A Control Mechanism for Sales Associates in High-End Retail

by

Shiou Lin Sam

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the

degree

of

Doctor of Philosophy in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2007
[June 2007]

Author ....
Department of Electrical Engineering and Computer Science
May 8, 2007

Certified by ....
Gabriel R. Bitran
Nippon Telegraph and Telephone Professor of Management Science
Thesis Supervisor

Accepted by ....
Arthur C. Smith
Chairman, Department Committee on Graduate Students

A Control Mechanism for Sales Associates in High-End Retail

by

Shiou Lin Sam

Submitted to the Department of Electrical Engineering and Computer Science
on May 8, 2007, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Computer Science and Engineering

Abstract

The strength of the high-end retail industry has traditionally been in marketing and branding while management and operational efficiency lagged behind. In the face of changing client demographics and increased competition, improving operations through better data management and utilization could prove promising. In this work, we attempt to do so by focusing on a neglected area in high-end retail - management of the sales associates. Our first finding is the existence of a disconnect between the data collected and the data required for better control of the associates. Recognizing the gap, we sidestep it by tapping the knowledge of many experienced sales associates through field work. This knowledge is then funnelled back to assist in modeling client behavior. The dynamics between an associate and his clients are modeled using an evolution model with stochastic client behavior. We show that under certain conditions, the optimal policy for an associate is a quasi-concave policy. In addition, we provide a methodology that would enable the associates to capture the full potential of a client while at the same time, allow management to reduce the variability in customer service within the store. The computational results indicate that such a mechanism, when compared to the commonly practiced policies, can achieve a substantial lift in revenue generated. In addition, the results also provide managerial insights and expose some common misconceptions.

Thesis Supervisor: Gabriel R. Bitran
Title: Nippon Telegraph and Telephone Professor of Management Science

Acknowledgments

I am grateful

To Professor Gabriel Bitran, for guiding, supporting and believing in me throughout this long process.

To Professors Duane Boning and Munther Dahleh, for their helpful insights and feedback.

To the OM group, for making me feel at home.

To Professor Hal Abelson, Professor Gerry Sussman and the EECS department, for enabling me to explore during the initial stages.

To the HB crew, for the colorful weekends.

To David Deveau, for taking me on as a student and his subsequent enthusiasm towards my often subpar music.

To Tor and Martin, for our lunches, dinners, discussions on convexity, polytopes and the likes.

To Joyce, for our swimming sessions, shopping and dining excursions.

To my brother, for being my link to the outside world.

To my parents, for their worries, encouragement, love and support.

To QQ, for welcoming me home with a wagging tail and tucked back ears every night.

To Tinglea, for being there, always.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The distribution of the United States household earnings has undergone tremendous changes in the past 30 years. In the 1970s, income distribution clustered around the average level, with few households earning more than the average. Today, the steep drop of number of households past the average is much more attenuated [1]. The data in Figure 1-1 obtained from the U.S Census Bureau shows that the number of consumers with an annual income of $80,000 or greater has been increasing steadily from 1996 to 2005, with 10% of the population and 25.6% of households falling into this segment in 2005. In Figure 1-2, a similar trend is observed for the number of consumers with an annual income of $100,000 or greater, with 6.7% of the population and 17.2% of households falling into this segment in 2005. As such, the belief that there are only two distinct consumer groups - the mass and the high-end - is no longer true. An entirely new segment, known as the Mass Affluent and famously dissected by Paul Nunes and Brian Johnson in their book "Mass Affluence: Seven New Rules of Marketing to Today's Consumer" [2], has emerged. This segment, characterized by its ability and willingness to pay a premium for brand and quality, is quickly becoming the new darling of the retail industry. The attractiveness of the mass affluent, driven by its lucrative nature and steady growth, is further accentuated by the influence it casts over the purchase behavior of the general public [3]. This phenomenon, sometimes known as the snob and the follower effect, has been well studied in the economics literature ([4],[5]).

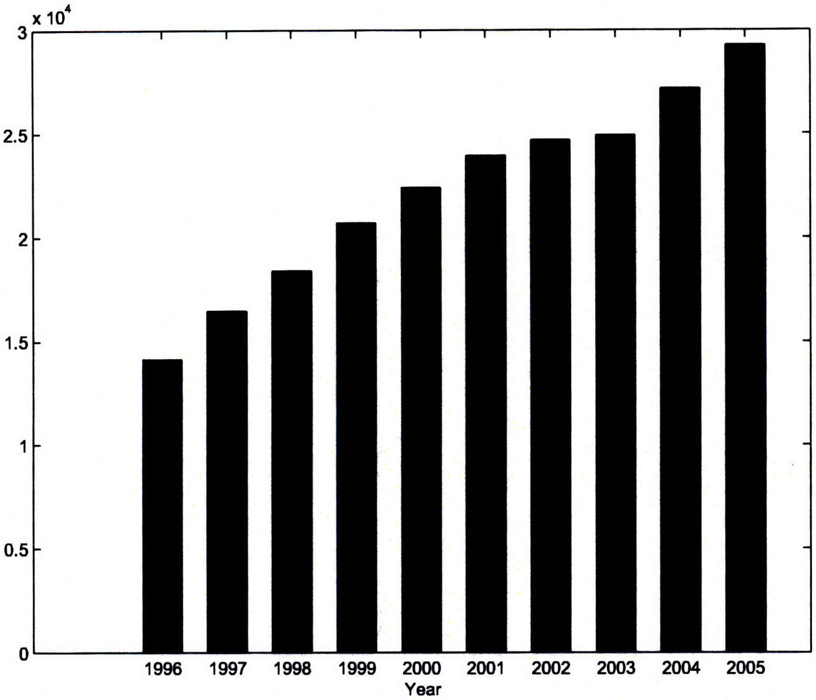

(a) Number of consumers with annual income of $80,000 or greater.

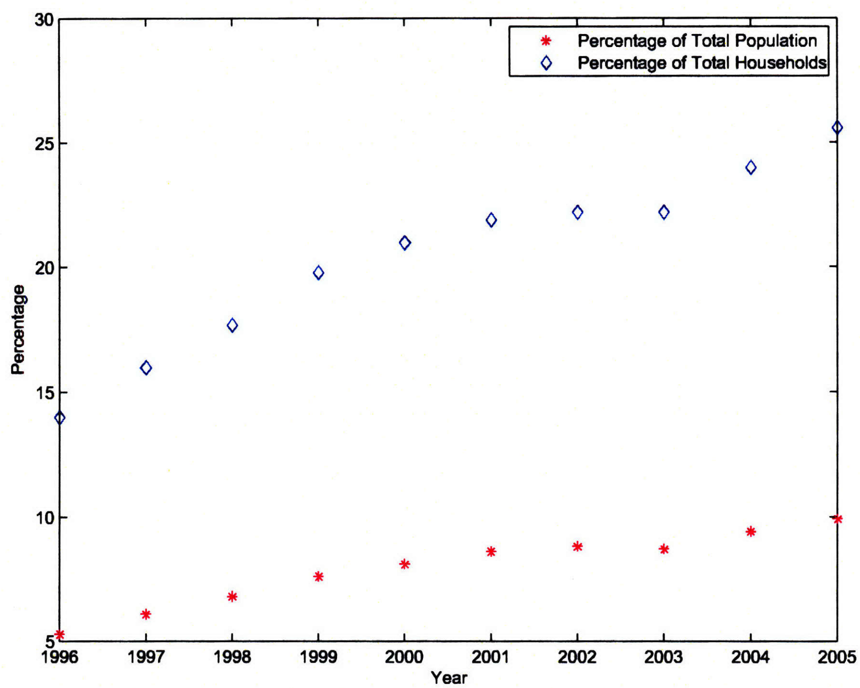

(b) Percentage of consumers with annual income of $80,000 or greater.

Figure 1-1: Income distribution of households in U.S. where annual income is greater than $80,000 (1996-2005).

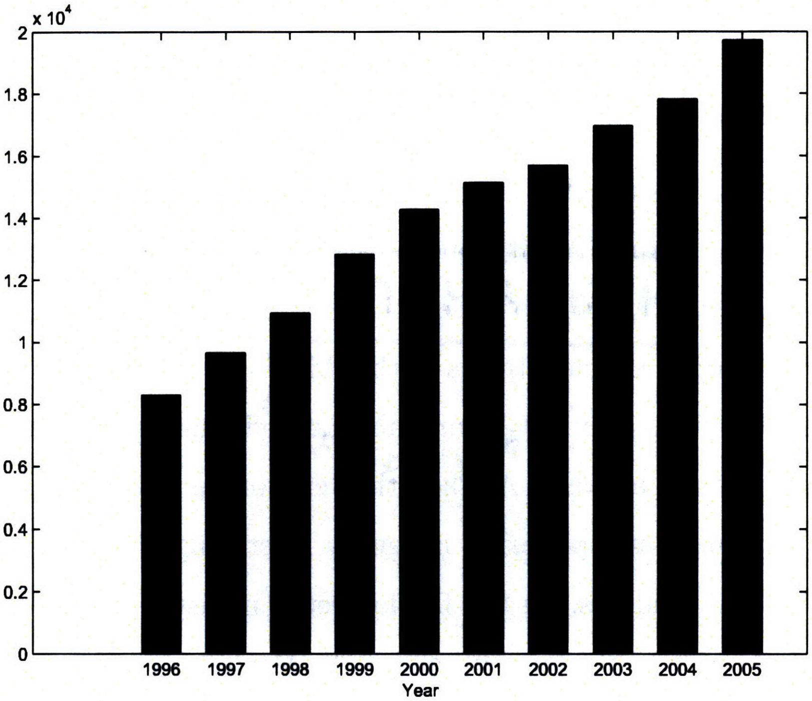

(a) Number of consumers with annual income of $100,000 or greater.

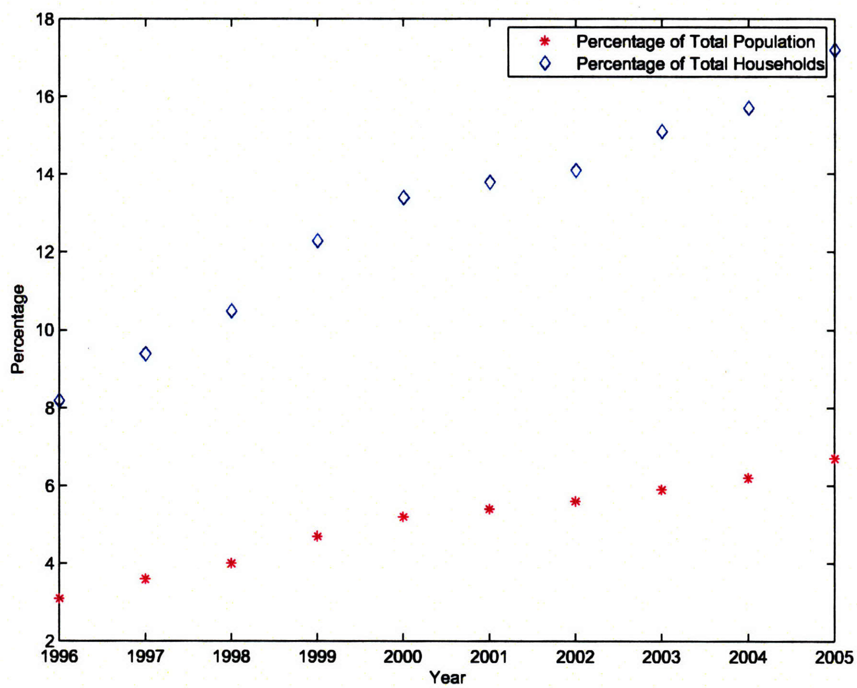

(b) Percentage of consumers with annual income of $100,000 or greater.

Figure 1-2: Income distribution of households in U.S. where annual income is greater than $100,000 (1996-2005).

One explanation suggests that the existence of this effect is driven by the desire to *not* be identified as the less-endowed. As such, the mass consumers tend to follow the consumption patterns of the well-endowed.

The emergence and subsequent recognition of the mass affluent segment have dramatically altered the retailing landscape. Traditional upscale retailers are attempting to diversify their product portfolio and to venture into price points more accessible to the mass affluent. We observe this in the luxury car industry as both BMW and Daimler now offer lower priced models. In the fashion industry, many luxury labels are offering separate sports and/or casual lines at lower price points as well. At the other end of the spectrum, the mass retailers are attempting to project an up-market image by leveraging third party brand equity in the creation of upscale products. In the fashion retail industry, this is exemplified by the relationships of H&M with Karl Lagerfeld[1] and Stella McCartney.[2] Target, often pronounced in faux French as "Tar-zhay", entered into high profile design partnerships across various product lines. In apparel, Target partnered with Isaac Mizrahi, while in kitchenware, it introduced Philippe Stark's line followed by Michael Graves' in quick succession.

These responses from existing and newly emerging players have resulted in a much more competitive high-end fashion retail industry, one that is projected to reach 50 billion USD in 2010 with an annual growth rate of 10% to 15%. The plethora of choices available further exacerbates the situation by reducing customer loyalty towards a particular brand. Many in high-end retail corporate upper management are turning to cultivating client relationships as a way to increase client loyalty and to differentiate themselves from the mass retailers. While the strategy seems sound, the implementation itself deserves scrutiny. Since client relationships are best built through sales associates, the role of an associate takes on increased importance in the face of the new landscape. Yet, has the management of these associates at the store level changed to reflect the added importance? Answers to this question hinge on the understanding of the current situation, which we delve into next in §1.1.

[1]Chief Executive of Design at House of Chanel, furs for Fendi and his eponymous label.
[2]Daughter of Paul McCartney, previously helming Chloé and presently her eponymous label.

1.1 Current Control and Relationship between the Retailer and Sales Associates

In most high-end retail stores, sales associates are hired on a commissioned basis. Some associates are fully commissioned while others receive an hourly rate in exchange for a lowered commission. Each associate usually maintains a client book containing detailed information about his clients - to be referred to when attempting to contact clients for generating sales. Figure 1-3 depicts a typical client page. The client book is built slowly over the years. A

Name	
Address	
Business	**Contact**
Occupation	Office
Look	Email
Family	**Events**
Spouse	Birthday
Children	Anniversary
Personal	**Sizes**
Vacation	Coat/Dress
Hobbies	Shirt/Blouse
Interests	Slack/Skirt
	Shoe

Figure 1-3: A generic client page.

walk-in client purchasing for the first time will usually be asked if he would like to be included in the central database. Concurrently, the associate selectively adds new clients to his book. It is worth noting that the client books and the central database are not synchronized; the information held in each are also very different. The central database consists of transaction histories - some of which are associated with names - of all customers who have bought from the store. Most of the data is not very user friendly since the transactions are often identified by stock keeping units (SKUs). As such, it is nearly impossible to backtrack simple information like size and color. The client books contain a subset of customers from the central database, and the information consists of mostly personal preferences. Clients stay active for a certain period of time, upon which, due to a variety of reasons, they stop visiting the store.

Management (both corporate and store level) offer little training to new and existing hires, and tend to focus their effort on marketing and client acquisition. Even though client relationships are acknowledged to be important at the corporate level, the actual task of cultivating these relationships are mostly delegated to the store level. Unfortunately, some store level managers believe that once a client is acquired, he will be enamored by the brand and its quality such that repeat business will follow naturally. The sales associates, due to their particular compensation scheme coupled with the distinct relationship they have with clients, often think of themselves as independent agents maintaining a store within a store. These associates are loosely managed, and are motivated by the amount of sales they push through the door. Under complete autonomy and little guidance, these associates naturally demarcate into two groups. The good associates generate a significant amount of revenue for the store, while the mediocre associates merely get by. As one drills for what makes a good associate, it appears that "good" is a misleading characterization. An associate in a high traffic store can do well by merely focusing on walk in clients. Such an associate is not necessarily considered "good" since he does very little client work, and consequently fails to capture the full potential of a client. On the other hand, an associate who makes the effort of maintaining and cultivating client relationships can sometimes fare worse, brought about by his little understanding of the nuances of client relationships. The different types of associate in the store, coupled with very little management control, result in high variability in customer service levels. Depending on pure luck, a new client could either have an associate who never bothers to contact or one who is overbearing. One cringes at the possibility of a high potential client being managed by a mediocre associate.

Having understood the status quo, and perhaps having overcome the initial shock, many questions follow. How much opportunity has been squandered? How much does the store stand to gain if the full potential of clients are realized? Is there a way to capture such potential systematically? Is such a system easily implementable? How does one reduce the variability in customer service level? Can one provide managerial insights?

Our research, then, is very much driven by providing answers to this series of questions. We begin by reviewing relevant literature in §1.2.

1.2 Literature Review

Of interest to us is literature relating to relationship management for sales associates. We find that much of the research relating to sales force focuses on compensation [6] and management [7] from a firm's perspective. The earlier works modeled sales force using a deterministic selling effort, while the later works evolved into using stochastic models. This line of research typically employs a game theoretic framework to find a compensation scheme such that both parties will not be incentivized to deviate from it. An interesting paper by Wernerfelt [8] looked at the function of sales assistance. In this paper, Wernerfelt models such function as providing a better match between the buyer with the product, and concludes that in certain industries, sales assistance does indeed prevail.

As direct literature search turns unfruitful, we take a step back and examine the mechanism used by associates for relationship building. In its simplest form, this mechanism is illustrated in Figure 1-4. In every period, the agent makes a contact decision based on the state of the customer. The customer observes the decision and takes his own action. The customer state is then updated, and the sequence of events is repeated for the next period. Presented this way, one observes that the dynamics of relationship building can be mapped easily to direct mailing and e-commerce industries. In the direct mailing industry, the agent is analogous to the catalog firm, the message to the catalog and the contact decision to the mailing decision. In the e-commerce industry, the agent is analogous to the firm, the message to the email and contact decision to the emailing decision. In the high-end retail industry, the agent is analogous to the sales associate, the message to the phone call and the contact decision to the calling decision.

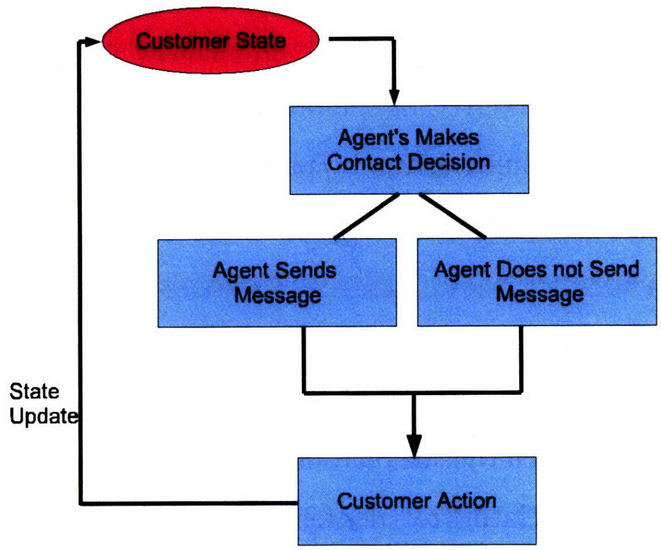

Figure 1-4: Mechanism used for relationship building.

Recognizing this allows us to leverage literature relating to the catalog industry. We begin by providing a brief industry overview. In the catalog sales industry, companies classify their customers into two categories: customers who belong to the *rental lists* and those who belong to the *house list*. Rental lists consist of names that have certain features in common and can be rented by the catalog companies. Thus, according to the products sold in the catalog, companies rent names from the market segments that are likely to buy them. The response rates for rental lists are generally low (0.5% to 1.5%). Once a customer from a rental list responds with a purchase, this customer becomes part of the house list, upon which companies can use the names without incurring a fee. In addition, house lists in general have a higher response rate of 3% to 8%.

Traditionally, catalog firms segment their customers along three dimensions - Recency, Frequency and Monetary value - otherwise known as the RFM classification. The most common definitions found in the industry for these three parameters are described in what follows. Recency of the last purchase is defined by the number of periods since the last purchase. Frequency is defined by the number of times a customer has responded with an order or by the fraction of mailings that the customer has responded to within a certain period of time. This definition is often times simplified further to consider only two values for frequency: one for customers that just entered the house list and the second for customers that have bought more than once from the company. Monetary value is defined by the average dollar amount per purchase, or by the dollar value that the customer spent in the last purchase or in all purchases to date.

The nature of the industry is rich enough that it has generated a wealth of research, roughly separable into catalog quality and appearance issues [9], customer returns [10], customer trust [11] and optimal mailing policies. Due to our interest in relationship management, in the following we will only focus on research relating to the optimal mailing policies.

Bitran and Mondschein [12] study optimal mailing and reordering policy in a stochastic setting. They focus on the cash flow and inventory constraints and their impact on optimal mailing strategies. They investigate the trade-off between sending catalogs to prospective customers and house customers using traditional RFM parameters to capture the state of the customers. They build a model to study how optimal solution changes with price, cost of goods, response rates and mailing costs. The computational results show that in the presence of cash flow constraints, it is better to send fewer catalogs to the recently responded customer and use the money to reach the less responsive ones. They also show that catalog companies should expect to spend a significant amount of resources in market tests to build their house lists.

Gönül and Shi [13] cast the problem in a game theoretic setting where a customer's utility is parameterized, thereby explicitly modeling the customer. They assume that the customers understand the firm's mailing policy and act in such a way that increases their chances of receiving a catalog. The firm and customer's problem are solved together to produce an optimal mailing policy. They make the strong assumption that a customer knows the cost of mailing and the mailing policy in simplified form, and that the information is publicly available. By doing this, they provide justification for not including budget constraints and inventory costs. They also assume that the firm has a very simplified cost structure which involves two variables: R (average revenue) and C (cost). Further, since parameters for the customer utility functions are estimated using past data, the catalog content is assumed not to vary substantially period to period. They find that it is not always optimal to mail to customers who have very recently purchased. It is better to save the mailing dollars for customers who have not responded recently. The catalog, in this case, serves as a reminder. Their result is similar to ones obtained by Bitran and Mondschein [12].

Elsner et al. [14] presented a multi-period optimization model of catalog mailings that allows for dynamic promotion or demotion of customers across a large number of customer segments. The model operates at three levels. They show that mailing to low-valued customers may be profitable even when these customers are not likely to respond immediately.

Simester, Sun and Tsitsiklis [15] tackle the problem by designing a discrete state space using customer histories and optimizing the mailing policy based on the state space. They calculate the transitional probabilities directly from the historical data and do not impose functional form assumptions. Thus, the model relies on stochasticity in the historical policy. In particular, if the historical policy is to mail only to customers who have recently purchased, they cannot estimate the effects of mailing to customers who have not purchased recently. They implemented the model using 13 variables to describe each customer's mailing and purchase histories. The model was validated in a field test with mixed results. The optimal policy recommended mailing less to customers who had recently purchased, but it turned out that there was insufficient stochasticity in the data to predict customer behavior under such scenario.

The existing literature in the catalog industry illustrates that two approaches could be used for modeling the associate. One is to use a dynamic programming formulation, while the other is to use a game theoretic framework. Similarly, to model the clients, one approach is to use historical data to predict future behavior; the other is to model the clients directly as solving their own utility maximization problem. In addition, the literature shows that the choice of methods does not change the fundamental result.

In the rest of this thesis we will use the ideas derived from the catalog industry literature to assist in our modeling. In tackling the problem we focus on the high-end retail environment and take the perspective of a sales associate, whom we define as follows.

Definition 1. *A sales associate is an employee who assists a client with his current and all future purchases for the entire duration he remains a client of the store.*

In Chapter 2 we formulate the problem of an associate and his clients. Chapters 3 and 4 analyze the single client problem in the infinite and finite horizon cases respectively. Chapter 5 analyzes the more realistic problem of an associate with multiple clients. We conclude with managerial insights and suggestions for future work in Chapter 6.

Chapter 2

Modeling An Associate and His Clients

The high-end retail industry is extremely seasonal and revolves around two distinct seasons - Spring/Summer and Fall/Winter. Retailers receive all merchandise in the beginning of each season and the assortment stays on the selling floor until close to the end of the season. As with all seasonal merchandise, the "leftovers" are discounted and eventually shipped out of the store. In recent years, high-end retailers are beginning to receive a small set of new assortment in the middle of the season. These refreshers are timed to coincide with the cruise and resort season. With this in mind, we model the sales associate problem as a multiperiod decision making process. In each period, the associate needs to choose a control policy in order to maximize his long term revenue stream. In this thesis we will focus on a single control variable which is binary in nature - the contact decision. The choice of control variable coupled with the periodicity of the industry implies that a sales associate will have approximately two opportunities to contact clients each season - in the beginning when the majority of the assortment arrives, and either in the middle of the season if updates are received, or close to the end when the retailer starts sale. We comment on other available control variables in §6.2.2.

To fully define the problem one also need to model the customer's response and behavior under different controls used by associates. There are multiple ways to obtain such a response model. Gönül and Shi [13] obtained client responses by directly solving each client's own utility maximization problem. In the initial stages of this work, an attempt was made to explicitly characterize the utility of a client, but the attempt was later shelved due to the intractability of the resulting model. We comment on this further in §6.2.3. An alternative method, employed by Simester et al. [15] and Bitran and Mondschein [12], involves using historical data to estimate transitional probabilities without imposing functional form assumptions. Such data is not readily available in the high-end retail industry, since an associate makes his own contact decisions using his client book and rarely keeps track of the policy and its impact on clients. Having exhausted the two common methods, we decide to take a different approach in modeling customer's responses in this work. The level of interaction between an experienced associate and his clients leads us to believe that such an associate could have a wealth of knowledge relating to customer responses. We attempt to harvest this by distilling conversations with experienced associates and managers such that the essence of client behavior is captured.

In this work we do not address competition directly, although one can argue that the effects of competition are embedded in the observed client responses. We also do not include the problem of client acquisition and assume that each associate already has a set of existing clients.

In the remainder of this chapter, we first formulate the simple case of an associate with a single client in §2.1. We then proceed to the realistic case of an associate with multiple clients in §2.2. Following this we present the model of a client in §2.3.

2.1 An Associate with a Single Client

In this simple case, we assume that the particular client remains in the associate's client book for $N-1$ periods. In the beginning of each period t, the client starts off in state i_t and the associate decides on a contact policy u_t. The client responds with a purchase decision and subsequently transitions into a new state j_{t+1} in the next period. The associate's problem can then be formulated as

$$V_t(i_t) = \max_{u_t \in \{0,1\}} \left[g(i_t, u_t) + \alpha \sum_j p_{iju_t} V_{t+1}(j_{t+1}) \right] \tag{2.1}$$

$$V_N(i_N) = 0 \tag{2.2}$$

where $V_t(i_t)$ denotes the optimal value function when the client is in state i_t, $u_t = 1$ denotes the associate contacting the client, $u_t = 0$ denotes the associate not contacting the client, α denotes the discount factor, and p_{iju_t} denotes the probability of the client transitioning from state i into state j under policy u_t. Since the client leaves the system after $N-1$ periods, the terminal cost is set to 0. The associate's single period expected reward is a function of his contact policy and the client's purchase decision, as expressed in Eq. (2.3).

$$g(i_t, u_t) = Rp_{ibu_t} - u_t c \tag{2.3}$$

where R is the revenue, p_{ibu_t} is the probability that receiver in state i in period t will make a purchase when policy u_t is adopted, and c is the cost associated with contacting the client, including the opportunity cost, cost of time and phone call. In this work we assume that the cost of contact is a constant. In reality, this cost varies by the day of the week and time of the day. Since an associate could potentially lose a walk-in sale when contacting clients, the cost of contact is highest when the store is busy.[1] We comment on this in §6.2.4.

2.2 An Associate with Multiple Clients

We now turn to the problem of an associate with multiple clients. In this more realistic case, the associate has a set of clients $\mathcal{S}$ (totaling $|\mathcal{S}|$) instead of a single client. In the beginning of each period, the set of clients start off in a vector of states $\mathbf{i_t}$, where $\mathbf{i_t} = [i_t^1, i_t^2, \cdots, i_t^{|\mathcal{S}|}]$. The associate makes a contact decision for each client, denoted by the vector

[1] In retail, the busiest days of the week are the weekends, while the busiest hours of the days vary depending on location.

$\mathbf{u_t} = [u_t^1, u_t^2, \cdots, u_t^{|S|}]$, where $u_t^s \in \{0,1\} \quad \forall \quad s \in \mathcal{S}$. Very often, the associate is limited in the number of clients that he can or will contact each period. This limitation arises naturally due to time constraints, or artificially as imposed by management[2] or inventory availability. The clients respond independently with their own purchase decisions and transition to a new vector of accessible states $\mathbf{j_{t+1}} = [j_{t+1}^i, j_{t+1}^2, \cdots, j_{t+1}^{|S|}]$. The problem can be formulated as follows, where $V_t(\mathbf{i_t})$ is the value function of the clients in states $\mathbf{i_t}$ in period t, M_t is the number of contacts to be made in that period and e is a column unit vector.

$$V_t(\mathbf{i_t}) = \max_{\mathbf{u_t}}\left[g(\mathbf{i_t}, \mathbf{u_t}) + \alpha \sum_{\mathbf{j_t}} p_{\mathbf{iju_t}} V_{t+1}(\mathbf{j_{t+1}})\right] \tag{2.4}$$

$$s.t. \qquad \mathbf{u_t}'e = M \tag{2.5}$$

$$V_N(\mathbf{i_N}) = 0 \tag{2.6}$$

2.3 Modeling the Client Behavior

We began this phase of the research by conducting interviews with associates and managers ([16], [17], [18], [19]) from high-end retail stores in Boston, including Hugo Boss, Hermès, Gucci and Barney's New York. The goals of the interviews are to understand how client work is currently performed, why certain methods are used, and to extract common knowledge about client behavior and purchase pattern. After distilling the conversations, we find that the experienced associates generally point to three factors that affect client behavior. The first factor, which we term the *purchase saturation effect*, is driven by the price points of the merchandise coupled with decreased loyalty among the clients. Associates point out that the majority of the clients get saturated and are not likely to purchase consecutively for many seasons. The second factor, which we term the *repeat purchase effect*, can be explained by the impact of freshness. Clients, either new or ones who have stopped buying for a couple of seasons, will return to the store and buy continuously for a couple of seasons in the beginning. The flurry of activity then dies down. The third factor, which

[2]Management sets targets periodically in an attempt to maintain some sort of service level to customers.

we term the *activation effect*, is less recognized than the other two. A small number of associates mentioned that if they call clients after having not seen them in the store for awhile, it is quite likely that these clients will end up making a purchase. Once the window of opportunity passes, it becomes more difficult to get these clients to purchase again.

In order to capture all three effects we choose to represent the state of a client i_t with a pair of variables (b_t, nr_t). b_t is defined as the total number of consecutive purchases by time t, while nr_t is defined as the total number of consecutive non-purchases by time t. Using this representation, the state transition of a client is depicted in Figure 2-1. Since the notion

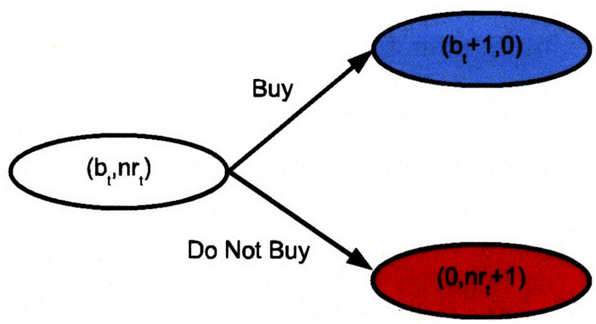

Figure 2-1: State transitions for a client at time t.

of "consecutive" is often confusing, we illustrate the state transitions with an example. A newly acquired client starts in the system with state $(1, 0)$. If the client makes a purchase in this period, he transitions into state $(2, 0)$. In the next period, if this client makes another purchase, he transitions into state $(3, 0)$. In the subsequent period, if this client does not make a purchase, he transitions into state $(0, 1)$. From this state, if he does not makes a purchase, he goes into state $(0, 2)$. In the next period, if the client makes a purchase, he returns to state $(1, 0)$. This process continues for a total of N periods, upon which the client will leave the system. Such a state representation, coupled with a time invariant transitional probability, assumes that a client is memoryless. In other words, we do not explicitly modeling the effect of the depth of a relationship. §6.2.1 provides suggestions to include this aspect into the model.

Our state definitions dictate that a particular client can either be in one of the buy states ($b_t \geq 1$) or one of the non-buy states ($nr_t \geq 1$) in any given period. Thus to capture the behavior, we look at two sets of transitional probabilities. The first set models the behavior of a client where $b_t \geq 1$, and the second models the behavior of a client where $nr_t \geq 1$. In Section 2.3.1, we propose transitional probabilities that capture the saturation effect. In Section 2.3.2, we propose transitional probabilities that capture all the aforementioned behavior. The transitional probabilities are stationary in both cases.

2.3.1 Monotonic Client Behavior

When a particular client exhibits monotonic behavior, we use monotonic transitional probabilities depicted in Figure 2-2 to capture his responses. A different client would exhibit the same monotonicity in his behavior, but the transitional probability would take on different values.

Figure 2-2(a) captures the behavior of this client in the buy states. Under monotonic behavior, the probability of purchase (thus transitioning into a higher buy state) is highest when the client has made one purchase and decreases monotonically to zero. Thus when a client has purchased many consecutive times, he is very unlikely to purchase again. In addition, when the client is being contacted ($u = 1$), his purchase response is uniformly higher than when he is not being contacted.

Figure 2-2(b) captures the behavior of this client in one of the non-buy states. Under monotonic behavior, the probability of non-purchase (thus transitioning into a higher non-purchase state) approaches one monotonically. Thus at very high non-purchase states, the client is very unlikely to purchase again. In addition, when the client is being contacted ($u = 1$), his non-purchase response is uniformly lower than when he is not being contacted.

2.3.2 Peaking Client Behavior

When a particular client exhibits all aforementioned behavior, we use transitional probabilities depicted in Figure 2-3 to capture his responses. Similarly, a different client would exhibit

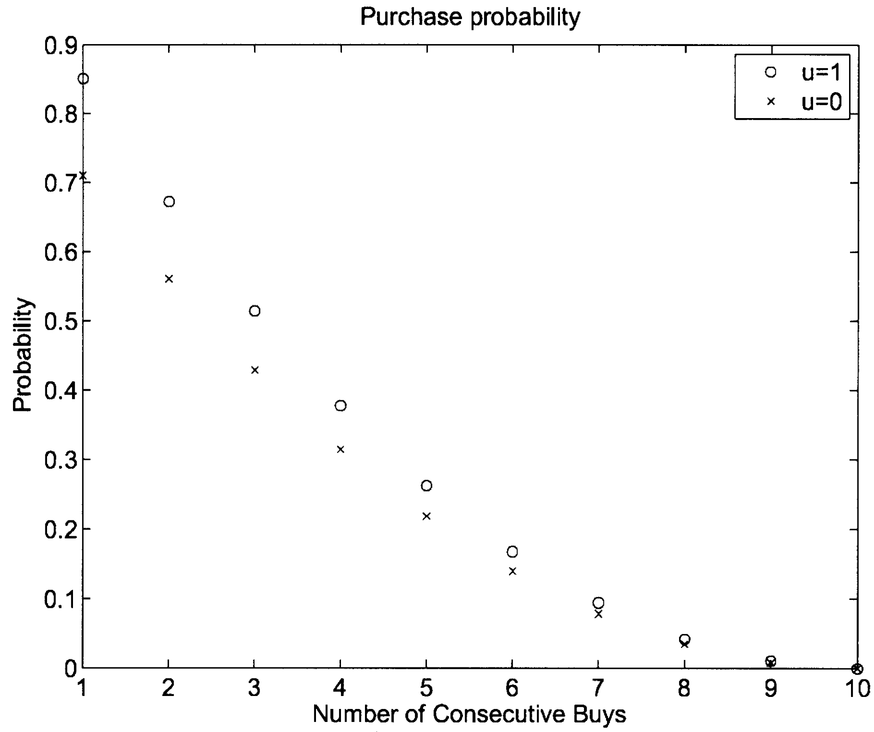

(a) Transitional probabilities for a client in buy states.

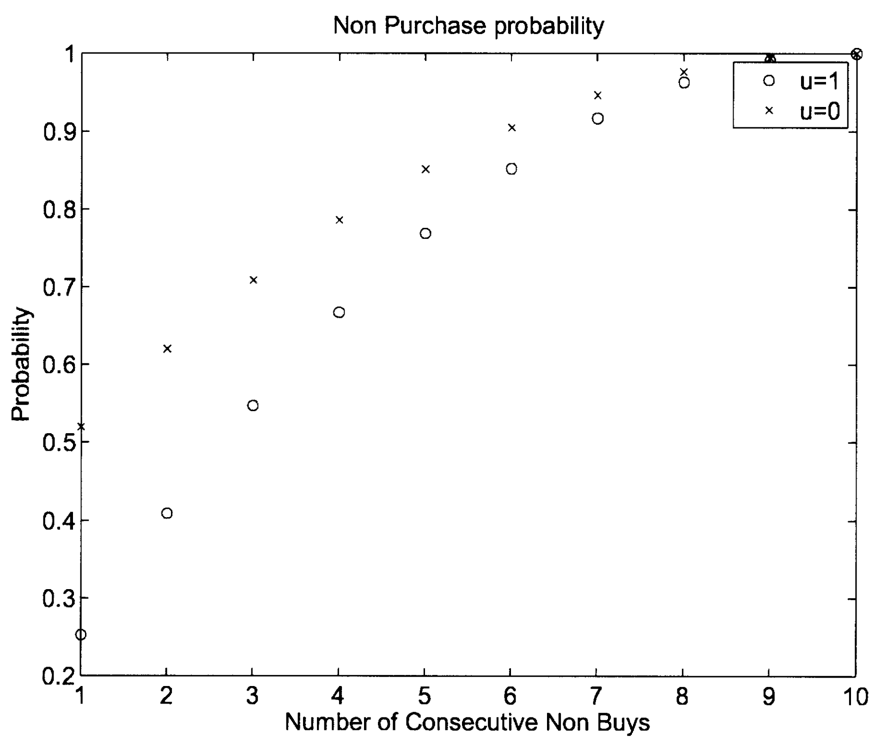

(b) Transitional probabilities for a client in non-buy states.

Figure 2-2: Transitional probabilities for a client who exhibits monotonic behavior.

the same peaking behavior, but with the peak shifted and the transitional probabilities taking on different values.

Figure 2-3(a) models this client in one of the buy states. In this figure, the probability of purchase is uniformly larger when the client is contacted ($u = 1$) than otherwise. The responses are not mononotic, but peaks when the number of consecutive buys equals to two. Note that the peak occurs under both policies. This is intentional, and is meant to capture the freshness effect, since a new client will tend to purchase consecutively in the beginning even without a contact.

Figure 2-3(b) models the same client in one of the non-buy states. The probability of non-purchase is uniformly larger when the client is not contacted than otherwise. When the client is not contacted ($u = 0$), the probability of non-purchase increases with the value of nr. On the other hand, when the client is contacted ($u = 1$), the non-purchase probability dips at $nr = 4$. In other words, the client is most likely to purchase again if contacted at $nr = 4$.

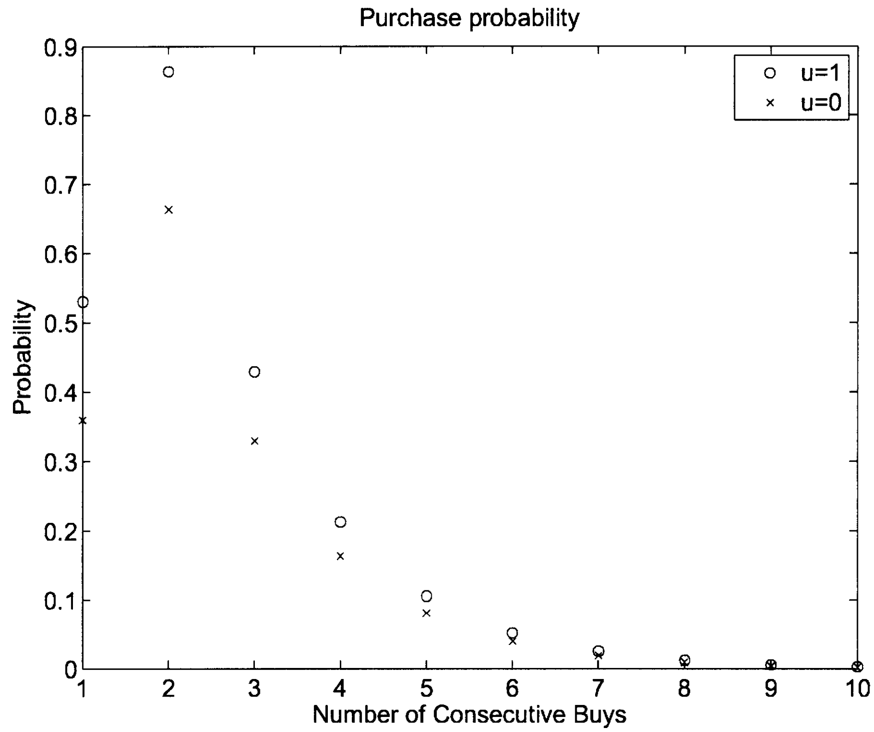

(a) Transitional probabilities for a client in buy states.

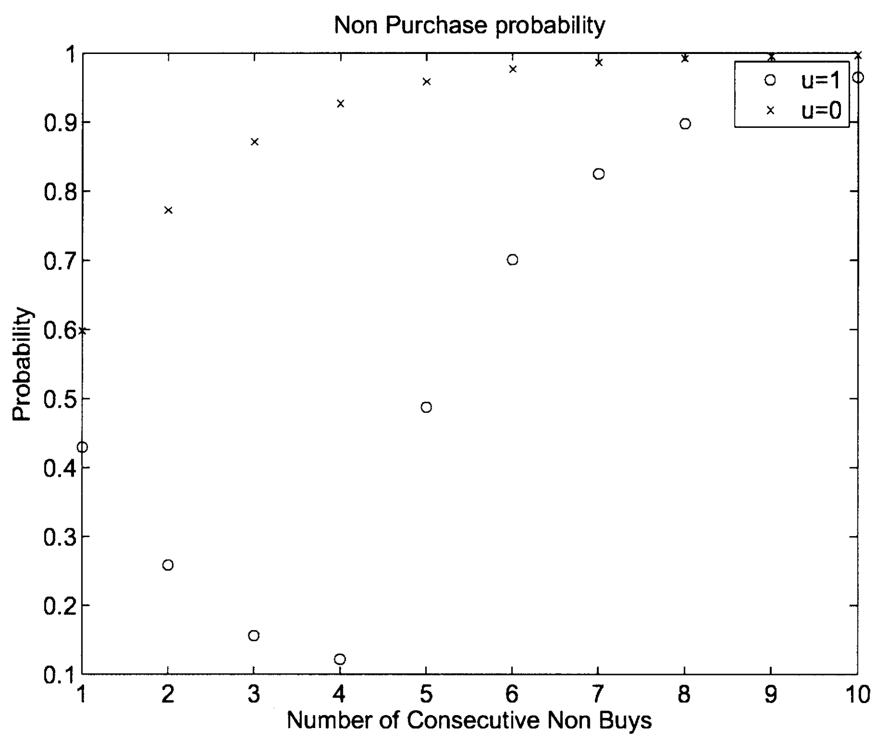

(b) Transitional probabilities for a client in non-buy states.

Figure 2-3: Transitional probabilities for a client who exhibits peaking behavior.

Chapter 3

A Single Client: Infinite Horizon Problem

In this chapter, we analyze the infinite horizon version of the problem formulated in Eq. (2.1)-Eq. (2.2). The motivation behind looking at the infinite horizon problem lies in its elegant and insightful analysis coupled with implementational simplicity.[1] Reformulated as an infinite horizon problem, the associate's problem is stated as in Eq. (3.1), where g_{iu} denotes the expected single stage cost in state i under policy u, and p_{iju} denotes the probability of transitioning from state i to j under policy u. Figure 3-1 depicts the admissible transitions for all states.

$$V(i) = \max_{u \in \{0,1\}} \left[g_{iu} + \alpha \sum_j p_{iju} V(j) \right] \tag{3.1}$$

Based on this formulation, we derive the structure of the optimal policy, where possible, for two types of client behaviors. The first behavior we look at is one where the client behaves monotonically, as illustrated in Figure 2-2. The second type of behavior we investigate is one where the client exhibits a peaking behavior, as illustrated in Figure 2-3. For each type of client behavior, we analyze the buy states and the non-buy states separately. Recall, by

[1] Optimal policy in infinite horizon problems are mostly stationary.

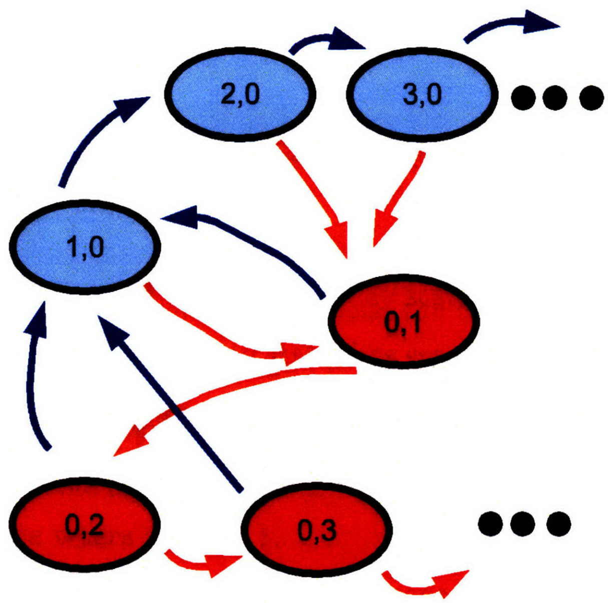

Figure 3-1: State transitions for the infinite horizon problem. The blue states represent consecutive buy states while the red states represent consecutive non-buy states.

our state definition, that a client in one of the buy states has bought at least one consecutive time while a client in one of the non-buy states has stopped buying for at least for one period.

When a client is in one of the buy states, the value of the state - denoted i - is uniquely determined by the value of b. As illustrated in Figure 3-2(a), upon a purchase, the client moves into a next state where $i = b + 1$. Otherwise, the client moves into the state $(0, 1)$. We thus use b (corresponding to the number of consecutive buys) to denote a unique buy state and $\bar{b}$ to denote the non-buy state $(0, 1)$. Figure 3-2(b) illustrates the revenue streams associated with the client actions using the simplified state notation.

Correspondingly, when a client is in one of the non-buy states, i is uniquely determined by the value of nr. As illustrated in Figure 3-2(c), upon a non-purchase, the client transitions into the next state where $i = nr + 1$. Otherwise, the client transitions into the state $(1, 0)$. We thus use nr (corresponding to the number of consecutive non-buys) to denote a unique non-buy state and b to denote the buy state $(1, 0)$. Figure 3-2(d) illustrates the revenue streams associated with the client actions using the simplified state notation.

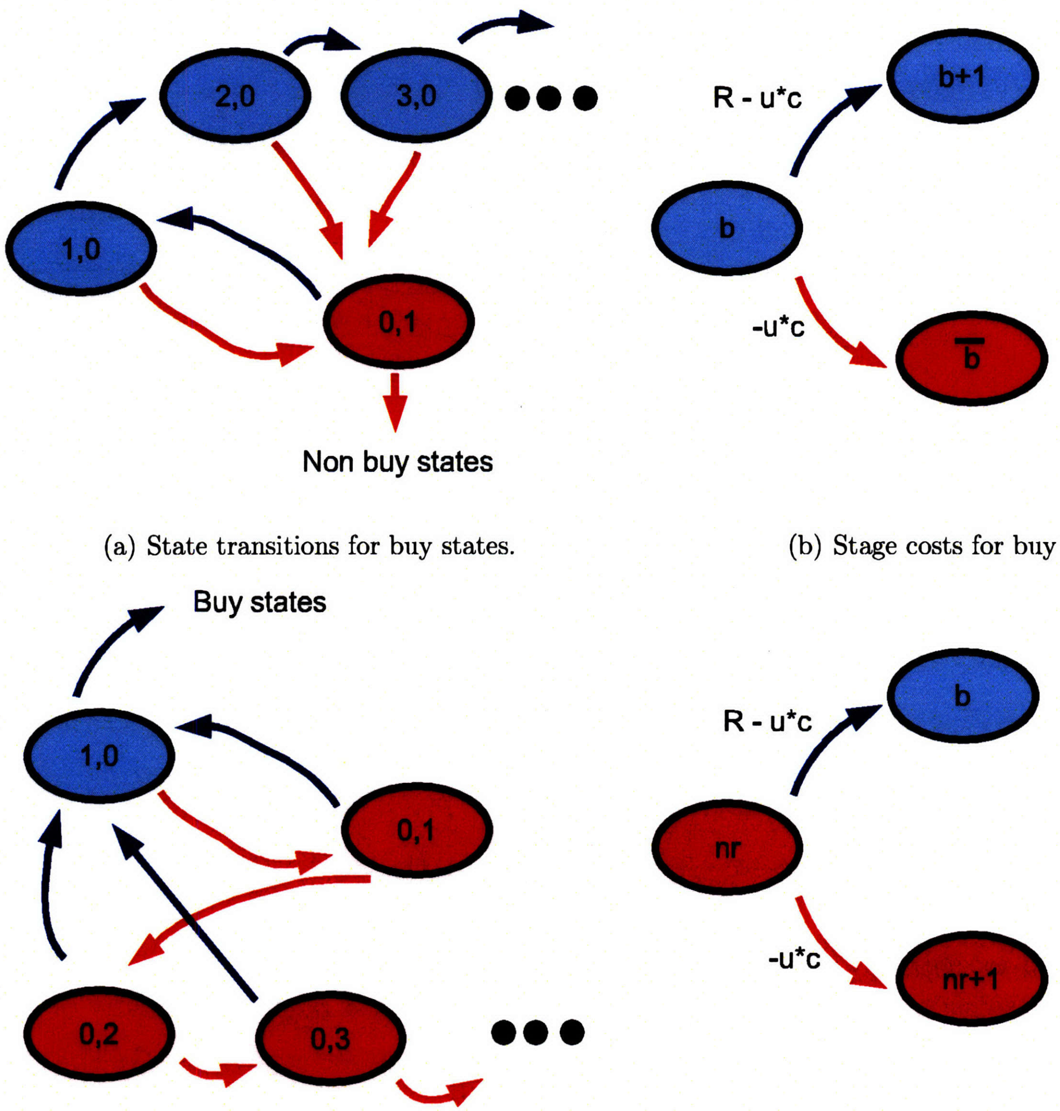

(a) State transitions for buy states.

(b) Stage costs for buy states.

(c) State transitions for non-buy states.

(d) Stage costs for non-buy states.

Figure 3-2: State transitions and costs.

The rest of the chapter is organized as follows. §3.1 and §3.2 focus on a monotonic client in the buy states and non-buy states respectively. §3.3 and §3.4 characterize the optimal policy for a non-monotonic client in the buy states and the non-buy states, respectively. We close this chapter by presenting computational results in §3.5.

3.1 Characterization of Optimal Policy for a Monotonic Client in Buy States

We begin by introducing the set of notation used in this section and by stating the assumptions invoked when a client is behaving monotonically.

$$
\begin{aligned}
p_u(b) &= \text{probability of transitioning from state } b \text{ into } b+1 \text{ under policy } u, \\
&\quad\ \text{where } u \in \{0,1\} \\
\Delta_{b,u} &= p_1(b) - p_0(b) \\
\bar{b} &\equiv \text{state (0,1)} \\
u^* &= \text{optimal policy}
\end{aligned}
$$

Using the above notation, the expected single stage cost g_{iu} is given by

$$
\begin{aligned}
g_{iu} &= (R - u \cdot c)p_u(b) + (-u \cdot c)(1 - p_u(b)) && (3.2) \\
&= Rp_u(b) - u \cdot c && (3.3)
\end{aligned}
$$

Assumption 1. *$p_1(b) \geq p_0(b) \quad \forall \quad b$ and $p_u(b)$ is convex and non-increasing.*

Assumption 2. *$p_1(b) - p_0(b)$ is monotonically non-increasing in b. In other words, the effect of a contact decreases as b increases.*

The first step in characterizing an optimal policy involves defining the notion of optimality. In other words, what is it that makes an associate better off when he contacts a client in a buy state? Since the associate has binary control, one would expect the optimality of the control to be determined by a ratio. Proposition 3.1.1 shows that such a ratio does exist.

Proposition 3.1.1. *If $\big(p_1(b) - p_0(b)\big) > \frac{c}{\big(R+\alpha V(b+1)-\alpha V(\bar{b})\big)}$, then $u^* = 1$ at b.*

Proof.

$$V(b)_{u=0} = Rp_0(b) + \alpha\big[p_0(b)V(b+1) + (1-p_0(b))V(\bar{b})\big]$$
$$V(b)_{u=1} = Rp_1(b) - c + \alpha\big[p_1(b)V(b+1) + (1-p_1(b))V(\bar{b})\big]$$

Since $u^* = 1$ when $V(b)_{u=1} > V(b)_{u=0}$, $u^* = 1$ when

$$Rp_1(b) - c + \alpha\big[p_1(b)V(b+1) + (1-p_1(b))V(\bar{b})\big] > Rp_0(b) + \alpha\big[p_0(b)V(b+1) + (1-p_0(b))V(\bar{b})\big]$$

Simplifying the expression gives

$$\big(R + \alpha V(b+1) - \alpha V(\bar{b})\big)\big(p_1(b) - p_0(b)\big) > c$$

□

The interpretation of Proposition 3.1.1 is rather intuitive. When the client is in state b, $R + \alpha V(b+1)$ is the resulting revenue stream if he makes a purchase, while $\alpha V(\bar{b})$ is the resulting revenue stream if he does not make a purchase. Since c is the cost of contact, the ratio $\frac{c}{\left(R+\alpha V(b+1)-\alpha V(\bar{b})\right)}$ can be interpreted as a weighted cost of contact. The weighted cost of contact is inversely proportional to the value function of the next buy state, $V(b+1)$. When $V(b+1)$ is high, the weighted cost of contact for the associate is low. On the other hand, a low $V(b+1)$ increases the weighted cost of contact for the associate. The term $p_1(b) - p_0(b)$ can be thought of as the impact of a contact when the client is in state b. Thus, when the impact is greater than the weighted cost at state b, it is optimal for an associate to contact the client.

Having defined the notion of optimality, we need to look at the behavior of the impact of a contact, $\Delta_{b,u}$, and the value function, $V(b)$, for different values of b. Assumption 2 states that $\Delta_{b,u}$ is monotonically non-increasing in b. It is then tempting to show that the expression

$\frac{c}{\left(R+\alpha V(b+1)-\alpha V(\bar{b})\right)}$ is monotonically non-decreasing in b, and that it approaches a limit.[2] Since R and $V(\bar{b})$ are constant in b, we can equivalently show that $V(b)$ is monotonically non-increasing in b, and that $V(b)$ approaches a limit as $b \to \infty$. The desired properties of $V(b)$ are shown in Propositions 3.1.2 and 3.1.3.

We now look at how the desired properties of $V(b)$ give rise to the existence of an optimal threshold policy. Figure 3-3 plots the two components that determine the optimality of a control. The dashed red line represents the impact of contact, $p_1(b) - p_0(b)$, while the solid black line represents the weighted cost of contact, $\frac{c}{\left(R+\alpha V(b+1)-\alpha V(\bar{b})\right)}$. It is apparent from the figure that the optimal policy for this particular set of curves is a threshold policy. It is also apparent that there will be cases where it is optimal to never contact a client. Propositions 3.1.4 and 3.1.5 establish the optimality of both, namely a never-contact policy and a threshold policy. In addition, we are also able to show in Proposition 3.1.6 that $V(b)$ is convex. This result will come in handy in subsequent analysis.

Proposition 3.1.2. $\lim_{b\to\infty} V(b) = \alpha V(\bar{b})$

Proof.

$$V(b) = \max\Big\{ Rp_0(b) + \alpha\big[p_0(b)V(b+1) + (1-p_0(b))V(\bar{b})\big],$$
$$Rp_1(b) - c + \alpha\big[p_1(b)V(b+1) + (1-p_1(b))V(\bar{b})\big] \Big\}$$

Since $\lim_{b\to\infty} p_0(b) = 0$ and $\lim_{b\to\infty} p_1(b) = 0$

$$\lim_{b\to\infty} V(b) = \max\Big\{ \alpha V(\bar{b}), -c + \alpha V(\bar{b}) \Big\}$$
$$= \alpha V(\bar{b})$$

□

[2]Such a combination will lead to the existence of an optimal threshold policy, as explained in the following paragraph.

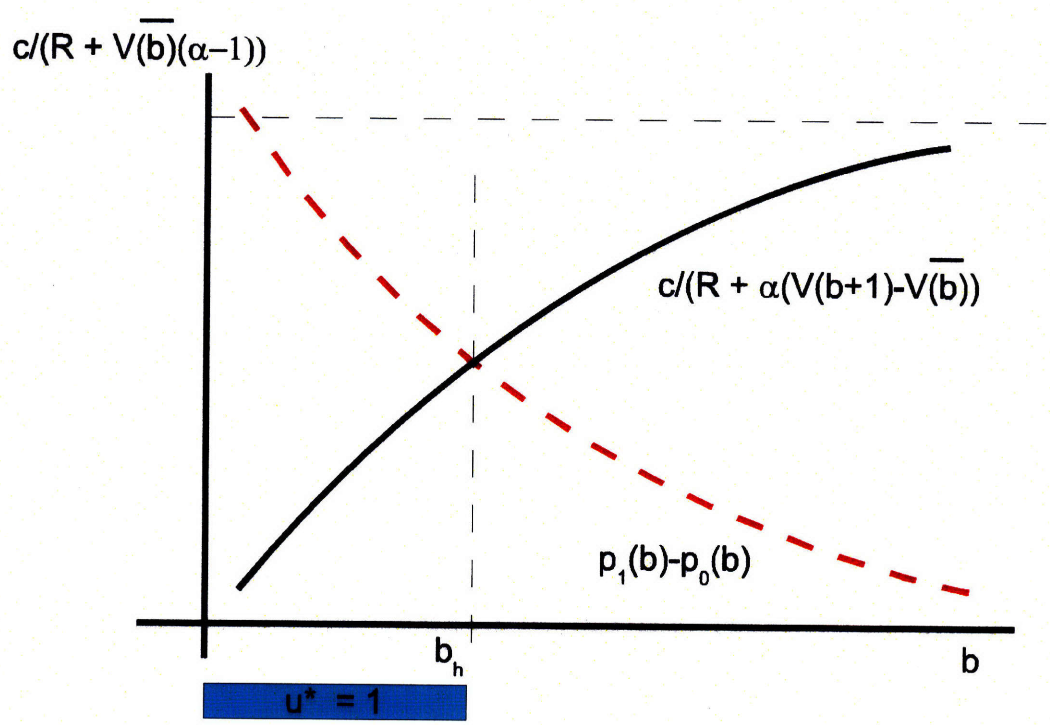

Figure 3-3: Characterization of optimal policy for the buy states.

Proposition 3.1.3. *$V(b)$ is monotonically non-increasing.*

Proof. Let $V_k(b)$ denote the value function of state b during iteration k. By the value iteration algorithm, $V(b) = \lim_{k->\infty} V_k(b)$ [20]. We assume, for induction, that $V_k(b) \geq V_k(b+1) \quad \forall \quad b$ and $R + \alpha V_k(b+2) \geq \alpha V_k(\bar{b})$. Then,

$$\begin{aligned} V_{k+1}(b) &= \max\Big\{ Rp_0(b) + \alpha\big[p_0(b)V_k(b+1) + (1-p_0(b))V_k(\bar{b})\big], \\ &\qquad Rp_1(b) - c + \alpha\big[p_1(b)V_k(b+1) + (1-p_1(b))V_k(\bar{b})\big] \Big\} \\ &= \max\Big\{ p_0(b)\big[R + \alpha V_k(b+1) - \alpha V_k(\bar{b})\big] + \alpha V_k(\bar{b}), \\ &\qquad p_1(b)\big[R + \alpha V_k(b+1) - \alpha V_k(\bar{b})\big] + \alpha V_k(\bar{b}) - c \Big\} \end{aligned}$$

By monotonicity of the transitional probabilities, $p_0(b) \geq p_0(b+1)$, and $p_1(b) \geq p_1(b+1)$

$$\begin{aligned} &\geq \max\Big\{ p_0(b+1)\big[R + \alpha V_k(b+1) - \alpha V_k(\bar{b})\big] + \alpha V_k(\bar{b}), \\ &\qquad p_1(b+1)\big[R + \alpha V_k(b+1) - \alpha V_k(\bar{b})\big] + \alpha V_k(\bar{b}) - c \Big\} \end{aligned}$$

By induction assumption, $V_k(b+1) \geq V_k(b+2)$

$$\begin{aligned} &\geq \max\Big\{ p_0(b+1)\big[R + \alpha V_k(b+2) - \alpha V_k(\bar{b})\big] + \alpha V_k(\bar{b}), \\ &\qquad p_1(b+1)\big[R + \alpha V_k(b+2) - \alpha V_k(\bar{b})\big] + \alpha V_k(\bar{b}) - c \Big\} \\ &= V_{k+1}(b+1) \end{aligned}$$

$\therefore V_{k+1}(b) \geq V_{k+1}(b+1)$. Since $V(b) = \lim_{k->\infty} V_k(b)$, we conclude that $V(b) \geq V(b+1)$, and that $V(b)$ is monotonically non-increasing in b. ☐

Proposition 3.1.4. *If $\frac{c}{R+\alpha V(2)-\alpha V(\bar{b})} > \Delta_{1,u}$, then it is never optimal to contact.*

Proof. $\frac{c}{R+\alpha V(b+1)-\alpha V(\bar{b})} \geq \frac{c}{R+\alpha V(2)-\alpha V(\bar{b})} \quad \forall \quad b \geq 1$ and $\Delta_{1,u} > \Delta_{b,u} \quad \forall \quad b \geq 1$.
$\therefore \frac{c}{R+\alpha V(b+1)-\alpha V(\bar{b})} > \Delta_{b,u} \quad \forall \quad b \geq 1$ and $u^* = 0 \quad \forall \quad b \geq 1$ ☐

Proposition 3.1.5. *If* $\exists$ b_h *such that* $u^* = 1$ *at* b_h *and* $u^* = 0$ *at* $b_h + 1$, *then* $u^* = 1 \;\; \forall \;\; b \leq b_h$ *and* $u^* = 0 \;\; \forall \;\; b \geq b_h + 1$

Proof. By optimality condition at b_h, we have $\frac{c}{R+\alpha V(b_h+1)-\alpha V(b)} < \Delta_{b_h,u}$. Further, $\forall \;\; b \leq b_h$, $\Delta_{b_h,u} \leq \Delta_{b,u}$ and $\frac{c}{R+\alpha V(b_h+1)-\alpha V(b)} \geq \frac{c}{R+\alpha V(b+1)-\alpha V(b)}$.
$\therefore \frac{c}{R+\alpha V(b+1)-\alpha V(b)} < \Delta_{b,u}$ and $u^* = 1 \;\; \forall \;\; b \leq b_h$.
Similarly, by optimality condition at $b_h + 1$, we have $\frac{c}{R+\alpha V(b_h+2)-\alpha V(b)} > \Delta_{b_h+1,u}$. Further, $\forall \;\; b \geq b_h + 1$, $\Delta_{b_h,u} \geq \Delta_{b,u}$ and $\frac{c}{R+\alpha V(b_h+2)-\alpha V(b)} \leq \frac{c}{R+\alpha V(b+1)-\alpha V(b)}$.
$\therefore \frac{c}{R+\alpha V(b+1)-\alpha V(b)} > \Delta_{b,u}$ and $u^* = 0 \;\; \forall \;\; b \geq b_h + 1$. □

Proposition 3.1.6. $V(b)$ *is convex*

Proof. The proof is rather lengthy, and relies on the fact that the optimal policy is either a never contact or a threshold policy. We separately show convexity in the different regions and at the boundary. Please refer to Appendix A.1 for details. □

3.2 Characterization of Optimal Policy for a Monotonic Client in Non-Buy States

Having characterized the optimal policy for the buy states, we now venture into the non-buy states. We begin by introducing the set of notation used in this section and by stating the assumptions invoked when a client is behaving monotonically in these states.

$$p_u(nr) = \text{probability of transitioning from state } nr \text{ into } nr+1 \text{ under policy } u, \text{ where } u \in \{0,1\}$$

$$\Delta_{nr,u} = p_0(nr) - p_1(nr)$$

$$b \equiv \text{state (1,0)}$$

Using the above notation, the expected single stage cost g_{iu} is given by

$$g_{iu} = (R - u \cdot c)(1 - p_u(nr)) + (-u \cdot c)(p_u(nr)) \tag{3.4}$$

$$= R(1 - p_u(nr)) - u \cdot c \tag{3.5}$$

Assumption 3. *$p_0(nr) \geq p_1(nr) \quad \forall \quad nr$ and $p_u(nr)$ is concave and non decreasing.*

Assumption 4. *$p_0(nr) - p_1(nr)$ is monotonically non-increasing in nr. In other words, the effect of a contact decreases as nr increases.*

We follow the flow in §3.1 by first defining the notion of optimality for the non-buy states. We expect, due to the binary nature of an associate's available control, that such notion will be determined by a ratio. Proposition 3.2.1 confirms our intuition.

Proposition 3.2.1. *If $\big(p_0(nr) - p_1(nr)\big) > \frac{c}{\big(R + \alpha V(b) - \alpha V(nr+1)\big)}$, then $u^* = 1$ at nr.*

Proof.

$$
\begin{aligned}
V(nr)_{u=0} &= R(1-p_0(nr)) + \alpha\big[p_0(nr)V(nr+1) + (1-p_0(nr))V(b)\big] \\
&= -p_0(nr)\big[R + \alpha V(b) - \alpha V(nr+1)\big] + \alpha V(b) + R \\
V(nr)_{u=1} &= R(1-p_1(nr)) - c + \alpha\big[p_1(nr)V(nr+1) + (1-p_1(nr))V(b)\big] \\
&= -p_1(nr)\big[R + \alpha V(b) - \alpha V(nr+1)\big] + \alpha V(b) + R - c
\end{aligned}
$$

Since $u^* = 1$ when $V(nr)_{u=1} > V(nr)_{u=0}$, $u^* = 1$ when

$$
\begin{aligned}
-p_1(nr)\big[R + \alpha V(b) - \alpha V(nr+1)\big] - c > \\
- p_0(nr)\big[R + \alpha V(b) - \alpha V(nr+1)\big]
\end{aligned}
$$

Rearranging terms,

$$
c < \big[R + \alpha V(b) - \alpha V(nr+1)\big]\big(p_0(nr) - p_1(nr)\big)
$$

□

Proposition 3.2.1 has the same form as Proposition 3.1.1, and thus offers the same intuitive interpretation. When the client is in state nr, $R + \alpha V(b)$ is the resulting revenue stream if he makes a purchase, while $\alpha V(nr+1)$ is the resulting revenue stream if he does not make a purchase. Since c is the cost of contact, the ratio $\frac{c}{\left(R+\alpha V(b)-\alpha V(nr+1)\right)}$ can be interpreted as a weighted cost of contact. This term, unlike when the client is in one of the buy states, is proportional to the value function of the next non-buy state, $V(nr+1)$. When $V(nr+1)$ is high, the weighted cost of contact is high. On the other hand, when $V(nr+1)$ is low, the weighted cost of contact is low as well. Since a contact increases the probability of purchase and causes the client not to go into the next non-buy state, the weighted cost of contact is high when the value function of the next non-buy state is high. $p_0(nr) - p_1(nr)$ can be thought of as the impact of a contact when the client is in state nr. Thus, when the impact is greater than the weighted cost at state nr, it is optimal for an associate to contact the client.

Since the expression for optimality hinges on the impact of a contact, $\Delta_{nr,u}$, and the value function, $V(nr)$, we proceed to look at how these two terms behave for different values of nr. Assumption 4 states that $\Delta_{nr,u}$ is monotonically non increasing in nr. Intuitively, $V(nr)$ should be non-increasing in nr as well. Recall that nr represents the number of consecutive non-buys. As such, a high value of nr indicates that the client has not purchased for many consecutive periods. It thus seems reasonable that the value function associated with such a client will be non increasing in nr, which we show in Proposition 3.2.2.

Proposition 3.2.2. *$V(nr)$ is monotonically non-increasing.*

Proof. Let $V_k(nr)$ denote the value function of state nr during iteration k. By the value iteration algorithm, $V(nr) = \lim_{k->\infty} V_k(nr)$. We assume, for induction, that $V_k(nr) \geq V_k(nr+1) \quad \forall \quad nr$ and $\alpha V_k(nr+1) - \alpha V_k(b) - R \leq 0$. Then,

$$\begin{aligned}
V_{k+1}(nr) = \max\Big\{ & R(1-p_0(nr)) + \alpha\big[p_0(nr)V_k(nr+1) + (1-p_0(nr))V_k(b)\big], \\
& R(1-p_1(nr)) - c + \alpha\big[p_1(nr)V_k(nr+1) + (1-p_1(nr))V_k(b)\big]\Big\} \\
= \max\Big\{ & -p_0(nr)\big[-\alpha V_k(nr+1) + \alpha V_k(b) + R\big] + \alpha V_k(b) + R, \\
& - p_1(nr)\big[-\alpha V_k(nr+1) + \alpha V_k(b) + R\big] + \alpha V_k(b) + R - c\Big\}
\end{aligned}$$

By monotonicity of the transitional probabilities and induction assumption

$$\begin{aligned}
p_0(nr) &\leq p_0(nr+1) \\
p_1(nr) &\leq p_1(nr+1) \\
-\alpha V_k(nr+1) + \alpha V_k(b) + R &< -\alpha V_k(nr+2) + \alpha V_k(b) + R \\
p_0(nr)\big[-\alpha V_k(nr+1) + \alpha V_k(b) + R\big] &\leq p_0(nr+1)\big[-\alpha V_k(nr+2) + \alpha V_k(b) + R\big] \\
p_1(nr)\big[-\alpha V_k(nr+1) + \alpha V_k(b) + R\big] &\leq p_1(nr+1)\big[-\alpha V_k(nr+2) + \alpha V_k(b) + R\big] \\
-p_0(nr)\big[-\alpha V_k(nr+1) + \alpha V_k(b) + R\big] &\geq -p_0(nr+1)\big[-\alpha V_k(nr+2) + \alpha V_k(b) + R\big] \\
-p_1(nr)\big[-\alpha V_k(nr+1) + \alpha V_k(b) + R\big] &\geq -p_1(nr+1)\big[-\alpha V_k(nr+2) + \alpha V_k(b) + R\big]
\end{aligned}$$

We then have the following inequality

$$\begin{aligned}
V_{k+1}(nr) \geq \max\Big\{ & p_0(nr+1)\big[\alpha V_k(nr+2) - \alpha V_k(b) - R\big] + \alpha V_k(b) + R, \\
& p_1(nr+1)\big[\alpha V_k(nr+2) - \alpha V_k(b) - R\big] + \alpha V_k(b) + R - c\Big\} \\
= V_{k+1}(nr+1) &
\end{aligned}$$

$\therefore V_{k+1}(nr) \geq V_{k+1}(nr+1)$. Since $V(nr) = \lim_{k->\infty} V_k(nr)$, we conclude that $V(nr) \geq V(nr+1)$, and that $V(nr)$ is monotonically non-increasing in nr. □

We now look at what the two components of Proposition 3.2.1 reveal about the form of optimal policy. Unlike the case where the client is in the buy state (§3.1), the ratio that determines optimality for the non-buy states, $\frac{c}{(R+\alpha V(b)-\alpha V(nr+1))}$, is monotonically non-increasing in nr due to the minus sign preceding the term $\alpha V(nr+1)$. Coupled with the behavior of $\Delta_{nr,u}$, the form of optimal policy is not apparent since both terms are monotonically non increasing in nr. Even so, there are two conjectures regarding optimal policy that can be made. One would guess that it is not optimal for an associate to contact a client if such client is in a state with very high nr. In addition, if the largest impact of contact is less than the smallest weighted cost, one would guess that it is never optimal to contact such a client. We show that these two conjectures are true in Propositions 3.2.3 and 3.2.4. Out of curiosity, in §3.2.1 we investigate the conditions under which an optimal threshold policy will result.

Proposition 3.2.3. *As $nr \to \infty$, $u^* \to 0$.*

Proof. As $nr \to \infty$, both $p_1(nr)$ and $p_0(nr)$ approach 1.

$$V(nr) = \max\{-p_0(nr)\big[R + \alpha V(b) - \alpha V(nr+1)\big] + \alpha V(b) + R \\ - p_1(nr)\big[R + \alpha V(b) - \alpha V(nr+1)\big] + \alpha V(b) + R - c\}$$

At large nr,

$$V(nr) = \max\{\alpha V(nr+1), \alpha V(nr+1) - c\}$$

$\therefore u^* = 0$ at large nr. □

Proposition 3.2.4. *If* $\min_{nr}\left(\frac{c}{(R+\alpha V(b)-\alpha V(nr+1))}\right) > \Delta_{1,u}$, $u^* = 0 \quad \forall \quad nr$

Proof. Since both $V(nr)$ and $\Delta_{nr,u}$ are monotonically non-increasing, we have $\forall \quad nr$

$$\frac{c}{\big(R + \alpha V(b) - \alpha V(nr+1)\big)} > \min_{nr} \frac{c}{\big(R + \alpha V(b) - \alpha V(nr+1)\big)}$$
$$\max_{nr} \Delta_{nr,u} = \Delta_{1,u} > \Delta_{nr,u}$$

□

3.2.1 On the Optimality of a Threshold Policy

Previously, we have shown that both $\Delta_{nr,u}$ and $\frac{c}{\big(R+\alpha V(b)-\alpha V(nr+1)\big)}$ are monotonically non-increasing. In addition, we have also shown that the optimal policy, as $nr \to \infty$, is to not contact. Thus a threshold policy is optimal if i) contacting a client in the first non-buy state (when $nr = 1$) is known to be optimal *and* ii) $V(nr)$ is convex. Under these two conditions, the terms $\Delta_{nr,u}$ and $\frac{c}{\big(R+\alpha V(b)-\alpha V(nr+1)\big)}$ will only crossover once, as illustrated by the red dashed line and black solid line respectively in Figure 3-4. We formalize this in Lemma 3.2.1.

Lemma 3.2.1. *If $u^* = 1$ at $nr = 1$ and $V(nr)$ is convex, then a threshold policy exists and it is optimal.*

Proof. We show by contradiction. Denote $h(nr) = \frac{c}{\big(R+\alpha V(b)-\alpha V(nr+1)\big)}$, $h(nr_\infty)$ the value of $h(nr)$ as $nr \to \infty$ and $\Delta_{nr_\infty,u}$ the value of $\Delta_{nr,u}$ as $nr \to \infty$. If there is no crossover (i.e,. threshold policy does not exist), $\Delta_{nr_\infty,k} \geq h(nr_\infty)$, which is a contradiction since at $\Delta_{nr_\infty,k} < h(nr_\infty)$ by Proposition 3.2.3. □

Motivated by Lemma 3.2.1, we now spend some time to establish the convexity of $V(nr)$. Unlike in the buy states, the optimal policy is not proven to be a threshold policy. As such, we cannot separate $V(nr)$ into two regions and show convexity for both regions and for the boundary. A more general approach for showing convexity of $V(nr)$ is needed. For exposition clarity we make a minor notation modification here and explicitly denote $V^*(nr)$ as the optimal value function for a client in state nr. We first define $\mathcal{V}$ to be a class of all

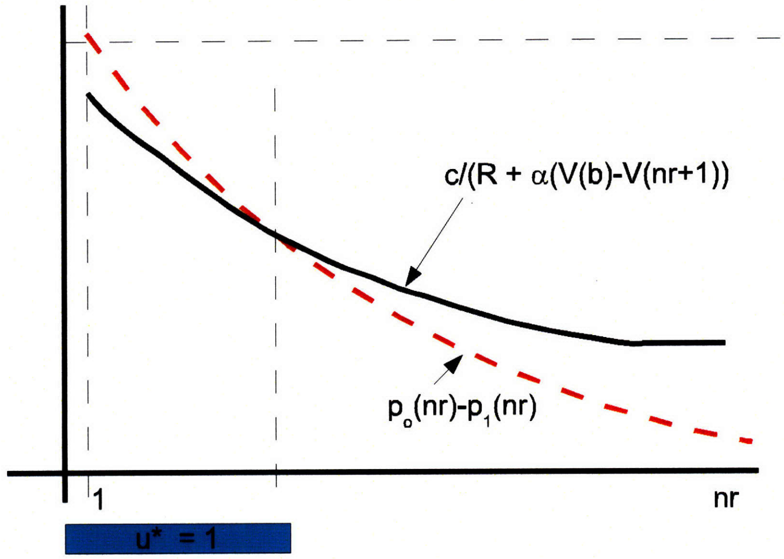

Figure 3-4: Optimal policy for non-buy states.

functions $V(nr)$ which are non-increasing and convex. Next, under the additional assumption that $p_u(nr)$ is such that $\bar{V}(nr) = -p_u(nr)\big[R + \alpha V(b) - \alpha V(nr+1)\big] \in \mathcal{V}$ if $V(nr+1) \in \mathcal{V}$, Lemma 3.2.2 establishes the convexity of $V^*(nr)$.

Lemma 3.2.2. $V^*(nr) \in \mathcal{V}$.

Proof. We show by value iteration. Define $V_0(nr) = R+\alpha V(b)-p_0(nr)\big[R+\alpha V(b)-\alpha V(nr+1)\big]$ and $V_1(nr) = R+\alpha V(b)-c-p_1(nr)\big[R+\alpha V(b)-\alpha V(nr+1)\big]$. Since convexity is closed under maximization, we need to show that $V_0(nr)$ and $V_1(nr)$ both $\in \mathcal{V}$.

We assume for induction that $V_k(nr+1)$ is convex and monotonically non-increasing and express $V_{k+1}(nr)$ as

$$V_{k+1}(nr) = \max\big\{R + \alpha V(b) - p_0(nr)\big[R + \alpha V(b) - \alpha V_k(nr+1)\big],$$
$$R + \alpha V(b) - c - p_1(nr)\big[R + \alpha V(b) - \alpha V_k(nr+1)\big]\big\}$$

Since $V_k(nr)$ is convex and monotonically non-increasing, $-p_0(nr)\big[R + \alpha V(b) - \alpha V_k(nr+1)\big]$ and $-p_1(nr)\big[R + \alpha V(b) - \alpha V_k(nr+1)\big]$ are both monotonically non-increasing and convex by assumption. Since convexity is closed under addition, $V_{k+1}(nr) \in \mathcal{V}$. Thus starting from any $V_0 \in \mathcal{V}$, we have $V_{k+1}(nr)$ convex and monotonically non-increasing. Since $\lim_{k\to\infty} V_{k+1}(nr) = V^*(nr)$, $V^*(nr)$ is monotonically non-increasing and convex. □

The assumption required for Lemma 3.2.2 to hold is quite strong and it is not obviously satisfied by any $p_u(nr)$. Driven by our inability to show that $V(nr)$ is convex, we turn to computational experiments. The computational results show that given the concavity in the response probabilities, $V(nr)$ is indeed not always convex. As a result, the optimal policy is also not guaranteed to be a threshold policy. Figure 3-5, in which the responses are modeled with a quadratic function, illustrates one such example.

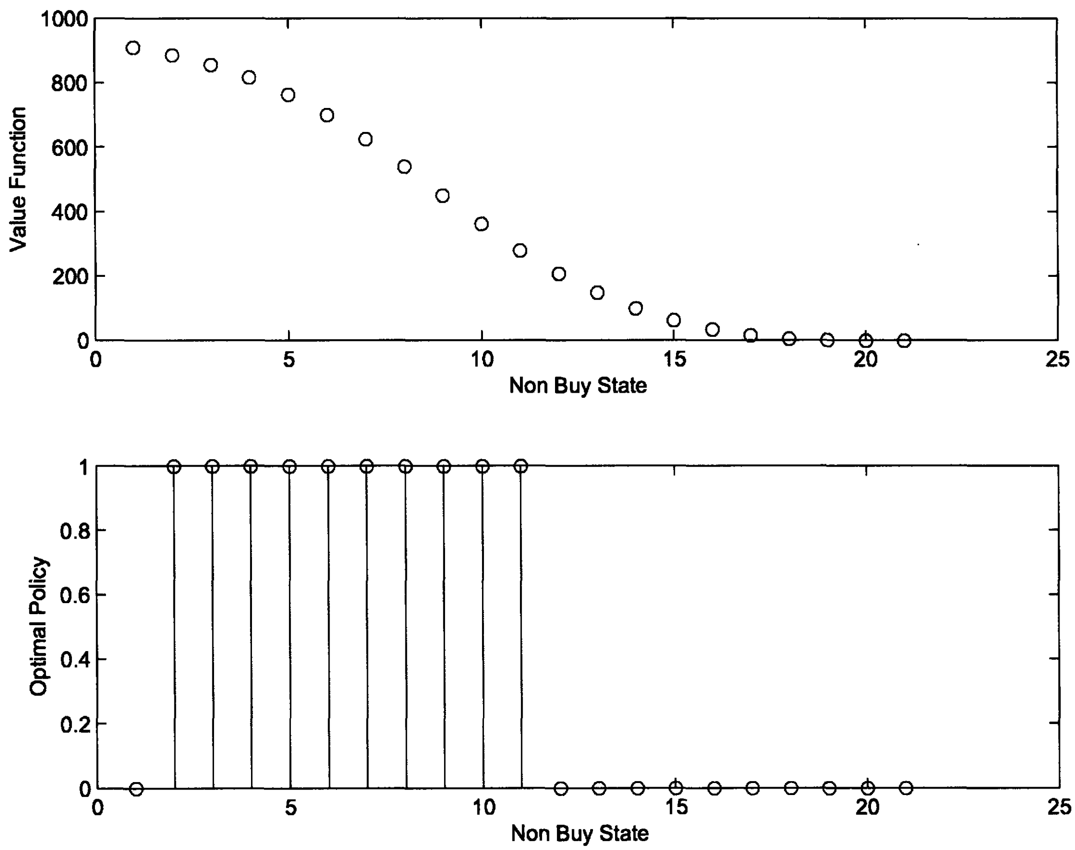

Figure 3-5: Example of a non-threshold optimal policy.

3.3 Characterization of Optimal Policy for a Non-Monotonic Client in Buy States

In §3.1 the optimal policy for an associate when the client behaves monotonically in the buy states is shown to be either a threshold policy or a never-contact policy. In this section, we look at how a non-monotonic client behavior affects the form of optimal policy. We begin by introducing the set of notations used and by stating the assumptions invoked when a client exhibits a peaking behavior in the buy states.

$$
\begin{aligned}
p_u(b) &= \text{probability of transitioning from state } b \text{ into } b+1 \text{ under policy } u, \\
&\quad \text{where } u \in \{0,1\} \\
\Delta_{b,u} &= p_1(b) - p_0(b) \\
\hat{b} &= \arg\max_b p_u(b) \\
\tilde{b} &= \arg\max_b V(b) \\
\bar{b} &\equiv \text{state } (0,1) \\
V(b)_u &= \text{value function in state b under the assumption that } u \text{ is optimal} \\
\mathcal{B}_c &= \text{set of buy states where } \tilde{b}+2 \leq b \leq \hat{b}+1
\end{aligned}
$$

Assumption 5. *$p_1(b)$ and $p_0(b)$ both exhibit peaking behaviors with a maximum at $\hat{b}$.*

Assumption 6. *$p_1(b) - p_0(b)$ is maximum at $\hat{b}$. In addition, $p_1(b) - p_0(b)$ is concave, non-decreasing $\forall \quad 1 \leq b \leq \hat{b}$ and non-increasing for $b \geq \hat{b}$. In other words, the effect of a contact increases towards the peak at $\hat{b}$ before decreasing with b.*

Since the optimality condition in Proposition 3.1.1 remains unchanged and $\Delta_{b,u}$ is dictated by Assumption 6, we proceed to look at the behavior of $V(b)$ for different values of b. Referring back to the transitional probabilities depicted in Figure 2-3(a), one would guess that due to the peaking in the transitional probabilities, it is unlikely that $V(b)$ will always be monotonically non-increasing. If so, will it be unimodal, or will it have multiple peaks?

We answer such questions in §3.3.1.

3.3.1 Property of the Value Function $V(b)$

We begin by noting that $\Delta_{b,u}$ behaves identically to the monotonic case for all values of $b \geq \hat{b}$. As such, we know that $V(b)$ will be monotonically non-increasing for $b \geq \hat{b}$. We thus focus on the region of $V(b)$ where $1 \leq b \leq \hat{b} - 1$. To establish unimodality of $V(b)$, we need to show that if there is a peak $\tilde{b}$, where $1 \leq \tilde{b} \leq \hat{b}$, then $V(b)$ is less than $V(\tilde{b}) \quad \forall \quad b \leq \tilde{b}$. When $\tilde{b} = 1$, $V(b)$ reduces to a monotonically non-increasing function. When $\tilde{b} = 2$, $V(b)$ can only be either monotonically non-increasing or unimodal. Thus we are only concerned about the case where the peak occurs anywhere in between $[3, \hat{b}]$. Lemma 3.3.1 shows that in that range, the value function $V(b)$ will be monotonically non-decreasing in b for $1 \leq b \leq \tilde{b}$. With that, Proposition 3.3.1 shows that $V(b)$ is unimodal.

Lemma 3.3.1. *If $V(b) \geq V(b-1)$, then $V(b-1) \geq V(b-2) \quad \forall \quad 3 \leq b \leq \hat{b}$.*

Proof.

$$\begin{aligned}
V(b-1) &= \max\{p_0(b-1)[R + \alpha V(b) - \alpha V(\bar{b})] + \alpha V(\bar{b}), \\
&\qquad\quad p_1(b-1)[R + \alpha V(b) - \alpha V(\bar{b})] + \alpha V(\bar{b}) - c\} \\
&\geq \max\{p_o(b-2)[R + \alpha V(b) - \alpha V(\bar{b})] + \alpha V(\bar{b}), \\
&\qquad\quad p_1(b-2)[R + \alpha V(b) - \alpha V(\bar{b})] + \alpha V(\bar{b}) - c\} \\
&\geq \max\{p_o(b-2)[R + \alpha V(b-1) - \alpha V(\bar{b})] + \alpha V(\bar{b}), \\
&\qquad\quad p_1(b-2)[R + \alpha V(b-1) - \alpha V(\bar{b})] + \alpha V(\bar{b}) - c\} \\
&= V(b-2)
\end{aligned}$$

The first inequality is true because $p_1(b) \geq p_1(b-1)$ and $p_0(b) \geq p_0(b-1) \quad \forall \quad b \leq \hat{b}$. The second inequality is true because $V(b) \geq V(b-1)$. □

Proposition 3.3.1. *$V(b)$ is unimodal in b with $\tilde{b} = \arg\max_b V(b)$, where $1 \leq \tilde{b} \leq \hat{b}$.*

Proof. We show that $V(b)$ is unimodal for i) $\tilde{b} = 1$, ii) $\tilde{b} = 2$ and iii) $3 \leq \tilde{b} \leq \hat{b}$.
$\tilde{b} = 1$ is trivial, $V(b)$ reduces to a monotonically non-increasing function. $\tilde{b} = 2$ is trivial as well. Lemma 3.3.1 takes care of case (iii). Starting from a peak value $\tilde{b}$, the value function preceding it will be monotonically non-decreasing in b. □

Knowing that $V(b)$ is unimodal, we now turn to look at when the peak $\tilde{b}$ occurs. Intuitively, the value function exhibits such peaking behavior due to the properties of the transitional probabilities. Take any triplet $V(b-1), V(b), V(b+1)$, where $V(b) \geq V(b+1)$ and $p_u(b-1) \leq p_u(b)$. The sufficient conditions for a peak to occur at b are

$$\frac{p_o(b-1)}{p_0(b)} \leq \frac{R + \alpha V(b+1) - \alpha V(\bar{b})}{R + \alpha V(b) - \alpha V(\bar{b})} \tag{3.6}$$

$$\frac{p_1(b-1)}{p_1(b)} \leq \frac{R + \alpha V(b+1) - \alpha V(\bar{b})}{R + \alpha V(b) - \alpha V(\bar{b})} \tag{3.7}$$

Thus when the ratio $\frac{p_u(b-1)}{p_u(b)}$ is small enough (caused by a relatively significant reduction in $p_u(b-1)$), there will be a peak at b. We analyze the necessary conditions for such peaking to occur in Appendix B.

It turns out that the unimodality of $V(b)$ greatly assists, but is not sufficient, in characterizing the optimal policies. For reasons that will become apparent in §3.3.2, we show that $V(b)$ is concave and non-increasing in $[\tilde{b}, \hat{b}+1]$. In Lemmas 3.3.2 to 3.3.5, we first show that a triplet of value functions always satisfy concavity when all states in the triplet belong to $\mathcal{B}_c$ regardless of the policy. The first triplet occurs when $b = \hat{b} + 1$, as shown by the blue boxes in Figure 3-6. As we move one state to the left, illustrated by the green boxes in the same figure, all possible states are also concave by the same set of lemmas. Thus $V(b)$ is concave in $\mathcal{B}_c$, as stated in Lemma 3.3.6.

Lemma 3.3.2. $V(b-2)_{u=1} - V(b-1)_{u=1} \leq V(b-1)_{u=1} - V(b) \quad \forall \quad \tilde{b}+2 \leq b \leq \hat{b}+1$

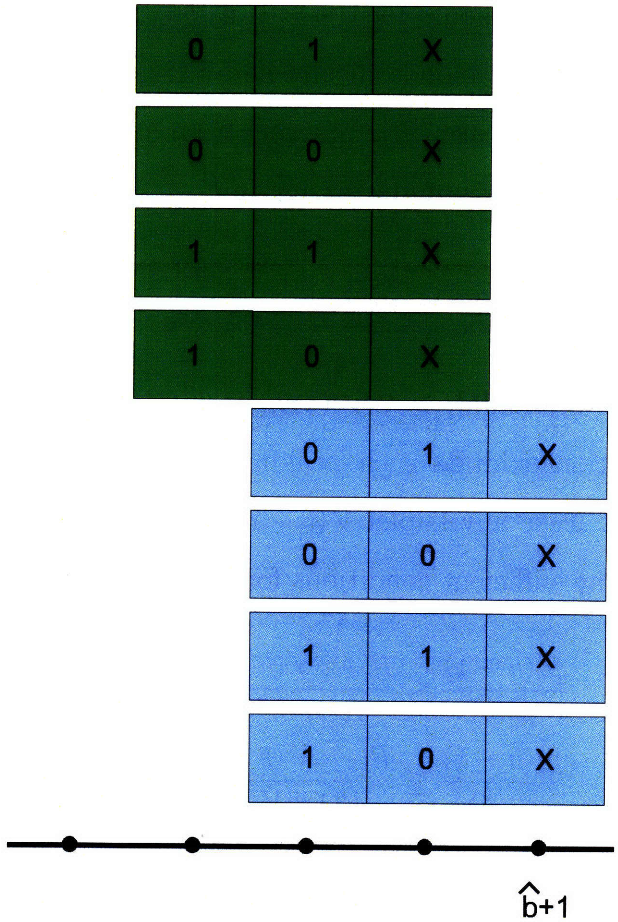

Figure 3-6: Partial concavity of $V(b)$.

Proof. Let $p_1(b-2) + \delta_1 = p_1(b-1)$. Then $\delta_1 \geq 0 \quad \forall \quad b \in \mathcal{B}_c$.

$$V(b-2)_{u=1} - V(b-1)_{u=1} = p_1(b-2)[R + \alpha V(b-1)_{u=1} - \alpha V(\bar{b})] - p_1(b-1)[R + \alpha V(b) - \alpha V(\bar{b})]$$

Substituting $p_1(b-2)$,

$$\begin{aligned}
&= (p_1(b-1)-\delta_1)[R+\alpha V(b-1)_{u=1}-\alpha V(\bar{b})]- \\
&\quad p_1(b-1)[R+\alpha V(b)-\alpha V(\bar{b})] \\
&= p_1(b-1)[\alpha V(b-1)_{u=1}-\alpha V(b)] - \delta_1[R+\alpha V(b-1)_{u=1}-\alpha V(\bar{b})] \\
&\leq p_1(b-1)[\alpha V(b-1)_{u=1}-\alpha V(b)] \\
&\leq V(b-1)_{u=1}-V(b)
\end{aligned}$$

The first inequality is true because $\delta_1[R+\alpha V(b-1)_{u=1}-\alpha V(\bar{b})] \geq 0$, the second inequality true because $p_1(b-1) \leq 1$ and $\alpha \leq 1$. □

Lemma 3.3.3. $V(b-2)_{u=1}-V(b-1)_{u=0} \leq V(b-1)_{u=0}-V(b) \quad \forall \quad \tilde{b}+2 \leq b \leq \hat{b}+1$

Proof.

$$\begin{aligned}
V(b-2)_{u=1}-V(b-1)_{u=0} &= p_1(b-2)[R+\alpha V(b-1)_{u=0}-\alpha V(\bar{b})]-c \\
&\quad - p_0(b-1)[R+\alpha V(b)-\alpha V(\bar{b})]
\end{aligned}$$

Since by assumption, $u^* = 0$ at $b-1$, we have $\frac{c}{R+\alpha V(b)-\alpha V(\bar{b})} > p_1(b-1)-p_0(b-1)$. In addition, let $p_1(b-2)+\delta_1 = p_1(b-1)$. Then $\delta_1 \geq 0 \quad \forall \quad b \in \mathcal{B}_c$.

$$\begin{aligned}
\therefore\ & p_1(b-2)[R+\alpha V(b-1)_{u=0}-\alpha V(\bar{b})]-c-p_0(b-1)[R+\alpha V(b)-\alpha V(\bar{b})] \\
&< p_1(b-2)[R+\alpha V(b-1)_{u=0}-\alpha V(\bar{b})]-p_1(b-1)[R+\alpha V(b)-\alpha V(\bar{b})] \\
&= (p_1(b-1)-\delta_1)[R+\alpha V(b-1)_{u=0}-\alpha V(\bar{b})]-p_1(b-1)[R+\alpha V(b)-\alpha V(\bar{b})] \\
&\leq p_1(b-1)[\alpha V(b-1)_{u=0}-\alpha V(b)] \\
&\leq V(b-1)_{u=0}-V(b)
\end{aligned}$$

The first inequality is true because $\delta_1[R+\alpha V(b-1)_{u=0}-\alpha V(\bar{b})] \geq 0$, the second inequality true because $p_1(b-1) \leq 1$ and $\alpha \leq 1$. □

Lemma 3.3.4. $V(b-2)_{u=0}-V(b-1)_{u=0} \leq V(b-1)_{u=0}-V(b) \quad \forall \quad \tilde{b}+2 \leq b \leq \hat{b}+1$

Proof. Let $p_0(b-2)+\delta_1 = p_0(b-1)$. Then $\delta_1 \geq 0 \quad \forall \quad b \in \mathcal{B}_c$.

$$V(b-2)_{u=0} - V(b-1)_{u=0} = p_0(b-2)[R + \alpha V(b-1)_{u=0} - \alpha V(\bar{b})] \\ - p_0(b-1)[R + \alpha V(b) - \alpha V(\bar{b})]$$

Substituting $p_0(b-2)$,

$$\begin{aligned} &=(p_0(b-1)-\delta_1)[R+\alpha V(b-1)_{u=0} - \alpha V(\bar{b})] - p_0(b-1)[R+\alpha V(b) - \alpha V(\bar{b})] \\ &=p_0(b-1)[\alpha V(b-1)_{u=0} - \alpha V(b)] - \delta_1[R + \alpha V(b-1)_{u=0} - \alpha V(\bar{b})] \\ &\leq p_0(b-1)[\alpha V(b-1)_{u=0} - \alpha V(b)] \\ &\leq V(b-1)_{u=0} - V(b) \end{aligned}$$

The first inequality is true because $\delta_1[R+\alpha V(b-1)_{u=0} - \alpha V(\bar{b})] \geq 0$, the second inequality true because $p_0(b-1) \leq 1$ and $\alpha \leq 1$. □

Lemma 3.3.5. $V(b-2)_{u=0} - V(b-1)_{u=1} \leq V(b-1)_{u=1} - V(b) \quad \forall \quad \tilde{b}+2 \leq b \leq \hat{b}+1$

Proof.

$$V(b-2)_{u=0} - V(b-1)_{u=1} = p_0(b-2)[R + \alpha V(b-1)_{u=1} - \alpha V(\bar{b})] \\ - p_1(b-1)[R + \alpha V(b) - \alpha V(\bar{b})] + c$$

Since by assumption, $u^* = 1$ at $b-1$, we have $\frac{c}{R+\alpha V(b)-\alpha V(\bar{b})} < p_1(b-1) - p_0(b-1)$. In addition, let $p_1(b-2)+\delta_1 = p_1(b-1)$. Then $\delta_1 \geq 0 \quad \forall \quad b \in \mathcal{B}_c$.

$$\begin{aligned} \therefore\ & p_0(b-2)[R+\alpha V(b-1)_{u=1} - \alpha V(\bar{b})] - p_1(b-1)[R + \alpha V(b) - \alpha V(\bar{b})] + c \\ &< p_0(b-2)[R+\alpha V(b-1)_{u=1} - \alpha V(\bar{b})] - p_0(b-1)[R+\alpha V(b) - \alpha V(\bar{b})] \\ &= (p_0(b-1)-\delta_1)[R+\alpha V(b-1)_{u=1} - \alpha V(\bar{b})] - p_0(b-1)[R+\alpha V(b) - \alpha V(\bar{b})] \\ &\leq p_0(b-1)[\alpha V(b-1)_{u=1} - \alpha V(b)] \\ &\leq V(b-1)_{u=1} - V(b) \end{aligned}$$

The first inequality is true because $\delta_1[R+\alpha V(b-1)_{u=1} - \alpha V(\bar{b})] \geq 0$, the second inequality is true because $p_0(b-1) \leq 1$ and $\alpha \leq 1$. ☐

Lemma 3.3.6. *$V(b)$ is concave for $b \in \mathcal{B}_c$*

Proof. Lemmas 3.3.2 to 3.3.5 establish the concavity of each triplet belonging to $\mathcal{B}_c$ for all combinations of policies. The concavity of the region follows directly. ☐

3.3.2 Optimal Policy for Unimodal $V(b)$

Having established the unimodality and concavity of $V(b)$ for $b \in \mathcal{B}_c$ when the client exhibits a peaking behavior in the buy states, we are now able to characterize the optimal policy for the associate. We begin by looking at when it is optimal to never contact a client. Since the maximum impact of contact occurs at $\hat{b}$ and the minimum weighted cost $\frac{c}{R+\alpha V(b+1)-V(b)}$ occurs at $b = \tilde{b}$, we expect that the optimal policy is to never contact when the maximum impact is smaller than the minimum weighted cost. Proposition 3.3.2 confirms this.

Proposition 3.3.2. *When* $\frac{c}{R+\alpha V(\tilde{b})-\alpha V(\bar{b})} > \Delta_{\hat{b},u}$, $u^* = 0 \quad \forall \quad b \geq 1$

Proof. $\frac{c}{R+\alpha V(b+1)-\alpha V(b)} \geq \frac{c}{R+\alpha V(\tilde{b})-\alpha V(\bar{b})} \ \forall \quad b \geq 1$ and $\Delta_{\hat{b},u} > \Delta_{b,u} \quad \forall \quad b \geq 1$.
$\therefore u^* = 0 \quad \forall \quad b \geq 1$. ☐

Given the property of the value function and the peaking in the client behavior, it seems unlikely that the optimal policy will always be threshold, as in §3.1. Instead, it could be that the optimal policy has two thresholds: a *lower* threshold, where the policy switches from a no-contact to a contact, and an *upper* threshold, where the policy switches from a contact back to no-contact. In order to show that such a quasi-concave policy is optimal, we separate the states into three regions, namely $b \leq \tilde{b}-1$, $\tilde{b}-1 \leq b \leq \hat{b}$ and $b \geq \hat{b}$. For $b \leq \tilde{b}-1$, we show in Propositions 3.3.3 and 3.3.4 that the optimal policy is either a never-contact policy or a threshold. Similarly, for $b > \hat{b}$, as shown previously in Propositions 3.1.4 and 3.1.5, the optimal policy is either a never-contact policy or a threshold policy. In the last region, where the states fall in $\tilde{b}-1 \leq b \leq \hat{b}$, we show in Proposition 3.3.5 that the optimal policy is either a never-contact policy, a threshold policy, or a quasi-concave policy.

Proposition 3.3.3. *If $u^* = 0$ at $\tilde{b}-1$, then $u^* = 0 \quad \forall \quad 1 \leq b \leq \tilde{b}-1$.*

Proof. $u^* = 0$ at $\tilde{b}-1$ dictates that $\frac{c}{R+\alpha V(\tilde{b})-\alpha V(nr)} > \Delta_{\tilde{b}-1,u}$. Since $\Delta_{\tilde{b}-1,u} \geq \Delta_{b,u} \forall \quad 1 \leq b \leq \tilde{b}-1$ and $\frac{c}{R+\alpha V(\tilde{b})-\alpha V(nr)} \leq \frac{c}{R+\alpha V(b)-\alpha V(nr)} \quad \forall \quad 1 \leq b \leq \tilde{b}-1$, we have $\frac{c}{R+\alpha V(b)-\alpha V(nr)} > \Delta_{b,u} \quad \forall \quad 1 \leq b \leq \tilde{b}-1$. Thus, $u^* = 0 \quad \forall \quad 1 \leq b \leq \tilde{b}-1$. The densely shaded region in Figure 3-7(b) illustrates this pictorially. □

Proposition 3.3.4. *If $\exists \quad b_l \leq \hat{b}-1$ such that $u^* = 1$ at b_l and $u^* = 0$ at $b_l - 1$, then $u^* = 0 \quad \forall \quad 1 \leq b \leq b_l - 1$ and $u^* = 1 \quad \forall \quad b_l \leq b \leq \tilde{b}-1$ and a threshold policy exists.*

Proof. At b_l, $\frac{c}{R+\alpha V(b_l+1)-\alpha V(nr)} < \Delta_{b_l,u}$. Since $\Delta_{b,u} \geq \Delta_{b_l,u} \forall \quad b_l \leq b \leq \tilde{b}$ and $\frac{c}{R+\alpha V(b_l+1)-\alpha V(nr)} \geq \frac{c}{R+\alpha V(b)-\alpha V(nr)} \forall \quad b_l < b \leq \tilde{b}-1$, we have $u^* = 1 \forall \quad b_l \leq b \leq \tilde{b}-1$. At $b_l - 1$, $\frac{c}{R+\alpha V(b_l)-\alpha V(nr)} > \Delta_{b_l-1,u}$. Since $\Delta_{b,u} \leq \Delta_{b_l-1,u} \forall \quad 1 \leq b \leq b_l - 1$ and $\frac{c}{R+\alpha V(b_l)-\alpha V(nr)} \leq \frac{c}{R+\alpha V(b+1)-\alpha V(nr)} \forall \quad 1 \leq b \leq b_l - 1$, we have $u^* = 0 \forall \quad l \leq b \leq b_l - 1$. The densely shaded region in Figure 3-7(a) illustrates this pictorially. □

Proposition 3.3.5. *The optimal policy where $\tilde{b}-1 \leq b \leq \hat{b}$ is one of the following:*

1. $u^* = 0 \quad \forall \quad \tilde{b}-1 \leq b \leq \hat{b}$
2. $\exists \, b_h$ *such that* $u^* = 1 \quad \forall \quad \tilde{b}-1 \leq b \leq b_h$ *and* $u^* = 0 \quad \forall \quad b_h + 1 \leq b \leq \hat{b}$
3. $\exists \, b_l$ *such that* $u^* = 0 \quad \forall \quad \tilde{b}-1 \leq b \leq b_l - 1$ *and* $u^* = 1 \quad \forall \quad b_l \leq b \leq \hat{b}$
4. $u^* = 1 \quad \forall \quad b_l \leq b \leq b_h$ *and* $u^* = 0$ *else*

Proof. We separate into two cases. In the first case, the optimal policy is to contact in state $\tilde{b}-1$. In the second case, the optimal policy is not to contact in state $\tilde{b}-1$.

Case 1: If $u^* = 1$ at $\tilde{b}-1$ When in this case, there are two possible forms of optimal policies as illustrated in the sparsely shaded region of Figure 3-7(a). The first policy, illustrated by the dotted red line, has a single crossover point $b_h \leq \hat{b}-1$ (due to the convexity of $\frac{c}{R+\alpha V(b)-V(nr)}$ and the concavity of $\Delta_{b,u}$ in this particular region) such that $u^* = 1 \quad \forall \quad b \leq b_h$ and $u^* = 0 \quad \forall \quad b \geq b_h + 1$. The second policy, illustrated by the dashed blue line, will not exhibit such a crossover point, and by the convexity of $\frac{c}{R+\alpha V(b)-V(nr)}$ and the concavity of $\Delta_{b,u}$ in this particular region, will have $u^* = 1 \quad \forall \quad \tilde{b}-1 \leq b \leq \hat{b}$.

Case 2: If $u^* = 0$ at $\tilde{b}-1$ When in this case, there are three possible forms of optimal policy as illustrated in the sparsely shaded region of Figure 3-7(b). The first policy, indicated by the dotted red line, occurs when there is no intersection between the curves $\Delta_{b,u}$ and $\frac{c}{R+\alpha V(b)-V(nr)}$, resulting in $u^* = 0 \quad \forall \quad \tilde{b}-1 \leq b \leq \hat{b}$. The second policy, indicated by the dashed blue line, has a single crossover point $b_l \geq \tilde{b}$, due to the convexity of $\frac{c}{R+\alpha V(b)-V(nr)}$ and the concavity of $\Delta_{b,u}$ in this particular region, such that $u^* = 1 \quad \forall \quad b_l \leq b \leq \hat{b}$. The last policy, indicated by the finely dashed green line, occurs when there are two crossover points, also due to the convexity of $\frac{c}{R+\alpha V(b)-V(nr)}$ and the concavity of $\Delta_{b,u}$ in this particular region, such that $u^* = 1 \quad b_l \leq b \leq b_h$ and $u^* = 1 \quad else$. □

Combining the characterization for all three regions results in either a never-contact policy or a quasi-concave policy for a unimodal $V(b)$.

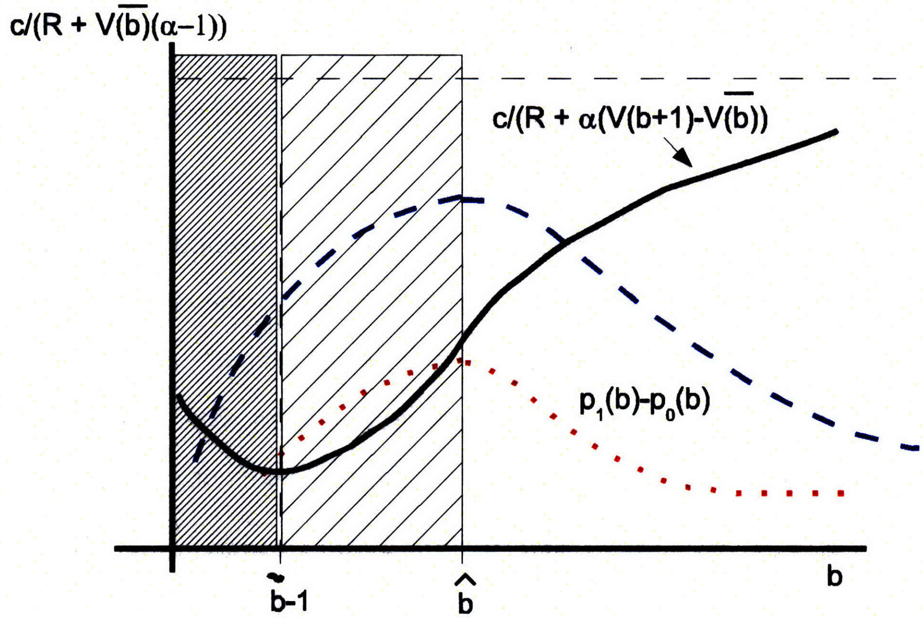

(a) Possible optimal policies when $u^* = 1$ at $\tilde{b} - 1$.

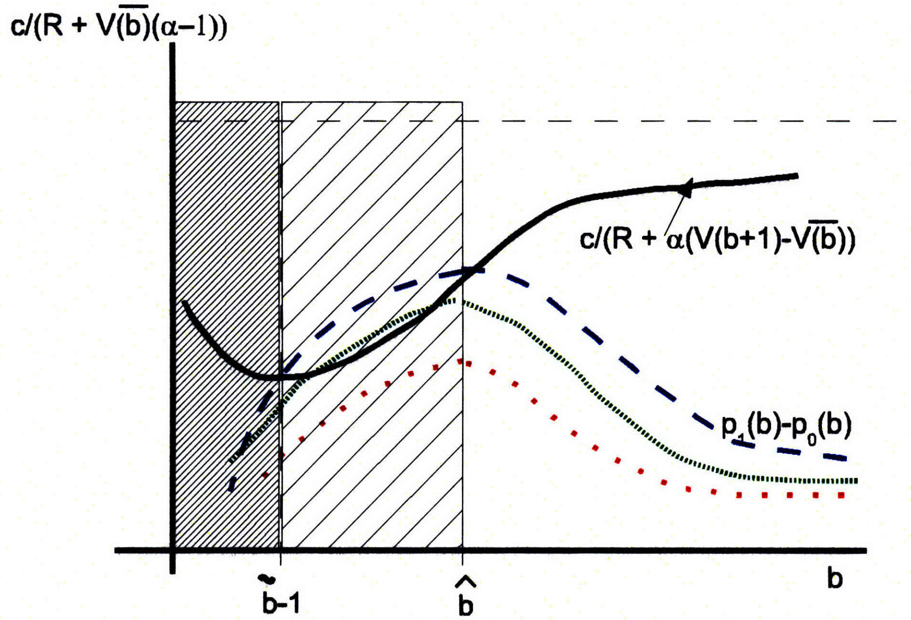

(b) Possible optimal policies when $u^* = 0$ at $\tilde{b} - 1$.

Figure 3-7: Possible optimal policies for buy states under peaking client behavior.

3.4 Characterization of Optimal Policy for a Non-Monotonic Client in Non-Buy States

In §3.2, we were unable to say much about the optimal policy when the client behaves monotonically in the non-buy states, other than that the threshold policy is not necessarily optimal. We suspect that the optimal policy when the client exhibits a peaking behavior will be more challenging to characterize. We begin, as usual, by introducing the set of notations used and by stating the assumptions invoked when the client exhibits a peaking behavior in the non-buy states.

$$
\begin{aligned}
p_u(nr) &= \text{probability of transitioning from state } nr \text{ into } nr+1 \text{ under policy } u, \\
&\quad \text{where } u \in \{0,1\} \\
\Delta_{nr,u} &= p_0(nr) - p_1(nr) \\
\hat{nr} &= \arg\max_{nr} p_u(nr) \\
\tilde{nr} &= \arg\max_{nr} V(nr) \\
b &\equiv \text{state (1,0)} \\
V(nr)_u &= \text{value function in state } nr \text{ under the assumption that } u \text{ is optimal}
\end{aligned}
$$

Assumption 7. *$p_0(nr)$ is concave and non-decreasing while $p_1(nr)$ exhibit peaking behavior with a minimum at $\hat{nr}$.*

Assumption 8. *$p_0(nr) - p_1(nr)$ is maximum at $\hat{nr}$. In addition, $p_0(nr) - p_1(nr)$ is non-decreasing $\forall \quad 1 \leq nr \leq \hat{nr}$ and non-increasing for $nr \geq \hat{nr}$. In other words, the effect of a contact peaks at $\hat{nr}$ before decreasing with nr.*

The optimality condition in Proposition 3.2.1 remains unchanged and $\Delta_{nr,u}$ is dictated by Assumption 8. Thus we turn to characterizing $V(nr)$ as the first step. To gain some insights to $V(nr)$, we refer back to the client behavior in Figure 2-3(b). The two responses corresponding to when the client is contacted ($u = 1$) and when the client is not contacted

($u = 0$) do not move in tandem. In other words, it is not obvious, unlike in §3.3, that the value function will be unimodal. We look at why this is so in §3.4.1.

3.4.1 Property of the Value Function $V(nr)$

Since we have shown in §3.2 that $V(nr)$ is monotonically non-increasing for $nr \geq \hat{nr}$, we focus on characterizing $V(nr)$ for $nr \leq \hat{nr}$. We first investigate the unimodality of $V(nr)$. In order for $V(nr)$ to be unimodal, we need to show, for all possible policy combinations, that if there is a peak at $\tilde{nr} \leq \hat{nr}$, the value function preceding it will be monotonically non-decreasing in nr. Unfortunately one cannot conclusively make such a statement. For a general state $nr \leq \hat{nr}$, the value function is given in Eq. (3.8).

$$\begin{aligned} V(nr) = \max\{&-p_1(nr)[R + \alpha V(b) - \alpha V(nr+1)] + \alpha V(b) + R - c, \\ &- p_0(nr)[R + \alpha V(b) - \alpha V(nr+1)] + \alpha V(b) + R\} \end{aligned} \tag{3.8}$$

If the peak $\tilde{nr}$ occurs at nr, we have $V(nr) \geq V(nr+1)$. Consider the case where $u^* = 0$ for both nr and $nr-1$. The value functions for both states are given by Eq. (3.9) and Eq. (3.10)

.

$$V(nr)_{u=0} = -p_0(nr)[R + \alpha V(b) - \alpha V(nr+1)] + \alpha V(b) + R \tag{3.9}$$

$$V(nr-1)_{u=0} = -p_0(nr-1)[R + \alpha V(b) - \alpha V(nr)] + \alpha V(b) + R \tag{3.10}$$

Since $p_0(nr-1) \leq p_0(nr) \quad \forall \quad nr \leq \hat{nr}$, $V(nr)_{u=0} \geq V(nr-1)_{u=0}$ is not always true. Computational results confirm that there are cases when $V(nr)$ will not be unimodal. Figure 3-8 illustrates one such example.

3.4.2 Optimal Policy Characterization

$V(nr)$ being non-unimodal makes the optimal policy very difficult to characterize. Proposition 3.2.4 still applies for determining the circumstances under which a never contact policy is optimal. It turns out we can also show that when the optimal policy at $\hat{nr}$ is not to

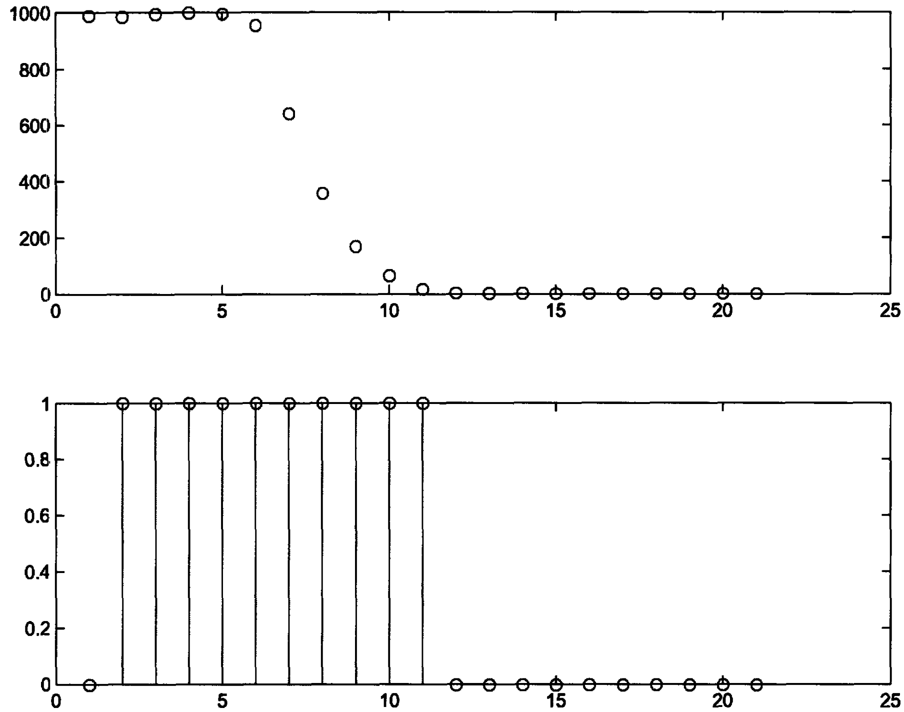

Figure 3-8: Computational example of a non-unimodal $V(nr)$.

contact, it is also optimal not to contact for all values of $nr < \widehat{nr}$ since the maximum impact of contact occurs at $\widehat{nr}$. We show this in Proposition 3.4.1 and Lemmas 3.4.1 to 3.4.2.

Lemma 3.4.1. *If* $u^* = 0$ *at* $\widehat{nr}$*, then* $u^* = 0$ *at* $\widehat{nr} - 1$

Proof. $u^* = 0$ at $\widehat{nr}$ implies that $\frac{c}{R+\alpha V(b)-\alpha V(\widehat{nr}+1)} > \Delta_{\widehat{nr},u}$. Since $V(\widehat{nr}) \geq V(\widehat{nr}+1)$, $\frac{1}{-V(\widehat{nr})} \geq \frac{1}{-V(\widehat{nr}+1)}$. In addition, $\Delta_{\widehat{nr},u} > \Delta_{\widehat{nr}-1,u}$. Thus, we have

$$\frac{c}{R+\alpha V(b)-\alpha V(\widehat{nr})} > \frac{c}{R+\alpha V(b)-\alpha V(\widehat{nr}+1)} > \Delta_{\widehat{nr},u} > \Delta_{\widehat{nr}-1,u}$$

and $u^* = 0$ at $\widehat{nr} - 1$. □

Lemma 3.4.2. $V(nr-1)_{u=0} > V(nr)_{u=0}$ *if* $V(nr) > V(nr+1)$.

Proof.

$$V(nr-1)_{u=0} = -p_0(nr-1)\big[R + \alpha V(b) - \alpha V(nr)\big] + \alpha V(b) + R \tag{3.11}$$

$$V(nr)_{u=0} = -p_0(nr)\big[R + \alpha V(b) - \alpha V(nr+1)\big] + \alpha V(b) + R \tag{3.12}$$

Since $V(nr) > V(nr+1)$ and $p_0(nr-1) < p_0(nr)$, we have $V(nr-1)_{u=0} > V(nr)_{u=0}$. □

Proposition 3.4.1. *If* $u^* = 0$ *at* $\hat{nr}$*, then* $u^* = 0 \quad \forall \quad nr < \hat{nr}$.

Proof. We showed in Lemma 3.4.1 that $u^* = 0$ at $\hat{nr}$ implies $u^* = 0$ at $\hat{nr}-1$. This means that $V(\hat{nr}-1) = V(\hat{nr}-1)_{u=0}$. We also showed in Lemma 3.4.2 that $V(\hat{nr}-1)_{u=0} \geq V(\hat{nr})_{u=0}$ is always true. Knowing that $V(\hat{nr}-1) = V(\hat{nr}-1)_{u=0} \geq V(\hat{nr})_{u=0}$, we have the following

$$\frac{c}{R + \alpha V(b) - \alpha V(\hat{nr}-1)} > \frac{c}{R + \alpha V(b) - \alpha V(\hat{nr})} > \Delta_{\hat{nr}-1,u} > \Delta_{\hat{nr}-2,u}$$

and $u^* = 0$ at $\hat{nr}-2$. Further, $V(\hat{nr}-2)_{u=0} > V(\hat{nr}-1)_{u=0}$ when $V(\hat{nr}-1) > V(\hat{nr})$ by Lemma 3.4.2. $u^* = 0$ at $\hat{nr}-3$ follows. This is true $\forall \quad k$ where $\hat{nr}-k > 0$. □

3.5 Computational Results

The size of the state space enables computation of the optimal policy for a single client when an absorbing state is defined. Figure 3-9 illustrates the state transitions with the absorbing state. A buy state with the number of consecutive buys of b_{max} will transition into the non-buy state $(0, 1)$ deterministically. A state with non-buys exceeding nr_{max} will transition probabilistically into the absorbing state upon a non-purchase. When this occurs, the client has not purchased for many consecutive periods and will not return in the future. We computed the optimal policy using value iteration algorithm [20]. Figures 3-10 to 3-

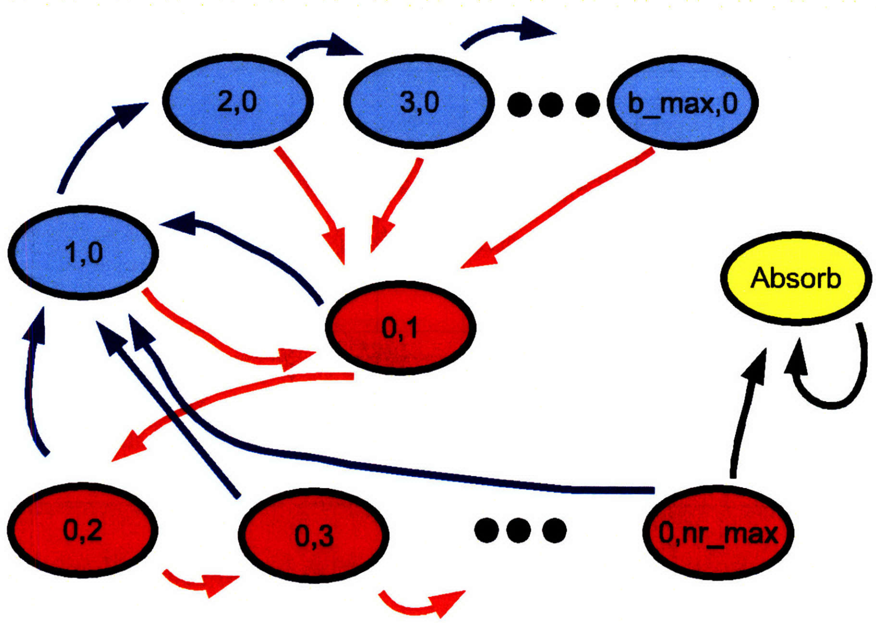

Figure 3-9: State transitions for the infinite horizon problem with an absorbing state. The blue states represent consecutive buy states while the red states represent consecutive non-buy states.

12 illustrate the optimal policy towards a client who exhibits a non-monotonic purchase behavior for different values of cost of contact, c. For example, if the cost of contact is relatively low, the associate should only contact when this particular client has purchased one or two consecutive times. The policy also points out that the associate should continue to activate the client if such client has not purchased for many consecutive periods. We note

that the forms of the optimal policies for the buy states are either threshold or never-contact, which coincides with the theoretical characterization of optimal policy. On the other hand, the optimal policies for the non-buy states are not always threshold.

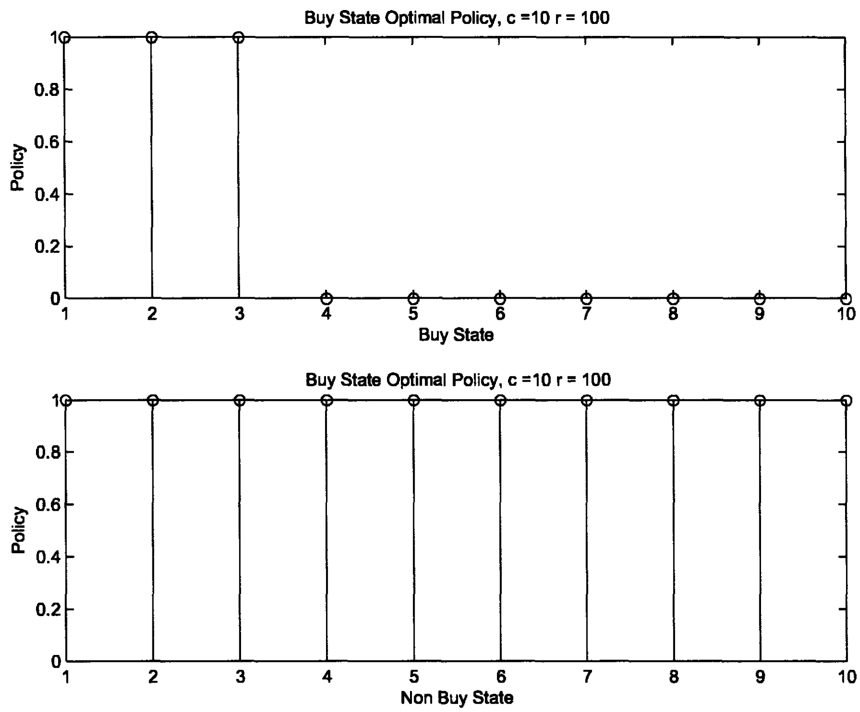

Figure 3-10: Optimal policy for a client when R = 100 and c = 10.

In the next set of computations we looked at the impact of a low discount factor - emulating a short-term oriented associate - on the optimal policy. Not surprisingly, we find that the low discount factor shifts the threshold of the optimal policy but does not change the form of the optimal policy. Figures 3-13 to 3-15 illustrate the optimal policy for the same client when $\alpha = 0.3$.

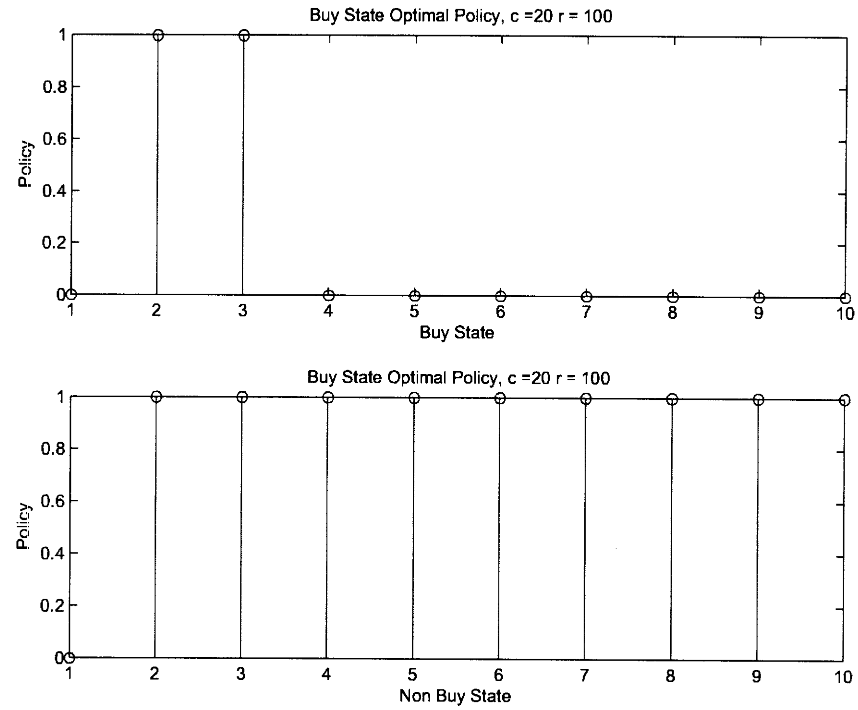

Figure 3-11: Optimal policy for a client when R = 100 and c = 20.

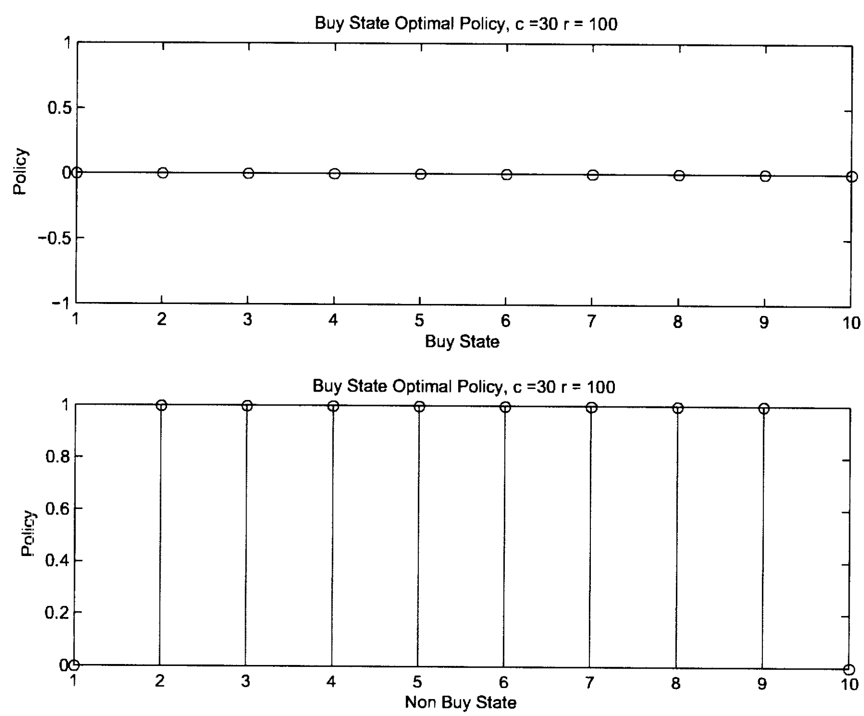

Figure 3-12: Optimal policy for a client when R = 100 and c = 30.

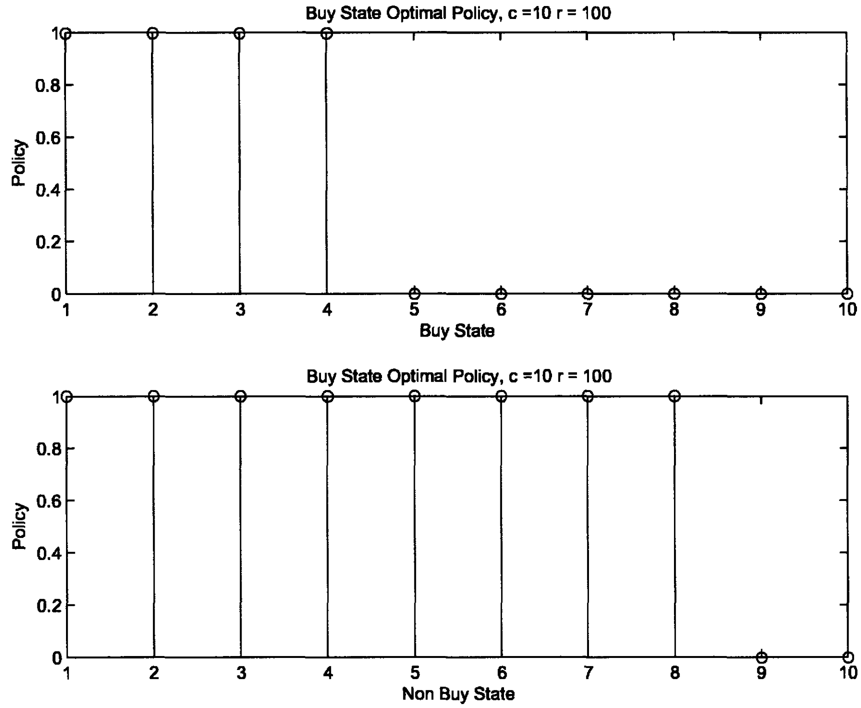

Figure 3-13: Optimal policy for a client when R = 100, c = 10 and $\alpha = 0.3$.

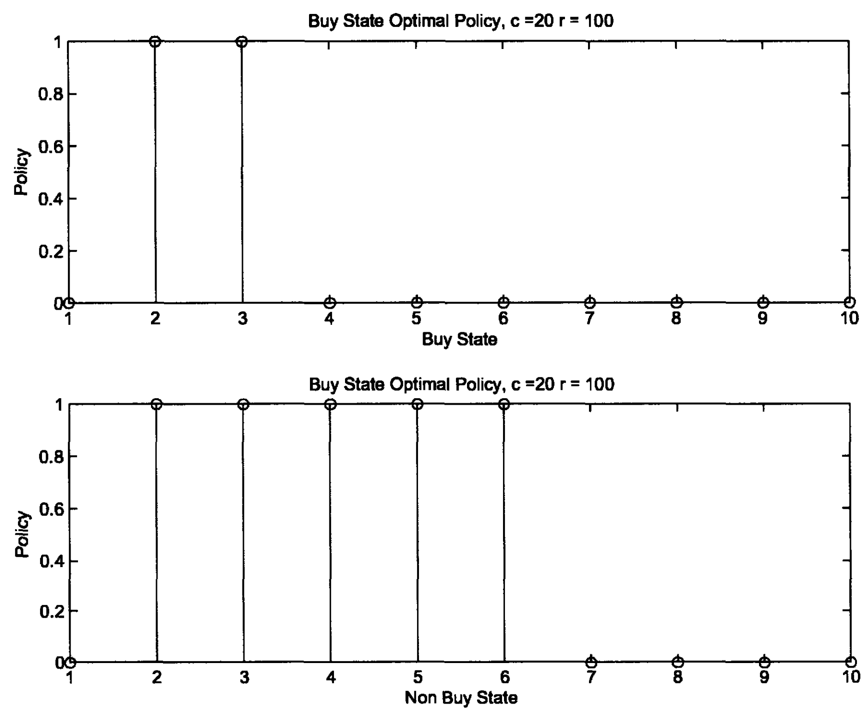

Figure 3-14: Optimal policy for a client when R = 100, c = 20 and $\alpha = 0.3$.

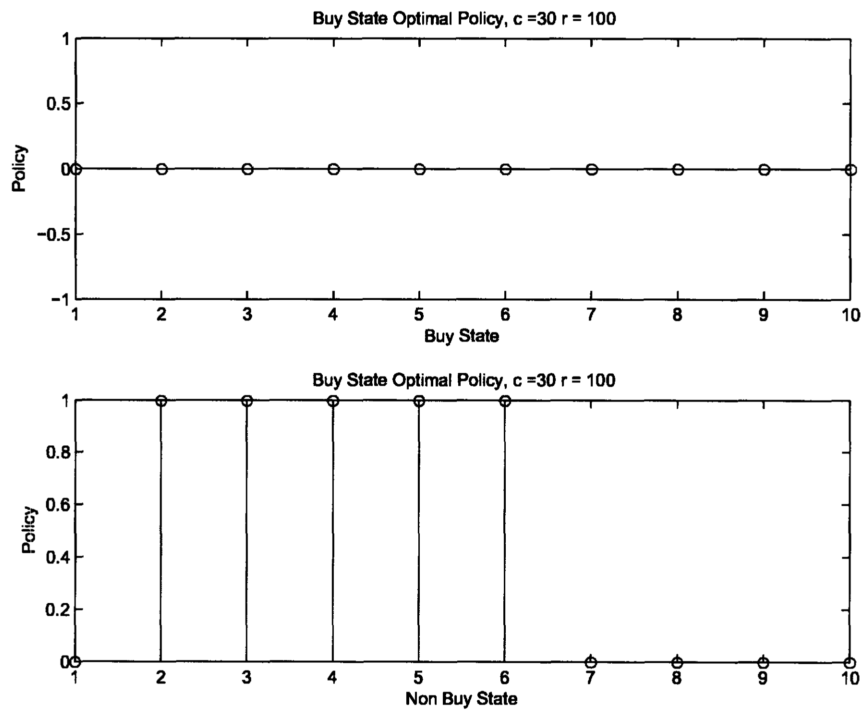

Figure 3-15: Optimal policy for a client when R = 100, c = 30 and $\alpha = 0.3$.

3.5.1 Insights

It is worthwhile to pause and understand the insights these results reveal. For each single client, we have provided a method for an associate to determine when it is optimal to contact a client. The optimal policy takes a very simple form. Towards a client who has not purchased for many consecutive periods, the policy guides the associate as to when to start reactivating and when to give up. On the other hand, towards a client who has been purchasing for consecutive periods, it guides the associate as to when to stop pestering. Interestingly, the policy also seems to suggest the importance of reactivating a client, as illustrated by the number of non-buy states in which it is optimal to contact. Lastly, the results confirm the inverse relationship between the cost and the number of contacts initiated. As the cost increases, the results indicate that the number of states where it is optimal to contact decreases.

The question that follows, naturally, is "What is the impact of the optimal policy?" To gauge such an impact, we compare the lifetime value of a client under some commonly practiced policies. For the extremely diligent associates, we assume that they will always call; for the extremely lazy associates, we assume that they never call and focus on walk-in clients instead; for the "smart" associates, we assume that they will chase after the repeat purchase clients; finally for the fickle associates, we assume that they use a random policy. Table 3.1 illustrates the relative performance of these policies under varying cost of contact. The best heuristic in each column is **boldfaced** while the worst is *italicized.* We first note

Policy	$\frac{c}{R} = .1$	$\frac{c}{R} = .2$	$\frac{c}{R} = .3$	$\frac{c}{R} = .4$
Contact all	**0.9921**	**0.9323**	**0.7536**	0.5360
Never Contact	*0.2450*	*0.2873*	0.3087	0.3273
Contact Buys Only	0.3240	0.3272	*0.2951*	*0.2529*
Random	0.7291	0.7644	0.7243	**0.6650**
Optimal Lifetime Value	1014.1790	865.0606	805.0540	759.2720

Table 3.1: Relative (to optimal) performance of various policies.

that when the cost of contact is very low, the extremely diligent associate performs rather well. Unfortunately, the performance decreases drastically with increasing cost. The lazy associate does not fare very well, for his performance is around 30% of the optimal policy. The "smart" associate does not fare well either since he fails to grasp the importance of reactivation. The fickle associate performs relatively well since a client is occasionally and randomly activated.

When the discount factor is low, a different picture emerges. Table 3.2 illustrates the relative performance for the same client when $\alpha = 0.3$. The best heuristic in each column is **boldfaced** while the worst is *italicized*. We first note that the optimal lifetime value is small. Its value is also close to the expected single stage cost since revenue in the future is heavily discounted.[3] Next, we note that the lazy associate does well, especially when the relative cost of contact is high. The "smart" associate and the extremely diligent associate both perform well under low relative cost of contact, but poorly under high relative cost of contact.

Policy	$\frac{c}{R} = .1$	$\frac{c}{R} = .2$	$\frac{c}{R} = .3$	$\frac{c}{R} = .4$
Contact all	**0.9999**	0.9426	*0.7126*	*0.4621*
Never Contact	*0.7982*	0.9504	**0.9749**	**0.9830**
Contact Buys Only	0.9575	*0.9258*	0.7300	0.5145
Random	0.8307	**0.9522**	0.9390	0.9087
Optimal Lifetime Value	68.6040	57.6210	56.1727	55.7083

Table 3.2: Relative (to optimal) performance of various policies when $\alpha = 0.3$

The contrasts between the relative performance of associates under two different values of α are rather interesting. If one were to assume that the future value of a client is insignificant, then the commonly practiced policy of *contacting clients when they have purchased* is actually close to optimal when the cost of contact is low. When the cost of contact is high, the associates are better off never contacting a client. Unfortunately, the future value of a client is often very significant. In following the commonly practiced policies, the associates are letting go of a significant opportunity.

[3] When α approaches 0, the problem reduces to a single stage problem.

Chapter 4

A Single Client: Finite Horizon Problem

Having analyzed the infinite horizon problem, we now look into the finite horizon problem as formulated in Eq. (2.1) - Eq. (2.2). Figure 4-1 depicts the evolution of the client as he moves through the system. A client enters the system in state $(1,0)$ at time $t = 1$. In the next period, depending on his action, he will either move into the next buy state or the non-buy state. For reasons similar to Chapter 3, we analyze the buy states and the non-buy states separately in §4.1 and §4.2. The same set of notation is used, with the exception that the value functions are now indexed by time t. Since the transitional probabilities are time invariant, Assumptions 1 to 8 remain valid as well. We close this chapter with computational results in §4.3.

4.1 Characterization of Optimal Policy for a Client in the Buy States

Rather than being time invariant, we expect the optimality condition of a policy to be indexed by time. In addition, we expect that such optimality be determined by a ratio. Proposition 4.1.1 illustrates this mathematically. When the client is in the final period, an

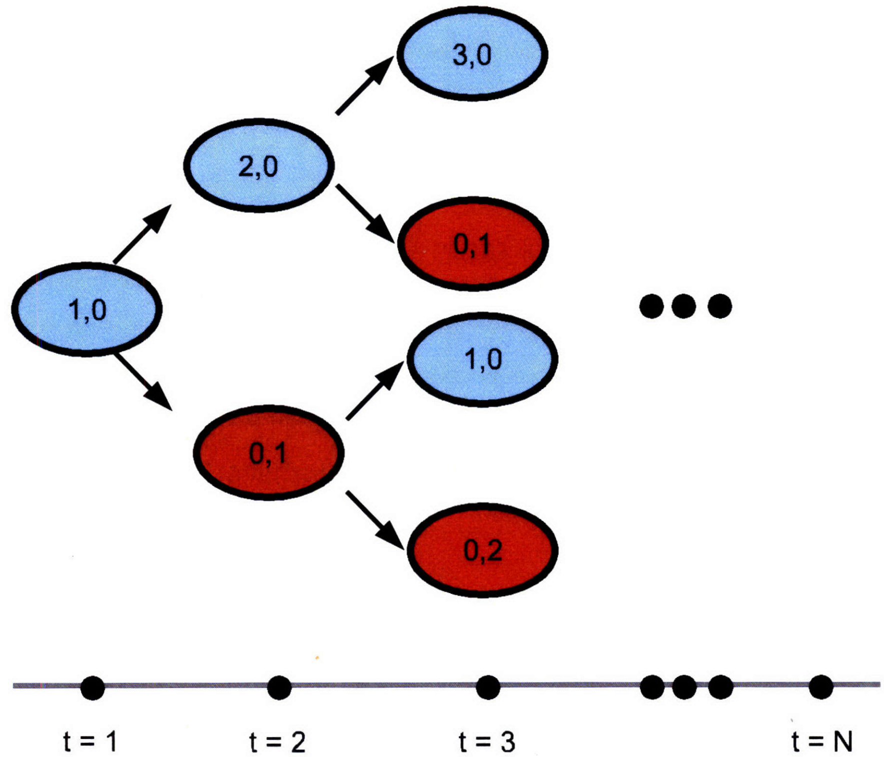

Figure 4-1: State transitions for a client for the finite horizon problem.

associate's policy should only be determined by the impact of a contact and the ratio of cost over revenue. Corollary 4.1.1 confirms this.

Proposition 4.1.1. *If* $\Delta_{b,u} > \frac{c}{R+\alpha V_{t+1}(b+1)-\alpha V_{t+1}(\bar{b})}$ *at time* t*, then* $u^* = 1$ *at* b*.*

Proof. Let $V_t(b)_{u=0}$ and $V_t(b)_{u=1}$ denote the value function at time t when the associate does not contact (u=0) and contact (u = 1) respectively in a buy state b.

$$V_t(b)_{u=0} = Rp_0(b) + \alpha\big[p_0(b)V_{t+1}(b+1) + (1-p_0(b))V_{t+1}(\bar{b})\big]$$
$$V_t(b)_{u=1} = Rp_1(b) - c + \alpha\big[p_1(b)V_{t+1}(b+1) + (1-p_0(b))V_{t+1}(\bar{b})\big]$$

Since $u^* = 1$ when $V_t(b)_{u=1} > V_t(b)_{u=0}$, $u^* = 1$ when

$$Rp_1(b) - c + \alpha\big[p_1(b)V_{t+1}(b+1) + (1 - p_0(b))V_{t+1}(\bar{b})\big] >$$
$$Rp_0(b) + \alpha\big[p_0(b)V_{t+1}(b+1) + (1 - p_0(b))V_{t+1}(\bar{b})\big]$$

$$p_1(b)\big[R + \alpha V_{t+1}(b+1) - \alpha V_{t+1}(\bar{b})\big] - c > p_0(b)\big[R + \alpha V_{t+1}(b+1) - \alpha V_{t+1}(\bar{b})\big]$$
$$p_1(b) - p_0(b) > \frac{c}{R + \alpha V_{t+1}(b+1) - \alpha V_{t+1}(\bar{b})}$$

□

Corollary 4.1.1. *If* $\Delta_{b,u} > \frac{c}{R}$, $u^* = 1$ *at time* $N - 1$.

Proof. Since $V_N(i) = 0 \quad \forall \quad i$, using Proposition 4.1.1 with $t = N - 1$ gives the result. □

The optimality condition bears strong resemblance to the infinite horizon characterization. As such we will not make additional comments on the intuition behind the condition. Since the optimality condition varies with time when the horizon is finite, one needs to look at how $V_t(b)$ changes with both time and the state b. We begin by looking at the effect of time on $V_t(b)$. A very natural result is that the value function associated with a particular state will decrease with time. We show this in Proposition 4.1.2. We also expect that when time t is fixed, the behavior of $V_t(b)$ will be identical to the behavior of $V(b)$. As such, the results for $V(b)$ in Chapter 3 can be applied here. For example, we show in Proposition 4.1.3 that $V_t(b)$ is monotonically non increasing for a fixed t when the client behaves monotonically. As such, the optimal policy for a fixed time t is either a never-contact or a quasi-concave policy. Thus, as we sweep t from 1 to N, the optimal policy will be time varying with different lower and upper thresholds.

Proposition 4.1.2. $\forall \quad t \leq N - 1,\ V_t(b) \geq V_{t+1}(b) \geq 0$

Proof. We show by induction.
We first show for $t = N - 1$, $V_{N-1}(b) \geq V_N(b) \geq 0$.

By definition, $V_N(b) = 0$. For $t = N - 1$

$$V_{N-1}(b) = \max\{Rp_1(b) - c, Rp_0(b)\} \geq 0$$

since $p_0(b) \geq 0$, $p_1(b) \geq 0$ and we are taking the maximum.

For $t \leq N - 2$, we invoke the following for induction hypothesis

$$V_{t+1}(\bar{b}) \geq V_{t+2}(\bar{b}) \geq 0$$

$$V_{t+1}(b) \geq V_{t+2}(b) \geq 0$$

For the buy states, we show $V_t(b) \geq V_{t+1}(b)$ by induction.

$$\begin{aligned}
V_t(b) &= \max\{R(p_1(b)) - c + \alpha p_1(b)V_{t+1}(b+1) + \alpha(1 - p_1(b))V_{t+1}(\bar{b}) \\
&\qquad R(p_0(b)) + \alpha p_1(b)V_{t+1}(b+1) + \alpha(1 - p_1(b))V_{t+1}(\bar{b})\} \\
&\geq \max\{R(p_1(b)) - c + \alpha p_1(b)V_{t+2}(b+1) + \alpha(1 - p_1(b))V_{t+2}(\bar{b}) \\
&\qquad R(p_0(b)) + \alpha p_1(b)V_{t+2}(b+1) + \alpha(1 - p_1(b))V_{t+2}(\bar{b})\} \\
&= V_{t+1}(b)
\end{aligned}$$

□

Proposition 4.1.3. $\forall \quad t \leq N-1$, *$V_t(b)$ is monotonically non increasing in b for fixed t.*

Proof. We show by induction that this is true.

For $t = N-1$

$$V_{N-1}(b) = \max\{Rp_1(b) - c, Rp_0(b)\}$$
$$V_{N-1}(b+1) = \max\{Rp_1(b+1) - c, Rp_0(b+1)\}$$
$$p_1(b) \geq p_1(b+1)$$
$$p_0(b) \geq p_0(b+1)$$
$$\therefore V_{N-1}(b) \geq V_{N-1}(b+1)$$

For $t \leq N-2$, we assume for induction that $V_{t+1}(b)$ is monotonically non-increasing and show that $V_t(b) \geq V_t(b+1)$

$$\begin{aligned}
V_t(b) = \max\{&Rp_1(b) - c + \alpha p_1(b)V_{t+1}(b+1) + \alpha(1-p_1(b))V_{t+1}(\bar{b}), \\
&Rp_0(b) + \alpha p_0(b)V_{t+1}(b+1) + \alpha(1-p_0(b))V_{t+1}(\bar{b})\} \\
= \max\{&p_1(b)(R+\alpha V_{t+1}(b+1) - \alpha V_{t+1}(\bar{b})) + \alpha V_{t+1}(\bar{b}) - c, \\
&p_0(b)(R+\alpha V_{t+1}(b+1) - \alpha V_{t+1}(\bar{b})) + \alpha V_{t+1}(\bar{b})\} \\
\geq \max\{&p_1(b+1)(R + \alpha V_{t+1}(b+1) - \alpha V_{t+1}(\bar{b})) + \alpha V_{t+1}(\bar{b}) - c, \\
&p_0(b+1)(R + \alpha V_{t+1}(b+1) - \alpha V_{t+1}(\bar{b})) + \alpha V_{t+1}(\bar{b})\}
\end{aligned}$$

By induction assumption

$$\begin{aligned}
\geq \max\{&p_1(b+1)(R + \alpha V_{t+1}(b+2) - \alpha V_{t+1}(\bar{b})) + \alpha V_{t+1}(\bar{b}) - c, \\
&p_0(b+1)(R + \alpha V_{t+1}(b+2) - \alpha V_{t+1}(\bar{b})) + \alpha V_{t+1}(\bar{b})\} \\
= V_t(b+1)&
\end{aligned}$$

□

4.2 Characterization of Optimal Policy for a Client in the Non-buy States

In this section we look at the optimality condition of a policy when the client is in the non-buy state. Similar to §4.1, we show in Proposition 4.2.1 that such condition is time dependent. Corollary 4.2.1 states the optimality condition of a client who is in the final period. We will not go into detail regarding the intuition of these results since they are identical to ones in §4.1.

Proposition 4.2.1. *If $\Delta_{nr,u} > \frac{c}{R+\alpha V_{t+1}(b)-\alpha V_{t+1}(nr+1)}$ at time t, then $u^* = 1$ at nr.*

Proof. Let $V_t(nr)_{u=0}$ and $V_t(nr)_{u=1}$ denote the value function at time t when the associate does not contact (u=0) and contact (u = 1) respectively in a non-buy state nr.

$$V_t(nr)_{u=0} = R(1 - p_0(nr)) + \alpha\big[p_0(nr)V_{t+1}(nr + 1) + (1 - p_0(nr))V_{t+1}(b)\big]$$
$$V_t(nr)_{u=1} = R(1 - p_1(nr)) - c + \alpha\big[p_1(nr)V_{t+1}(nr + 1) + (1 - p_0(nr))V_{t+1}(b)\big]$$

Since $u^* = 1$ when $V_t(nr)_{u=1} > V_t(nr)_{u=0}$, $u^* = 1$ when

$$R(1 - p_1(nr)) - c + \alpha\big[p_1(nr)V_{t+1}(nr + 1) + (1 - p_0(nr))V_{t+1}(b)\big] >$$
$$R(1 - p_0(nr)) + \alpha\big[p_0(nr)V_{t+1}(nr + 1) + (1 - p_0(nr))V_{t+1}(b)\big]$$

Simplifying the above gives

$$p_0(nr) - p_1(nr) > \frac{c}{R + \alpha V_{t+1}(b) - \alpha V_{t+1}(nr + 1)} \tag{4.1}$$

□

Corollary 4.2.1. *If $\Delta_{nr,u} > \frac{c}{R}$, $u^* = 1$ at time $N - 1$.*

Proof. Since $V_N(i) = 0 \quad \forall \quad i$, using Proposition 4.2.1 with $t = N - 1$ gives the result. □

Due to the time dependent nature of the optimality condition, we look at how $V_t(nr)$ changes with both time and state nr in order to fully characterize the optimal policy. We show in Proposition 4.2.2 that the value function associated with a particular state will decrease with time. For a particular value of t along the time axis, the behavior of $V_t(nr)$ becomes time invariant since the value of t is fixed. As such, for each value of t, the results in §3.2 and §3.4 apply and the threshold policy is not always optimal.

Proposition 4.2.2. $\forall \quad t \leq N-1,\ V_t(nr) \geq V_{t+1}(nr) \geq 0$

Proof. We show by induction.

By definition, $V_N(nr) = 0$. For $t = N-1$

$$V_{N-1}(nr) = \max\{R(1-p_1(nr)) - c, R(1-p_0(nr))\} \geq 0$$

For $t \leq N-2$, we invoke the following for induction hypothesis

$$V_{t+1}(nr) \geq V_{t+2}(nr) \geq 0$$

$$V_{t+1}(b) \geq V_{t+2}(b) \geq 0$$

We show $V_t(nr) \geq V_{t+1}(nr)$ by induction.

$$\begin{aligned}
V_t(nr) &= \max\{R(1-p_0(nr)) + \alpha\big[p_0(nr)V_{t+1}(nr+1) + (1-p_0(nr))V_{t+1}(b)\big], \\
&\qquad R(1-p_1(nr)) - c + \alpha\big[p_1(nr)V_{t+1}(nr+1) + (1-p_0(nr))V_{t+1}(b)\big]\} \\
&\geq \max\{R(1-p_0(nr)) + \alpha\big[p_0(nr)V_{t+2}(nr+1) + (1-p_0(nr))V_{t+2}(b)\big], \\
&\qquad R(1-p_1(nr)) - c + \alpha\big[p_1(nr)V_{t+2}(nr+1) + (1-p_0(nr))V_{t+2}(b)\big]\} \\
&= V_{t+1}(nr)
\end{aligned}$$

□

Policy	$\frac{c}{R} = .1$	$\frac{c}{R} = .2$	$\frac{c}{R} = .3$	$\frac{c}{R} = .4$
Contact all	541.65	434.31	326.97	219.63
Never Contact	231.31	231.31	231.31	231.31
Contact Buys Only	298.26	256.50	214.75	173.00
Random	427.26	351.29	358.74	276.83
Optimal Lifetime Value	545.42	464.18	432.68	409.69

Table 4.1: Lifetime value function for various policies when N = 16

4.3 Computational Results

Our state definition allows the optimal policy under finite time horizon to be computed by brute force recursion from the very last time period. The computational results align with our theoretical characterizations. For clients in the buy states, the optimal policy is quasi-concave with time dependent upper and lower thresholds. For clients in the non-buy states, the threshold policy is not necessarily optimal. An example of the optimal policy can be seen in Figure 4-2, which plots the policy against the time period and the state variable. The results are obtained using a cost of $c = 20$ and a revenue of $R = 100$. A contact policy is indicated by value one while a non contact policy is indicated by value zero. For the buy states (Figure 4-2(a)), we see that the optimal policy is a time dependent quasi-concave policy. For the non-buy states (Figure 4-2(b)), we see that the optimal policy is also time dependent and quasi concave.

To gauge the impact of an optimal policy, we compare the differences in lifetime value[1] of a client under various policies. Table 4.1 illustrates the lifetime values of various heuristics while Table 4.2 illustrates their performance relative to the optimal policy. When the relative cost is low, a blanket contact policy works best (.99 relative to the optimal policy), followed by the "contact buys only" policy and the "never contact" policy. When the relative cost is high, the "never contact policy" outperforms the other two. For both "contact buys only" and "never contact" policies, the performance deteriorates as the relative cost increases.

To gauge the viability of the infinite horizon approximation, we compared the perfor-

[1]Defined as the value function of a client for the entire period of time he remains in the system. In this case, it is given by $V_{16}(1, 0)$.

Policy	$\frac{c}{R} = .1$	$\frac{c}{R} = .2$	$\frac{c}{R} = .3$	$\frac{c}{R} = .4$
Contact all	0.9931	0.9357	0.7557	0.5361
Never Contact	0.4241	0.4983	0.5346	0.5646
Contact Buys Only	0.5468	0.5526	0.4963	0.4223
Random	0.7834	0.7568	0.8291	0.6757

Table 4.2: Comparison of lifetime value under different policies and costs.

Policy	$\frac{c}{R} = .1$	$\frac{c}{R} = .2$	$\frac{c}{R} = .3$	$\frac{c}{R} = .4$
u^* for ∞ Horizon, N=10	0.9992	0.9992	0.9997	0.9999
u^* for ∞ Horizon, N=16	0.9996	0.9996	0.9998	0.9999
u^* for ∞ Horizon, N=20	0.9997	0.9997	0.9999	0.9999

Table 4.3: Comparison of infinite horizon optimal policy under different costs.

mance of the infinite horizon optimal policies to the finite horizon optimal policies. Table 4.3 shows that the infinite horizon optimal policies perform to within 0.999 of the finite horizon optimal policy. In addition, the results were quite insensitive to time horizon N, as modifying it did not significantly change the performance.

We conclude this chapter by noting that the time varying optimal policy does not perform significantly better than the corresponding infinite horizon optimal policy. As such, it is much less attractive compared to its time invariant counterpart due to implementation and computational difficulties. The computational complexity further increases in the more realistic case where an associate has multiple clients.

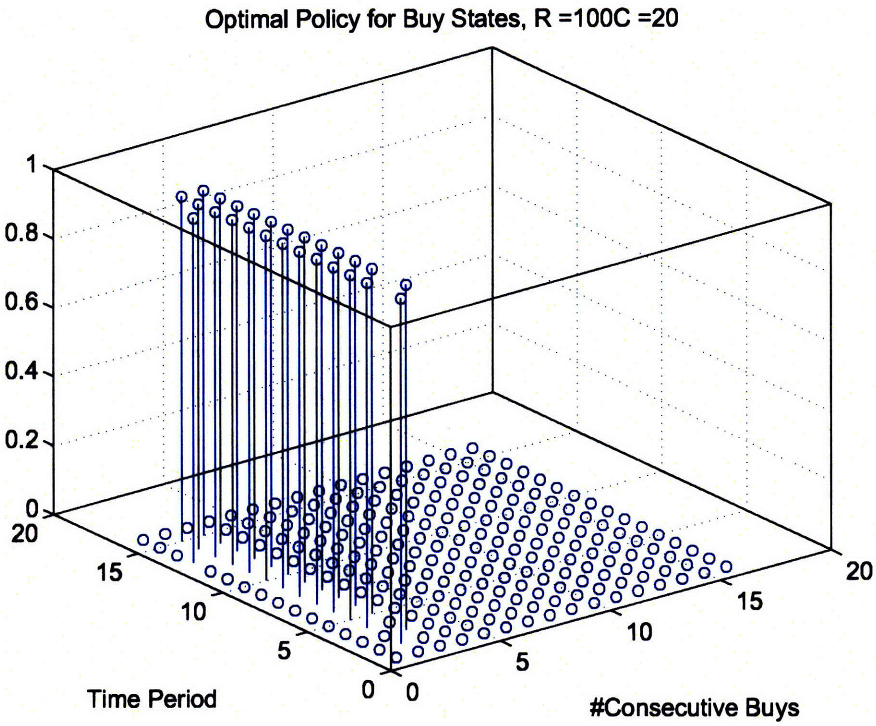

(a) Optimal contact policy for buy states.

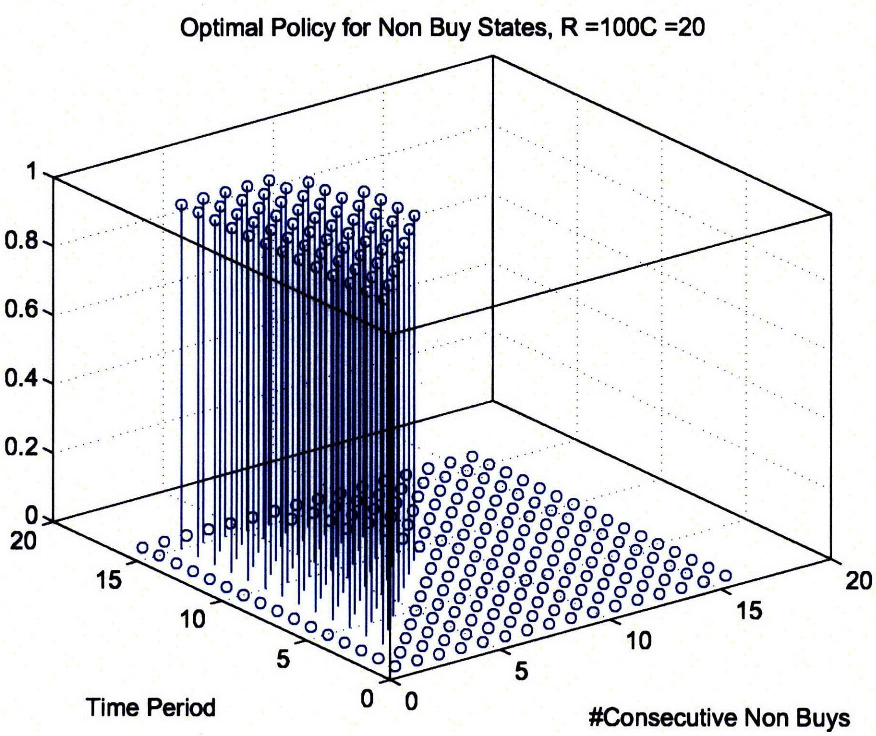

(b) Optimal contact policy for non-buy states.

Figure 4-2: Optimal policies for a client with peaking behavior under finite time horizon.

Chapter 5

Multiple Clients: Infinite Horizon Problem

In previous chapters we discussed the control mechanisms of an associate who only has a single client and analyzed this scenario under both infinite and finite time horizons. We now turn to analyzing the problem where an associate has multiple clients. The infinite horizon version of the problem originally formulated in §2.2 is restated as follows, where $V(\mathbf{i})$ is the value function of the clients in states $\mathbf{i}$ and M is the number of contacts to be made in that period.

$$V(\mathbf{i}) = \max_{\mathbf{u}}[g(\mathbf{i}, \mathbf{u}) + \alpha \sum_{\mathbf{j}} p_{\mathbf{iju}} V(\mathbf{j})] \tag{5.1}$$

$$s.t. \qquad \mathbf{u}'e = M \tag{5.2}$$

The problem in Eq. (5.1) - Eq. (5.2) falls into the category of a restless bandit problem, which is a generalization of the classical multiarmed bandit problem. In the classical multiarmed bandit problem, one out of n projects can be worked on in any time period. The project selected will return an expected reward and its state will evolve accordingly, while the passive projects remain static. In the restless bandit problem, the passive projects do not remain static. Rather, they evolve according to a set of passive transitional probabilities. The analogy between the restless bandit problem and the associate's problem is clear. Each client corresponds to a project, and in each period, the associate selects M clients to work on. The ones selected (or active) will evolve according to a set of *active* transitional probabilities while the ones not selected (or passive) will evolve according to a set of *passive* transitional probabilities. This problem was first investigated by Whittle in his seminal paper in 1988 [21]. Papadimitriou and Tsitsiklis [22] studied the complexity of the problem in 1999 and established that it is PSPACE-hard even with deterministic transition rules. In the remainder of this chapter, we first establish an upper bound for $V(\mathbf{i})$ in Section 5.1. Following that we propose heuristics in Section 5.2 and provide computational results to judge their performance in Section 5.3.

Before delving into the analysis, we introduce a set of simplified notation to be used throughout this chapter. For each client s, we define the following:

$$\begin{aligned}
i^s &= \text{state of client s, equivalent to } s_{th} \text{ element of } \mathbf{i} \\
V^s(i^s) &= \text{Value function of client s in state } i^s \\
r^s &= \text{single period reward yielded by client s, dependent on state and action} \\
g_1^s(i^s) &= \text{expected single period reward when client in state } i^s \text{ is contacted} \\
g_0^s(i^s) &= \text{expected single period reward when client in state } i^s \text{ is not contacted} \\
(P_{s1}V^s)(i^s) &= \text{expected value function of the next period when client in state } i^s \text{ is contacted} \\
(P_{s0}V^s)(i^s) &= \text{expected value function of the next period when client in state } i^s \text{ is not contacted} \\
\alpha &= \text{discount factor}
\end{aligned}$$

5.1 Upper Bound of Value Function

The upper bound derivation closely follows Whittle's approach [21]. In his paper, an upper bound for the undiscounted restless bandit problem with $M = 1$ was derived by replacing the per-period constraint with an averaged version, facilitating the introduction of an Lagrangian multiplier. We take an identical approach in deriving the upper bound for the discounted bandit problem with $M \geq 1$. The constraint of contacting M clients every period is relaxed to the following:

$$\mathbf{E}[\sum_{t=0}^{\infty} \alpha^t m(t)] = \frac{M}{1-\alpha} \tag{5.3}$$

In other words, we only need the discounted long run average of clients being contacted to be $\frac{M}{1-\alpha}$. The problem, in which the constraint is relaxed as in Eq. (5.3), can be reformulated as follows:

$$\max \quad \mathbf{E}[\sum_{s\in\mathcal{S}} \sum_{t=0}^{\infty} \alpha^t r^s] \tag{5.4}$$

$$s.t \quad \mathbf{E}[\sum_{s\in\mathcal{S}} \sum_{t=0}^{\infty} \mathbb{I}_s(t)\alpha^t] = \frac{|\mathcal{S}|-M}{1-\alpha} \tag{5.5}$$

where $\mathbb{I}_s(t) = 0$ when the client is contacted and $\mathbb{I}_s(t) = 1$ when the client is not contacted. Denote the solution to the relaxed constraint problem in Eq. (5.4) - Eq. (5.5) as $V_R(\mathbf{i})$. We solve the problem by introducing a Lagrangian multiplier κ.

$$\max_{\pi,\kappa} \quad \mathbf{E}[\sum_{s\in\mathcal{S}} (\sum_{t=0}^{\infty} \alpha^t (r^s + \kappa \mathbb{I}_s(t))) - \kappa(\frac{|\mathcal{S}|-M}{1-\alpha})] \tag{5.6}$$

Next, we maximize the primal variable π (the policy):

$$\max_{\pi} \quad \mathbf{E}[\sum_{s\in\mathcal{S}} (\sum_{t=0}^{\infty} \alpha^t (r^s + \kappa \mathbb{I}_s(t)))] \tag{5.7}$$

When $\kappa = 0$, for a particular client s, Eq. (5.7) reduces to

$$\max_{\pi} \sum_{t=0}^{\infty} \alpha^t r^s \tag{5.8}$$

The solution to Eq. (5.8) can be obtained by solving the equivalent problem stated as follows.

$$V^s(i^s) = \max\big[g_1^s(i^s) + \alpha(P_{s1}V^s)(i^s), g_0^s(i^s) + \alpha(P_{s0}V^s)(i^s)\big] \quad \forall \quad s \in \mathcal{S} \tag{5.9}$$

For a general κ, obtaining the solution to the formulation in Eq. (5.7) for a client s is equivalent to solving

$$V^s(i^s, \kappa) = \max\big[g_1^s(i^s) + \alpha(P_{s1}V^s)(i^s),\ \kappa + g_0^s(i^s) + \alpha(P_{s0}V^s)(i^s)\big] \quad \forall \quad s \in \mathcal{S} \tag{5.10}$$

It follows that the solution to the primal problem in Eq. (5.7) for a given κ is

$$\sum_{s \in \mathcal{S}} V^s(i^s, \kappa) = \max_{\pi} \quad \mathbf{E}\big[\sum_{s \in \mathcal{S}}(\sum_{t=0}^{\infty} \alpha^t(r^s + \kappa \mathbb{I}_s(t)))\big] \tag{5.11}$$

We now minimize the dual variable κ in order to solve the original problem in Eq. (5.6). For a given M, the solution is given by

$$V_R(\mathbf{i}, M) = \min_{\kappa} \quad \big[\sum_{s \in \mathcal{S}} V^s(i^s, \kappa) - \kappa(\frac{|S| - M}{1 - \alpha})\big] \tag{5.12}$$

The expression in Eq. (5.12) is quite intuitive. The term $\sum_{s \in \mathcal{S}} V^s(i^s, \kappa)$ gives the maximum possible reward for all clients. Since the associate is constrained to maintain a certain average, the total reward thus needs to be reduced accordingly by $\kappa(\frac{|S|-M}{1-\alpha})$.

Proposition 5.1.1. $V^*(\mathbf{i}) \leq V_R(\mathbf{i})$.

Proof. The upper bound immediately follows since the constraints for $V_R(\mathbf{i})$ is a relaxed version of the original constraint. □

5.2 Heuristics

In the absence of the optimal solution we turn to investigating various heuristics in order to solve the multiple client problem. The first heuristics we consider are myopic in nature. We follow by looking at a one-step lookahead policy and a policy based on linear programming relaxation.

5.2.1 Myopic Policies

We consider two simple myopic policies - the greedy policy (GR) and a modified version of the greedy policy (V^s). The greedy policy dictates that the associate picks M clients with the largest active single-stage reward every period. The ability to compute $V^s(i^s)$,[1] the unconstrained value function of a particular client, allows for a modified myopic policy. Instead of looking at the single-stage reward of each client, the associate selects M clients with the highest value of $V^s(i^s)$. Mathematically, this heuristic solves the following problem in each period.

$$\max_{\mathbf{u}} \sum_{s \in \mathcal{S}} V^s(i^s) \cdot u^s \tag{5.13}$$

$$s.t. \qquad \mathbf{u}'e = M \tag{5.14}$$

where $u^s \in \{0, 1\} \quad \forall \quad s \in \mathcal{S}$ is the s^{th} element of vector $\mathbf{u}$.

5.2.2 One Step Lookahead Policy (1_{sla})

We improve on the myopic policy by looking at one-step lookahead policies. Since we are able to calculate the value functions of each individual client, we approximate $V(\mathbf{i})$ by a separable function and let $V(\mathbf{i}) \approx \tilde{V}(\mathbf{i}) = \sum_{s=1}^{|\mathcal{S}|} V^s(i^s)$. Let $\mathcal{M}$ denote the set of clients selected. The one-step lookahead policy selects M clients that maximizes the following:

$$\sum_{s \in \mathcal{M}} g_1^s(i^s) + \sum_{t \notin \mathcal{M}} g_0^t(i^t) + \alpha \sum_{s \in \mathcal{M}} (P_{s1} V^s)(i^s) + \alpha \sum_{t \notin \mathcal{M}} (P_{t0} V^t)(i^t) \tag{5.15}$$

[1]See §3.5

Let

$$\sum_{t \notin \mathcal{M}} g_0^t(i^t) = \sum_{t \in \mathcal{S}} g_0^t(i^t) - \sum_{s \in \mathcal{M}} g_0^s(i^s) \tag{5.16}$$

$$\sum_{t \notin \mathcal{M}} (P_{t0}V^t)(i^t) = \sum_{t \in \mathcal{S}} (P_{t0}V^t)(i^t) - \sum_{s \in \mathcal{M}} (P_{s0}V^s)(i^s) \tag{5.17}$$

A simple substitution shows that Eq. (5.15) can be rewritten as

$$\begin{aligned} &\sum_{s \in \mathcal{M}} \left[g_1^s(i^s) - g_0^s(i^s) + \alpha\big((P_{s1}V^s)(i^s) - (P_{s0}V^s)(i^s)\big) \right] \\ &\quad + \sum_{t \in \mathcal{S}} \left[g_0^t(i^t) + \alpha(P_{t0}V^t)(i^t) \right] \end{aligned} \tag{5.18}$$

If we define an index $m(i^s)$ as in

$$m(i^s) = g_1^s(i^s) - g_0^s(i^s) + \alpha\big((P_{s1}V^s)(i^s) - (P_{s0}V^s)(i^s)\big) \tag{5.19}$$

then the solution to Eq. (5.18) is to pick the top M clients according to index $m(i^s)$. Note that when $M = 1$, this one step lookahead policy reduces to the policy derived by Bertsekas in [23]. By Proposition 1.3.7(b) in [23], the performance of the one-step lookahead policy is to within $\frac{2\alpha\epsilon}{1-\alpha}e$ of the optimal solution, where $V(\mathbf{i}) - \epsilon e \leq \tilde{V}(\mathbf{i}) \leq V(\mathbf{i}) + \epsilon e$.

The index $m(i^s)$ ranks each client according to a value which makes contacting and not contacting a client equivalent. Interestingly, a similar idea was proposed by Whittle in his seminal paper on restless bandits [21]. In this paper, Whittle proposed an index based on a project independent subsidy (v) pitched at the right level to ensure that M projects are active on average. Projects with subsidy values greater that v will be made passive while projects with subsidy values less than v will be made active. Whittle's index, formulated for the average cost problem,[2] recovers the classical Gittin's index, and was conjectured by Whittle and later shown by Weber and Weiss [24] to be asymptotically optimal for the exactly-m

[2]The equivalent for the discounted problem is derived in Eq. (5.12), where κ is the project independent subsidy.

problem under an additional technical condition. This index, although very intuitive, is only limited to the class of problems which is *indexable*, defined as follows.

Definition 2. *(From [21]) A project is* indexable *for a given discount factor* $0 < \alpha < 1$ *if the set of states where it is optimal to take passive action increases monotonically from* $\emptyset$ *to a full state space as* κ *grows from* $-\infty$ *to* $+\infty$*.*

Intuitively, the multiple client problem is indexable. For a given client, it is possible to rank all the states according to their respective break-even point, κ. As we sweep κ from $-\infty$ to $+\infty$, the set of states where is it optimal to not contact increases monotonically.

5.2.3 Primal-Dual Heuristic

Bertsimas and Niño-Mora [25], in an attempt to work around the indexability requirement, proposed a different index policy based on the primal and dual solutions to an equivalent linear programming (LP) formulation of Whittle's relaxation.

We define the following to facilitate the restating of their formulation:

$x_1^{i^s}$ = total expected discounted time that client s is in state i^s and active

$x_0^{i^s}$ = total expected discounted time that client s is in state i^s and passive

$\mathbf{x_s}$ = column vector of $x_1^{i^s}$ and $x_0^{i^s}$

$$\beta^{i^s} = \begin{cases} 1 & \text{if client } s \text{ is in state } i^s \text{ at time } t = 0 \\ 0 & \text{else} \end{cases}$$

$\mathcal{I}_s$ = finite state space of client s

$$\mathcal{Q}_s^1 = \{\mathbf{x_s} \in \mathcal{R}_+^{|\mathcal{I}_s \times 0,1|} | x_1^{j^s} + x_0^{j^s} = \beta^{j^s} + \alpha \sum_{i^s \in \mathcal{I}_s} p_1^{i^s j^s} x_1^{i^s} + p_0^{i^s j^s} x_0^{i^s}, j^s \in \mathcal{I}_s\}$$

The primal problem can now be stated as

$$(P) \qquad \max \sum_{s\in\mathcal{S}} \sum_{i^s\in\mathcal{I}_s} \left[g_1^s(i^s) x_1^{i^s} + g_0^s(i^s) x_0^{i^s} \right] \tag{5.20}$$

$$s.t \quad \mathbf{x_s} \in \mathcal{Q}_s^1, \quad s \in \mathcal{S} \tag{5.21}$$

$$\sum_{s\in\mathcal{S}} \sum_{i^s\in\mathcal{I}_s} x_1^{i^s} = \frac{M}{1-\alpha} \tag{5.22}$$

and the corresponding dual as

$$(D) \qquad \min \sum_{s\in\mathcal{S}} \sum_{j^s\in\mathcal{I}_s} \beta^{j^s} \kappa^{j^s} + \frac{M}{1-\alpha} \kappa \tag{5.23}$$

$$s.t \quad \kappa^{i^s} - \alpha \sum_{j^s\in\mathcal{I}_s} p_0^{i^s j^s} \kappa^{j^s} \geq g_0^s(i^s), \qquad i^s \in \mathcal{I}_s, s \in \mathcal{S} \tag{5.24}$$

$$\kappa^{i^s} - \alpha \sum_{j^s\in\mathcal{I}_s} p_1^{i^s j^s} \kappa^{j^s} + \kappa \geq g_1^s(i^s), \qquad i^s \in \mathcal{I}_s, s \in \mathcal{S} \tag{5.25}$$

$$\kappa \geq 0 \tag{5.26}$$

Let $\bar{x}_1^{i^s}, \bar{x}_0^{i^s}$ and $\bar{\kappa}^{i^s}, \bar{\kappa}$ denote the optimal solutions to the primal and dual problems respectively. Further, let the optimal reduced cost coefficients be as defined in Eq. (5.27) and Eq. (5.28).

$$\bar{\gamma}_0^{i^s} = \bar{\kappa}^{i^s} - \alpha \sum_{j^s\in\mathcal{I}_s} p_0^{i^s j^s} \bar{\kappa}^{j^s} - g_0^s(i^s) \tag{5.27}$$

$$\bar{\gamma}_1^{i^s} = \bar{\kappa}^{i^s} - \alpha \sum_{j^s\in\mathcal{I}_s} p_1^{i^s j^s} \bar{\kappa}^{j^s} + \bar{\kappa} - g_1^s(i^s) \tag{5.28}$$

Then the heuristic is as follows:

1. Given the current states of the $|\mathcal{S}|$ clients, compute the indices $\delta_{i^s} = \bar{\gamma}_1^{i^s} - \bar{\gamma}_0^{i^s}$

2. Contact M clients who have the smallest indices. In case of ties, contact the clients with $\bar{x}_1^{i^s} > 0$

This particular index is not as intuitive as Whittle's. In order to gain some insight, we note that $\bar{\gamma}_1^{i^s}$ and $\bar{\gamma}_0^{i^s}$ can be thought of as the rate of decrease in objective value of the primal problem per unit increase in $x_1^{i^s}$ and $x_0^{i^s}$ respectively. Thus, the smaller $\bar{\gamma}_1^{i^s}$, the better it is for a client to spend more active time in state i^s. In deciding which client to contact (or equivalently to make active), one would pick a client when $\bar{\gamma}_1^{i^s} < \bar{\gamma}_0^{i^s}$. This leads to a natural interpretation of the heuristic proposed - an associate would rank the clients according to $\bar{\gamma}_1^{i^s} - \bar{\gamma}_0^{i^s}$ in increasing order and pick the first M to contact.

5.3 Impact of the Control Mechanism

In this section we attempt to establish the impact of implementing a control mechanism. We do so by simulating the various heuristics and comparing their performance to the established upper bound in §5.1. In addition, we compare them to some "policies" commonly practiced by sales associates, which we elaborate on in the following paragraph.

We find that managers and associates alike often rank clients by the total amount of money spent. In addition, when not offered guidance, associates tend to call clients who are fresher in their minds. As they flip through their client lists, they naturally select clients who they recently helped. We use the following heuristics to capture these "policies".

1. TopRev - Rank clients according to the total amount of money spent. Pick the top M clients.

2. JustBought - Rank clients according to how recent and frequent they purchase. For clients in the buy states, rank according to the value of b. If needed, rank clients in the non-buy states according to their average revenue.

5.3.1 Methodology

Figure 5-1 depicts our methodology for simulating the heuristics. Prior to simulation, an initial set of clients is generated. For each heuristic, the same set of initial clients is used to simulate its performance. A single run consists of simulating all the heuristics for a predetermined T periods. Depending on the particular heuristic, clients will have different evolution processes. Next state realization is simulated using a draw from a uniform distribution in unit interval. If the result of the draw is less than or equal to the transitional probability of the client, the transition occurs.[3] The period revenue for each heuristic is then calculated, and the period counter is incremented. This process is repeated for T periods, upon which the simulation for a single run is complete. We stop the simulation at run l if the percentage difference between the mean of the first $l-1$ runs and l runs is less than 10^{-4}. In order to gauge the performance of these heuristics, we computed the upper bound using the formulation in Eq. (5.20) - Eq. (5.22). The result was checked using the alternative formulation in Eq. (5.12).

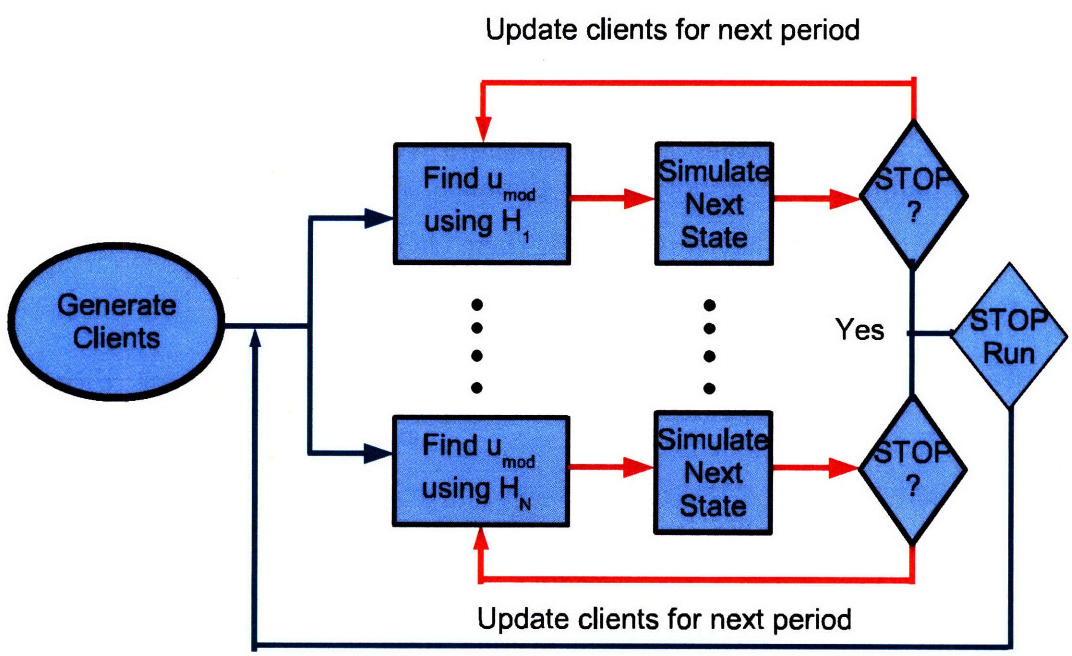

Figure 5-1: Simulation methodology.

[3]For example, if the probability of transitioning from a buy state of value 1 to 2 is .8, and if the value of the random draw is .7, the transition occurs.

5.3.2 Results

In this section we present the relative performance of the heuristics as compared to the upper bound.[4] The number of periods T is set such that terms where $\alpha^t \geq 10^{-4}$ are truncated.[5] Tables 5.1 and 5.3 break down the heuristics by the total revenue generated from all clients. The corresponding standard errors are tabulated in Tables 5.2 and 5.4. Each column tabulates the performance of various heuristics for a given value of constraint M. The best heuristic in each column is **boldfaced** while the worst is *italicized.* Figures 5-2 and 5-3 plot the performance of various heuristics relative to the upper bound.

We first note that the commonly practiced policies are the two worst performing heuristics. Among the two, the policy of going after clients who just purchased is worse. Next, we observe that the two myopic policies perform better than the two commonly practiced policies, with the greedy policy being the better of the two. The top two performers are the Primal-Dual and the One-step Lookahead policies, with both achieving $> 95\%$ of the upper bound. The lift achieved by using the best performing heuristic rather than one of going after clients who just purchased is significant. Simulation results indicate that the lift ranges from 65% to as high as 95% for a randomly generated group of clients.

5.3.3 Insights

In addition quantifying the promising lift achieved when the associates practice a better policy in contacting clients, the results confirm that the commonly practiced policies are indeed a misconception towards client work. Chasing after repeat buyers not only is suboptimal, it is one of the worst performing policies. In addition, the guideline of "calling as many as you can" often given out by management is vague and suboptimal. The methodology proposed can potentially assist management in determining the number of clients that should

[4]Weber and Weiss [24] showed that the upper bound is asymptotically optimal for the average cost per stage problem. We do not investigate the strength of the upper bound for the discounted problem since our goal is to provide a way to judge the *relative* performance of the heuristics and commonly practiced policies.

[5]Reducing the value of T does not change the relative ranking of the heuristics. The absolute performance of the heuristics, however, are further from the upper bound. This is because the upper bound is derived based on the infinite horizon assumption. Increasing the value of T improves the performance of some of the heuristics (P-D and 1_{sla}).

Heuristic	M=5	M=10	M=15	M=20
TopRev	28849.12	31386.02	33316.02	35099.47
Rank by b	*26264.18*	*26130.08*	*26738.67*	*27858.34*
V^s	31805.00	35568.68	37932.31	40055.35
Greedy	34281.28	38495.91	41474.78	43135.64
1_{sla}	39655.16	44054.00	46081.41	**46961.25**
Primal-Dual	**39686.90**	**44108.73**	**46114.95**	46568.91
Upper Bound	41196.23	45738.30	47492.26	48160.59

Table 5.1: Performance of heuristics for $|S| = 50$. The best performing heuristic is **boldfaced** while the worst performing heuristic is *italicized.*

Heuristic	M=5	M=10	M=15	M=20
TopRev	3.77	4.03	4.07	4.36
Rank by b	3.88	3.98	4.16	4.21
V^s	3.67	3.47	3.65	3.69
Greedy	3.63	3.96	4.31	4.09
1_{sla}	3.86	4.46	4.43	4.54
Primal-Dual	3.97	4.24	4.53	4.63

Table 5.2: Standard error for $|S| = 50$.

Heuristic	M=10	M=20	M=30	M=40
TopRev	54847.30	59016.72	62948.09	66583.35
Rank by b	*49538.49*	*49849.25*	*50704.48*	*52909.51*
V_s	60321.31	67310.76	72806.83	76916.82
Greedy	65804.51	74380.72	80432.77	84938.60
1_{sla}	**79170.50**	**88233.22**	**92249.02**	93453.69
Primal-Dual	79139.42	88039.57	91889.42	**93531.81**
Uppper Bound	81653.52	90247.17	94123.54	95587.07

Table 5.3: Performance of heuristics for $|S| = 100$. The best performing heuristic is **boldfaced** while the worst performing heuristic is *italicized.*

Heuristic	M=10	M=20	M=30	M=40
TopRev	7.08	7.16	7.99	7.66
Rank by b	7.39	7.70	8.12	8.32
V_s	6.81	7.14	7.18	7.20
Greedy	6.86	7.16	7.79	8.11
1_{sla}	7.80	8.20	7.96	8.66
Primal-Dual	7.61	8.15	8.46	8.19

Table 5.4: Standard error for $|S| = 100$.

be contacted each period. In addition, it can also assist sales associates in determining which clients should be selected. Given these tools, the variability in customer service level within a store could be reduced.

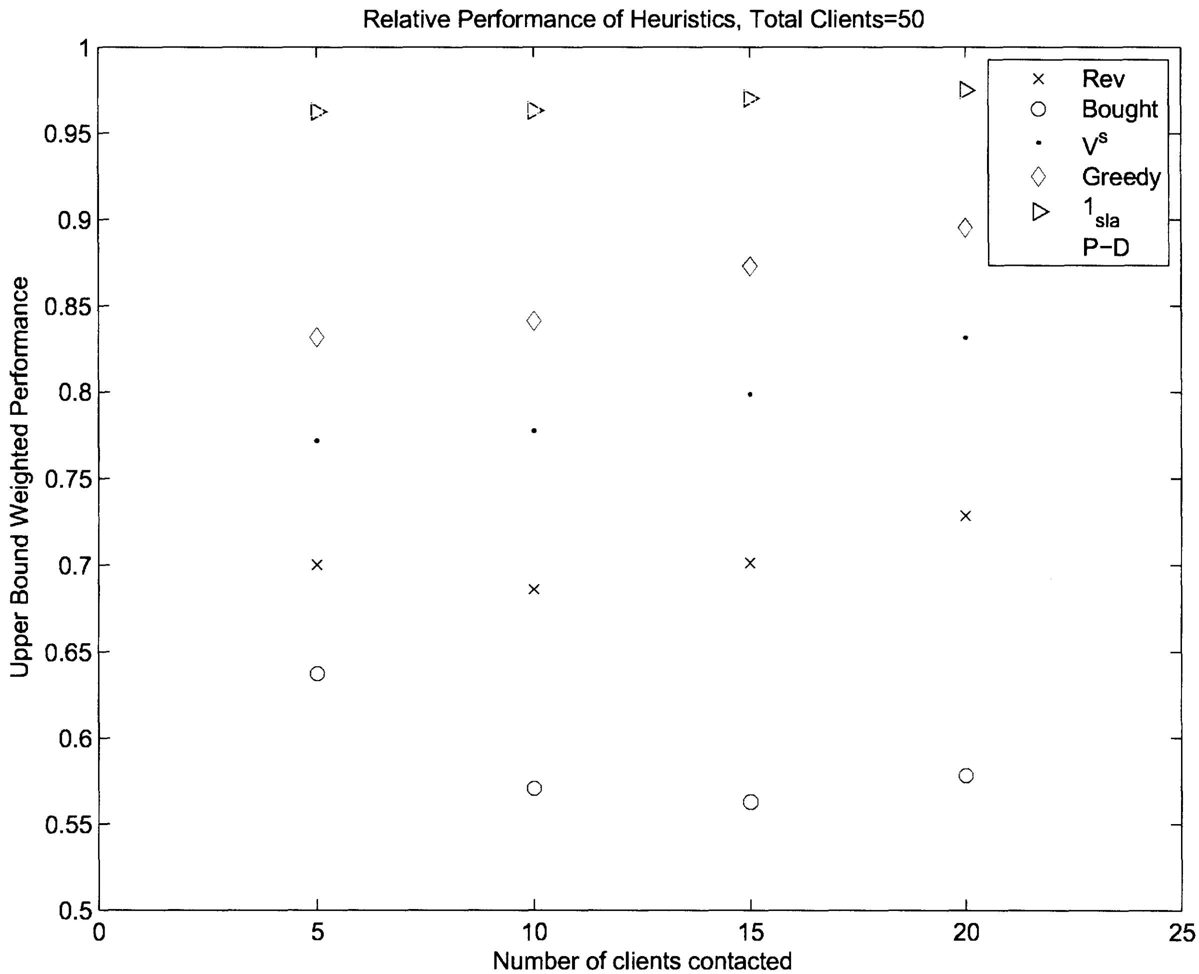

Figure 5-2: Relative performance of various heuristics when total number of clients = 50.

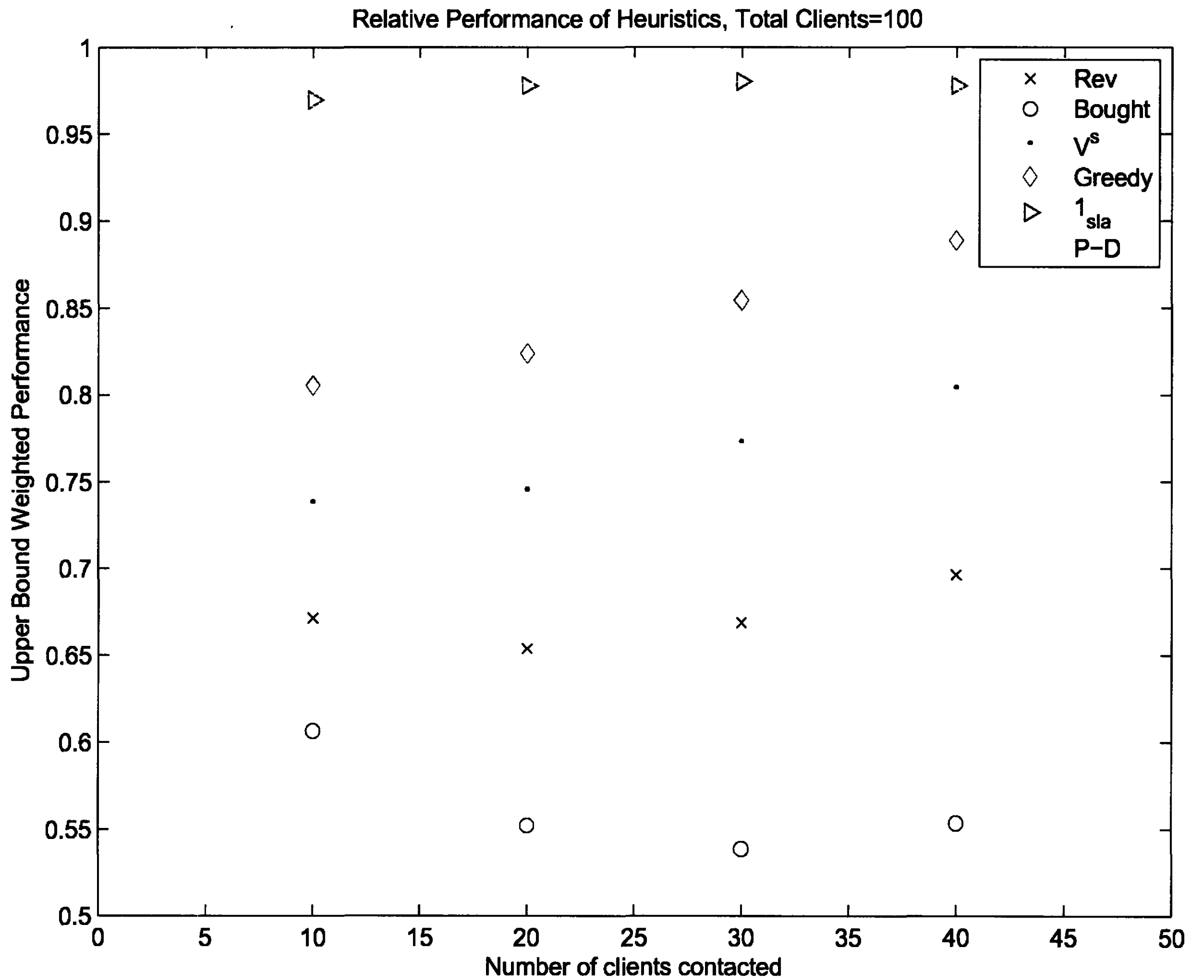

Figure 5-3: Relative performance of various heuristics when total number of clients = 100.

Chapter 6

Conclusions and Extensions

In this final chapter, we provide some concluding remarks and offer suggestions for extensions and future work.

6.1 Concluding Remarks

In this work we looked at a neglected problem in high-end retail - the control and management of sales associates. As we further researched the area, the existence of a disconnect between data currently collected - which has limited usefulness - and the data required for better control of the associates became apparent. We sidestepped the data unavailability issue by tapping the knowledge of many experienced sales associates through field work. This knowledge was then funneled back to assist in modeling client behavior. We proceeded to frame the control problem of an associate using dynamic programming. In analyzing the problem, we began by focusing on the case where an associate only has a single client and subsequently extended to the case where an associate has multiple clients. In both cases, the analysis was separately performed for clients in the buy states and the non-buy states.

When the associate only has a single client, we showed that the optimal policy under monotonic client behavior in the buy states is a threshold policy. Further, we showed that the optimal policy under peaking client behavior in the buy states is a quasi-concave policy.

On the other hand, when the client is in the non-buy states, the characterization of the optimal policy was hindered by the non-concavity and non-unimodality of the optimality condition. Computational results for a single client showed that the commonly used policy of *contacting clients who recently purchased* performed to within approximately 30% of the optimal policy.

When the associate has multiple clients, the problem was formulated as a restless bandit. Due to the complexity of such problems, we did not attempt to characterize the optimal policy. Instead, we provided an upper bound based on Lagrangian relaxation and proceeded to look at various heuristics that could be used to solve the problem. Computational results showed that the best heuristic performed to within approximately 96% of the upper bound, while the commonly practiced policy of *contacting clients who recently purchased* performed to within approximately 55% of the upper bound.

The qualitative data obtained through extensive field work coupled with the mathematical model in this work provided answers to many of the questions initially raised. We showed that current operations in the high-end retail stores are far from optimal and that associates are squandering opportunities to capture long-term potential of clients. The proposed mechanism is simple, implementable and will generate a substantial lift in revenue. With better control and management of the sales associates, the variability in customer service level will be reduced. In addition, the results also provide managerial insights and shed light on some common misconceptions.

In this work we have made several simplifying assumptions to assist in modeling and to maintain tractability. As a result, the model serves as a basis from which extensions could be built. In the following section we will address these assumptions and extensions.

6.2 Model Extensions and Future Work

6.2.1 Incorporating Updating Mechanism

In this work we have used a set of time invariant transitional probability to model the behavior of a particular client. Coupled with the choice of variables used for the state representation, the following assumptions were implicitly made:

1. The client behavior, due to the time invariant transitional probability, does not change with time.

2. The client behavior is independent of the total number of past interactions made between the client and the associate.

Both of these assumptions are quite limiting. One would expect the transitional probability to be non-stationary. With passage of time, the associate would have a better estimate of the client behavior. In addition, one would also expect a dependency between the client behavior and the total number of past interactions. The current state definition does not differentiate between a completely new client in state $(1,0)$ and an old client who happens to return back to state $(1,0)$.[1] Future work should look at ways to address these assumptions. One straightforward method is to enlarge the state space at the expense of computational complexity. In an effort to maintain computational ability, an alternative is to use an updating mechanism to capture the changing client behavior and client-associate relationship. For each new client, we propose an initial set of transitional probabilities as depicted in Figure 2-2. Each point on the transitional probability curve represents the response probability of a particular client in the given state i under policy u, which we model as a Bernoulli random variable with parameter $p_{i,u}$. The initial priors are set such that the mean of the Beta distribution, $\frac{\alpha}{\alpha+\beta}$, is equal to the estimated response. In each period, the updating occurs as depicted in Figure 6-1 for buy states and in Figure 6-2 for non-buy states.

[1]This could be justified if we ignore the impact of human relationships, but the main motivation behind using sales associates is to take advantage of the relationship bond between two parties.

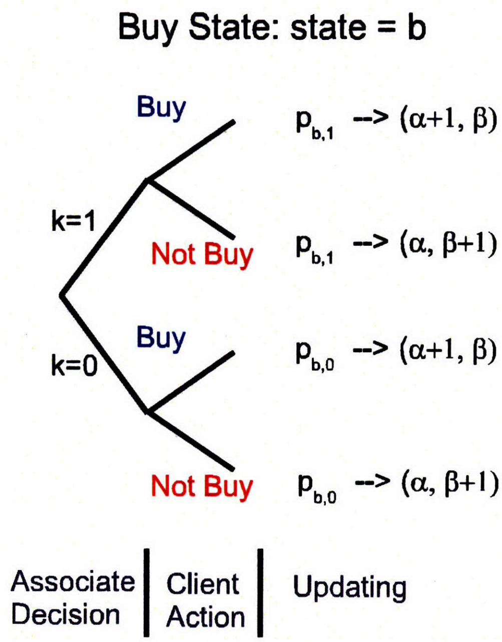

Figure 6-1: Updating for buy states. A *success* corresponds to a purchase and a *failure* corresponds to a non-purchase.

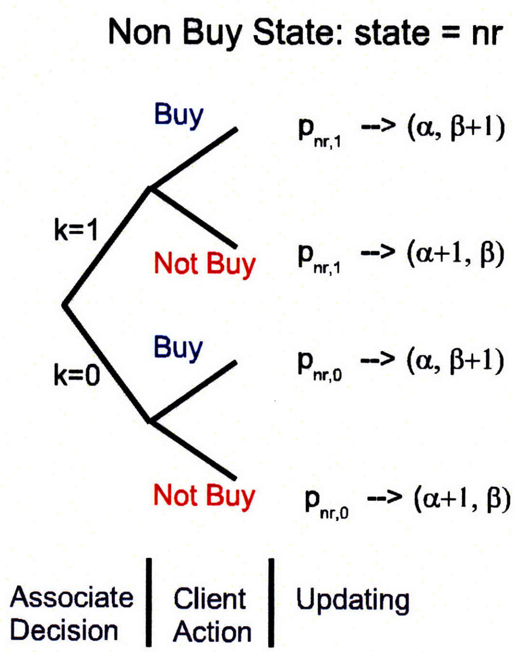

Figure 6-2: Updating for non-buy states. A *success* corresponds to a non-purchase and a *failure* corresponds to a purchase.

Unfortunately, the slow speed of learning in the high-end retail environment suggests that more research into the impact, practicality and efficiency of such an updating mechanism is needed. The time varying transitional probabilities also make the optimal policy characterization more difficult.

6.2.2 Expanded Control Structure

In this work we have modeled the associate as having a binary control of contacting or not contacting a client. Future work could look at other control variables available to the associate. One example of such a control variable is the type of message sent by the associate. In high-end retail, associates could contact clients for various reasons, including following up on a sale and calling to inform clients of special in-store charity or sales events. Incorporating this aspect into the model would increase the state space of the associate policy - instead of the decision of contacting or not contacting, the policy will also dictate the type of message to send during a contact.

Another control implicitly available to the associates is pricing. Since the retailing industry goes into sales season twice a year, the associate can potentially choose to call a particular client only during sales season. The ability to make this decision optimally hinges on knowing the type of client one has. Some clients value the assortment and size availability at the beginning of the season, thus are more likely to purchase at full price.[2] Others value the discounted price points, and will only purchase during sales season. Successful modeling using the expanded control structure will require a substantial amount of data to be gathered such that customer responses to various different types of control can be gauged.

6.2.3 Game Theoretic Framework

In this work, we used a dynamic programming framework to model the problem of an associate and his clients. Alternatively, one could have used a game theoretic framework. Such

[2]High-end retailers do not carry many units of each assortment. For example, a particular style and fabric of men's suit will come in various sizes ranging from 38R to 46R, but a store will only carry 2 units for each size.

a framework requires one to explicitly model the utility of a client, and could potentially provide some interesting results since it can capture the interaction between the associate and the client more accurately. Instead of a simple two-step interaction where the associate contacts the client and the client responds with a purchase decision, a more complete sequence of interaction is illustrated in Figure 6-3. In each period the associate makes a contact decision based on the client's state. The client then decides whether or not to visit the store. If the client chooses to visit, he will inspect the item and subsequently make a purchase decision. At the end of the period, the client updates the state of the associate by either upgrading, downgrading or maintaining his old state. The incorporation of a customer's feedback mechanism could be interesting for future work. In addition, one can also enrich the model to capture two stages of decision making from the client's perspective - the first is the decision to visit the store, the second is the decision to purchase.

Using a game theoretic framework also allows one to study the incentive misalignment in retailing. One obvious example of such misalignment occurs between the firm and the associate since the firm is focused on profit while the associate is focused on pure revenue. In addition, the game theoretic framework also lends itself well to modeling the in store competition between the sales associates and general competition between firms. These problems might have been studied in the economics literature, but it is worthwhile to look at them again under different a light.

6.2.4 Management of Sales Associate

This work looked at the management of sales associates at a tactical level. Future work could look into both strategic and operational levels. At the strategic level, one wonders if forming a one-to-one relationship between an associate and his clients is the best method for managing client relationships. Very often, a departing associate will bring his clients away with him. Would a gating process be a better way to manage these relationships? In a gating process, a client would evolve through different associates in his lifetime, each with increasing level of capability. At the operational level, one could look at scheduling of

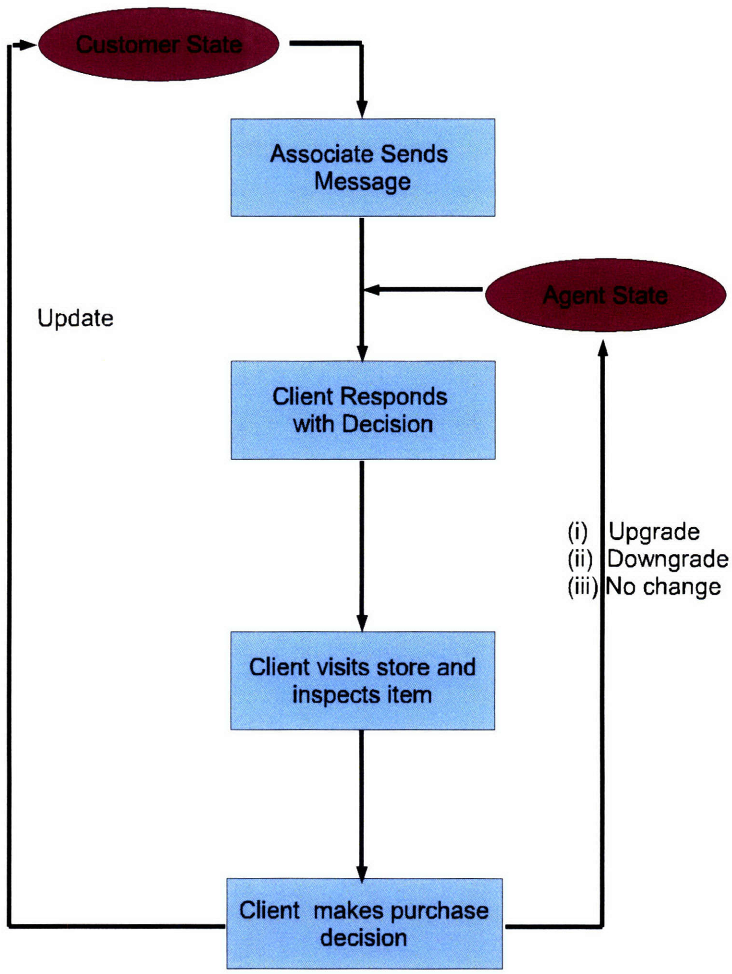

Figure 6-3: A detailed sequence of associate and client interaction.

associates to better match supply and demand in the services sector. From the associate's perspective, one could also look at the impact of the changing cost of contact on optimal policy.

Appendix A

Proofs

A.1 Proof for Proposition 3.1.6

Proof. We show for two cases, i) $u^* = 0 \quad \forall \quad b$ and ii) $u^* = 1 \quad \forall \quad b \leq b_h$, $u^* = 0 \quad \forall \quad b > b_h$.

Case 1: $u^* = 0$

In this case, we know that $V(b) = p_0(b)\big[R + \alpha V(b+1) - \alpha V(\bar{b})\big] + \alpha V(\bar{b}) \quad \forall \quad b$. Define b_{max} as the maximum state for $V(b)$. We construct $\tilde{V}(b)$ such that

$$\tilde{V}(b) = \begin{cases} V(b) & b \in [1, b_{max}] \\ V(b_{max}) & b = b_{max}+1, b_{max}+2 \end{cases}$$

We show by induction that $\tilde{V}(b) - \tilde{V}(b+1) \geq \tilde{V}(b+1) - \tilde{V}(b+2)$. Since $\tilde{V}(b) = V(b)$ for $b \in [1, b_{max}]$, $V(b)$ is also convex.

Base Case:

We show that $\tilde{V}(b_{max}) - \tilde{V}(b_{max}+1) \geq \tilde{V}(b_{max}+1) - \tilde{V}(b_{max}+2)$. This is true by construction.

(Induction Hypothesis) $\tilde{V}(b) - \tilde{V}(b+1) \geq \tilde{V}(b+1) - \tilde{V}(b+2)$.

We now show that $\tilde{V}(b-1) - \tilde{V}(b) \geq \tilde{V}(b) - \tilde{V}(b+1)$

$$\tilde{V}(b-1) = p_0(b-1)\big[R + \alpha\tilde{V}(b) - \alpha V(\bar{b})\big] + \alpha V(\bar{b}) \tag{A.1}$$

$$\tilde{V}(b) = p_0(b)\big[R + \alpha\tilde{V}(b+1) - \alpha V(\bar{b})\big] + \alpha V(\bar{b}) \tag{A.2}$$

$$\tilde{V}(b-1) - \tilde{V}(b) = p_0(b-1)\big[R + \alpha\tilde{V}(b) - \alpha V(\bar{b})\big] - p_0(b)\big[R + \alpha\tilde{V}(b+1) - \alpha V(\bar{b})\big] \tag{A.3}$$

Substituting $p_0(b-1) - p_0(b) = \delta_1$, we have

$$\tilde{V}(b-1) - \tilde{V}(b) = \delta_1[R + \alpha\tilde{V}(b) - \alpha V(\bar{b})] + p_0(b)(\alpha\tilde{V}(b) - \alpha\tilde{V}(b+1)) \tag{A.4}$$

Similarly, substituting $p_0(b) - p_0(b+1) = \delta_2$, we have

$$\tilde{V}(b) - \tilde{V}(b+1) = \delta_2[R + \alpha\tilde{V}(b+1) - \alpha V(\bar{b})] + p_0(b+1)(\alpha\tilde{V}(b+1) - \alpha\tilde{V}(b+2)) \tag{A.5}$$

By convexity of $p_0(b)$ and monotonicity of $V(b)$, $\delta_1 \geq \delta_2$ and $\alpha\tilde{V}(b) \geq \alpha\tilde{V}(b+1)$. By monotonicity of $p_0(b)$ and induction hypothesis, $p_0(b) \geq p_0(b+1)$ and $(\alpha\tilde{V}(b) - \alpha\tilde{V}(b+1)) \geq (\alpha\tilde{V}(b+1) - \alpha\tilde{V}(b+2))$. So we have shown that $\tilde{V}(b-1) - \tilde{V}(b) \geq \tilde{V}(b) - \tilde{V}(b+1)$ and that $V(b)$ is convex for this case.

Case 2: $u^* = 1 \quad \forall \quad b \leq b_h,\ u^* = 0 \quad \forall \quad b > b_h$

The proof for Case 1 can be duplicated here to show that $V(b)$ is convex in the region $b \in [b_h + 1, b_{max}]$. Since $V(b)$ is convex in this region, we have

$$\begin{aligned} &p_0(b_h)\big[R + \alpha V(b_h+1) - \alpha V(\bar{b})\big] - p_0(b_h+1)\big[R + \alpha V(b_h+2) - \alpha V(\bar{b})\big] \\ &\quad\geq p_0(b_h+1)\big[R + \alpha V(b_h+2) - \alpha V(\bar{b})\big] - p_0(b_h+2)\big[R + \alpha V(b_h+3) - \alpha V(\bar{b})\big] \\ &\quad= V(b_h+1) - V(b_h+2) \end{aligned} \tag{A.6}$$

Further, since $u^* = 1$ at b_h, we have

$$V(b_h) = p_1(b_h)\big[R + \alpha V(b_h+1) - \alpha V(\bar{b})\big] - c \geq p_0(b_h)\big[R + \alpha V(b_h+1) - \alpha V(\bar{b})\big] \tag{A.7}$$

and

$$V(b_h) - V(b_h+1) \geq V(b_h+1) - V(b_h+2) \tag{A.8}$$

We now need to show that $V(b_h-1) - V(b_h) \geq V(b_h) - V(b_h+1)$. Since $u^* = 1$ at $b_h - 1, b_h$ and $u^* = 0$ at $b_h + 1$, we have

$$V(b_h-1) = p_1(b_h-1)\big[R + \alpha V(b_h) - \alpha V(\bar{b})\big] + \alpha V(\bar{b}) - c \tag{A.9}$$

$$V(b_h) = p_1(b_h)\big[R + \alpha V(b_h+1) - \alpha V(\bar{b})\big] + \alpha V(\bar{b}) - c \tag{A.10}$$

$$V(b_h+1) = p_0(b_h+1)\big[R + \alpha V(b_h+2) - \alpha V(\bar{b})\big] + \alpha V(\bar{b}) \tag{A.11}$$

We want to show that

$$\begin{aligned} p_1(b_h-1)\big[R + \alpha V(b_h) - \alpha V(\bar{b})\big] - p_1(b_h)\big[R + \alpha V(b_h+1) - \alpha V(\bar{b})\big] \\ \geq p_1(b_h)\big[R + \alpha V(b_h+1) - \alpha V(\bar{b})\big] - c - p_0(b_h+1)\big[R + \alpha V(b_h+2) - \alpha V(\bar{b})\big] \end{aligned} \tag{A.12}$$

Since $u^* = 0$ at $b_h + 1$, we have the following inequality

$$-p_1(b_h+1)\big[R + \alpha V(b_h+2) - \alpha V(\bar{b})\big] > -c - p_0(b_h+1)\big[R + \alpha V(b_h+2) - \alpha V(\bar{b})\big] \tag{A.13}$$

We now need to show

$$p_1(b_h-1)\big[R+\alpha V(b_h)-\alpha V(\bar{b})\big]-p_1(b_h)\big[R+\alpha V(b_h+1)-\alpha V(\bar{b})\big]$$
$$\geq p_1(b_h)\big[R+\alpha V(b_h+1)-\alpha V(\bar{b})\big]-p_1(b_h+1)\big[R+\alpha V(b_h+2)-\alpha V(\bar{b})\big] \tag{A.14}$$

We showed in Eq. (A.8) that

$$V(b_h)-V(b_h+1)\geq V(b_h+1)-V(b_h+2) \tag{A.15}$$

Let $p_1(b_h)+\delta_1=p_1(b_h-1)$, then

$$p_1(b_h-1)\big[R+\alpha V(b_h)-\alpha V(\bar{b})\big]-p_1(b_h)\big[R+\alpha V(b_h+1)-\alpha V(\bar{b})\big] \tag{A.16}$$
$$=\delta_1\big[R+\alpha V(b_h)-\alpha V(\bar{b})\big]+p_1(b_h)(\alpha V(b_h)-\alpha V(b_h+1)) \tag{A.17}$$

Similarly, let $p_1(b_h+1)+\delta_2=p_1(b_h)$, we have

$$p_1(b_h)\big[R+\alpha V(b_h+1)-\alpha V(\bar{b})\big]-p_1(b_h+1)\big[R+\alpha V(b_h+2)-\alpha V(\bar{b})\big] \tag{A.18}$$
$$=\delta_2\big[R+\alpha V(b_h+1)-\alpha V(\bar{b})\big]+p_1(b_h+1)(\alpha V(b_h+1)-\alpha V(b_h+2)) \tag{A.19}$$

By convexity of $p_1(b)$ and monotonicity of $V(b)$, $\delta_1\geq\delta_2$ and $\alpha V(b)\geq\alpha V(b+1)$. By monotonicity of $p_1(b)$ and Eq. (A.15), $p_0(b)\geq p_0(b+1)$ and $(\alpha V(b_h)-\alpha V(b_h+1))\geq(\alpha V(b_h+1)-\alpha V(b_h+2))$. So we have shown that $V(b_h-1)-V(b_h)\geq V(b_h)-V(b_h+1)$. To get the base case for induction in region $b\in[1,b_h]$, we use $V(b_h-1)-V(b_h)\geq V(b_h)-V(b_h+1)$ to show

$$\begin{aligned}&p_1(b_h-2)\big[R+\alpha V(b_h-1)-\alpha V(\bar{b})\big]-p_1(b_h-1)\big[R+\alpha V(b_h)-\alpha V(\bar{b})\big]\\&\quad\geq p_1(b_h-1)\big[R+\alpha V(b_h)-\alpha V(\bar{b})\big]-p_1(b_h)\big[R+\alpha V(b_h+1)-\alpha V(\bar{b})\big]\end{aligned} \tag{A.20}$$

The induction step is shown in Eq. (A.14) - Eq. (A.19) and $V(b)$ is convex. □

Appendix B

Necessary Conditions for Peaking at b

Since $\hat{b} = \arg\max_b p_u(b)$, $V(b)$ is monotonically non increasing $\forall \quad b \geq \hat{b}$. Let $V(\hat{b}+1) + \delta_{\hat{b},\hat{b}+1} = V(\hat{b})$, where $\delta_{\hat{b},\hat{b}+1} > 0$. For $V(\hat{b}-1) \geq V(\hat{b})$ to hold, one of the following conditions need to hold.

1. $V(\hat{b}-1)_{u=0} \geq V(\hat{b})_{u=0}$
2. $V(\hat{b}-1)_{u=1} \geq V(\hat{b})_{u=1}$
3. $V(\hat{b}-1)_{u=0} \geq V(\hat{b})_{u=1}$
4. $V(\hat{b}-1)_{u=1} \geq V(\hat{b})_{u=0}$

We proceed to elaborate on each of the conditions.

1. $V(\hat{b}-1)_{u=0} \geq V(\hat{b})_{u=0}$

$$V(\hat{b}-1)_{u=0} = p_0(\hat{b}-1)[R + \alpha V(\hat{b}) - \alpha V(\bar{b})] + \alpha V(\bar{b}) \tag{B.1}$$

$$= p_0(\hat{b}-1)[R + \alpha V(\hat{b}+1) - \alpha V(\bar{b})] + \alpha p_0(\hat{b}-1)\delta_{\hat{b},\hat{b}+1} + \alpha V(\bar{b}) \tag{B.2}$$

$$V(\hat{b})_{u=0} = p_0(\hat{b})[R + \alpha V(\hat{b}+1) - \alpha V(\bar{b})] + \alpha V(\bar{b}) \tag{B.3}$$

The inequality holds when

$$p_0(\hat{b}-1)[R+\alpha V(\hat{b}+1)-\alpha V(\bar{b})]+\alpha p_0(\hat{b}-1)\delta_{\hat{b},\hat{b}+1} \geq p_0(\hat{b})[R+\alpha V(\hat{b}+1)-\alpha V(\bar{b})] \tag{B.4}$$

$$\delta_{\hat{b},\hat{b}+1} \geq [R+\alpha V(\hat{b}+1)-\alpha V(\bar{b})]\left[\frac{p_0(\hat{b})}{p_0(\hat{b}-1)}-1\right]\frac{1}{\alpha} \tag{B.5}$$

2. $V(\hat{b}-1)_{u=1} \geq V(\hat{b})_{u=1}$

$$V(\hat{b}-1)_{u=1} = p_1(\hat{b}-1)[R+\alpha V(\hat{b})-\alpha V(\bar{b})]+\alpha V(\bar{b})-c \tag{B.6}$$

$$= p_1(\hat{b}-1)[R+\alpha V(\hat{b}+1)-\alpha V(\bar{b})]+\alpha p_1(\hat{b}-1)\delta_{\hat{b},\hat{b}+1}+\alpha V(\bar{b})-c \tag{B.7}$$

$$V(\hat{b})_{u=1} = p_1(\hat{b})[R+\alpha V(\hat{b}+1)-\alpha V(\bar{b})]+\alpha V(\bar{b})-c \tag{B.8}$$

The inequality holds when

$$\delta_{\hat{b},\hat{b}+1} \geq [R+\alpha V(\hat{b}+1)-\alpha V(\bar{b})]\left[\frac{p_1(\hat{b})}{p_1(\hat{b}-1)}-1\right]\frac{1}{\alpha} \tag{B.9}$$

3. $V(\hat{b}-1)_{u=0} \geq V(\hat{b})_{u=1}$

A simple substitution shows that the inequality holds when

$$\delta_{\hat{b},\hat{b}+1} \geq [R+\alpha V(\hat{b}+1)-\alpha V(\bar{b})]\left[\frac{p_1(\hat{b})}{p_0(\hat{b}-1)}-1\right]\frac{1}{\alpha}-\frac{c}{p_0(\hat{b}-1)\alpha} \tag{B.10}$$

4. $V(\hat{b}-1)_{u=1} \geq V(\hat{b})_{u=0}$

A simple substitution shows that the inequality holds when

$$\delta_{\hat{b},\hat{b}+1} \geq [R + \alpha V(\hat{b}+1) - \alpha V(\bar{b})] \left[\frac{p_0(\hat{b})}{p_1(\hat{b}-1)} - 1 \right] \frac{1}{\alpha} - \frac{c}{p_1(\hat{b}-1)\alpha} \tag{B.11}$$

If none of the conditions hold, $V(\hat{b}-1) < V(\hat{b})$ and a peak occurs at $\hat{b}$. The above conditions can be generalized for any $b \leq \hat{b}$.

Appendix C

Table of Symbols

Symbol	Definition
i_t	State of a client in period t
$V_t(i_t)$	Optimal value function in period t when the client is in state i_t
$g(i_t, u_t)$	Single period expected reward
R	Average revenue an associate derives from a client if there is a purchase
c	Associate's cost of contact
α	Discount factor
$\mathcal{S}$	Set of clients belonging to an associate
$\mathbf{i_t}$	Vector of states corresponding to the set of clients
$\mathbf{u_t}$	Vector of policies corresponding to the set of clients
M_t	Number of clients to be contacted each period
$V_t(\mathbf{i_t})$	Optimal value function in period t for the set of clients
b	Number of consecutive buys
nr	Number of consecutive non-buys
u	Associate's policy, where $u \in \{0, 1\}$
$p_u(b)$	Probability of transitioning from state b into $b + 1$ under policy u
$\Delta_{b,u}$	$p_1(b) - p_0(b)$
$\bar{b}$	Maps to state (0,1)

Symbol	Definition
$p_u(nr)$	Probability of transitioning from state nr into $nr+1$ under policy u
$\Delta_{nr,u}$	$p_0(nr) - p_1(nr)$
b	Maps to state $(1,0)$ when used in the context of non-buy states
$\hat{b}$	$\arg\max_b p_u(b)$
$\tilde{b}$	$\arg\max_b V(b)$
$V(b)_u$	Value function in state b under the assumption that u is optimal
$\mathcal{B}_c$	Set of buy states where $\tilde{b}+2 \leq b \leq \hat{b}+1$
$\hat{nr}$	$\arg\max_{nr} p_u(nr)$
$\tilde{nr}$	$\arg\max_{nr} V(nr)$
$V(nr)_u$	Value function in state nr under the assumption that u is optimal
i^s	State of client s, equivalent to s_{th} element of i
$V^s(i^s)$	Value function of client s in state i^s
r^s	Single period reward yielded by client s, dependent on state and action
$g_1^s(i^s)$	Expected single period reward when client s in state i^s is contacted
$g_0^s(i^s)$	Expected single period reward when client s in state i^s is not contacted
$(P_{s1}V^s)(i^s)$	Expected value function of the next period when client s in state i^s is contacted
$(P_{s0}V^s)(i^s)$	Expected value function of the next period when client s in state i^s is not contacted

Bibliography

[1] Paul F. Nunes, Brian A. Johnson, and R. Timothy S. Breene. Selling to the moneyed masses. Harvard Business Review, 2004.

[2] Paul Nunes and Brian Johnson. *Mass Affluence: Seven New Rules of Marketing to Today's Consumer.* Harvard Business School Press, 2004.

[3] David Ellis. Are you a member of the mass affluent? CNNMoney.com, 2005.

[4] Laurie Simon Bagwell and B. Douglas Bernheim. Veblen effects in a theory of conspicuous consumption. *The American Economic Review*, 86(3):349–373, 1996.

[5] Giacomo Corneo and Olivier Jeanne. Conspicuous consumption, snobbism and conformism. *Journal of Public Economics*, 66:55–71, 1997.

[6] A. T. Coughlan and S. K. Sen. Salesforce compenstaion: Theory and implications. *Marketing Science*, 8:324–342, 1989.

[7] Dan Horsky and Paul Nelson. Evaluation of salesforce size and productivity through efficient frontier benchmarking. *Marketing Science*, 15(4):301–320, 1996.

[8] Birger Wernerfelt. On the function of sales assistance. *Marketing Science*, 13(1):68–82, 1994.

[9] Ann Marie Fiore and Hong Yu. Effects of imagery copy and product samples on responses toward the product. *Journal of Interactive Marketing*, 15(2):36–46, 2001.

[10] James D. Hess and Glenn E. Mayhew. Modeling merchandise returns in direct marketing. *Journal of Direct Marketing*, 11(2):20–35, 1997.

[11] Denise D. Schoenbachler and Geoffrey L. Gordon. Trust and customer willingness to provide information in database driven relationship marketing. *Journal of Interactive Marketing*, 16(3):2–16, 2002.

[12] Gabriel R. Bitran and Susana V. Mondschein. Mailing decisions in the catalog sales industry. *Management Science*, 42(9):1364–1391, 1996.

[13] Füsun Gönül and Meng Ze Shi. Optimal mailing of catalogs: A new methodology using estimable structural dynamic programming models. *Management Science*, 44(9):1249–1262, 1998.

[14] Ralf Elsner, Manfred Krafft, and Arnd Huchzermeier. Optimizing rhenania's mail-order business through dynamic multilevel modeling. *Interfaces*, 33(1):50–66, 2003.

[15] Duncan I. Simester, Peng Sun, and John N. Tsitsiklis. Dynamic catalog mailing policies. *Management Science*, 52(5):683–696, 2006.

[16] Nunzio Bellofatto. General manager of a retail store. Personal Conversation, 2006.

[17] Terri Davis. Assistant manager of a retail store. Personal Conversation, 2006.

[18] Cynthia M. Assistant manager of a retail store. Personal Conversation, May 2006.

[19] Elizabeth Knox. Buyer of a retail store. Personal Conversation, May 2006.

[20] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control, Vol I.* Athena Scientific, 2000.

[21] P. Whittle. Restless bandits: Activity allocation in a changing world. *Journal of Applied Probability*, 25A:301–313, 1988.

[22] Christos H. Papadimitriou and John N. Tsitsiklis. The complexity of optimal queuing network control. *Mathematics of Operations Research*, 24(2):293–305, 1999.

[23] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control, Vol II*, chapter 1.5. Athena Scientific, 2000.

[24] R. R. Weber and G. Weiss. On an index policy for restless bandits. *Journal of Applied Probability*, 27:637–648, 1990.

[25] Dimitris Bertsimas and Jose Nino-Mora. Restless bandits, linear programming relaxations, and a primal-dual index heuristic. *Operations Research*, 48(1):80–90, 2000.