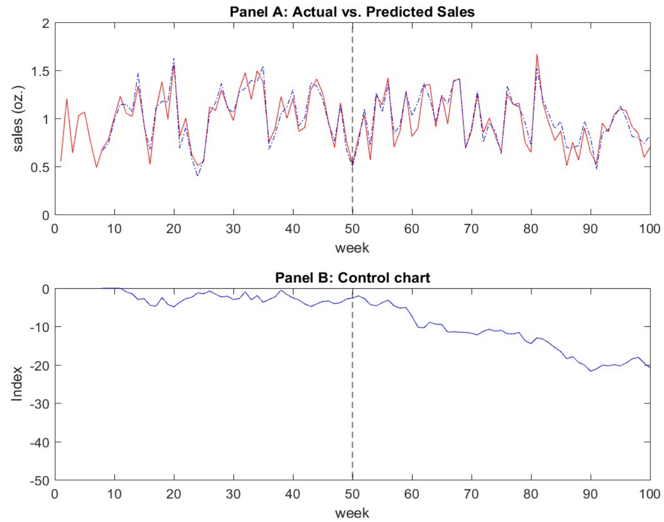

(c) The Response to Price Change at Period 50

Before change: $Sales = 1 - 0.5 \times Price + 0.6 \times Promotion + \varepsilon$
After change: $Sales = 1 - 0.7 \times Price + 0.6 \times Promotion + \varepsilon$

Figure 13: Monitoring Downward Trend in Sales

The control chart indexes in the three scenarios serve as clear indicators of changes. In scenario (a) when no change take place, the index fluctuates around zero (mean = -1.50; std = 1.44) and the minimum value of the index is -4.85. In contrast, when a change take place in the intercept of the model in scenario (b) at period 50, the index falls below -10 within seven periods and it continues to fall to -47.69 at period 100. Similarly, when a change takes place in the coefficient of the price variable, the index drops below -10 within ten periods and it continues to drop to -20.76 at period 100. The descriptive statistics of the indexes for the three scenarios are summarized in Table 5, which shows clear differences in the value of the control chart index before and after change.

An examination of the control charts in scenario (b) and (c) show that the control chart indexes start to descent around period 50. This is a very nice property of the control chart index. In cases where we set the threshold to set off the alert at -20 and we detect the change in 10 periods, the control chart plot allows us to trace back the time when the change first take place. This time point can be managers' first clue in identifying possible causes of the change, and managers can search for relevant events taking place around the time point and examine their impacts on the market.

Table 5: Descriptive Statistics of the Control Chart Index in the Three Scenarios

	Before Period 50			After Period 50		
	Mean	Std	Min	Mean	Std	Min
(a) No Change	-2.47	1.38	-4.85	-0.62	0.78	-2.98
(b) Change in the Intercept	-2.52	1.36	-4.85	-30.37	12.47	-47.69
(c) Change in the Price Coefficient	-2.50	1.36	-4.85	-12.94	5.52	-21.64

3.5 Empirical Study

We tested the effectiveness of the method in the context of a new product introduction, which has been shown to have significant impacts on the incumbents in the market. Because new product introduction usually increases the competition in the marketplace, it generally results in decrease in the sales of the incumbents' sales. Thus, we focus on detecting downward trends in the sales of the brands after the new production. Similar to the simulation study, we focus on the ability of the method to detect downward trends once then take place and facilitate us to trance the cause of the change back to the event of the new product introduction. We state the research questions formally as following.

(a) Is the proposed method able to identify the change in market's response due to the introduction of the new product?

(b) Is the proposed method able to identify the time when the change first take place in the market? Can the proposed method identify a time approximate to the time when the new products are introduced into the market?

Data Description

In this study, we use the Dominick's Finer Foods database to demonstrate the ability of the proposed model to identify systematic changes in the market place. This data set is made publically available by the Kilts Center for Marketing, University of Chicago. It spans 400 weeks from September of 1989 to May of 1997 and contains weekly observations of sales and shelf prices of the products sold by Dominick's Finer Foods in the Chicago area[4]. For the purpose of this study, we aggregate the data to the market level by brand and we obtain the market-level price of each brand by averaging the unit price of the brand across stores and UPCs. From the data set, we also obtain the entry dates of the new products or new product variants into the market, and we confirm these entries using news reports and articles extracted from Nexis and Factiva data base. We apply the proposed model on the toothpaste category. This category is chosen because they are characterized by intensive competition among brands, which often changes the market landscape. During the observation period, these categories has experienced one or several events of new product entries, including new brands or new product variants, which has been shown to have consistent impact on the market.

Toothpaste Category

---

[4] Although a nationwide data including sales from other supermarket and other types of grocery stores would be helpful in examining the changes brought by new product introduction, researchers in the past conclude that the data and sales pattern obtained from one chain or one region resembles that of the nationwide data (Sudhir 2001; Chintagunta, Dube, and Singh 2003). Because these change patterns are also shown in the sales of the Dominick's Finer Food chain, it is sufficient for our demonstration purpose.

We analyze seven major brands of the toothpaste category, which account for 78.45% of UPCs and 97% of sales in the category. Among them, Crest and Colgate are the leading brands, which account for 36.32% and 29.74% of the market share, respectively. The toothpaste category features increasing competition during the observation period. This intensified competition is evident in the fact that the number of UPCs carried by the Dominik's chain has increase from 150 in September 1989 to 248 in May 1997, representing a 65.33% increase in the number of UPCs.

The Introduction of Mentadent into the Toothpaste Category

While there are several new brand entries during this period[5], only one brand, Mentadent, has gained considerable amount of market share during the observation period. Mentadent is a brand of Unilever and was later sold to Church & Dwight in 2003. Mentadent toothpaste is positioned as a product that protects the gum rather than a cure for dental caries. It is known for its dual chamber pump, which is a uniquely designed toothpaste dispenser that store two ingredients of the toothpaste separated. For the majority of the Mentadent toothpaste products, these two ingredients are baking soda and peroxide. During the brushing process, the two ingredients are mixed and can react with each other. to create oxygen bubbles, which, claimed by Mentadent, has the ability to clean, whiten and refresh the mouth.

Mentadent entered the Chicago market in September 1993 (208th week of the observation period) with the baking soda and peroxide formulation. In the first year of its entry, Mentadent gained 3.89% of the market share. This number jump to 6.27% in the second year. Other brands entering the market during the observation period are Rembrand,

[5] Arm & Hammer entered the market with its baking soda formulation on January 1990 (19th week of the observation period). Because this entry happens at the beginning of our observation window, we do not monitor the impacts of this entry on other brands.

Sensodyne, Topol, Pluswhite, Oral-B, Rapid and Tom's. By the end of the observation period, none of these brands gain more than 1% of the market share.

Operationalization of the proposed model

We monitored the weekly sales volume in the toothpaste category. To calculate the sales volume, we sum up the total volume (in ounce) of the toothpaste sold by brand i at week t and conduct a log transformation of the overall sales figure. The price variable is computed by averaging the price across all UPCs among all Dominick's stores. For price promotion, Dominick's store offers promotion by either applying direct discount of the original price or giving away addition bonus points to its loyalty program.

$$Sales_{it} = \beta_0 + \beta_1 \times Price_{it} + \beta_2 \times Discount_{it} + \beta_3 \times BonusBuy_{it} + \varepsilon_{it}$$

In this paper, we try to keep the specification of the baseline model as simple as possible. This is because the focus of this paper is not to improve the prediction on the sales figure, but to quickly identify of any deviations from the baseline model that we estimated using past data.

Results

We focus on two qualities when assessing the effectiveness of the method. The first quality that we looked for in the proposed model is whether the method is able to identify changes in the market response due to introduction of new products. The second quality is whether the proposed method is able to identify the time when the change first take place in the market. Ideally, the time point identified by the model should close to the time when the new products were introduced into the market.

We monitored the performance of the toothpaste category and we presented the results in Figure 16. Figure 16 contains seven subsets of plots, each representing the result of a brand. Panel A in each subset presents the comparison of the actual and predicted sales. A review of these plots reveals that predicted sales approximate the actual sales and it is difficult to identify the pattern of the differences between the predicted sales and actual sales. These plots illustrate the importance of a tool that can help managers detect systematic changes in market responses. Panel B in each subset presents the result of the control chart index constructed using the method proposed in this paper, and it measures the downward deviation from the original model.

(a) Pepsodent

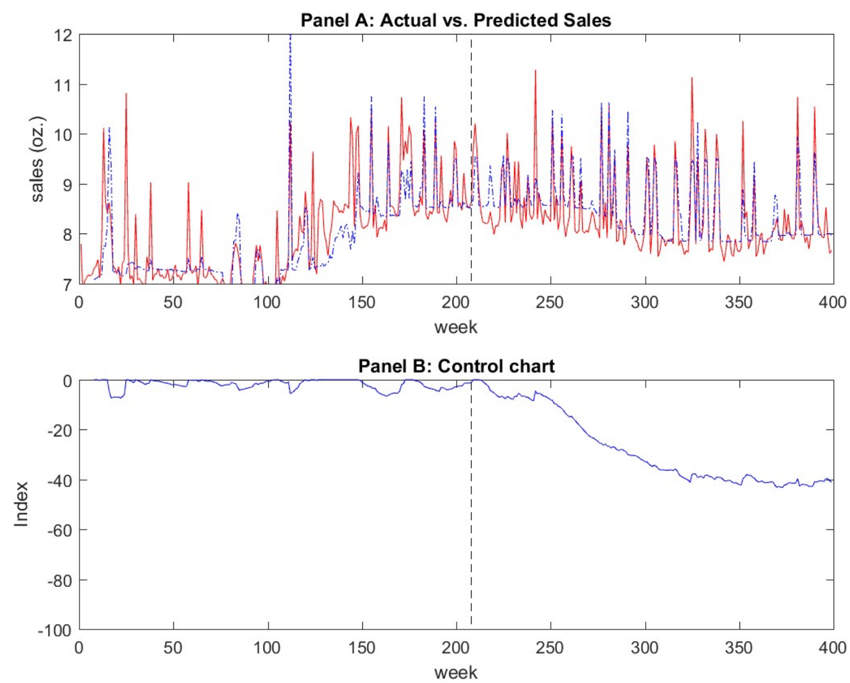

(b) AIM

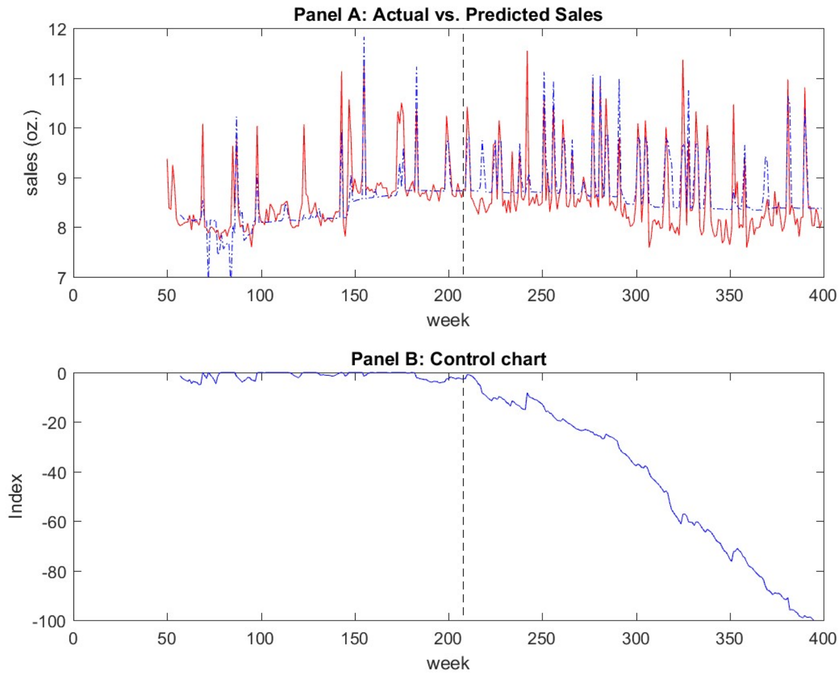

(c) Cologate

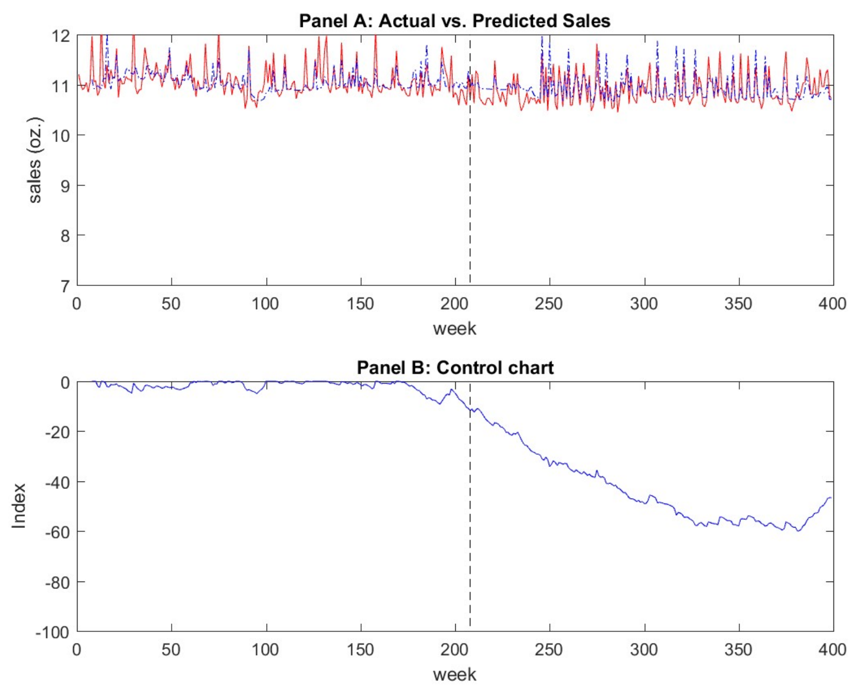

(d) Aqua

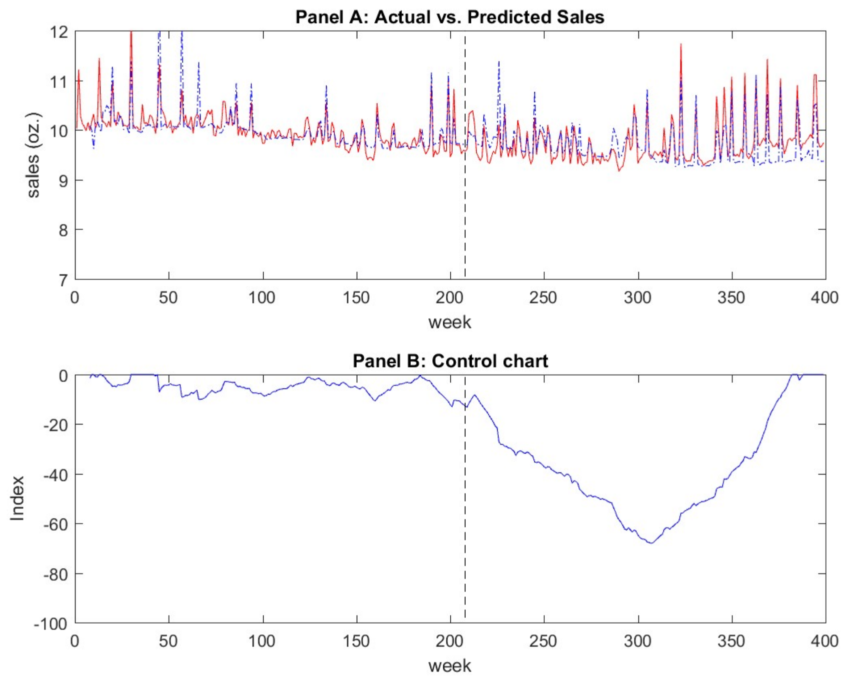

(e) Arm & Hammer

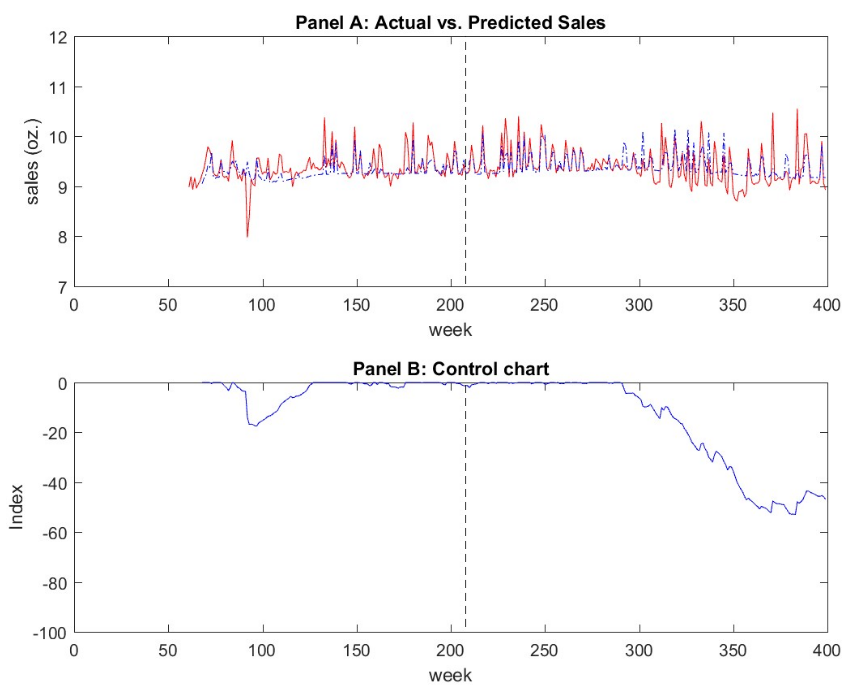

(f) Closeup

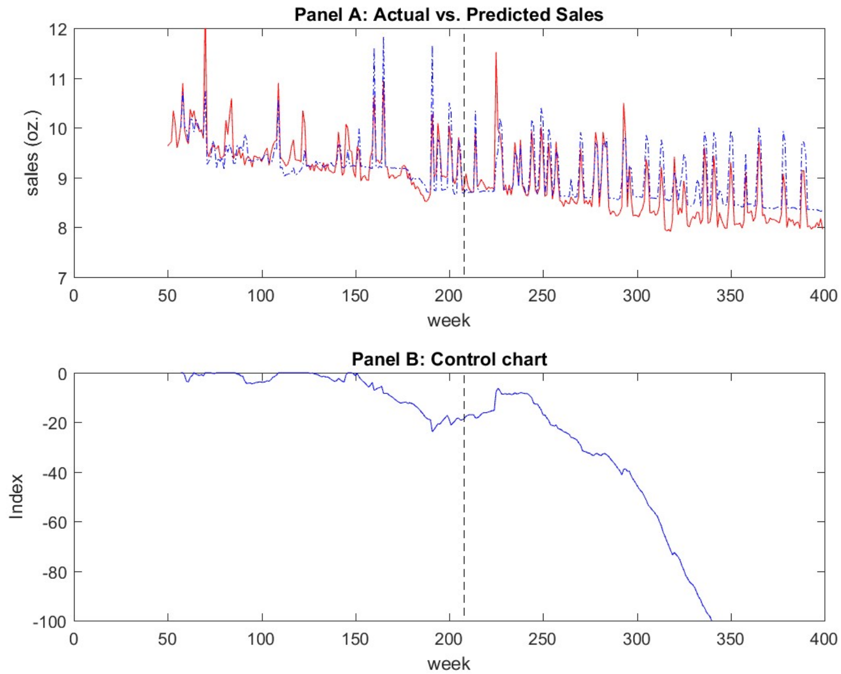

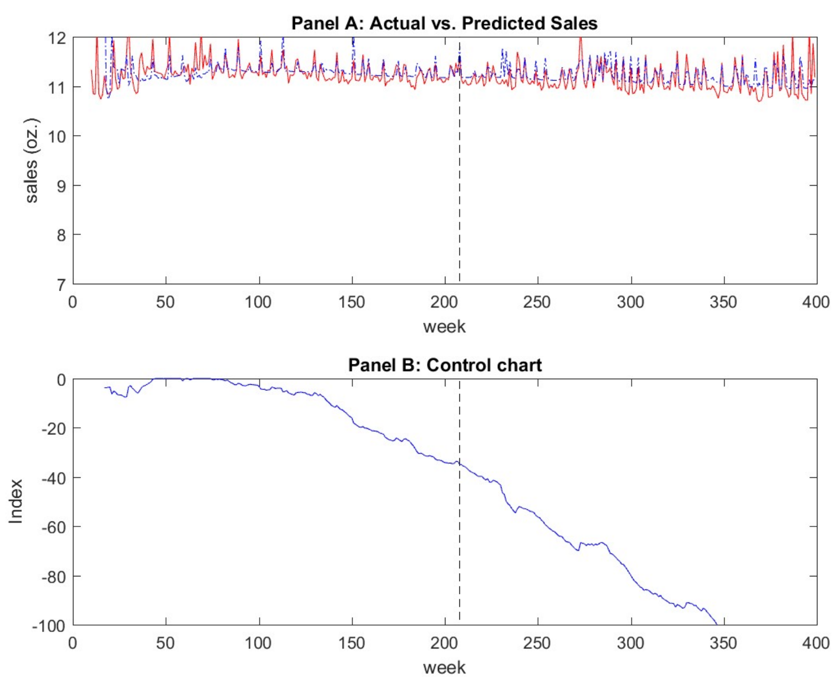

(g) Crest

Figure 14: Monitoring the Downward Trend in the Toothpaste Category

Using the control chart index, we can now clearly identify the systematic differences between the actual and the predicted sales, which is originally hard to detect. In these plots, when the market is stable, the chart index fluctuates around zero; when the actual sales is systematically lower than the predicted sales, the index becomes smaller and smaller. For example, Pepsodent's sales were systematically lower than the prediction after the entry of Mentadent. In Pepsodent's case, the mean of the index before the change is -1.83 (std = 1.95). After the entry of Mentadent, the index drop below -10 within 38 weeks, serving as a strong signal for market turbulence. Because the impact of Mentadent's entry persists over the years, this index continues to drop and reaches -46.83 at the end of the observation period (week 399, which is 191 periods after Mentadent's entry).

Based on the results from the control chart, we find that four brands show systematic downward trend after the entry of Mentadent, and these brands are Pepsodent, AIM, Cologate and Aqua (Figure 2(a) – (d)). We also find that the sales of Arm & Hammer remain stable after the entry of Mentadent (Figure 2(e)). We validate the results by conducting a regression analysis to confirm the regime change after the new product entry. The model is as follows.

$$
\begin{aligned}
Sales_{it} = {} & \beta_0 + \beta_1 \times Price_{it} + \beta_2 \times Discount_{it} + \beta_3 \times BonusBuy_{it} + \beta_4 \times Change \\
& + \beta_5 \times Change \times Price_{it} + \beta_7 \times Change \times Discount_{it} + \beta_8 \\
& \times Change \times BonusBuy_{it} + \varepsilon_{it}
\end{aligned}
$$

In this model, $Change$ is the indicator variable with "0" representing the periods before the entry of Mentadent and "1" representing the periods after the entry. The results are represented in Table 6. The result of this analysis confirms that the results from the control chart index. The coefficients of the models for Pepsodent, AIM, Cologate and Aqua are significantly different after the entry of Mentadent. In contrast, there is no significant difference for Arm & Hammer.

Table 6: Analyzing the Changes in the Sales after the Entry of Mentadent

	Pepsodent	**AIM**	**Cologate**	**Aqua**	**Arm & Hammer**
Intercept	9.83*	24.29*	13.46*	11.45*	8.62
	(0.21)	(1.65)	(1.84)	(0.39)	(1.32)
Price	-0.95*	-12.06*	-1.13	-0.73*	0.21
	(0.12)	(1.28)	(0.81)	(0.16)	(0.46)
Discount	0.78*	-2.00*	-0.04	0.52*	0.83
	(0.14)	(0.38)	(0.38)	(0.14)	(0.33)
Bonus	1.62*	-1.92*	1.76*	2.26*	0.95
Buy	(0.25)	(0.46)	(0.55)	(0.30)	(0.72)
Change	1.26*	-7.15*	-2.29	-0.02	1.34
	(0.50)	(1.74)	(1.90)	(0.51)	(1.37)
Change×	-1.13*	5.33*	0.92	-0.05	-0.42
Price	(0.35)	(1.35)	(0.83)	(0.21)	(0.48)
Change×	-0.32	1.53*	1.13*	0.05	-0.29
Discount	(0.23)	(0.40)	(0.44)	(0.16)	(0.34)
Change×	-0.54	2.02*	-0.12	-1.24*	0.43
BonusBuy	(0.38)	(0.50)	(0.60)	(0.34)	(0.74)

Note: The coefficients are shown in the first line and the standard deviation is shown in the bracket. The coefficients significantly different from zero are marked with "*".

The control chart index allows managers to monitor the market continuously and detect changes once they take place. At times, managers are not aware of all events that might cause changes in the market, and this is when the control chart becomes helpful. As we can observe in Figure 2(f) and (g), the sales of Closeup and Crest began to decrease long before the entry of Mentadent. This pattern of the control chart index indicates that other events might play a role in the decrease in the sales of Closeup and Crest. If managers only analyze the event of Mentadent's entry, they might overrate the impact of Mentadent and ignore other influential factors.

The panels Bs of each set of the plots also show us another great feature of the control chart index, which is its ability to allow us to trace back to the original time point when the change first take place. For example, if we set the threshold for alarm at -20, we can discover a systematic downward trend at week 258, which is 50 weeks after the occurrence of the event. Based on the control chart index, we would be able to trace back to the first time point when the systematic change starts to accumulate by looking for the latest time point when the index is reset to zero. In the case of Pepsodent, we discovered that the downward trend starts at period 216, which is very close to the time point of Mentadent's entry. The proximity of the two time points provides managers a solid ground to speculate Mentadent's entry as a likely cause of the downward trend in Pepsodent's sales. Table 7 listed the detected change points for the four brands that showed systematic downward trend after Mentadent's entry, and these detected change points are all close to Mentadent's entry.

Table 7: The Detected Change Points

Brand	**Detected Change Point**	**Differences from Pepsodent's entry**
Pepsodent	216	+8 periods
AIM	204	-4 periods
Cologate	163	-45 periods
Aqua	177	-31 periods

3.6 Conclusion

Given the turbulent and the dynamic market today, managers need tools to assist the early detection of changes in the market. In this paper, we propose a solution based on Cusum control chart to help managers detect trends in the sales data. In the proposed solution, we incorporate market response model into the control chart index and we solve the problem of unknown baseline model by dynamically estimating the baseline model as new data arrives. We also propose a new use and interpretation of the control chart. We point out that the time when the control chart index starts to descend or raise is the time when changes in the market take place. In this way we can narrow the range of the time of the event that causes the changes in the market.

In order to validate the effectiveness of the proposed method in sales management, we evaluate the performance of the proposed method in simulation experiments and using empirical evidence. To empirically assess the effectiveness of the proposed solution, we applied the method in the context of new product introduction, which has been shown to have significant impacts on the incumbents in the market. In both simulation experiment and empirical data, the proposed method has demonstrated its power in both detecting systematic changes from the baseline model and identifying the time when change first takes place.

The proposed solution is an efficient solution for detecting the market's deviation from existing market response model using a data-driven approach. Upon the discovery of the trend, the proposed method is limited in providing further information about the change. For example, the control chart index does not reveal how coefficients in the model are different from the baseline model. Further analysis is required to obtain a deeper understanding of market after change and how it differs from the previous model.

Chapter 4

Essay 3

Detection of Customers' Life Change Using Control Chart Approach

4.1 Introduction

Consumers are embracing loyalty programs with rising enthusiasm. Starting from 973 million in 2000, the number of loyalty memberships in US jumped to 1.8 billion in 2009 and further leapt to 3.3 billion in in 2015 (Colloquy, 2015); it is estimated that 75% of the shoppers in the US today have at least one loyalty card in their wallets (Berry, 2013). In another study of global respondents by AC Nielsen, nearly 60 percent of said that loyalty programs were available where they shopped, and of those, 84 percent said they were more likely to visit those retailers (Nielsen, 2013).

Accompanying consumers' rising appetite for loyalty programs is the intensified competition among the loyalty programs. This intensified competition manifests in the escalated efforts to acquire and retain the customers, using tactics such as new-customer bonus and one-shot deals layered over the loyalty program. What intensifies the competition landscape even more is the expansion of loyalty programs to various economic sectors such as finial services, retail, travel & Hospitality, entertainment, telecom and Internet. As a result, loyalty programs across industries are competing for the scare resource of customer attention. With limited time and energy, customers choose to only actively engage in a small number of loyalty programs. Studies showed that in 2013, while US households participate in an average of 21.9 loyalty programs, they stay active in only 9.5 of those programs (Berry, 2013). Thus, it becomes more and more challenging for loyalty programs to engage customers. In response to this challenge, companies look to the coalition loyalty programs as a solution.

A Coalition loyalty program is a loyalty card platform or system that allows customers to earn rewards from two or more merchants. A typical coalition loyalty program uses loyalty points as a common currency in the program and all eligible purchases are translated to points that can eventually be redeemed at partners. This mechanism provides the convenience to the customers in redeeming their rewards and thus has quickly gained popularity among the customers. Some of the world's leading coalition loyalty programs include Air Miles (70% penetration in Canada), Nectar (68% penetration in U.K.), FlyBuys (60% penetration in Australia) and Payback (60% penetration in Germany) (SLI, 2013).

In spite of its rising popularity among consumers, doubts exist on whether the merchants can benefit from the coalition loyalty programs. Merchants participating in the coalition loyalty program partner with other merchants and thus relinquish the sovereignty of their reward programs and share the platform with other merchants. The concerns of the coalition loyalty program arise from the comparison with the traditional one-vendor loyalty program. With in-house loyalty program, merchants can use rewards as a way to encourage repetitive purchase and cultivate customer loyalty. It is believed that this link between reward and patronage is diluted by the presence of other merchants in the program and the rewards largely foster the loyalty to the reward program, instead of the merchants (SLI, 2013). In this way, merchants participating in the coalition loyalty programs are constraint on the utilization of reward to cultivate loyalty. If the coalition program and its merchants focus solely on the economic value that they can return to the customers, the coalition program will soon be dragged into an indulgence of unprofitable promotions. To avoid this mistake, a successful coalition program, like other forms of alliances, should allow partners to make effective use of the resources that other partners bring to the program and collaborate with each other to expand the market (Das & Bing-Sheng, 2000; Sheth & Parvatiyar, 1992). To be more specific, the most valuable resource, which is also the focus of this study, is the customers brought by the merchants into the coalition program. For merchants to be successful in a coalition program, they should develop a unique proposition in the loyalty program, effectively target and cross-sell to customers of other partners and collaborate with each other to unleash the full potential of the joint customer base.

Thus, an objective of the study is to understand whether merchants can benefit from partnering in the coalition program. We examine this issue by comparing customers with different shopping portfolios: Do customers who cross-buy at multiple merchants yield higher value for the company? Furthermore, we are interested in comparing the benefits gained by two merchants when they are paired: for customers whom cross-buy in two stores, do they yield more value for the two companies respectively? Is the gain for one company higher than the other? Even more importantly: Are there situations in which one of the partners gains and the other partner losses? This is critical because the feasibility of partnership lies not only in the benefit of one partner but a win-win situation for both of the partners (Reichheld & Teal, 2001).

If there is value in partnering with other merchants, the following question arises: which partner should a merchant partner with? In another word, what types of cross-buying customers are valuable in terms of their shopping portfolio? Being selective in choosing the customers to serve is a critical issue. When a merchant has access to a whole pool of customer information in the coalition loyalty programs, it is a common and tempting mistake to market to all customers. However, untargeted cross selling is not beneficial. According to a study by Shah and Kumar (2012), one in five customers who cross buy are unprofitable and they account for the 70% of customer loss. Kumar (2012) pointed out that these unprofitable customers are those who are attracted to promotions, who have a strict limit of their spending or who demand more than other customers. In addition to these reasons, in a coalition loyalty program with a diversified customers base attracted by all walks of business, an unprofitable customer can simply caused by sending marketing contacts to uninterested customers. A way to identify the valuable partners and conduct effective marketing contacts is to mine the data to understand the underlying associations among merchants in their customer resources.

To verify the feasibility of partnerships within loyalty program and facilitate the selection of partners, we created an index called Partner Leveraged Value (PLV). What this index measures is the increment in a customer's lifetime value when the customer is also actively shopping at other partner stores. With this index, we also tested three hypotheses on the characteristics of merchants that make good partners, providing some rules of thumb to search for partners.

What is different in this study from other studies on loyalty programs and what this study contributes to the loyalty program literature is the emphasis on the concept of cooperation. The pervasive view in the existing literature is that loyalty programs are promotional tools to accelerate customer purchase and marketing tools to fend off the competitors. In a coalition loyalty program, merchants who share a same customer base, which often turns out to be merchants that compete with each other, can in fact work together to create benefits and rewards unique for their common customers. Through this collaboration, merchants participating in the loyalty program can deliver unique value proposition to attract customers who fall out of love with merchants who bombard them with advertising and promotional schemes.

An obvious benefit from this coalition is the sharing of the cost of operating a loyalty program. A sizeable amount of investment is required for managing a loyalty program. For example, it takes 500 employees for the Tesco's loyalty program, Clubcard, to manage customer communication, which includes customer call center, direct mailing and club magazine that issues four times a year. By participating in a coalition loyalty program, merchants with little expertise and resources in customer management are able to outsource the tasks such as maintaining database, call center and member website.

4.2 Background and Theory

Loyalty Program

A Coalition loyalty program is an extension of the traditional one-vendor loyalty programs. In general, a loyalty program is a tool of customer relationship management that aims to reward the loyalty customers to build a connection between these customers and the company. The marketing literature generally considers loyalty programs as long-term-oriented promotional tools that allow customers to accumulate their rewards through repetitive patronage and redeem the reward in the future (Byung-Do, Mengze, & Srinivasan, 2001; Lewis, 2004). The coalition loyalty programs inherit this delayed reward structure while adding a distinct feature of partnership: merchants participating in the coalition loyalty programs works with other merchants on the same platform – as both collaborators and competitors. From this view, the coalition loyalty program is a type of business alliance that joins the force of existing competitors and indirect competitors for the purpose of efficient use of customer resource and future growth opportunity (Sheth & Parvatiyar, 1992).

However, current literature has given little attention to this partnering side of the loyalty program. In fact, the majority of the studies on loyalty program today focus on the single-vendor programs because they have simpler structure and present themself as a well-defined problem to the researchers. As a result, the context of coalition has received limited attention. In the following section, we first review the existing literature on loyalty program.

Existing studies on loyalty program have extensive investigation on the factors that influence the effectiveness of the coalition loyalty program. Based on the framework proposed by Varadarajan and Jayachandran (1999), the outcome of marketing strategies in general depends on the interplay between the internal environment and the external environment. From the perspective a merchant participating in a loyalty program, the internal controllable environment is the merchants' capability in managing the program, and the external uncontrollable factors are the consumers and the competitors, which include competitors in the same industry and competing loyalty programs (Liu & Yang, 2009). Besides these two factors, merchants participating in the coalition loyalty programs are also influenced by intermediary factors, which includes the platform of the coalition loyalty program and the other partners on the same platform. Merchants have partial control of these two factors: they can actively communicate with other partners to learn their best practices in the program and join forces with them to create promotions. However, a merchant cannot dictate other partners' promotions or alter the design of the loyalty program at their wish. At the same time, the merchants in loyalty programs will experience the spillover effect of other partners' actions – either an increase in service satisfaction (Lemon & Wangenheim, 2009) or service failure (Schumann, Wünderlich, & Evanschitzky, 2014). We propose that merchants in a coalition loyalty program need to actively manage their relationships with their partners and consider the partners as useful resources in customer management.

In the following section of this chapter, we first review the existing studies that examine how loyalty programs are influenced by external factors, intermediary factors and internal management. From the review, we note the importance of collaboration in the loyalty program and the relative void of studies in this field.

Consumers

Existing studies have discovered two patterns in consumers' interactions with the loyalty programs that persist across the different programs. The first behavior pattern is forward-looking (Kopalle, Sun, Neslin, Sun, & Swaminathan, 2012; Lewis, 2004). The majority of the loyalty programs are designed as a delayed reward program in which customers shop and accumulate points to obtain a reward in the future. This future orientation is accounted using the dynamic programing method in which the expectations of future rewards are accounted when making a purchase decision. Lewis (2004) leads the research in this field and found an improvement in model performance when this future orientation is accounted. Closely related to the forward-looking behavior is the second and the most prominent pattern termed point pressure effect, which describes the acceleration in customers' purchase frequency when they are close to the reward (Kivetz, Urminsky, & Zheng, 2006). This phenomenon roots in a general human behavior called goal gradient effect, which is first discovered in the experiments of the animals. In general, when human and animals are approaching the goal, the prospects of achieving the goal promote more efforts towards the goal (Anderson, 1933). The classic experiment of this effect in marketing is the coffee shop experiment by Kivetz et al. (2006). The researchers set up the experiment in a coffee shop where consumers are rewarded for every ten cups of coffee that they purchased. The researchers found that when consumers are moving towards the tenth free coffee, the time lapsed between the two coffees gets shorter. This pulling effect of reward on purchase is discovered in loyalty programs across industries and with different types of structures (Drèze & Nunes, 2011; Kivetz et al., 2006; Zhang & Breugelmans, 2012).

In addition to these patterns, researchers are also interested in whether the loyalty program can have persistent effects in changing customer behavior in the long run. Beyond obtaining the reward, researchers found a partial reset effect when consumers start a new pursuit of the reward (Drèze & Nunes, 2011). This reset refers to the slowing down of purchase frequency after the customers redeem the points at the reward threshold and the customers are distant from the reward again. In this second pursuit, although consumers' purchase frequency is lower than that before consumers are about to achieve the goal, this frequency is still higher than that when consumers first start to collect the reward in the loyalty program. Thus, this reset is only partial, indicating the experience of accumulating and redeeming the points increase the baseline rate of consumption. In a longer term, loyalty programs are able to increase the purchase level of the light buyers (Liu, 2007) and eliminate cherry-picking (Lal & Bell, 2003).

Since the coalition loyalty programs maintain the basic structure of a traditional loyalty program, similar behavioral patterns are expected to be found in the coalition program and benefit the merchants. We note that these patterns also encourage consumers to conduct cross-buying across partners in the coalition network. The ample opportunities offered by the merchants allows consumers not only the convenience to accumulate their loyalty points but also faster progress towards the reward. Furthermore, when consumers obtain the reward in the form of vouchers, they have the flexibility to use the vouchers at a large number of merchants in network. In this way, the coalition loyalty programs make cross-selling/cross-buying a win-win strategy for both the merchants and consumers. This win-win situation is the base for a healthy merchant-customer relationship (Reichheld & Teal, 2001).

Competition

The competition landscape shapes the loyalty programs. Loyalty program is applied as a defensive tool to shield the company from the competitors. Studies showed that consumers participated in a point-based loyalty programs are less responsive to competitors' promotions (Zhang & Breugelmans, 2012). Vice versa, consumers participating in competitors' loyalty program showed decrease in the share of wallet as well as lifetime value at the focal company (Mägi, 2003; Meyer-Waarden, 2007). Because the threat to lose customers is high in a competitive market, it is more likely to observe company's adoption of loyalty program in such a setting (Leenheer & Bijmolt, 2008b). This in turn crowed the space of loyalty programs. When competitors are all equipped with the loyalty program, loyalty programs themselves as a promotional tool no longer serves as a point of differentiation. This reduction in the effectiveness of the loyalty programs due to market saturation is confirmed in a study of the airline industry (Liu & Yang, 2009). The researchers found that with the increase in the number of frequent-flier programs in which a consumer enrolls, the effectiveness of the loyalty program in increasing the frequency of flying with the companies decreases. This is a reason often cited by critics of the loyalty programs. The critics argue that the loyalty programs create debt to the companies without creating long-term sustainable competitive advantages (Dowling & Uncles, 1997).

Loyalty program is more than a tactical tool that trade reward with future patronage. In fact, it should be considered as a useful resource within the firm (Liu & Yang, 2009). While loyalty cards provide a more accurate way to track customer activities, this piece has its own value. Instead of the intangible connection with the firm, the card is a tangible presence. It transforms the abstract idea of loyalty to the company to a concrete action of holding the card, keeping it in the wallet and swiping the card for transactions. It also serves as the ambassador of the company with its omnipresence around the customer that reminds the customers of the company. Another example is the loyalty currencies. These currencies provide an alternative way of promotions other than direct price discounts which is proved to be more be more responsive and more cost efficient for the retailers (Zhang & Breugelmans, 2012). These currencies can also be combined with real currency in pricing to reduce customers' price sensitivity (Drèze & Nunes, 2004). More importantly, these currencies can be traded to other companies as an additional source of revenue as what airline companies do with their miles.

From this resource point of view, loyalty programs need to be complemented with other resources to fully realize its potential. For example, a loyalty program requires customer resource from which a company can draw loyal customers. Another example is product resource. A company with diversified product portfolio and widespread product distribution can do a better in relating their loyalty programs to consumers. In a traditional one-vendor loyalty program, these complementary resources can only come within the company. As a result, companies with higher market shares generally have more effective loyalty programs because the high-share companies are better equipped with the complementary resources. (Liu & Yang, 2009; Nako, 2004). In contrast, a coalition loyalty program allows merchants to utilize the resources from the partners and make up for their disadvantage in these complementary resources. For example, the coalition program pooled together customers from various companies, giving partners a large customer base to recruit new customers. The diversified product and availability of the service provide consumers the convenience to participate in the loyalty program. We propose that partners can go beyond this most basic form of the coalition. Utilizing the knowledge from the coalition loyalty program, merchants can identify the valuable customers through their shopping behavior at the merchants and together with the partners, merchants can create unique value proposition for their joint customers.

Program Design

Extensive studies have been conducted in this area to configure the best practice of program design for loyalty programs. The studies have focused on two key components of the loyalty programs, which are the point structure and rewards.

On the side of point structure, researchers found that the threshold to obtain the rewards and the design of tiers influence the attractiveness of a loyalty program to a consumer. While a low threshold is not challenging enough to motive consumers to increase their purchase frequencies (Kivetz et al., 2006), a high threshold can be difficult to reach and thus making the loyalty program unappealing for consumers (O'Brien & Jones, 1995). When deciding whether to join a loyalty program, consumers would evaluate their effort advantage related to the program requirement (Kivetz & Simonson, 2003). Thus, the threshold in a loyalty program can be used as a selection mechanism for what the company considers as a valuable customer.

On the side of the tiers, researchers found adding the design of tiers in the program can improve the performance of a loyalty program (Kopalle et al., 2012). While loyalty programs typically offer monetary rewards (e.x., cash back and free drinks) as frequency rewards, the tier benefits typically come in the form of soft benefits that improves customers’ experience (e.x., access to lounges in the airport, a personal shopper and invitations to special events). Researchers found that unlike frequency rewards that are only relevant to the price sensitive shoppers, these tier benefits attract both price-sensitive shoppers and service-oriented shoppers. The attractiveness of the tier benefit make tier another source of motivation for consumers to shop using the loyalty program, thus adding additional point pressure effect to the original frequency rewards structure (Kopalle et al., 2012). Moreover, the tiers can also create a psychological benefit of social status, allowing consumers with a high tier feel a sense of superiority (Drèze & Nunes, 2009). This benefit of social status is often independent of the exact benefit offered in a tier. Researchers found that for customers who are eligible for gold status, adding a silver tier beneath the gold tier increase the attractiveness of the program even when these gold-tier customers do not perceive a substantial difference between the benefits offered at a silver tier and gold tier (Drèze & Nunes, 2009).

The studies on rewards have emphasized the integration and synthesis between the program and the rewards. One key feature of the reward is its aspirational level. The match between the point structure and the rewards are important because different point structures can stimulate different preferences for the aspirational value of the reward (Kivetz & Simonson, 2002). Using experimental studies, researchers have found that customers prefer necessity rewards with low aspirational value when the efforts required to obtain the rewards are low and opportunities to receive the reward is high. In contrast, when the requirement of effort is high, consumers would prefer more luxury rewards with high aspirational value. In addition to the congruency between reward and point structure, the congruency between reward and the focal brand is also crucial. In general, the loyalty program is more effective when the reward is congruent with the brand. Meanwhile, companies need to consider other environmental factors and adjust the type of rewards offered according to consumers' preferences, such as involvement (Roehm, Pullins, & Roehm Jr, 2002), promotional reactance (Kivetz, 2005) and price sensitivity (Byung-Do et al., 2001).

The key takeaway from these past studies of loyalty program design is the importance of integration among different components of the loyalty program. Consumers care more than the monetary value of the reward. Consumers participating in the loyalty program also care about better service and the aspirational value of the reward. Different components of the loyalty programs - the point structure and the reward - need to orchestrate together to deliver a central value proposition. In the loyalty programs, the merchants have one additional tool to juggle: The partners.

Internal management

Program management makes a great difference in the success of a loyalty program. The effectiveness of a loyalty program depends on a company's technical capability to comprehend the data, extract customer insights and transform the knowledge into executable actions (Leenheer & Bijmolt, 2008a). Besides the modification in the design of loyalty programs, the actions that managers can take to influence the customers are in the form of one-shot promotions. Examples of the one-shot promotions are: "Spend $100 in July and get a $10 voucher" or "Triple points when shopping next week". Past studies found that one-shot promotions within the loyalty programs enhanced the program by increasing consumers' purchase probability and spending (Lewis, 2004). The powers of these one-shot promotions are enhanced by the ability of the loyalty program to capture consumers' past behaviors, which are good indications of future behavior, as has been proved in numerous studies. Before the invention of the loyalty program, companies who wish to conduct direct marketing need to rent lists of addresses and emails from list brokers and try to figure out the value of the customers on the list based on limited demographic information provided by the brokers, or coupling the list with commercial segmentation systems such as PRIZM (Verhoef et al., 2010). This method is still popular today even for companies with in-house loyalty program when companies are acquiring new customers. Merchants within a coalition loyalty program have the luxury to know the customers better before customers' patronage thus can customize the promotion based on a more complete view of customers' purchase behavior.

While promoting to the right customer at the right time can enhance the relationships with customers (Kumar, Venkatesan, & Reinartz, 2006), untargeted promotions can do a great harm to the companies beyond the waste of marketing budget (Shah & Kumar, 2012). Undisciplined promotions can break the routine habits formed in the past without building a more frequent consumption pattern, while stimulating consumers' sensitivity to price and promotions (Liu-Thompkins & Tam, 2013). For example, for customers routinely using a loyalty program credit card to pay for gas, a promotion such as "extra 10% off" can increase the customers' sensitivity of such promotions and lead them to anticipate such offers in the future and look for similar offers from competing programs. Furthermore, frequently receiving uninterested promotions can exhaust customers' interest in the loyalty program, fatigue customers for future promotions, and raise their resentment of the loyalty program, if not totally abandon the program.

Partners

In the previous section, we review of the existing literature on one-vendor program. We show that the effectiveness of the loyalty program is shaped by several factors, such as competition, consumer, program design and company's internal management effort. These factors are still influential in the context of the multi-vendor loyalty program. In addition, the partners within the same coalition become a crucial factor for the merchants participating in the coalition loyalty program. However, few studies have been conducted to investigate the context of the coalition loyalty program and little is known about the value and impacts of partnerships in loyalty program.

Through a search of the literature, we found three studies on this topic. The common theme of these studies is to investigate whether customers' interaction with other partners in the coalition loyalty program can have impact on the relationship between the customer and the focal merchant and the results are mixed. One study is conducted by Dorotic, Fok, Verhoef and Bijmolt (2011), investigating a coalition loyalty program of five vendors from different industries (Grocery, Electronics, DIY shop, Fuel and department store) using aggregate data at the merchant level. In this study, the researchers did not find significant impacts of joint promotions by multiple partners or spillover effect of partners' promotion. This result casts doubt on the effectiveness of joint-promotion in the coalition program. Another study by Schumann et al. (2014) investigated how the service failure by a partner influence consumers' loyalty of the coalition program. The researchers found that while a customer's loyalty to a partner in the coalition can contribute to the overall attractiveness of the coalition loyalty program, it can also lead to more damage to the program when a service failure of the partner take place. However, this study did not investigate the impact of the service failure on a specific merchant in the coalition. Both of these studies do not provide evidence for the benefits for a merchant to participate in a coalition loyalty program. Not only that the synergies among the partners are not found in the coalition, but also that merchants can suffer from possible damage due to partners' misbehavior.

The only evidence that showed the benefit of a partnership in loyalty program is a study conducted by Lemon & Wangenheim (2009). They investigated the partnership between a European airline company and companies providing serve in car rental, hotel booking and credit card. The researchers found that consumers perceive partners to have different degrees of compatibility between the additional services offered by the partners and the core service of the loyalty program. In their study, car rental and hotel booking are considered to have a stronger fit with the core service of airline than credit card service. Positive relationships can be observed between the usage of airline service and the two addition services with high level of fit. The more customers use the airline service, the more that the customers will use these two additional services in the future; In turn, the more customers book the car rental and hotel through the coalition loyalty program, the more the customers will use the airline service in the future. However, no significant relationships exist between the core airline service and credit card usage. This result indicates whether a merchant can benefit from the synergies of the coalition program depends on the fit between the focal merchants and the partners in the coalition program.

We point out here that there exist two types of coalition loyalty programs that are vastly different in how much the merchants have control over the design of the program. One type of the coalition loyalty program is a proprietary loyalty program that offers partnering opportunities to non-competing companies, such as the airline loyalty program studied by Lemon & Wangenheim (2009). These proprietary loyalty programs are inherently one-vendor programs because the companies that own the programs have the power to alter the design of the loyalty program. These programs by nature attract the customers who are loyal to the founding company of the loyalty program. Thus, for a merchant to join partnership in this type of coalition loyalty program, the major consideration is the fit and potential synergy with the core service provider. In this study, we focus on another type of coalition loyalty programs – the ones run by a third party. Examples of these coalition loyalty programs include Nectar in U.K. and Payback in Germany. These coalition programs are not structured around any core service or company. Instead, they attract partners from all types of categories. Thus, for the customers, the selling point of these programs is the ample opportunities to accumulate rewards and redeem the rewards.

This difference in structure brings in two differences in the operation of the loyalty program for the merchants. First, merchants face a more diversified customer base in these coalition programs because the customers are attracted by a wide array of partners. Second, merchants have more flexibility and options in creating joint promotions with other partners. These differences provide more flexibility in merchants' customer management, but it also poses challenges. One challenge is that the diversified partners and customer base make it hard for the merchants to identify whether there is a natural fit with the coalition. Another challenge is to identify the partners that provide more value to work close with.

4.3 Research Question and Hypotheses

The purpose of forming the alliance is for the partners to reinforce each other's effort in customer development and customer retention. Participating merchants are bounded by the membership responsibilities to share the cost of the loyalty program, the transaction information of the customers and the direct marketing channels. In return, the merchants can benefit from the potential synergies effect of the coalition program.

Marketing to the customers in a coalition program is similar to the promotion of cross buying. Cross buying refers to the behavior of buying several products from the same provider (Ngobo, 2004). In the context of the coalition program, cross buying refers to customers' patronage of several in-network partners within the coalition program. Cross buying is advertised as the key benefit of the coalition loyalty program because the large coalition network provides customers with ample opportunities to save and they can obtain the rewards more quickly if they choose to shop with more in-network merchants. Cross buying is considered as a reflection of the depth of relationships between customers and the companies. This is because customers with higher level of cross buying tend to have longer relationship with the company, purchase more frequently and pay higher margin. In this way, customers with higher level of cross buying have more value for the company. Transferring this knowledge to the coalition loyalty program, we know that customers with higher level of cross buying are more valuable for the program.

However, for the coalition loyalty program to succeed, the program and the merchants need to have congruency in their goals. That is, the increased engagement with the program should also prove to be beneficial for the merchants participating the coalition. In the traditional cross-buying problem, customers purchase multiple products from one company. Thus, for every one additional product and service that customers buy, the company generally makes more profit. The situation is different for the merchants in the coalition loyalty program. If a cross buying customer is only buying occasionally from the merchant, and the profit margin of the customer is low, this customer is not a loyalty customer for the merchant. The coalition loyalty program requires the merchant to pay the loyalty reward for the low value customers, because these customers can reach the reward threshold by shopping widely in the network. In contrast, this cost can be saved if the merchant choose to not participating in the loyalty program or set up its own loyalty program. If the merchant's in-house loyalty program set a threshold for customers to reach the reward, it would take a long time for the customer to get the reward or no reward at all. A coalition loyalty program that brings low value customers creates only debts to the merchants.

The practitioners are divided in their views of the relationship between customers' cross-buying behavior and the value to the merchant. In a recent article on Fast Company, Bran Pearson(2014), the CEO of LoyaltyOne, recognized that "Shared customers are unfaithful customers" is a common concern for the merchants when choosing to participate in the loyalty program. However, he also argued that when customers are cross buying in the coalition network, it creates network effect that "inspires a higher rate of cumulative purchasing, for increased spending and basket size, among the individual brands" and "builds and maintains top-of-mind awareness". Given this divided view of the coalition loyalty program, it is worthwhile to analyze the relationship between the degree of cross buying in the coalition program and the value of the customers to the merchants. Does the number of high value customers outweigh the number of low value customers? Does the coalition loyalty program attract a large mass of low value customers that warrant action to quit the coalition?

Research Question: Is Cross-Buying in Coalition Loyalty Program Beneficial for the Merchants?

In this study, we argue that merchants should take an active role in managing the collaborations with other partners and incorporate the information of customers' relationship with the coalition program and other partners in their customer management effort.

In this section, we develop hypotheses about the subsets of the merchants within the loyalty program that creates value when collaborate with each other. These hypotheses are based on the theoretical and empirical knowledge from studies in retailing and business alliance. Given the void in the analysis of partnerships within the loyalty program, we draw on the evidence from current research in consumers' shopping behavior that identifies the underlying association among retailers. These associations are the fundamental sources that create the synergy among the retailers when they work together.

Collaboration between Full-line Store and Specialty Store

Retailers in general can be grouped into three categories, general merchandisers, broad-line specialists and limited-line specialist using the criterion of consistency in the product line (Miller, Reardon, & McCorkle, 1999). A retailer has high consistency in the product line when its assortments are all closely related in end use. For example, LEGO stores have a high consistency in product line because its offerings are all focused on the constructional toys and their accessories. Based on this criterion of consistency, a limited-line specialist is defined as a retailer that provides highly consistent offers in a specific product category. The LEGO store mentioned above is an example of limited-line specialist. Comparing with a limited-line specialist, a broad-line specialist carries a wider range of offerings that satisfy more generic needs in the product category. An example from the children product category is the Toys "R" Us, which carries a broader line ranging from toys to children's clothing, including the LEGO construction toys. The third category, the general merchandiser, offers the broadest product lines with little consistency among its offerings. Mass merchandisers such as Wal-Mart, Kmart and Target all belong to this category.

Although competition exists among different types of retailers, studies in the retailing industry suggest that the relationships among retailers of different types are characterized as mutually beneficial. This mutual benefit is evident in the typical structure of shopping centers that we commonly observe today: The shopping center is anchored by a few full-line generalists, while a number of small specialty stores locates around them (Sheth & Sisodia, 2002). Studies have showed that being in an agglomeration formed by a collection of generalists and specialists has positive effect on retailers' traffic and profit.

One reason for this positive agglomeration effect is that the proximity of multiple retailers allows consumers to combine multiple shopping trips into one multi-purpose shopping trip, saving both time and energy in traveling among different shopping locations. For example, during a visit to the department store, consumers can also get a haircut while picking up toys and picture books for their children. In this way, the shopping center as a whole becomes an attractive shopping destination to the consumers, which in turn increase the traffic to each of the retailers.

Another reason is that the agglomeration also allows consumers to acquire rich product information and compare different products in one location. This benefit the consumers by reducing the uncertainty associated with the purchase, making it easier for them to make the purchase decision. In another word, consumers are more likely to make the purchase on the site, rather than postponing the purchase or making the purchase somewhere else.

The third and the most fundamental reason is the complementary in the offerings provided by the full-line store and the specialty store. Even with some overlap in the product offerings, the generalist and specialists have very different value positioning in terms of their collection of assortment. Generalists often offer a few standard versions of the products in the category, while specialists provide a wide range of selection within the category, usually with slightly higher price. In this way, the generalists and specialists complement each other in the product offerings, instead of a direct competition.

A coalition loyalty program with both a full-line retailer and category specialist can capture this synergy due to geographic proximity and the complementarity in the product category. Thus, we propose the following hypothesis.

H1: Partners who are geographically approximate to the anchored merchant benefits more from the coalition loyalty program.

Collaborations between Companies within the Same Category

Customers who are cross buying with companies in the same category are valuable customers. This is cross-buying behaviors is a possible indication that the customer has a higher and expandable demand in this category. It is demonstrated in studies using analytical models and empirical analysis that for customers with expandable demand in the category, loyalty program serves as an effective tool for firms to compete with alternatives from other categories that satisfy the same need (Kopalle & Neslin, 2003; Liu & Yang, 2009). In the context of the coalition loyalty program, the coalition allows the focal merchants to join force with merchants in the same category to compete with both offering from other industries and those in the same industry but provided by merchants outside of the coalition.

While competition exists within the coalition among merchants in the same category, there are several benefits to join the same coalition. Having several merchants in the same category allow the coalition loyalty program satisfy customers' need of variety seeking. Variety seeking refers to the phenomenon of consumers selecting different brands at different purchase occasions. In the retailer setting, this means that consumers choose to visit different retailers when there is need to shop. Variety seeking can stem from consumers need to obtain the optimal mix of information (Farquhar & Rao, 1976). For example, they visit different store to form a better knowledge of a category's price and alternative offerings. Variety seeking can be also caused by satiation (McAlister, 1982). Consumers can get bored with a retailer when they visit the same place over and over again. Variety seeking can also be a reflection of consumers' desire for change (Bawa, 1990). In this case, it is just a random decision to shop in a different place. Variety seeking is a crucial aspect of customer behavior that merchants should factor in their strategic planning. In cases where consumers have need of variety seeking, the customer can still be captured with the coalition loyalty program.

Another benefit is that when several merchants in a same category rally in the coalition program, a strong association is built between the category and the loyalty program. Therefore, when consumers with the loyalty program card decide to make a purchase in the category, they are more likely to use the loyalty card and choose a retailer in the category. When the habit of using the card to shop in the category is fostered, consumers are less likely to become active in other loyalty program network, shielding merchants from more fierce competition outside of the coalition. In this way, partners with similar product category form a "value net" to capture the customers.

While retailers in the same category are often in the head-to-head competition with each other, coalition loyalty program provides these merchants an opportunity to grow the customers together. Therefore, we propose the following hypothesis.

H2: Partners who with similar product/service offerings benefits more from the coalition loyalty program.

4.4 Data and Method

Data

Data from a European coalition loyalty program is used in this study. This coalition program offers customers the opportunity to accumulate loyalty points at more than 370 merchant partner stores that participated in the coalition program. When signed into the program, customers would receive a loyalty credit card with Visa function, which can be used like any other Visa credit card. When customers use this loyalty credit card in the coalition network, customers can obtain 1 or 0.5 loyalty points for every dollar that they spent. Once a customer accumulated 500 loyalty points, the loyalty program will issue the customer a voucher, which is valid for redemption at a number of partner stores for two years. Before September 2009, 500 loyalty points were equivalent to a 15-dollar voucher that can be used for 30% of purchase. After September 2009, the coalition loyalty program reduces the limit on voucher usage and allows vouchers to pay for 100% of the purchase. At the same time, the value of the loyalty points is devalued: 500 loyalty points are equivalent to 5-dollar vouchers. The merchants participating in the program are from all types of industries, including both full-line retailers such as department stores, broad-line specialist such as electronic retailers and limited-line specialist such as hair-salons and wine-stores. We do not disclose the name of the company and its partners because the sponsor of the dataset would like to remain anonymous.

The data set records all transaction activities of the selected customers when they use the loyalty program card and the window of observation covers the period from January 2000 to November 2011. The analysis of this study is based on consumers participating in the coalition loyalty program form January 2000 to December 2008. This means that we observe these consumers' shopping behavior since the first day that they participated in the program. Therefore, left censoring is not a problem in this study. The customers in the data set are selected based on a stratified sampling method: First, customers are grouped into different cohorts based on their year of registration in the program, with customers registered in the same year designated to the same cohort. Then, roughly one thousand customers are selected from each cohort. This type of cohort data is common for studies in customer management area. We provide a more detailed description of the sample by cohorts in Table 8. In the data set, customers signed in the program in the same year are grouped in one cohort and named by the year of registration. For example, customers signed into the program in year 2000 is called cohort 2000. This table shows that customers across different cohorts present similar card usage pattern in the frequency of card uses, monthly spending using the card and cross buying. We note that there is increase in card uses and monthly spending for cohort 2006, 2007 and 2008. This is because the loyalty program allows customers to use the card to also shop out-of-network like other Visa credit card. For these three cohorts, the in-network consumption frequency and spending level is comparable to the previous cohorts.

Table 8: Summary Statistics of Cohorts

Cohort	N	Average Number of Card Uses	Average Monthly Spending	Average Number of Merchants Visited
Cohort 2000	754	40.21	558.80	2.68
Cohort 2001	594	28.85	533.03	2.93
Cohort 2002	495	27.0	531.05	2.88
Cohort 2003	552	32.41	579.08	2.63
Cohort 2004	508	21.97	616.30	2.96
Cohort 2005	440	38.26	594.33	2.80
Cohort 2006	515	71.93	657.92	2.68
Cohort 2007	662	102.75	750.28	2.19
Cohort 2008	612	99.13	709.52	2.62

The data set contains information of 13010 customers. Among them, 4342 of customers (33.14% of the sample) have never used the card after signing into the program. We exclude these customers from our analysis because these customers have not been successful attracted to the loyalty program, leaving a sample of 8759 customers. Among the consumers who ever used the card, the number of times that consumers use the card ranges from 1 to 1381, with a median number of 22 uses. Given the customer has used the card in the month, the average monthly spending is 618.02 with standard deviation 977.29. The number of in-network merchants that consumers visited ranges from 1 to 42, with a media of 5 merchants.

On the side of the merchants, we observe customers making purchase in 144 partners in the data. Among them, we selected 16 merchants for the analysis in this study because they own sizeable customer bases that make them meaningful for other merchants participating in the coalition. One of the merchants, department store 101, has a special place in this coalition. Department store 101 is the oldest and most well-known department store in the city. It has created its own loyalty program, which was then merged into the coalition loyalty program, making department store 101 the biggest sponsor of the coalition loyalty program. Thus, our study focus on the customers of department store 101 within the coalition loyalty program.

Customer Lifetime Value

Customer lifetime value (CLV) is a metric that evaluates the profitability of a customer throughout history of the customer's transaction with the company. This metric is widely adopted in both academia and practice. In the industry, companies such as IBM, Harrah's, Capital One, LL Bean and ING have applied CLV as a tool to evaluate and manage their customer management efforts. Financial institutes also apply this method to evaluate the market value of the customer companies, whose most valuable assets are the large customer bases. In the academic field, researchers confirmed that CLV as a tool is more effective than other customer metrics in selecting the valuable customers and allocating marketing resources. Numerous studies have been conducted using CLV as the most important customer level outcome.

Gupta et al. (2006) summarized three reasons for the popularity of the CLV metric. First, CLV meets the needs of the practitioners to identify the return of marketing actions such as promotion and advertising. Using CLV, the positive impacts of marketing actions purchase incident and retention can be quantified by monetary value, which is better than other abstract constructs such as awareness and attitude. Second, CLV is a disaggregate customer level metrics that is proved to have great diagnostic value in allocating marketing resources. Using CLV, marketers can effectively segment customers by their profitability and identify the most valuable customers. Marketers can also evaluate marketing efforts at individual level using CLV, making it possible to customize marketing policy at individual level.

In this study, we use CLV as the criteria to evaluate the value of a customer for a merchant participating in the coalition loyalty program. In particular, we follow Reinartz & Kumar's (2000) approach, which limit the calculation of CLV to the net present value of cash flow provided by customers over the first 36-month period since they first shop at the merchant. Constraining the calculation within a 36-month period allow us to equally judge the value of the customers with different length of observation period due to the different timing of their participation in the program. The 36-month window is a reasonable time frame because a large proportion of customer lifetime value is captured within three-year period. In the example given by Kumar, Venkatesan, Bohling, & Beckmann (2008), with a retention rate of 75% and an annual discount rate of 20%, 86% of CLV is captured within three years' time. Furthermore, customers' behavior is usually more versatile when they first have contact with a new product or service (Fader, Hardie, & Chun-Yao, 2004). The first 36-month period can capture most of the variances among the consumers in terms of their usage behavior of the coalition loyalty program.

Therefore, we calculate the net present value of customer i for merchant j during the first 36 months starting from the month when customer i made the first purchase at merchant j.

$$CLV_{ij} = \sum_{t=1}^{36} \frac{I(shop\ at\ merchant\ j_{it}) \cdot GC_{it}}{(1+r)^t} - \frac{MT_{it}\ \cdot \overline{MC}}{(1+r)^t}$$

$I(shop\ at\ merchant\ j_{it})$ = 1 if customers have shopped at merchant j at time t

= 0 if customers have not shopped at merchant j at time t

GC_{it} = The gross contribution by customer i at time t, which is computed as amount of spending by customer i at time t × .1, assuming the profit margin is 10%

MT_{it} = Number of marketing contacts delivered to customer i at time t

$\overline{MC}$ = Average cost of a single marketing contact for merchant i. The cost is assumed to be $0.1 in this case

r = Monthly discount factor, this is assumed to be 0.0125 in this case, which is equivalent to 15% annual rate

In this calculation, we make assumptions on discount rate, profit margin and the cost of marketing contacts. The monthly discount rate is set at 0.0125, which is same as the discount rate set in Reinartz & Kumar's (2000) study. The gross contribution of customer i at time t is calculated as the multiplication of customer i's spending at merchant j and average profit margin. In this study, we set the profit margin of monthly revenue to be 10% based on the annual report of the merchants participated in the coalition program. This is a conservative figure compared to the 30% rate assumed in other studies of retailer. The cost component is calculated by multiplying the number of marketing contacts received by consumer i with the average cost of each marketing that is responsible by merchant j, which is assumed to be 0.1 dollar. Because the direct marketing materials are sent to the consumers with monthly credit card bill or through email, we expect a low cost of these marketing contacts. The cost is even lower for each merchant because all merchants participating in the program share the marketing cost. We note that the one-time acquisition cost for each customer is not included because the information is not available. If this information is available, then it can be easily incorporated in the analysis by subtracting the acquisition cost from the current function.

Steps of Analysis

To understand the nature of cross buying, we examine the following questions step by step.

1. Is there positive correlation between Consumers' Cross-Buying Behavior and CLV at focal merchant's store?

Before a deeper engagement in the coalition loyalty program, it is crucial for the focal merchant to understand the nature of the relationship between their business and the coalition program: What is the relationship between the customers' engagement in the coalition program and the value of the customer at a specific merchant? In this study, we measure the customers' engagement with the coalition program by the number of in-network merchants that the customers have visited. This behavior of customers shopping at multiple merchants within the network is also called as customers' "cross-buying" within the network. The more merchants visited by the customers, the more engaged the customer is with the loyalty program. If the merchant can benefit from the network effect from the coalition loyalty program, then we should observe that with the increase in the number of in-network merchants visited by the customer, the value of the customer to the merchant would also increase. To explore the existence of this positive relationship between the customers' cross-buying behavior and the value of the customer for a merchant, we first calculated a Pearson correlation between the number of merchants visited by the customers and the CLV of the customers. A positive correlation indicates positive synergy between the program and the focal merchant, whereas a negative correlation indicates a lack of synergy between the program and the focal merchant.

To gain a better understanding of the relationship between the cross-buying behavior and CLV of the customers, we then plot a 3 by 3 cross table between the customers' CLV and the cross buying behavior. We categorized the customers on the CLV dimension into three groups: (a) a low CLV group containing customers with CLV in the lower 50 percentile; (b) a high CLV group containing customers with CLV between the 50 and 75 percentile; (c) a super CLV group containing customers with CLV that is higher than the 75 percentile. We categorized the customers on the cross-buying dimension using the same method.

Profiling the customers using the cross-tab gives the managers a better understanding of the integration between the merchant and the coalition. If a positive relationship exists between cross-buying behavior and the CLV of the customers to the merchant, then the we shall observe a concentration of the sample in the diagonal cells of the cross table. We are also interested in the off-diagonal cells of the table. Customers categorized into the cells in the lower right cells are the customers with low lifetime value but conduct a lot of cross-buying activities. These customers are the ones that are engaged in the program but do not generate much profit for the merchants. Having a large proportion of customers in this cell is a bad sign for the merchant: while the merchant contributes to the coalition by providing customers the convenience to shop and cumulate loyalty points in one additional place, the merchant fails to benefit from the network effect from these customers. Customers categorized into the upper left cells are the customers with high lifetime value but low cross-buying behavior. These are the customers who patronage the merchants but they do not care about the benefits of the coalition program. For these customers, the merchants' participation in the coalition loyalty program has no or little benefits for them. Having sizeable customers in these off-diagonal cells indicates lack of synthesis between the merchants and the coalition loyalty program, which can be caused by either a lack of fit between the coalitions or poor management.

2. Which partners create positive synergy for the focal merchant through consumers' cross buying?

Not all partnerships work. To fully utilize the resources of the coalition program, it is important for the focal store to decompose the positive synergy at the program level to the partner level and identify those partners that contribute this positive synergy. Knowing these partners allow the focal merchant to create more efficient joint promotions that leverage the value of the customers.

Calculate PLV: In this step, we can use the CLV as a diagnostic metric to evaluate the value of the partners. More specifically, we compare the CLV of the customers who only shop at the focal merchant and the CLV of the customers who also actively shop at partner k. The difference in the CLV between the two groups of customers represents the value of a partner to the merchant. We conduct this analysis only on the customers of the merchant j, that is, customers who ever shop in merchant j. The ordinal least square model is applied in the analysis of the value of partners to a merchant. The model is as following.

$$\mathrm{CLV}_{ij} = \alpha + \beta \cdot Partners_i + \varepsilon_i$$

In this model, the dependent variable is CLV_{ij}, which is the net present value of customer i for merchant j. The explanatory variables are the customers' relationship status with other merchants' in the coalition program during the 36 months period, which is $Partners_i$ in the above model. $Partners_i$ is a 15×1 vector. Each element in this vector is a dummy variable, with 1 representing that customer i is an active customer at partner k's store. That is, the customer needs to make at least one purchase during the three. The intercept, α, represents the average CLV of a customer. The coefficients, β, represent the amount of CLV leveraged by the corresponding partner. Thus, this coefficient is the partner-leveraged value (PLV). We conduct these three steps of analysis for the 16 merchants in the coalition program.

3. Test of hypothesis

To test the hypothesis, we identify several pairs of merchants that satisfy the descriptions in the hypothesis. To test hypothesis 1, we identify a small set of merchants in the data set that locate in the same shopping center. This set of merchants contains a department store (a generalist), a restaurant (specialist), a wine store (specialist) and a hair salon (specialist). To test hypothesis 2, we identify two other department stores: department store 507 and department store 525. We are interested in two results for a specific merchant: (a) As the focal merchant, does other partners leverage CLV in my customer base? (b) As a partner to other merchant, do I leverage the CLV is my partner' customer base? If the answers to both of these questions are yes, then it makes a good partnership within the coalition loyalty program.

4.5 Empirical Findings

Part I: What is the relationship between customers' engagement with the coalition and their profitability for a specific merchant?

Analysis of Department Store 101

To test the strength of the relationship between customers' engagement in the coalition loyalty program and their lifetime value for department store 101, we calculate the bivariate Pearson correlation between the number of merchants visited by the customer and the lifetime value of the customers. The correlation coefficient r is 0.188 ($p<.001$), which means that moderate linear association exists between customers' degree of cross buying and their profitability at the department store 101. This positive association indicates that customers' engagement with the coalition loyalty program helps to improve their lifetime value at department store 101. This moderate positive association is further confirmed when customers are categorized into nine segments by their cross buying behavior and lifetime value to the department store. We find that 563 out of 1824 customers fall in the diagonal cells, which accounts for 30.86% of the total population. The results are summarized in Table 9.

Table 9: Customer Segments by CLV and Cross Buying Behavior of Department Store 101

	No Cross Buying			Low Cross Buying			High Cross Buying		
	N	CLV	Cross-buying	N	CLV	Cross-buying	N	CLV	Cross-buying
	Segment 1			Segment 4			Segment 7		
Low CLV	74	9.45	0	286	10.11	2.96	552	10.83	8.47
	Segment 2			Segment 5			Segment 8		
High CLV	22	40.11	0	134	39.20	3.01	299	40.41	8.58
	Segment 3			Segment 6			Segment 9		
Super CLV	8	123.59	0	94	109.63	3.13	355	168.66	9.07

Beyond the linear association between cross buying and profitability, we are also interested in the contribution of each segment to the merchant's overall profitability. We discovered several remarkable patterns in Table 9. The first finding is that the customers who are actively engaged with the coalition loyalty program constitutes the core of the business for department store 101. This is evident in value of customers in segment 9. While segment 9 only contained 19.45% of the customers who visited the store in this sample, it contributes 60.50% of the customer equity to the store. Customers in segment 9 also have the highest average net-present lifetime profit per customer ($168.66). This value is significantly higher than that of segment 7 and 8, which are the other two segments that also contain customers from the top 25% most valuable customers of the sample ($p<0.001$). In contrast, for customers in high CLV group and low CLV group, the increase in the level of cross buying do not accompany a significant increase in customers' value ($p= 0.293$ for High CLV group; $p = 0.0571$ for Low CLV group). This result reveals that the benefit of cross buying is most useful in leveraging the high value customers for department store 101, but not for less valuable customers. In addition, for customers with high level of cross buying, the level of cross buying is significantly higher for customers in the super CLV group ($p=0.0123$). This indicates that these most valuable customers for department store 101 also more capable in utilizing the benefits offered by the coalition loyalty program.

The second finding is that a considerable size of customers is accumulated in segment 7. These customers shop wildly within the coalition loyalty program and have actively engaged with the loyalty program. However, these customers generate very little profit for the company. While this segment contains 30.26% of the customers visited the store, it only contributes 6.04% of the total customer equity to the store. These customers can be loyal customers to other partners in the coalition program and enjoy accumulating extra points by shopping at department store101. Another possibility is that these customers are low value customers for all merchants. The customers are attracted by the money return provided by the coalition program and use it as a way to save money. For either case, to department store 101, these customers compose a burden brought by the coalition loyalty program. These customers can consume the resources of the merchants, leaving less resource for the high value customers. For example, the flock of low value customers attracted to the department store can make the space crowded and occupy the staffs' time in serving them.

The third finding is that the customers in the off-diagonal cells do not equally distributed across the cells. The customers are concentrated on the upper left cells (segment 4, 7 and 8) and very few of them are in the lower left cells (segment 2, 3 and 6). The number of customers in segment 4, 7 and 8 is 9.16 times that in segment 2, 3 and 6. Furthermore, Customers with no cross buying accounts for only 1.75% of the super CLV customers and 4.83% of the high CLV customers; customers with low level of buying accounts for 20.56% of the super CLV customers and 29.45% of the high CLV customers. In all, the results reveal that actively cross buying within the coalition network is a necessary characteristic of the valuable customers for the department store 101. Although high level of cross buying does not guarantee a high value customer, a low level of cross buying is generally associated with low lifetime value for the department store 101. Thus, encouraging cross buying is beneficial for department store 101.

Analysis of Other Partners

For a coalition to thrive, it requires that the majority of the partners in the coalition benefit from the coalition network, instead of a few anchoring firms. Thus, we also examine whether the partners sharing the same coalition loyalty program with department store 101 can also benefit from consumers' increasing engagement with the program. That is, whether there is positive relationship between the number of in-network partners the customer visited and the lifetime value of the customers. Table 10 shows the Pearson correlation between customers' cross buying and lifetime value for the rest of the 15 merchants.

From this result, we find mixed relationship between customers' engagement with the coalition program and customers' lifetime value for the merchants participating in the coalition program. Only two of the merchants, department store 507 and supermarket 994, have positive associations between the cross buying behavior and their lifetime value at the merchant. The Pearson's correlation is 0.08 for department store 507 and 0.091 for supermarket 994. These associations are much weaker than that for department store 101. Eight merchants observe no significant association. This result shows that partners within the same coalition benefit in different degree from customers' engagement with the coalition program and the amount of benefits that the partners obtain depends on the position of the partner in the coalition.

Table 10: Correlation between cross buying behavior and CLV

	Correlation	p-value
Department store 507	0.080	0.005
Supermarket 525	-0.053	0.270
Electronic Retailer 106	-0.131	<.001
Supermarket 994	0.091	0.028
Sportswear 836	-0.084	0.003
Sportswear 828	0.040	0.161
Sportswear 835	-0.080	0.000
Clothing store 440	0.005	0.895
Clothing store 855	0.074	0.052
Wine Store 840	-0.109	0.056
Restaurant 108	0.007	0.882
Travel Agency 105	-0.088	0.016
Hair Salon 816	0.008	0.901
Auto Maintenance 431	-0.086	0.021
Gas Station 430	0.020	0.704

Part II: Who are the Good Partners?

We then go beyond the number of in-network partner stores that a customer visited and evaluate the relationship between the specific partners that a customer visited and the customer's lifetime value. Since the previous study revealed that department store 101 is the one partner that benefit from the synergies of cross buying, we will focus on the analysis of department store 101. To achieve this goal, we calculated the slopes and significance of the coefficients from the regression analysis of customer lifetime value as a function of customers' purchasing status in other partners' stores. The result is presented in Table 11.

Table 11: Analysis of Partners of Department Store 101

Variable	Estimate		t value
Intercept	37.79	*	11.93
Department store 507	13.35	*	2.97
Department store 525	55.15	*	8.08
Electronic Retailer 106	20.52	*	4.84
Supermarket 994	-6.06		-0.77
Sportswear 836	-19.47	*	-2.96
Sportswear 828	-16.82	*	-2.89
Sportswear 835	-13.76	*	-3.02
Clothing store 440	-15.37		-1.54
Clothing store 855	23.39	*	3.98
Wine Store 840	50.69	*	6.20
Restaurant 108	-8.10		-1.21
Travel Agency 105	9.74		1.43
Hair Salon 816	45.94	*	5.85
Auto Maintenance 431	-10.62		-1.37
Gas Station 430	-6.97		-0.63
	Adjusted R^2		0.13

We found that customers' status in other partners' store explain a considerable proportion of the variance of CLV for department store 101 (R-square = 0.1). This result is consistent with the previous analysis using the crosstab and correlation. It also indicates that customers' cross buying behavior is an influential factor for CLV of department store 101.

We termed the slopes of the coefficients as PLV. Using PLV, we can compare the contribution of the partners to the CLV at department store 101. Based on the results from the regression, we can categorize the merchants into three groups. The First group contains the high value partners. The partners in this group are: department store 507 (PLV = 13.35, p=0.003), supermarket 525 (PLV = 55.15, p<0.001), electronic Retailer 106 (PLV = 20.52, p

<0.001), Wine Store 840 (PLV= 50.68, p<0.001) and Hair Salon 816 (PLV = 45.94, p<0.001). Customers who were also visiting these partners have a significantly higher CLV at department store 101 than those customers who were not. This shows that these high value partners help to attract customers with high CLV to department store 101. The second group contains the low value partners. The partners in this group are sportswear 836 (PLV= -19.47, p=0.003), sportswear 828 (PLV = -16.82, p = 0.004) and sportswear 835 (PLV = -13.76, p=0.003). Customers who were also visiting these partners tend to have a significantly lower CLV at department store 101 than those were not. This shows that these low value partners attract customers who are not interested in frequently shopping in department store 101 and have little value to department store 101.The third group is the ones that do not have significant impact on the CLV at department store 101. The partners in this group are supermarket 994, Clothing store 440, Restaurant 108, Travel Agency 105, Auto Maintenance 431 and Gas Station 430. This result shows that partners vary in their ability to attract high value customers for department store101.

Based on these results, we specifically examine the partners related to the two hypotheses. The results show that customer visiting clothing store 855, wine store 840 and hair salon 816 have significantly higher value than customers who did not, suggesting that customers who also shop at the specialists at the same location have significantly higher lifetime value for department store 101 (See Table 12), thus providing support for hypothesis 1. We found that customers who visited department store 507 and department store 525 also tend to have significantly higher value than customers who did not, suggesting that customers who also shop at other department stores have significantly higher lifetime value for department store 101, thus providing support for hypothesis 2.

Table 12: Value of Partners in the Same Shopping Mall

Partner	focal Store				
	Department store 101	Electronic Retailer 106	Clothing store 855	Wine Store 840	Hair Salon 816
Department store 101	--	-18.16 (-3.37)	2.15 (1.09)	-15.93 (-1.56)	-3.84 (-0.74)
Electronic Retailer 106	18.26* (4.21)	--	3.22 (1.69)	18.71 (1.89)	-2.34 (-0.58)
Clothing store 855	18.90* (3.66)	-17.27 (-2.32)	--	-5.19 (-0.46)	1.02 (0.27)
Wine Store 840	38.21* (5.32)	7.33 (0.75)	0.51 (0.17)	--	0.97 (0.09)
Hair Salon 816	37.74* (5.14)	-3.99 (-0.42)	3.39 (1.28)	1.25 (0.08)	--

To further illustrate this point, we compare segment 7 (customers with high level of cross buying but low CLV) and segment 9 (customers with high level of cross buying and high CLV) on the percentage of customers being active members in the partner stores. The results are presented in Table 13.

Table 13: Comparing Segment 7 and Segment 9

Probability of Shopping at Different Partners		
	Segment 7 (N=552) High cross buying, Low CLV	Segment 9 (N=355) High cross buying, high CLV
Department store 507	0.27	0.44
Department store 525	0.09	0.24
Electronic Retailer 106	0.53	0.65
Supermarket 994	0.12	0.08
Sportswear 836	0.23	0.08
Sportswear 828	0.28	0.14
Sportswear 835	0.53	0.34
Clothing store 440	0.10	0.04
Clothing store 855	0.16	0.23
Wine Store 840	0.08	0.12
Restaurant 108	0.13	0.12
Travel Agency 105	0.11	0.15
Hair Salon 816	0.04	0.17
Auto Maintenance 431	0.14	0.10
Gas Station 430	0.05	0.05

We found that the percentage of customers being member in the high value partner stores in segment 9 is all higher than that in segment 7; On the other hand, the percentage of customers shopping in low value partners in segment 9 is all lower than that in segment 7. This result shows that a distinct difference exists between these two segments of customers in the partners that they visited. It indicates that for department 101 to benefit from customers' engagement with the coalition, it is not just the level of engagement that matters, but also the partners that attract the customers to the network.

Table 14: Value of Department Store 101 to Other Partners in the Coalition Loyalty Program

Variable	Estimate		t value
Department store 507	4.89		1.58
Department store 525	-1.98		-1.29
Electronic Retailer 106	-18.16	*	-3.37
Supermarket 994	4.76		1.60
Sportswear 836	-0.57		-0.39
Sportswear 828	-0.18		-0.17
Sportswear 835	-2.27		-1.05
Clothing store 440	-2.54		-1.28
Clothing store 855	2.15		1.09
Wine Store 840	-15.93		-1.56
Restaurant 108	-1.15		-0.58
Travel Agency 105	18.17		0.57
Hair Salon 816	-3.84		-0.74
Auto Maintenance 431	-3.43		-0.53
Gas Station 430	-6.79		-0.75

For a partnership to sustain, both parties need to benefit from the relationship. Therefore, department store 101 should also care about the value that it can bring to other partners in the coalition. The same regression analysis is conducted for the other 15 partners in the coalition to evaluate the value of the partners. Using the results, these15 partners can evaluate the partners following the same steps that we have done for department store 101.

In this part, we do not provide the same detailed description of the results for the 15 partners. Instead, we will focus on the value of department store 101 to the partners and the results are shown in table 6 and the result is presented in Table 14. Unfortunately, based on this analysis, customers who also shop at department store 101 do not show a significant increase in the lifetime value to the partners. That is, partners do not gain benefits by having customers who also shop at department store 101.

4.6 Discussion

Implications for Marketing Literature

Given the popularity of the coalition loyalty programs today, little is known about whether merchants can benefit from such type of program. This study shows that not all the partners in the coalition loyalty program are able to achieve the synergy with the partners in the loyalty program. For many partners in the coalition loyalty program, customers' active engagement with the program does not transform into high customer lifetime value. Furthermore, the synergies between a pair of partners are asymmetric. The result shows that as a major sponsor of the coalition loyalty program, a general merchandiser can benefit from the leverage by other specialty partner stores. However, the study also shows that the partners do not have significant leverage in customer lifetime value by having customers who also patronage the general merchandiser. Finally, this study gives us a better understanding of the customers attracted by the coalition loyalty program. We found that the benefit of crossing buying in the coalition loyalty program attracts two groups of vastly different customers. One group is most favorable customers for the merchants – customers with the high lifetime value and high level of cross buying. These customers account for only a small proportion of the customer base but contribute more than half of the total customers values to the merchant. The other group is the least favorable – customers with low lifetime value and high level of cross buying. The number of these customers is large and they have very little contribution to the customer equity of the merchant.

Implications for Marketing Practice

When the coalition loyalty program is gaining popularity in the market, it is easy for the merchants to be tempted to follow other merchants and join the program, with the hope to gain access to the vast customer base. This study cautions against the urge to participate in the coalition loyalty program simply because of the popularity of the coalition loyalty program. Instead, firm should carefully evaluate the potential synergies between the existing partners and themselves to understand whether they can benefit from such a coalition. Such synergies can come from the proximity to other partners in the coalition network and the similarity with the existing partners in the product offerings. Without such a synergy, coalition loyalty program does not provide the merchants with the desirable segment of customers.

Managers need to be aware of the hidden cost in participating in such a loyalty program. A tempting reason to participate in such a loyalty program is the lower cost of managing the loyalty program because the costs are shared by other partners in the coalition. We argue that being able to share the cost does not necessarily reduce the overall cost of the program. There is hidden cost for a merchant to pay in a coalition loyalty program – rewards to the customers with low loyalty. Our analysis shows that the coalition loyalty program brings a flock of low value customers to the merchants. These customers are loyal to the cash back rewards provided by the loyalty program instead of the merchant. They shop wildly in the network and thus are able to reach the reward threshold. As a result, merchants pay for the cost of the rewards of these low value customers, which could be avoided in the merchant's own in-house loyalty program. These low value customers can also create a crowding effect in the store and consume other resources of the retailer, such as the staff's time in serving these customers, resulting in lower satisfaction of other customers shopping at the store. We point out that there exist other alternatives to create a loyalty program at a lower cost. Today, there emerges a flock of consulting companies or software companies that provide loyalty program solutions. Merchants can customize the loyalty program based on their own needs or choose to adopt the standardized modules. These companies include 500 friends, 123 social media, Ariesy, Crowad Factory and Pluck.

Once participated in the coalition loyalty program, the merchants need to ensure that they commit the critical resources to the management of the program in order to fully exploit the benefit of the coalition program. One critical resource is the analytical capability to extract the managerial insights from the big data of customer information provided by the coalition loyalty program. Based on this data, merchants can have a more comprehensive understanding of the customers' preferences, habits and lifestyle from their transaction history at other partners. To acquire a deeper understanding of consumer behavior, the merchants need to consistently invest in the customer research units and acquire the talent to analyze the big data. Another resource is the managerial actions to enhance collaborations with other partners. For example, managers need to actively initiate collaborations with the valuable partners in order to create effective joint promotions to the customers.

Finally, for the partners that gain synergies from the coalition loyalty program, we suggest that it is beneficial for them to collaborate with other partners and develop joint promotions and marketing actions that increase customer values for both parties. An alliance with only a few partners benefiting is not sustainable. The drop-off of partners not benefiting from the coalition can hurt customers' confidence and satisfaction in the coalition loyalty program, resulting in a damage of those partners who are benefiting from the coalition.

Future Research

Several limitations in this study present opportunities for future studies. First, this study looked into only one coalition loyalty program in Europe. Through this one example, we show that synergies among partners are not guaranteed, even in a nation-wide coalition loyalty program with a large customer base and partner network. The pervasiveness of this lack-of-synergy program gives us an alert of the effectiveness of coalition loyalty programs. To access the severity of this problem in general, additional studies should be conducted on other coalition loyalty programs. With data from other coalition loyalty programs, researchers will also be able to explore how governance of the coalition loyalty programs, such as the design of the loyalty programs, the criteria of partner selection and different ways of information and cost sharing structures, influence the value of the program for their partners.

Second, this study did not factor the marketing actions taken by the merchants in the analysis. It is possible that the distinct differences between the values of the customers are a result of the marketing campaigns created by the merchants based on the coalition platform. It is possible that the strong position of department store 101 in the coalition program is a result of its aggressive and effective marketing actions. Unfortunately, detailed information on the partners' marketing campaigns is often not recorded by the coalition program. With such information, future studies can examine how partners' own actions influence their gains from the coalition loyalty program. We also urge managers of the coalition loyalty program to keep track of the partners' campaign efforts because it allows coalition program to give suggestions to its partners about effective campaigns based on the coalition platform. This knowledge transferring can become a valuable competitive advantage of the coalition platform.

Chapter 5

Conclusion

The access to rich customer data have provide many new opportunities for research in the area of customer dynamics. In this dissertation, I look into three problems that managers frequently encounter and provide generalizable frameworks that managers can apply to analyze the data and transform the data to actionable insights. From a managerial standpoint, this dissertation discusses customers' dynamics under three types of influences, which are major life events, entry of new products and the participation in a coalition loyalty program. The solution provided in this dissertation can help managers improve their targeted marketing efforts by selecting the right message for the right customer at the right time. From a methodological standpoint, the solutions proposed in this dissertation harness two characteristics of customer data. One is the richness of the data. I integrated records of various behaviors to generate deep sights on customers' underlying life states and evaluate the effectiveness of a coalition loyalty program. Another is the streaming nature of the data. To allow data to be processed as they arrive, I draw from the CUSUM control chart method and alter the method to accommodate the nature of marketing problems.

Yet, many exciting research opportunities awaits. The adopted framework of CUSUM can be applied to other business problems, such as monitoring product and service performance, brand association, value of an asset or a company, salespeople performance using data from sources like news reports, social media, online reviews, customer surveys and communication logs. The essence of these problems are the same. Using data from various sources, we can model the desirable pattern or the current state of a consumer, a product or a company. Because a change in the underlying states (i.e., the quality of the products, perceived image of a brand, the productivity of a salesperson) can result in change in observed behaviors (i.e., volume of sales, comments and reviews and detailed content of service requests), we can then use the CUSUM control chart to quickly detect any deviation from the desirable or current pattern.

Extending the solution proposed in this dissertation to other contexts and monitoring data of different formats can be challenging. Information from the new media today are presented in forms such as texts, pictures and videos. However, the Cusum method, as well as the majority of the statistical tools are designed to process numerical data. In order to apply the existing models, records in the form of text, picture and video need to be transformed into a numerical format. The way to conduct this transformation is to extract meaningful features, and then for each feature, rate the level of the feature or group the attributes into different categories.

Analyzing non-numerical data can be challenging in two particular ways. One is the extraction of meaningful features for the problem at hand. Studies today have only scratched minimal features from non-numerical data. Take the twitter data, a popular and widely available source of data, for example. Massive information from the twitter data is presented in the text format. However, the commonly used attributes of the twitter data are number of twitts, sentiment of the twits (i.e., positive, negative and neutral). Development in text mining techniques (e.x., topic modeling) can lead to the extraction of more complex and meaningful features. Future research can examine different ways to extract meaningful features from non-numerical data.

Another challenge is the complexity of computation when analyzing non-numerical data. To represent graphic and video data in numerical format, we might need many variables. The increase in the number of variables can lead to the increase in computation difficulty, making some tools infeasible, such as Bayesian estimation method. In this paper, we propose the CUSUM framework can be one solution of the problem. By developing a iterative equation, we can easily update results as new data arrive. Future research can look into more ways to simplify the computation through approximation approach or parallel computing.

Given the rich data at individual customer level and the advanced analytical tools, managers are able to gain deeper insights on customers' behaviors than ever before. However, managers need to employ these tools with great cautions. Most statistical tools are only effective for a specific type of problem or under a set of assumptions. The violation of the assumptions can lead to bias in the results. Managers need to combine their experience and intuition when interpreting the results. Furthermore, targeted marketing efforts, such as one-to-one marketing, can arouse consumers' concerns on privacy issues. In these cases, managers need to watch for possible backfires of a targeted campaign.

Bibliography

Ailawadi, K., Lehmann, D., & Neslin, S. (2003). Revenue Premium as an Outcome Measure of Brand Equity. *Journal of Marketing*, 67(4), 1-17.

Anderson, A. (1933). Runway Time and the Goal Gradient. *Journal of Experimental Psychology*, 16(3), 423-428.

Andreasen, A. (1984). Life Status Changes and Changes in Consumer Preferences and Satisfaction. *Journal of Consumer Research*, 11 (3), 784-794.

Baek, E. & Hong, G. (2004). Effects of Family Life-Cycle Stages on Consumer Debts. *Journal of Family and Economic Issues*, 25 (3), 359-385.

Bawa, K. (1990). Modeling Inertia and Variety Seeking Tendencies in Brand Choice Behavior. *Marketing Science*, 9(3), 263-278.

Benzies, K., Tough, S., Tofflemire, K., Frick, C., Faber, A. & Newburn-Cook, C. (2006). Factors Influencing Women's Decisions About Timing of Motherhood. *Journal of Obstetric, Gynecologic, & Neonatal Nursing*, 35 (5), 625-33.

Berry, J. (2013). Bulking Up: The 2013 Colloquy Loyalty Census. Retrieved March 29, 2016 from https://www.loyalty.com/research-insights/whitepapers/the-2013-colloquy-loyalty-census-us-edition

Brown, S. & Venkatesh, V. (2005). Model of Adoption of Technology in Households: A Baseline Model Test and Extension Incorporating Household Life Cycle. *MIS Quarterly*, 29 (3), 399-426.

Byung-Do, K., Shi, M., & Srinivasan, K. (2001). Reward Programs and Tacit Collusion. *Marketing Science*, 20(2), 99.

Cannon, J. & Perreault Jr., W. (1999). Buyer-Seller Relationships in Business Markets. *Journal of Marketing Research*, 36(4), 439-460.

Chandola, V, Sukumar S. & Schryver J. (2013). Knowledge discovery from massive healthcare claims data. *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, Chicago, Illinois, USA: ACM.

Cocco, J., Gomes F. & Maenhout, P. (2005). Consumption and Portfolio Choice over the Life Cycle. *Review of Financial Studies*, 18 (2), 491-533.

Colloquy. (2015). U.S. Customer Loyalty Program Memberships Top 3 Billion for First Time, 2015 Colloquy Census Shows. Retrieved March 29, 2016, from https://www.colloquy.com/latest-news/2015-colloquy-loyalty-census/

Columbus, L. (2015). Retrieved March 29, 2016 from http://www.forbes.com/sites/louiscolumbus/2015/03/15/data-analytics-dominates-enterprises-spending-plans-for-2015/#685ef59c3eb4

Cornfield, D. (1983). Chances of Layoff in a Corporation: A Case Study. *Administrative Science Quarterly*, 28 (4), 503-20.

Cosenza, R. & Davis, D. (1981). Family Vacation Decision Making Over the Family Life Cycle: A Decision and Influence Structure Analysis. *Journal of Travel Research*, 20 (2), 17-23.

Das, T. K., & Bing-Sheng, T. (2000). A Resource-Based Theory of Strategic Alliances. *Journal of Management*, 26(1), 31-62.

Dillon, W., Böckenholt, U., Borrero, M., Bozdogan, H., Sarbo, W., Gupta S., Kamakura, W., Kumar, A., Ramaswamy, B., & Zenor, M. (1994). Issues in the Estimation and Application of Latent Structure Models of Choice. *Marketing Letters*, 5 (4), 323-334.

Dorotic, M., Fok, D., Verhoef, P., & Bijmolt, T. (2011). Do vendors benefit from promotions in a multi-vendor loyalty program? *Marketing Letters*, 22(4), 341-356.

Dowling, G. R., & Uncles, M. (1997). Do Customer Loyalty Programs Really Work? *Sloan Management Review*, 38(4), 71-82.

Drèze, X., & Nunes, J. (2004). Using Combined-Currency Prices to Lower Consumers' Perceived Cost. *Journal of Marketing Research*, 41(1), 59-72.

Drèze, X., & Nunes, J. (2009). Feeling Superior: The Impact of Loyalty Program Structure on Consumers' Perceptions of Status. *Journal of Consumer Research*, 35(6), 890-905.

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

Du, R. & Kamakura, W. (2006). Household Life Cycles and Lifestyles in the United States. *Journal of Marketing Research*, 43 (1), 121-32.

Du, Y., Ye, H., & Damangir, S. (2015). Leveraging Trends in Online Searches for Product Features in Market Response Modeling. *Journal of Marketing*, 79 (1), 29-43.

Erdem, T. & Keane, M. (1996). Decision-Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets. *Marketing Science*, 15(1), 1-20.

Fader, P., Hardie, B., & Chun-Yao, H. (2004). A Dynamic Changepoint Model for New Product Sales Forecasting. *Marketing Science*, 23(1), 50-65.

Farquhar, P., & Rao, V. (1976). A Balance Model for Evaluating Subsets of Multiattributed Items. *Management Science*, 22(5), 528-539.

Fernández-Villaverde, J. & Krueger, D. (2007). Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data. *Review of Economics & Statistics*, 89(3), 552-65.

Fritzsche, D. (1981). An Analysis of Energy Consumption Patterns by Stage of Family Life Cycle. *Journal of Marketing Research*, 18 (2), 227-232.

George, M., Rowlands D., Price, M., & Maxey J. (2005). *The Lean Six Sigma Pocket Toolbook*. New York, NY: McGraw-Hill.

Gobillon, L., Magnac, T., & Selod, H. (2011). The Effect of Location on Finding a Job in the Paris Region. *Journal of Applied Econometrics*, 26 (7), 1079-112.

Gourinchas, P. & Parker, J. (2002). Consumption over the Life Cycle. *Econometrica*, 70(1), 47–89.

Gupta, S., Hanssens, D., Hardie, B., Kahn, W., Kumar, V., Lin, N., & Sriram, N. (2006). Modeling Customer Lifetime Value. *Journal of Service Research*, 9(2), 139-155.

Haenlein, M. (2013). Social Interactions in Customer Churn Decisions: The Impact of Relationship Directionality. *International Journal of Research in Marketing*, 30(3), 236-48.

Hanley, A. (2005). Receiver Operating Characteristic (ROC) Curves. In *Encyclopedia of Bio-statistics*. Retrieved March 29, 2016 from http://onlinelibrary.wiley.com/doi/10.1002/0470011815.b2a04043/abstract

Haughton, D. & Oulabi, S. (1997). Direct marketing modeling with CART and CHAID. *Journal of Interactive Marketing*, 11(4), 42-52.

Heilman, C., Bowman, D. & Wright G. (2000). The Evolution of Brand Preferences and Choice Behaviors of Consumers New to a Market. *Journal of Marketing Research*, 37(2), 139-155.

Helsen, K., & Schmittlein, D. (1999). Analyzing Duration Times in Marketing: Evidence for the Effectiveness of Hazard Rate Models. *Marketing Science*, 11(4), 395-414.

Homburg, C., Hahn, A., Bornemann, T., & Sandner, P. (2014). The Role of Chief Marketing Officers for Venture Capital Funding: Endowing New Ventures with Marketing Legitimacy. *Journal of Marketing Research,* 51(5), 625-44.

Jones-Farmer, F., Megahed M. & Allison, L. (2015). Statistical Perspectives on "Big Data". In Knoth, S. (Eds.), *Frontiers in Statistical Quality Control* (pp. 29-47), Wolfgang: Springer.

Jung, J. & Tran, C. (2014). Medical consumption over the life-cycle. *Empirical Economics*, 47 (3), 927-57.

Kivetz, R. (2005). Promotion Reactance: The Role of Effort-Reward Congruity. *Journal of Consumer Research*, 31(4), 725-736.

Kivetz, R., & Simonson, I. (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.

Kivetz, R., & Simonson, I. (2003). The Idiosyncratic Fit Heuristic: Effort Advantage as a Determinant of Consumer Response to Loyalty Programs. *Journal of Marketing Research*, 40(4), 454-467.

Kivetz, R., Urminsky, O., & Zheng, Y. (2006). The Goal-Gradient Hypothesis Resurrected: Purchase Acceleration, Illusionary Goal Progress, and Customer Retention. *Journal of Marketing Research*, 43(1), 39-58.

Kopalle, P., & Neslin, S. (2003). The Economic Viability of Frequency Reward Programs in a Strategic Competitive Environment. *Review of Marketing Science*, 1, 1-39.

Kopalle, P., Sun, Y., Neslin, S., Sun, B., & Swaminathan, V. (2012). The Joint Sales Impact of Frequency Reward and Customer Tier Components of Loyalty Programs. *Marketing Science*, 31(2), 216-235.

Kreyenfeld, M. (2010). Uncertainties in Female Employment Careers and the Postponement of Parenthood in Germany. *European Sociological Review*, 26 (3), 351–66.

Kumar, V., Venkatesan, R., & Reinartz, W. (2006). Knowing What to Sell, When, and to Whom. *Harvard Business Review*, 84(3), 131-137.

Kumar, V., Venkatesan, R., Bohling, T., & Beckmann, D. (2008). The Power of CLV: Managing Customer Lifetime Value at IBM. *Marketing Science*, 27(4), 585-599.

Lal, R., & Bell, D. (2003). The Impact of Frequent Shopper Programs in Grocery Retailing. *Quantitative Marketing & Economics*, 1(2), 179-202.

Lansing, J. & Leslie, K. (1957). Family Life Cycle as an Independent Variable. *American Sociological Review*, 22 (5), 512–19.

Leenheer, J., & Bijmolt, T. (2008). Which retailers adopt a loyalty program? An empirical study. *Journal of Retailing and Consumer Services*, 15(6), 429-442.

Lemon, K., & Wangenheim, F. (2009). The Reinforcing Effects of Loyalty Program Partnerships and Core Service Usage: A Longitudinal Analysis. *Journal of Service Research*, 11(4), 357-370.

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

Li, S., Sun, B., & Montgomery, A. (2011). Cross-Selling the Right Product to the Right Customer at the Right Time. *Journal of Marketing Research*, 48(4), 683-700.

Liao, T. (2005). Clustering of Time Series Data: A Survey. *Pattern Recognition*, 38(11) 1857-1874.

Liechty, J., Pieters, R. & Wedel M. (2003). Global and local covert visual attention: Evidence from a Bayesian Hidden Markov Model. *Psychometrika*, 68 (4), 519-41.

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

Liu, Y., & Yang, R. (2009). Competing Loyalty Programs: Impact of Market Saturation, Market Share, and Category Expandability. *Journal of Marketing*, 73(1), 93-108.

Liu-Thompkins, Y., & Tam, L. (2013). Not All Repeat Customers Are the Same: Designing Effective Cross-Selling Promotion on the Basis of Attitudinal Loyalty and Habit. *Journal of Marketing*, 77(5), 21-36.

Lu, W., & Tong, H. (2009). Detecting Network Anomalies Using CUSUM and EM Clustering. In Zhihua Cai and Zhenhua Li and Zhuo Kang and Yong Liu (Eds.), *Advances in Computation and Intelligence* (pp. 297-308). Berlin, Heidelberg: Springer.

Luo, A. & Kumar, V. (2013). Recovering Hidden Buyer–Seller Relationship States to Measure the Return on Marketing Investment in Business-to-Business Markets. *Journal of Marketing Research*, 50(1), 143-160.

Machado, F., & Sinha, R. (2007), Smoking Cessation: A Model of Planned vs. Actual Behavior for Time-Inconsistent Consumers. *Marketing Science*, 26 (6), 834-850.

Mägi, A. (2003). Share of Wallet in Retailing: The Effects of Customer Satisfaction, Loyalty Cards and Shopper Characteristics. *Journal of Retailing*, 79(2), 97-106.

Mathur, A., Moschis, G., & Lee, E. (2008). A Longitudinal Study of the Effects of Life Status Changes on Changes in Consumer Preferences. *Journal of the Academy of Marketing Science*, 36 (2), 234-246.

McAlister, L. (1982). A Dynamic Attribute Satiation Model of Variety-Seeking Behavior. *Journal of Consumer Research*, 9(2), 141-150.

McLeod, P., & Ellis, J. (1982). Housing Consumption Over the Family Life Cycle: An Empirical Analysis. *Urban Studies*, 19 (2), 177-185.

McWilliams, A., & Siegel, D. (1997). Event Studies in Management Research: Theoretical and Empirical Issues. *Academy of Management Journal*, 40(3), 626-657.

Megahed, F., Woodall, W., & Camelio, J. (2011). A Review and Perspective on Control Charting with Image Data. *Journal of Quality Technology*, 43 (2), 83-98.

Menken, J., Trussell, J., Stempel, D., & Babakol, O. (1981). Proportional Hazards Life Table Models: An Illustrative Analysis of Socio-Demographic Influences on Marriage Dissolution in the United States. *Demography*, 18 (2), 181-200.

Metz, C. (1978). Basic principles of ROC analysis. *Seminars in Nuclear Medicine*, 8 (4), 283-298.

Meyer-Waarden, L. (2007). The effects of loyalty programs on customer lifetime duration and share of wallet. *Journal of Retailing*, 83(2), 223-236.

Michael, R., & Brandon, T. (1985). Entry into Marriage and Parenthood by Young Men and Women: The Influence of Family Background. *Demography*, 22 (4), 515-44.

Miller, C., Reardon, J., & McCorkle, D. (1999). The Effects of Competition on Retail Structure: An Examination of Intratype, Intertype, and Intercategory Competition. *Journal of Marketing*, 63(4), 107-120.

Montgomery, D. (2007), *Introduction to statistical quality control*. Hoboken, NJ: John Wiley & Sons.

Moon, S., Kamakura, W., & Ledolter, J. (2007). Estimating Promotion Response When Competitive Promotions Are Unobservable. *Journal of Marketing Research*, 44 (3), 503–15.

Morrison, D. (1969). On the Interpretation of Discriminant Analysis. *Journal of Marketing Research*, 6 (2), 156–63.

Nako, S. (2004). Frequent Flyer Programs and Business Travelers: An Empirical Investigation. *Logistics and Transportation Review*, 28(4), 394-414.

Netzer, O., Lattin, J., & Srinivasan, V. (2008). A Hidden Markov Model of Customer Relationship Dynamics. *Marketing Science*, 27 (2), 185–204.

Ngobo, P. (2004). Drivers of Customers' Cross-buying Intentions. *European Journal of Marketing*, 38(9), 1129-1157.

Nielsen. (2013). How loyal are your customers? A View of Loyalty Sentiment around the World. Retrieved March 29, 2106 from http://www.nielsen.com/content/dam/nielsenglobal/eu/nielseninsights/pdfs/Nielsen%20Global%20Report%20of%20Loyalty%20Sentiment%20November%202013.pdf

O'Brien, L., & Jones, C. (1995). Do Rewards Really Create Loyalty? *Harvard Business Review*, 73(3), 75-82.

Page, E. (1954). Continuous Inspection Schemes. *Biometrika*, 41 (1/2), 100-15.

Page, E. (1961). Cumulative Sum Charts. *Technometrics*, 3 (1), 1-9.

Pearson, B. (2014). Brand - Loyalty Myths, and Why Teaming up Beats Going it Alone. Retrieved March 29, 2016 from http://www.fastcompany.com/3007449/brand-loyalty-myths-and-why-teaming-beats-going-it-alone

Perron, P., (1989). The great crash, the oil price shock and the unit root hypothesis. *Econometrica*, 57, 1361-1401.

Punj, G., & D. Stewart (1983). Cluster Analysis in Marketing Research: Review and Suggestions for Application. *Journal of Marketing Research*, 20 (2), 134–48.

Pyzdek, T., & Keller, P. (2003). *The Six Sigma handbook: a complete guide for green belts, black belts, and managers at all levels*. New York, NY: McGraw-Hill.

Reeves, M., & Deimler, M. (2011), Adaptability: The New Competitive Advantage. *Harvard Business Review*, 89 (7), 134-41.

Reichheld, F., & Teal, T. (2001). *The Loyalty Effect: The Hidden Force Behind Growth, Profits, and Lasting Value*. Brighton, MA: Harvard Business Publishing.

Reinartz, W., & Kumar, V. (2000). On the Profitability of Long-Life Customers in a Noncontractual Setting: An Empirical Investigation and Implications for Marketing. *Journal of Marketing*, 64(4), 17-35.

Roehm, M., Pullins, E., & Roehm Jr, H. (2002). Designing Loyalty-Building Programs for Packaged Goods Brands. *Journal of Marketing Research*, 39(2), 202-213.

Schmittlein, D., & Mahajan, V. (1982). Maximum Likelihood Estimation for an Innovation Diffusion Model of New Product Acceptance. *Marketing Science*, 1 (1), 57.

Schmittlein, D., & Morrison, D. (1983). Modeling and estimation using job duration data. *Organizational Behavior and Human Performance*, 32 (1), 1-22.

Schumann, J., Wünderlich, N., & Evanschitzky, H. (2014). Spillover Effects of Service Failures in Coalition Loyalty Programs: The Buffering Effect of Special Treatment Benefits. *Journal of Retailing*, 90(1), 111-118.

Schwartz, E., Bradlow, E., & Fader, P. (2014). Model Selection Using Database Characteristics: Developing a Classification Tree for Longitudinal Incidence Data. *Marketing Science*, 33 (2), 188–205.

Schweidel, D., Bradlow, E. & Fader, P. (2011). Portfolio Dynamics for Customers of a Multi-Service Provider. *Management Science*, 57(3) 471-486.

Schweidel, D., Young-Hoon, P., & Jamal, Z. (2014). A Multiactivity Latent Attrition Model for Customer Base Analysis. *Marketing Science*, 33 (2), 273-86.

Shah, D., & Kumar, V. (2012). The Dark Side of Cross-Selling. *Harvard Business Review*, 90(12), 21-23.

Sheth, J., & Parvatiyar, A. (1992). Towards a Theory of Business Alliance Formation. *Scandinavian International Business Review*, 1(3), 71-87.

Sheth, J., & Sisodia, R. (2002). *The Rule of Three: Surviving and Thriving in Competitive Markets*. New York, NY: The Free Press.

SLI. (2013). Is the U.S. the Afghanistan of Coalition Loyalty? Overcoming the Challenges of the U.S. Market. Retrieved March 29, 2016, http://www.sli21.com/?whitepaper=is-the-us-the-afghanistan-of-coalition-loyalty

Sood, A., & Tellis, G. (2009). Do Innovations Really Pay Off? Total Stock Market Returns to In-novation. *Marketing Science*, 28(3) 442-456.

Sriram, S., Balachander, S. & Kalwani, M. (2007). Monitoring the Dynamics of Brand Equity Using Store-Level Data. *Journal of Marketing*, 71 (2), 61-78.

Su, X. & Bozeman, B. (2009). Dynamics of Sector Switching: Hazard Models Predicting Changes from Private Sector Jobs to Public and Nonprofit Sector Jobs. *Public Administration Review*, 69 (6), 1106-14.

The CMO Survey. (2015). Retrieved March 29, 2016, from http://cmosurvey.org/results/ http://cmosurvey.org/results/

Thomaz, F., & Swaminathan, V. (2015). What Goes Around Comes Around: The Impact of Marketing Alliances on Firm Risk and the Moderating Role of Network Density. *Journal of Marketing*, 79(5), 63-79.

Tsui, F., J. Espino, Dato, V., Gesteland, P., Hutman, J., & Wagner, M. (2003). Technical De-scription of RODS: A Real-time Public Health Surveillance System. *Journal of the American Medical Informatics Association*, 10 (5), 399–408.

Varadarajan, P., & Jayachandran, S. (1999). Marketing Strategy: An Assessment of the State of the Field and Outlook. *Journal of the Academy of Marketing Science*, 27(2), 120-143.

Verhoef, P., Venkatesan, R., McAlister, L., Malthouse, E. C., Krafft, M., & Ganesan, S. (2010). CRM in Data-Rich Multichannel Retailing Environments: A Review and Future Research Directions. *Journal of Interactive Marketing*, 24(2), 121-137.

Wald, A. & Wolfowitz, J. (1948). Optimum Character of the Sequential Probability Ratio Test. *The Annals of Mathematical Statistics*, 19 (3), 326-39.

Wald, A. (1945). Sequential Tests of Statistical Hypotheses. *The Annals of Mathematical Statistics*, 16 (2), 117-86.

Webb, M. (2006). *Sales and marketing the six sigma way*. Fort Lauderdale, Florida: Kaplan.

Wedel, M. (2000), *Market segmentation: Conceptual and methodological foundations*. New York, NY: Springer.

Wells, W. & Gubar, G. (1966). Life Cycle Concept in Marketing Research. *Journal of Marketing Research,* 3 (4), 355-63.

Wiles, M. A., & Danielova, A. (2009). The Worth of Product Placement in Successful Films: An Event Study Analysis. *Journal of Marketing*, 73(4), 44-63.

Wiles, M., Jain, S., Mishra, S., & Lindsey C. (2010). Stock Market Response to Regulatory Reports of Deceptive Advertising: The Moderating Effect of Omission Bias and Firm Reputation. *Marketing Science*, 29(5) 828-845.

Wilkes, R. (1995). Household Life-Cycle Stages, Transitions, and Product Expenditures. *Journal of Consumer Research*, 22 (1), 27–42.

Zhang, J., & Breugelmans, E. (2012). The Impact of an Item-Based Loyalty Program on Consumer Purchase Behavior. *Journal of Marketing Research*, 49(1), 50-65.