4-6-2007

Value of Learning and Acting On Customer Information

H. Henry Cao
Cheung Kong School of Business

Bao-Hong Sun
Carnegie-Mellon University, bsun@andrew.cmu.edu

Preliminary Version Comments Welcome

Value of Learning and Acting On Customer Information

H. Henry Cao Baohong Sun *

Current Version: April 6, 2007

*H. Henry Cao is Professor of finance with the Cheung Kong Graduate School of Business, Beijing, 100738; e-mail: hncao@ckgsb.edu.cn. Sun is Associate Professor of Marketing with the Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA 15213, bsun@andrew.cmu.edu. We thank seminar participants at the University of Chicago for their valuable comments.

Value of Learning and Acting on Customer Information

Abstract

The rapid growth in demand and supply of sophisticated data mining and analytical decision tools calls for research to understand the value of learning, as well as how learning interacts with firms' day-to-day marketing strategies.

In this paper, we consider a market in which the firm can reduce its service cost when it matches the right product with the right customer. Facing uncertainty about customer types, the firm can gradually learn using observed service costs realized from recent interactions as noisy signals. On the basis of the most updated information, the firm makes matching and selection decisions to maximize its long-term profit. By solving the closed form solution to the fully dynamic optimization problem with infinite horizon, we analytically investigate the dynamic and endogenous nature of learning processes, the interaction between learning and decision making, and the evolution of profit over time. We also examine how optimal decision paths, market size in steady states and the evolution of profit are affected by parameters such as discount rate and the precision of information. Extending the model to a setting in which the precision of signals is proportional to the units of goods, we study how firms can endogenize the speed of learning. Our results shed light on the value of learning and acting on information, the time point to discontinue service to a customer, the duration and amount of short-term financial losses before learning pays off, and how the firm can modify its decisions to speed up the learning process.

Keywords: customer information management; customer heterogeneity; uncertainty; learning; customer-centric CRM; forward-looking; optimization; stochastic dynamic programming; continuous time; data mining; analytical decision making

1 Introduction

Today, marketing managers recognize that customer knowledge has become a promising means of building comparative advantages, defending competition, and growing revenue (Sawhney, et al. 2004). Customers have come to represent firm assets, and firms' marketing focus has shifted from “product-centric” or ”campaign-centric” to “customer-centric.” The traditional process of mass marketing is being challenged by a new interactive marketing approach (Blattberg and Deighton 1991; Haeckel 1998) or by one-to-one marketing that aims to profit from relationship building (Peppers et al. 1999; Wind and Rangaswamy 2001; Rust et al. 2004). This trend has been fueled by the rapid development of the Internet and digital technology, which tremendously increase the amount of detailed customer information available and create highly interactive environments for marketing communications (Shugan 2004).

Given the central role of individualization in moving toward customer-centricity, companies from all types of industries, ranging from manufacturing to online retailing, explore innovative ways to investigate customer information. Sophisticated tools and data warehousing and data mining technologies offer promising means of gaining detailed knowledge about each customer. Realizing that understanding customers provides only a foundation but that acting on information represents the ultimate step towards customer-centric marketing, firms also seek analytical marketing solutions that enable them to customize their marketing decisions, such as promotion, advertising, and product design, according to their most updated customer information. These data mining and analytical tools let them retrieve the most recent customer information, improve their knowledge on the preferences of individual customers, respond directly to customer requests, and provide customers with highly customized intervention decisions (Winer 2001). A new report by the independent market analyst Datamonitor forecasts global enterprise investments in data mining and analytical marketing tools will grow from an estimated \$2.3 billion in 2005 to more than \$3 billion in 2009.[1]

For example, the recommendation system adopted by Amazon.com allows the online re-

[1]See http://www.crm2day.com/news/crm/115241.php.

tailer to analyze customer preferences using their recent purchase histories and automatically recommends other products for the purpose of improving customers' shopping experiences and future cross-sell opportunities. Credit card companies use real-time data mining to screen credit applicants, predict credit fraud, and customize their offers. Ameritrade develops long-term, one-to-one marketing communication plans to improve the lifetime value of each customer by considering the benefits of customer acquisition over a long period of time. Department store Bloomingdale's adopted a proprietary system which accesses two years worth of customer data and delivers messages to the sales associates' personal digital assistants (PDAs) at the point of sales. Sales associate then informs customer of a special offer they qualify for. Casino Harrah's merged its customer loyalty and reservation stystem to identfy best customers and create targeted offers and rewards.[2]

In response to the growing demand for integrated data mining and analytical decision tools, increasing software applications, under the names of customer relationship management (CRM) or Web analytics (WA), have made significant inroads among commercial firms, which allows the adoptors to automate their data mining and decision making processes. For instance, the recently released On-Demand Customer Relationship Management by Saleforce.com claims that the software provides "point-and-click customization that fits the way you sell" and that "real-time analytics empower your business to make better decisions." Microsoft Dynamics CRM 3.0 Professional Datasheet claims that it "is a complete customer relationship management solution that provides all of the tools and capabilities needed to create and easily maintain a clear picture of customers from first contact through purchase and post-sales." Similarly, the concept and technology of e-business on-demand recently advocated by IBM is to host the CRM solutions for small and medium size companies.

Conceptually, the core idea behind the practice of customer-centric marketing is to identify customer heterogeneity (e.g., demand, taste, cost to serve, sensitivity to firm's decision variables) and customize marketing interventions that are relevant to the status and preferences of each customer. The goal is to acquire, nurture, and retain customers to maximize customer lifetime contribution to the company's profit, while controlling costs. Intuitively, three components characterize customer-centric marketing: Customer knowledge,

[2]See "The Customer-Centric Store: Deliverting the total experience, IBM Institute for Business Value." on web site http://www-935.ibm.com/services/us/imc/pdf/g510-4027-the-customer-centric-store.pdf.

which identifyies customers, follows their evolving demand, and/or tracks changes in customer preference (Rossi et al. 1996); foresight about customers' future reaction to current marketing interventions and hence their long-term profit contributions (Berger et al. 2002); and a sequence of customized and intertemporally related marketing treatments (Anderson 2006).

Then, implementing customer-centric marketing takes two iterative steps: (1) The firm continuously learns about each individual customer by analyzing results from the recent interactions, and then (2) the firm adapts its decisions according to its recent knowledge about each customer. The first step involves learning; the firm analyzes customer information on the basis of revealed customer reactions to the firm's most recent actions. The second step pertains to acting on information; the firm incorporates the updated knowledge into its marketing decisions, such as price, promotion, advertising, and product design. During these iterative processes, updated knowledge continuously adjusts firm decisions, and the resulting customer reactions again inform the learning process. Thus, learning and decision making are inter-dependent. In the rest of this paper, we term this increasingly observed business practice of continuous learning and acting on customer information for the purpose of maximizing long-term profit "customer information management" (CIM).

The rapid growth of demand for and supply of sophisticated data mining and analytical decision tools calls for research to understand the value of learning and how learning interacts with a firm's day-to-day marketing strategies. Unfortunately, data mining and analytical decision making followed separate paths in academic literature. On the one hand, marketing researchers develop statistical approaches to identify consumer heterogeneity (e.g., Kamakura and Russell 1989; Rossi et al. 1996) without explicitly deriving firm decisions. This line of research results in a snap shot segmentation of customers and scorerankings of consumers based on relevant variables. These segmentation methods are useful tools for campaign-centric marketing, which maximizes the return of investment of each campaign event by treating them independently over time. On the other hand, financial and information economists have explicitly treated firms as decision makers that optimally act on information. However, by focusing on information at the aggregate demand or product level, this stream of literature ignores firms' incomplete information about individual consumers.

Only recently has some sparse discussion about the idea of integrating learning and decision making appeared (e.g., Berger et al. 2002; Rust and Verhoef 2005; Venkatesan and Kumar 2004; Kamakura et al. 2006; Rust and Chung 2006; Sun et al. 2006). Some empirical works demonstrate how learning can be incorporated into marketing decisions for the purpose of maximizing customers' long-term profit contributions in the application of service matching (Sun and Li 2005). However, to our knowledge, no research systematically investigates the profit implications of firm learning or how the process and results of learning interact with a firm's day-to-day marketing decisions. The fundamental managerial and research question therefore is whether and how CIM enables companies to serve customers more effectively, reduce the cost of marketing and communication, and translate better customer knowledge into long-term profitability. Many specific issues pertaining to the benefits of CIM remain open.

First, how can firms better utilize individual customer information that proliferates in the modern highly interactive marketing and communication environment? Will learning and acting on customer information help firms move toward customer-centric decision making? An understanding of these issues will shed some light on how to improve the static and exogenous segmentation approaches commonly adopted by industry practice. It also will help companies planning to adopt CIM systems evaluate the value of their investment.

Second, how do we characterize the nature of the dynamic learning? How do the processes and results of learning interact with companies' marketing decisions? A firm must understand how learning makes its day-to-day marketing strategies different from those it would apply without learning. Alternatively, if the results of executing marketing decisions serve as input for the learning process, how can a firm alter its decisions to facilitate learning?

Third, how does the process of learning affect the steady-state market size in the long run? What are the implications for the intertemporal flow of profit? In other words, learning and acting upon information leads to better customer relationships and/or increased operating efficiency, which implies greater profitability? An understanding of this issue helps companies estimate the short-term costs of learning and the point at which learning and acting on customer information will start to pay off.

In this paper, we formulate a firm's learning and decision making as integrated solutions to a stochastic dynamic programming problem. This framework provides good approximations of those increasingly observed situations in which a firm, as a decision maker, learns about individual customer preferences, predicts future consequences associated with each possible intervention, calculates the long-term profit implications of current interventions, and chooses a sequence of optimal intervention that maximizes the sum of discounted future expected profits.[3] In particular, we assume a monopolistic firm provides a limited degree of customization in terms of two kinds of continuous services, A and B, to two types of customers, A and B. Assuming that both products have the same value to both customers, we let customers differ in terms of service cost – if they are provided with the right kind of service, expected service costs should be low, but otherwise, costs will be high. We also assume the firm remains uncertain about customer types at the beginning. However, the firm can rely on the observed service costs it realizes from recent interactions as noisy signals and gradually learn about customer types. Using the most updated knowledge about each customer, the firm makes matching (when service costs are not too high) and discontinuation (when service costs are too high) decisions to maximize long-term profit. Adopting a continuous time approach, we provide closed-form solutions to the fully dynamic optimization problem with an infinite horizon. We analytically investigate the properties of dynamic learning processes, the interaction between learning and decision making, and the evolution of profit over time. Furthermore, we examine how parameters such the discount rate and the precision of information affect firm decisions and profit. We also extend the model to a setting in which signal precision is proportional to the units of goods sold to shed light on how a firm's pricing strategies may be affected by learning.

Our analysis generates many interesting insights on how learning and decisions interact and the value of CIM. First, we find that learning enables a firm to improve its decisions over time by better matching services to customer types and discontinuing services to unprofitable customers. Under some parameters, the firm initially serves all customers. As the firm starts to learn, the perceived likelihood of the customer belonging to one particular type becomes heterogeneous. On the basis of its heterogeneous perceptions about individual customers, the firm may selectively decline to offer services to those deemed too costly to serve. As a result, the market size gradually decreases and converges to a lower bound, at which the customer

[3] In the rest of this paper, we use the terms products and services interchangeably.

types of all remaining customers are known. This approach is very different from the optimal strategy of a firm without learning, which mandates that it offer the same product to all customers at all time periods. As a result, a firm without learning denies all customers with negative profit from the very beginning.

Second, when the firm discontinues services to customers with negative expected profit flows, the firm with learning is less likely to deny customers with expected negative profits at the beginning, because it recognizes the possible misclassification and the ability to correct this error in the future. The firm must find the right balance between making a type I error and a type II error. If the firm discontinues services to those customers whose expected profit is only slightly negative, it may suffer too much type II errors, because it turns away many customers who are potentially profitable customers. However, if the firm waits too long and keeps providing services to those customers who are expected to be highly unprofitable, it could be making too much type I errors in providing services to those to whom the firm should discontinue its offers. The optimal point thus depends on the intricate trade-offs of these two types of errors. We show that firms are willing to make more type I errors when the discount rate is low and information precision is high.

Third, we show that dynamic learning and acting on customer information improves long-term profits because the firm is better able to match customers with the right services and discontinue services to unprofitable customers, which lowers its service costs. A learning firm will incur a short-term profit loss, because being able to identify customers in later periods enables it to recoup all its losses in the long run, which contrasts with the constant profit flow of firms without learning. We analyze the intertemporal pattern of expected profit and thus determine the duration, maximum, and accumulated financial loss before learning and acting on information pays off for the firm. Such analyses of the duration and the maximum accumulated loss allows the firm to make better financial planning. We further demonstrate that the improvement of profit (or value of information) decreases with the discount rate and increases with information precision.

Fourth, we illustrate that firm decisions can be modified to facilitate learning. When the precision of learning per unit of time is proportional to the units sold to customers, the firm will undercut prices to induce the customer to buy more units of services for the purpose of speeding up the learning process. We demonstrate that it is beneficial for the firm to treat learning as part of the decision process and view the speed of learning itself as an endogenous variable.

The rest of the paper is organized as follows: In section 2, we review the related literature. In section 3, we develop a dynamic model with gradual learning and then extend the model to multiple units of goods in which the precision of learning per unit of time is proportional to the units sold to customers. In section 4, we discuss the managerial implications. Finally, in section 5, we conclude the paper and discuss some limitations and future research directions.

2 Literature Review

Substantial economics research deals with learning about (aggregate) consumer demand or product quality uncertainty. For example, Grossman et al. (1977), McLennan (1984), and Trefler (1993) study the pricing decision of a monopolist with demand uncertainty. Harrinton (1995) considers the case of a duopoly, and Judd and Riordan (1994) and Bergemann and Valimake (1997) study pricing strategies when both buyers and sellers are uncertain about a product's quality. Focusing on learning at the aggregate level, the literature on information economics analyzes aggregate demand uncertainty or product quality. Recently, some marketing research papers have discussed whether firms should share knowledge about individual customers with its competitors. In a static competitive setting, Chen et al. (2001) analytically study whether it is optimal to share customer information with competitors in a state of imperfect targetability, that is, when the firms' ability to predict the purchase behaviors of individual consumers for the purpose of customizing prices or product offers, is imperfect. They show that the competing firms can improve their ability to identify price-sensitive switchers and price-insensitive loyal customers by sharing information.

Another related stream of research proposes empirical models to determine more personalized levels of marketing interventions and thereby manage long-term customer value (e.g., Bult and Wansbeek 1995; Schmittlein and Peterson 1994; Gonul and Shi 1998; Venkatesan and Kumar 2004; Rust and Verhoef 2005; Lewis 2005; Netzer et al. 2005; Sun and Li 2005; Sun et al. 2006). For example, controlling for customer heterogeneous characteristics, Gonul and Shi (1998) study the optimal direct mail policy in a dynamic environment that enables customers to maximize utility and the direct mailer to maximize profit. Assuming a myopic company, Rust and Verhoef (2005) derive the optimal marketing intervention mix for intermediate-term customer relationship management (CRM). Lewis (2005) adopts a dynamic programming based approach to derive the optimal pricing policy of newspaper subscriptions that adjusts discounts as the customer relationship evolves. Hitsch (2006) estimates an empirical model of optimal product launch and exit under demand uncertainty. In a static setting, Netzer et al. (2005) construct and estimat a hidden Markov model that dynamically segments customers; they demonstrate how a firm can alter consumer buying behavior. Narayanan and Manchanda (2006) model heterogeneity across physicians in terms of their rates of learning about new drugs and run simulations to demonstrate that firms can improve their resource allocations for marketing communication decisions by taking into account heterogeneity across consumers and over time. Although some of these articles adopt optimization, none of them focus on the integrated learning, forward-looking, and optimization components that are essential for CIM.

Recently, Sun et al. (2005) and Sun (2006 a, b) suggested a conceptual framework of a two-step procedure (adaptive learning and proactive marketing decisions) with three components for customer-centric decision making (adaptive learning of customer individual preferences, forward-looking into future marketing consequences of current marketing interventions, and optimization to balance cost and benefit). In an empirical application to service allocations, Sun and Li (2005) formulate service allocation decisions of a call center as a solution to a stochastic dynamic programming problem. They empirically demonstrate that learning and acting on long-term marketing consequences prompts optimal allocation decisions that improve customer retention, reduce service costs, and enhance profit. With the applications to service allocation, these papers only provide empirical evidence about learning and acting on information; therefore, further research must explicitly investigate the learning process and systematically establish how dynamic and endogenous learning interact with firm decisions and improve profit in a general setting.

Methodologically, our work relates to dynamic structural models with learning developed to examine consumers' dynamic decisions regarding brand, quantity, and purchase timing (e.g., Erdem and Keane 1996), stockpiling behavior (e.g., Assuncao and Meyer 1993; Krishna 1992; Gonul and Srinivasan 1996; Sun et al. 2003; Erdem et al. 2003), and consumption (Sun 2005). Chintagunta et al. (2006) offer a good summary of research in this area. Unlike prior research, which attempts to establish that consumers are sophisticated decision makers, we treat firms as decision makers that learn about customers, take into account the effect of current marketing interventions on future customer reactions, and optimally adjust their marketing interventions to maximize customer lifetime value.

3 The Dynamic Model with Gradual Learning

We consider a gradual process of information revelation in a dynamic model with continuous time and an infinite horizon. While the continuous time model is only an approximation to the real world setting, it does provide us with tractability and closed form solutions, which allows us to investigate the inter-temporal nature of gradual learning and acting upon customer information. In the continuous time model, the monopolistic firm maximizes discounted expected customer life time profit contribution with an infinite horizon by making optimal decisions on whether and what service to provide and what price to charge to each customer. In the next few subsections, we describe the information structure and learning, followed by studying the firm's decisions, and then analyzing firm's steady state customer size and profit in the long-run.

3.1 Information and Learning Process

Consider a market in which customers receive a service continously over time. Suppose a monopolistic firm provides two types of services, $S = A, B$, to a unit mass of heterogeneous customers. Customers are indexed by $i \in [0, 1]$, and they are of two types $T_i = A, B$. Each customer has one unit of demand. The value flow of each unit of service is v for these customers.

The firm makes the following decisions: whether or not to provide the service, which service to provide if it decides to provide service, what price to charge after the firm has decided which service to provide. Let D_{it} be the dummy variable indicating whether to provide service, and S_{it} be the type of service the firm provides at time t to customer i. Let p_{it} be the price the firm charges, and $q_{it} = 0, 1$ denote the units of service customer i purchases at time t.

The cost of serving customer i is $dc_i(S_{it})$ between time t to time $t + dt$ when service S_{it} is offered, and we assume,

$$dc_i(S_{it}) = c(T_i, S_{it})dt + \sigma d\hat{W}_{it}, \tag{1}$$

where

$$c(T_i, S_{it}) \equiv \begin{cases} c_l & \text{for } T_i = S_{it} \\ c_h & \text{for } T_i \neq S_{it} \end{cases} . \tag{2}$$

The function $c(T_i, S_{it})$ defined in Equation (2) is the service cost flow without the noise term. The process $\hat{W}_{it}$ is a standard Brownian motion and is independent across customers. Equation (1) indicates that the firm does not have perfect information about service cost, which is the customer type-specific service cost $c(T_i, S_{it})dt$ plus a noise term $\sigma d\hat{W}_{it}$. The coefficient σ measures the size of the noise in the service cost, such that when σ is large, the firm learns little from observing the cost. We assume that $c_h > v > c_l$, so the firm will match S_{it} with T_i if it can identify customer i's type.

Let λ_{i0} be the prior belief the firm has about the probability that customer i is of type A. We assume that the indicator variable that customer i is of type A is i.i.d across customers, which implies that λ_{i0} is the same, and $\lambda_{i0} = \lambda_0$. Following the law of large numbers, the proportion of type A customers in the population is λ_0.

Let λ_{it} denote the posterior belief of the firm at time t that customer i is of type A. We then have

$$\lambda_{it} = Pr(i \in A | F_{it}),$$

where F_{it} is the information set the firm has about customer i at time t. We let the realized cost observed from the past serve as the information source. Thus, F_{it} consists of the history of observed service costs generated from all the past interactions with customer i.

Over time, the firm observes the actual serving costs realized from past interactions and updates its belief about the type of the customer λ_{it} in a Bayesian fashion. Let $\bar{c}_i(S_{it})$ denote the expected cost flow of customer i at time t, such that $\bar{c}_i(S_{it}) = E[c(T_i, S_{it})|F_{it}]$, and $d\lambda_{it}$ denote the change of λ_{it} over time from t to $t + dt$. Following Lipster and Shiryaev (1977), the updating process of λ_{it} using Kalman filtering of the realized cost follows a process

$$\begin{aligned}
d\lambda_{it} &= [-\frac{\lambda_{it}(1-\lambda_{it})(c_h - c_l)}{\sigma^2}[c(T_i, S_{it}) - \bar{c}_i(S_{it})]dt - \frac{\lambda_{it}(1-\lambda_{it})(c_h - c_l)}{\sigma}d\hat{W}_{it}] \\
&\quad \times [I(S_{it} = A) - I(S_{it} = B)] \\
&= \frac{\lambda_{it}(1-\lambda_{it})(c_h - c_l)}{\sigma}dW_{it},
\end{aligned}$$

where $I(E)$ denotes the indicator function for event E, such that $I(E) = 1$ if E occurs and 0 otherwise, and

$$dW_{it} = [I(S_{it} = B) - I(S_{it} = A)]\{[dc_i(S_{it}) - \bar{c}_i(S_{it})dt]/\sigma\}$$

is a Brownian motion process from the firm's perspective. The process of the posterior belief λ_{it} has zero drit and the difference between realized cost and expected cost $dc_i(S_{it}) - \bar{c}_i(S_{it})dt$ serves as the new information.

We define $\sigma(\lambda_{it})$ as the instantaneous standard deviation of λ_{it},

$$\sigma(\lambda_{it}) \equiv \frac{\lambda_{it}(1-\lambda_{it})(c_h - c_l)}{\sigma}. \tag{3}$$

The variable $\sigma(\lambda_{it})$ measures how fast λ_{it} is updated and is linear in the ratio of signal to noise $(c_h - c_l)/\sigma$. When σ is smaller, there is less noise in the observed cost, and the firm learns more about the customer. In addition, λ_{it} gets updated faster, and the corresponding instantaneous volatility in the change of λ_{it} is higher. Similarly, when $c_h - c_l$ is larger, the signals coming from different types of customers are more distant from one another on average, so we obtain more informative signals. As a result, updating of λ_{it} is faster as well. Note that as λ_{it} moves toward either 0 or 1, the standard deviation of $d\lambda_{it}$ and the updating moves to zero. In the long run, the customer type can be identified as time reaches infinity.

3.2 Firm's Action

We assume that the firm is risk neutral and maximize the present value of discounted expected profits with a discount rate of r. We assume that the firm can only provide one service at a time. The expected profit contributed by consumer i over the time interval t to $t + dt$ is:

$$d\pi_{it} = [p_{it} - \bar{c}_i(S_{it})]D_{it}q_{it}(p_{it})dt.$$

Because the reservation price for the customer is v, the customer will purchase the service at time t if and only if $p_{it} \leq v$. Consequently, the firm always chooses $p_{it} = v$ if it wants to sell to the customer. Moreover, customer i buys one unit of service from the firm, $q_{it}(p_{it}) = 1$, at $p_{it} = v$. The expected profit $d\pi_{it}$ reduces to:

$$d\pi_{it} = [v - \bar{c}(S_{it})]D_{it}dt.$$

For any given adapted strategy path $\{S_{it}D_{it}\}$, the life time profit contribution of customer i starting from the time 0 to infinity is

$$\Pi_i(\{S_{it}, D_{it}\}, \lambda_0) = E\int_0^{\infty} e^{-rt}[v - \bar{c}(S_{it})]D_{it}dt.$$

The firm's problem thus is to find an optimal path (S_{it}, D_{it}) adapted to the firm's information set at time t (F_{it}) to maximize customer i's long-term expected profit contribution to the firm. The firm's posterior belief about the customer at time t (λ_{it}) serves as the state variable that the firm can act on, and the firm's decisions are driven by λ_{it}.

Given the symmetry of the cost function between serving A and B, the optimal strategy of service matching is

$$S_{it} = S(\lambda_{it}) \equiv \begin{cases} A, \text{ for } \lambda_{it} \geq 0.5 \\ B, \text{ for } \lambda_{it} < 0.5. \end{cases} \tag{4}$$

As a result, the expected cost flow under the optimal service strategy, denoted by $c(\lambda_{it})$

given service $S(\lambda_{it})$, is

$$\begin{aligned} c(\lambda_{it}) \equiv \bar{c}_i(S(\lambda_{it})) &= \begin{cases} \lambda_{it}c_l + (1-\lambda_{it})c_h, \text{for} \lambda_0 \geq 0.5 \\ (1-\lambda_{it})c_l + \lambda_{it}c_h, \text{for} \lambda_0 < 0.5 \end{cases} \\ &= (0.5 + |\lambda_{it} - 0.5|)c_l + (0.5 + |\lambda_{it} - 0.5|)c_h. \end{aligned} \tag{5}$$

Next, we determine the optimal strategy for D_{it} using the Bellman equation, which states that the value function (or maximized cumulative profit on the optimal path) satisfies sequential optimality. To derive the Bellman equation, we consider the value function at time t with updated belief λ_{it}:

$$V(\lambda_{it}) \equiv E_t \int_t^{\infty} e^{-r\tau}[v - c(\lambda_{i\tau})]D^*_{i\tau} d\tau.$$

On the optimal path where $D^*_{it} = 1$ over the interval $[t, t+dt]$, we have

$$V(\lambda_{it}) = [v - c(\lambda_{it})]dt + (1 - rdt)E_t[V(\lambda_{i(t+dt)})]. \tag{6}$$

Using Ito's lemma, as shown in the Appendix, the Bellman equation reduces to the following inhomogeneous ordinary differential equation:

$$V(\lambda_{it}) = \frac{v - c(\lambda_{it})}{r} + \frac{\sigma^2(\lambda_{it})}{2} V''(\lambda_{it}).$$

The value function has two terms that can be understood as follows: The first term $(v - c(\lambda_{it}))/r$ is the present value of the profit if the firm does not implement CIM. The firm receives a constant expected profit flow of $v - c(\lambda_{it})$ at each period of time, and the present value for such a perpetuity is $(v - c(\lambda_{it}))/r$. The second term is the difference of the expected profit on the optimal path using information and the expected profits without information. Consequently, the second term represents the gains from acting on information on the optimal path, which is proportional to the instantaneous volatility of λ_{it} $(\sigma(\lambda_{it}))$ as defined in Equation (3) and V''. When the firm uses the information optimally, it must receive more profits, so the gains from information on the optimal path must be positive. Consequently, $V'' \geq 0$, and V is convex in λ_{it}. Note that λ_{it} is a martingale and from

Jensen's inequality, the mean of the value function is greater than the value function of the mean. Consequently, the expected value function must increase over time.

Convexity occurs because the firm has both switching and termination options. With learning, the firm can reduce the mismatch between services and customers and discontinue services to unprofitable customers. This option is more valuable when λ_{it} is more volatile, such that the firm learns more quickly about customer types.

The solution for the value function is given by the following theorem:

Theorem 1 *The value function is convex and symmetric around $\lambda_{it} = 0.5$, i.e. $V(\lambda_{it}) = V(1 - \lambda_{it})$. Moreover, if $\lambda_{it} \geq 0.5$ and the firm is providing service to customer i, we have*

$$V(\lambda_{it}) = \frac{v - c(\lambda_{it})}{r} + b\lambda_{it}^{-(\gamma-1)/2}(1 - \lambda_{it})^{(\gamma+1)/2}, \tag{7}$$

where

$$\gamma \equiv \sqrt{1 + \frac{8r\sigma^2}{(c_h - c_l)^2}} > 1.$$

The coefficient b in the general solution of the value function will be determined subsequently. Before we proceed, note that V is a convex function and symmetric around $\lambda_{it} = 0.5$. Consequently, the value function must achieve its minimum at $\lambda_{it} = 0.5$. If $V(\lambda_{it})$ reaches 0, the best strategy of the firm is to provide no more service for this customer. As a result, there are two possible cases. In the first case, the value function is positive for all λ_{it}, and the firm provides service for all λ_{it}. Specifically, it provides service A for $\lambda_{it} \geq 0.5$ and service B for $\lambda_{it} < 0.5$. We term the first case service matching, for which the coefficient b can be determined by the continuity of the first-order derivative at the switching boundary $\lambda_{it} = 0.5$. In the second case, there exists an interval $[1 - \hat{\lambda}, \hat{\lambda}]$ around 0.5 such that the value function is 0 in the interval. When $\lambda_0 > \hat{\lambda}$, the firm starts with service A but stops whenever $\lambda_{it} = \hat{\lambda}$. Similarly, when $\lambda_0 < 1 - \hat{\lambda}$, the firm starts with service B and stops when $\lambda_{it} = 1 - \hat{\lambda}$. We term the second case customer selection, because the firm will gradually discontinue service to unprofitable customers and the market size dwindles to a constant of less than 1. In the second case, the coefficient b and the optimal stopping point $\hat{\lambda}$ can be jointly determined by the continuity of the value function and its first-order derivative at the stopping point $\hat{\lambda}$.

Service Customization

We first consider the service customization case in which the firm provides service for all λ_{it}. Because V is a convex, symmetric function around $\lambda_{it} = 0.5$, it must achieve its minimum at $\lambda_{it} = 0.5$ and thus $V'(0.5) = 0$. Alternatively, by symmetry, we have $V(\lambda_{it}) = V(1 - \lambda_{it})$, and it follows that $V'(\lambda_{it}) = -V'(1 - \lambda_{it})$. At $\lambda_{it} = 0.5$, the continuity of the first-order derivative implies that $V'(0.5) = -V'(0.5)$, which gives $V'(0.5) = 0$. From the condition $V'(0.5) = 0$, we can determine the coefficient b:

$$b = \frac{(c_h - c_l)}{r\gamma}.$$

For the customization policy to be optimal, the minimum point of the value function, $V(0.5)$, must be positive. Given

$$V(\lambda_{it}) = \frac{v - c(\lambda_{it})}{r} + \frac{(c_h - c_l)}{r\gamma}\lambda_{it}^{-(\gamma-1)/2}(1 - \lambda_{it})^{(\gamma+1)/2},$$

we have,

$$V(0.5) = \frac{v - (c_h + c_l)/2}{r} + \frac{(c_h - c_l)}{2r\gamma}$$

The condition for the customization policy to be optimal, $V(0.5) > 0$, reduces to

$$\frac{(c_h - v)(\gamma - 1)}{(c_h - c_l)(\gamma - 1) + 2(v - c_l)} < 0.5.$$

The value function of the service customization case is as given in Proposition 1.

Proposition 1 *When*

$$\frac{(c_h - v)(\gamma - 1)}{(c_h - c_l)(\gamma - 1) + 2(v - c_l)} < 0.5,$$

Figure 1: Value Function under Service Customization

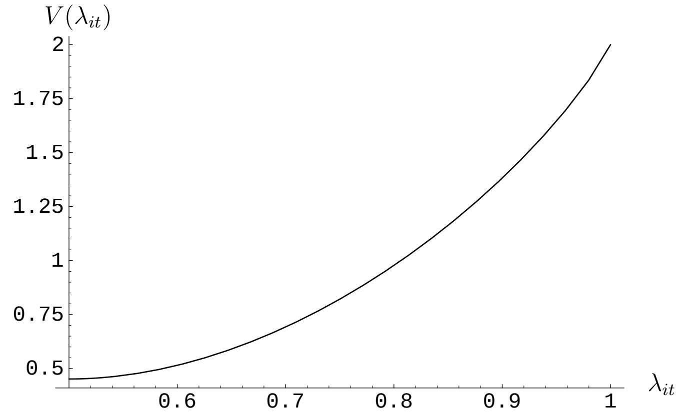

The value function with information vs λ_{it}. We have $v = 0.1, c_l = 0, c_h = 1, r = 0.05$, and $\sigma = 1$.

the customization strategy is optimal, and the firm will switch between services A and B as λ_{it} crosses 0.5. The value function is

$$V(\lambda_{it}) = \frac{v - c(\lambda_{it})}{r} + \frac{(c_h - c_l)}{r\gamma}\lambda_{it}^{-(\gamma-1)/2}(1-\lambda_{it})^{(\gamma+1)/2}$$

for $\lambda_{it} \geq 0.5$ *and* $V(\lambda_{it}) = V(1-\lambda_{it})$ *for* $\lambda_{it} < 0.5$.

In Figure 1, we plot the value function for $\lambda_{it} > 0.5$. Notice that $V(0.5) > 0$, which implies that the firm provide service to all customers at all times. Moreover, the value function increases with λ_{it} and is convex.

Customer Selection

In line of the preceding discussion, when

$$\frac{(c_h - v)(\gamma - 1)}{(c_h - c_l)(\gamma - 1) + 2(v - c_l)} \geq 0.5,$$

the cumulative profit is negative around $\lambda_{it} = 0.5$ with the customization strategy. Therefore, the customization strategy is not optimal, and the firm will not provide services for λ_{it} in the neighborhood of 0.5. Due to the symmetry of services, there exists an interval $[1-\hat{\lambda}, \hat{\lambda}]$, such that the firm will not provide services for any customer whose λ_{it} falls in this stop region. Specifically, $\hat{\lambda}$ and $1-\hat{\lambda}$ are stopping points for services A and B, respectively.

Note that $1-\hat{\lambda} \leq \hat{\lambda}$ implies that $\hat{\lambda} \geq 0.5$. In this case, suppose that $\lambda_0 \geq \hat{\lambda}$, in which case the firm will initially provide service A and stop providing services to customer i when λ_{it} reaches $\hat{\lambda}$. Similarly, when $\lambda_0 < \hat{\lambda}$, the firm will initially provide service B and then stop at the point $\lambda_{it} = 1-\hat{\lambda}$. We next determine the stopping point $\hat{\lambda}$ and the coefficient b in the value function.

In the stop region $[1-\hat{\lambda}, \hat{\lambda}]$, no service is provided, and the value function is 0. Outside the stop region, the value function is given by Equation (7). At point $\lambda_{it} = \hat{\lambda}$, the value function must satisfy the value matching and smooth pasting conditions (see Lipster and Shiryayev 1979), such that the value function and its first-order derivative must be continuous at the stopping point $\hat{\lambda} \geq 0.5$:

$$V(\hat{\lambda}) = 0, V'(\hat{\lambda}) = 0.$$

Plugging these conditions back to Equation (7), we get:

$$\hat{\lambda} = \frac{(c_h - v)(\gamma - 1)}{(c_h - c_l)(\gamma - 1) + 2(v - c_l)}$$

and

$$b = \frac{2(v - c_l)}{r(\gamma - 1)} \left(\frac{\hat{\lambda}}{1 - \hat{\lambda}} \right)^{(\gamma+1)/2}$$

Note that for the solution to be consistent, we must have $\hat{\lambda} \geq 0.5$, as shown previously. The value function for $\hat{\lambda} < 0.5$ can be obtained using the symmetry argument that $V(\lambda_{it}) = V(1-\lambda_{it})$, and we derive the following:

Proposition 2 *In the customer selection scenario, the optimal stopping point for service A is $\lambda_{it} = \hat{\lambda}$, where*

$$\hat{\lambda} = \frac{(c_h - v)(\gamma - 1)}{(c_h - c_l)(\gamma - 1) + 2(v - c_l)} \geq 0.5,$$

Figure 2: Value Function under Customer Selection

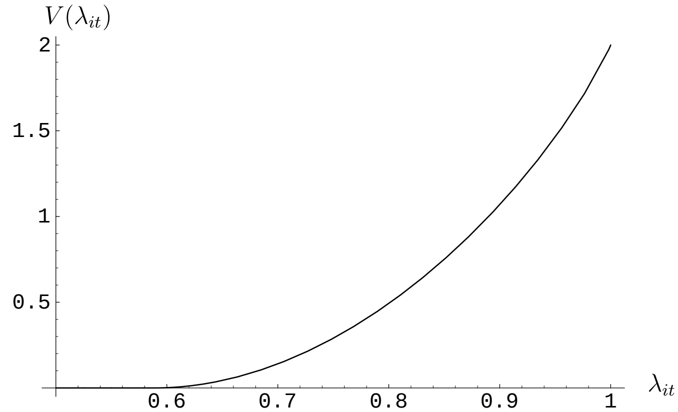

The value function with information vs λ_{it}; $v = 0.1, c_l = 0, c_h = 1, r = 0.05$, and $\sigma = 1.5$.

and the optimal stopping point for service B is $1 - \hat{\lambda}$. *The value function for* $\lambda_0 > \hat{\lambda}$ *is*

$$V(\lambda_{it}) = \frac{v - c(\lambda_{it})}{r} + b\lambda_{it}^{-(\gamma-1)/2}(1 - \lambda_{it})^{(\gamma+1)/2}.$$

The value function for $\lambda_0 < 1 - \hat{\lambda}$ *is* $V(\lambda_{it}) = V(1 - \lambda_{it})$. *Finally, the value function for* λ_{it} *in* $[1 - \hat{\lambda}, \hat{\lambda}]$ *is 0.*

In Figure 2, we plot the value function in the customer selection case with respect to λ_{it} for $\lambda_{it} \geq 0.5$. The value function is increasing and convex for $\lambda_{it} \geq \hat{\lambda}$ and 0 for λ_{it} between 0.5 to $\hat{\lambda}$.

The constant $\hat{\lambda}$ represents the optimal stopping point at which the firm can give up on a customer and discontinue service. The determination of the stopping point is very intriguing. We can show that $\hat{\lambda} < \lambda^*$, where λ^* is the profit flow breakeven point of the firm. Thus, the firm loses money on average serving a customer with $\hat{\lambda}$. To determine the optimal stopping point, the firm faces a trade-off about which point to discontinue service. If the firm chooses a $\hat{\lambda}$ too high, while reducing short-term losses, it will make too much type II error. Many A type customers will be discontinued, and long-term profits will decline. If the firm chooses a $\hat{\lambda}$ too low, it will make too much type I error. Many B type customers

Figure 3: Stopping Point as a Function of the Discount Rate

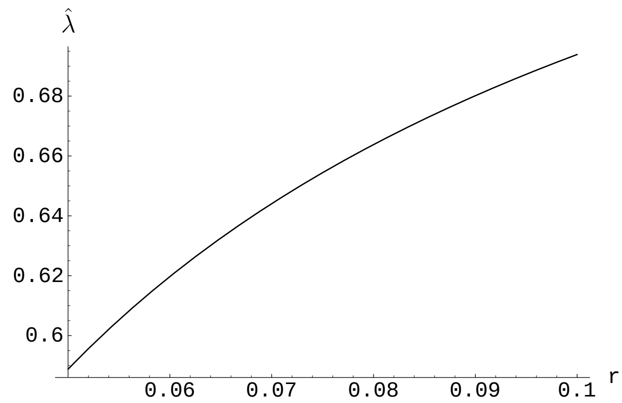

$v = 0.1, c_l = 0, c_h = 1, and \sigma = 1.$

Figure 4: Stopping Point as a Function of the Noise in the Information

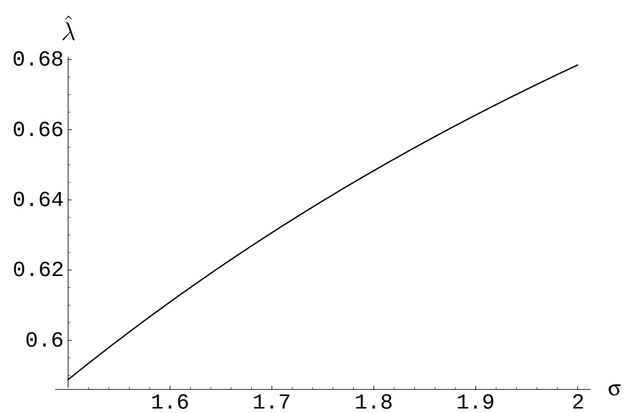

The stopping point vs. the standard deviation of noise σ; $v = 0.1, c_l = 0, c_h = 1$, and $r = 0.05$.

will receive service A, and the firm will be losing too much money to these customers in the short run, even though the firm increases its long-term profit flows. The optimal $\hat{\lambda}$ has to balance the intertemporal trade-off between type I versus type II error.

We note that $\hat{\lambda}$ is an increasing function of γ that increases with discount rate r and imprecision of the signal, σ. In Figures 3 and 4, we plot $\hat{\lambda}$ as a function of r and σ, respectively. To summarize, we have the following proposition:

Proposition 3 *The stopping point $\hat{\lambda}$ increases with the imprecision of information σ and the discount rate r.*

Intuitively, when the discount rate r is smaller, future profits become more important, and thus, it becomes more important for the firm to identify customers. In this case, type II error becomes more important, and the firm is less likely to give up on the customer, which means $\hat{\lambda}$ is smaller. When σ is small, the firm obtains more precise information and is willing to learn by keeping the customer. Type I error is less likely to happen in the future, and as a result, $\hat{\lambda}$ is smaller.

3.3 Value of Information

We now determine the value of information, defined as the difference between the profit functions with and without information. When the firm does not learn, it only provides the service when $v - c(\lambda_{it}) > 0$. The firm's expected profits at time t will be a constant $(v - c(\lambda_{it}))^+$.

Thus, the value of information is

$$VI(\lambda_{it}) = V(\lambda_{it}) - (v - c(\lambda_{it}))^+/r. \tag{8}$$

The comparative statics of the value of information with respect to λ_{it}, σ, and r appear in Proposition 4.

Proposition 4 *In both the customer selection and the service customization scenarios, the value of information increases with λ_{it} for $\lambda_{it} < \lambda^*$ but decreases for $\lambda_{it} > \lambda^*$. The value of information decreases with σ and r.*

In Figures 5 and 6, for $\lambda_{it} > 0.5$, we plot the value of information in the customization and selection scenarios as a function of λ_{it}, respectively. Both value functions with and without learning increase with λ_{it}. However, for $\lambda_{it} < \lambda^*$, the value function without learning is 0, so the value of information increases with λ_{it} for $\lambda_{it} < \lambda^*$. For $\lambda_{it} > \lambda^*$, the expected profit remains the same at all times for the firm that does not learn. The dependence of its expected profits on λ_{it} also remains the same for all time points. In contrast, the firm with the ability to learn can better match the customer with the right service or discontinues service to unprofitable customers. As a result, the value function with learning is less dependent on the initial value of λ_{it}. Because both value functions are increasing functions of λ_{it}, for $\lambda_{it} > \lambda^*$, the value of information declines with λ_{it}. The maximum value of information is achieved at λ^*.

Proposition 4 also states that the value of information decreases with r and σ. Intuitively, when r is smaller, the future is more important. Learning also is more valuable because the firm can make more informative decisions in the future, so the value of information is higher. Similarly, when σ is small, the information is more precise, and the firm can learn faster, which implies that the value of information is higher.

3.4 Steady-State Market Size in the Case of Customer Selection

In the customer selection case, customer i is discontinued once λ_{it} hits the stopping point $\hat{\lambda}$. Therefore, we next consider a few interesting issues regarding the properties and outcomes of the learning process. Will the market size dwindle to 0 or converge to a positive size? How fast does convergence occur? How might the proportion of customer types of remaining customers change over time?

Again, we focus on the case of $\lambda_0 > 0.5$, so the firm serves service A to all customers. Let $H(t)$ denote the probability that customer i will hit the stopping point $\hat{\lambda}$ before time t.

Figure 5: Value of Information under Service Customization

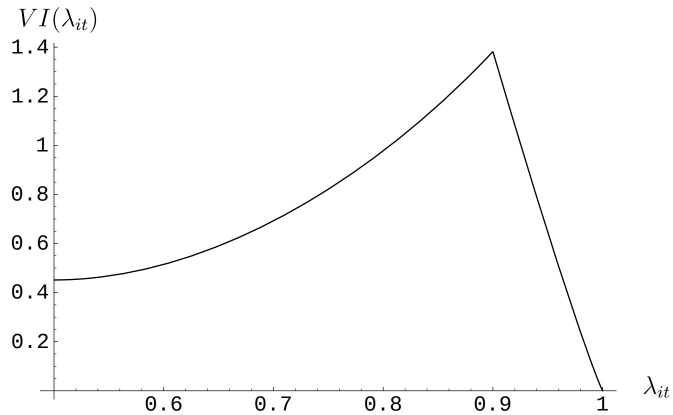

The value of information vs λ_{it}; $v = 0.1, c_l = 0, c_h 1, r = 0.05$, and $\sigma = 1$.

Figure 6: Value of Information under Customer Selection

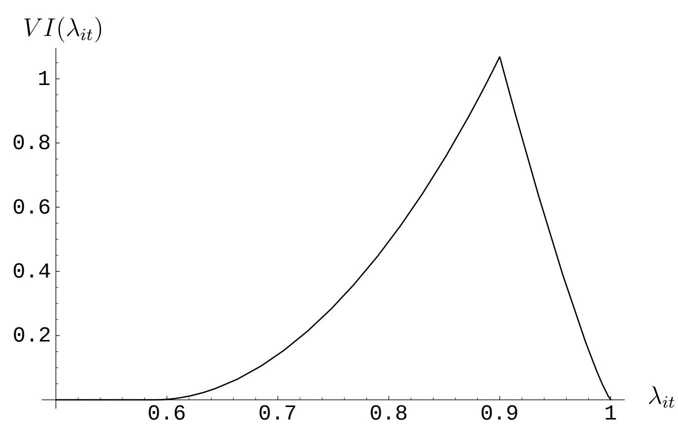

The value of information vs λ_{it}; $v = 0.1, c_l = 0, c_h 1, r = 0.05$, and $\sigma = 1.5$

By the law of large numbers, $H(t)$ is also the proportion of customers in the population that hits $\hat{\lambda}$. As a result, $1 - H(t)$ is the size of the remaining customers at time t. Determination of $H(t)$ allows the firm to analyze how fast the market stabilizes and how the average λ_{it} of the remaining customers changes over time. The expression of $H(t)$, which we derive in the Appendix, is as follows:

$$H(t) = \frac{\lambda_0}{\hat{\lambda}}[N(\frac{(\hat{y} - y_0) - \sigma_s^2 t/2}{\sigma_s \sqrt{t}}) + \exp(\hat{y} - y_0) N(\frac{(\hat{y} - y_0) + \sigma_s^2 t/2}{\sigma_s \sqrt{t}})],$$

where

$$\hat{y} = \ln\left(\frac{\hat{\lambda}}{1 - \hat{\lambda}}\right)$$

and

$$y_0 = \ln\left(\frac{\lambda_0}{1 - \lambda_0}\right).$$

As approaches infinity, $H(t)$ converges to a constant:

$$\lim_{t \to \infty} H(t) = \frac{1 - \lambda_0}{1 - \hat{\lambda}}.$$

Thus, the higher the $\hat{\lambda}$, the higher $H(\infty)$ and the lower the long-run market size $1 - H(\infty)$. Moreover, $H(\infty) < 1$, and thus, the long-run market size is strictly above 0 for $\lambda_0 > \hat{\lambda}$, such that the firm provides service to customers initially. Let $\bar{\lambda}(t)$ denote the proportion of A type customers among the remaining customers. By the law of large numbers, $\bar{\lambda}(t)$ is the expectation of λ_{it}, conditional on $\lambda_{it} > \hat{\lambda}$. We must have

$$(1 - H(t))\bar{\lambda}(t) + H(t)\hat{\lambda} = \lambda_0,$$

because λ_{it} is a martingale for all i. Thus,

$$\bar{\lambda}(t) = \frac{\lambda_0 - H(t)\hat{\lambda}}{1 - H(t)}.$$

Because $H(t)$ is increasing in t, $\bar{\lambda}(t)$ must also increase in t. We have Proposition 5.

Figure 7: Customer Size

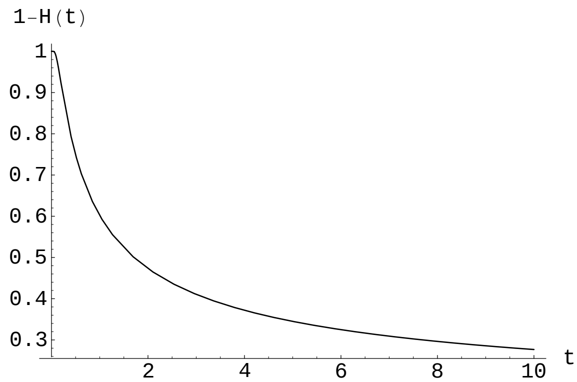

The market size over time; $v = 0.1, c_l = 0, c_h = 1, r = 0.05, \sigma = 1.5$, and $\lambda_0 = 0.7$.

Proposition 5 *Suppose that $\hat{\lambda} > 0.5$. As t approaches infinity, the market size converges to a constant, $\frac{\lambda_0 - \hat{\lambda}}{1-\hat{\lambda}}$. The market size decreases with σ and r. The proportion of the A type of customers , $\bar{\lambda}(t)$, moves to 1, as t approaches infinity.*

In Figure 7, we plot market size as a function of t. The remaining customers decrease but converge to a constant. The market size decreases with $\hat{\lambda}$. When $\hat{\lambda}$ is large, each customer is more likely to hit, $\hat{\lambda}$ and thus, $H(t)$ is higher, and the steady-state market size is lower. As we discussed previously, the stopping point $\hat{\lambda}$ increases with σ and r, and the market size decreases with σ and r. Intuitively, with smaller σ and r, the firm is able to learn faster and has a stronger long-term perspective, so the steady-state market size increases.

In Figure 8, we plot the $\bar{\lambda}(t)$ of the remaining customers. As the firm gets rid of its perceived unprofitable customers, the $\bar{\lambda}(t)$ of retained customers increases monotonically with time. Determining $\bar{\lambda}(t)$ allows the firm to recognize approximately how fast it will be able to identify customers. For example, if the firm needs to know the time point at which 95% of the remaining customers will be of A type, that can do so by simply inverting $\bar{\lambda}(t)$ to get $\bar{\lambda}^{-1}(0.95)$.

Figure 8: Average Lambda of Remaining Customers

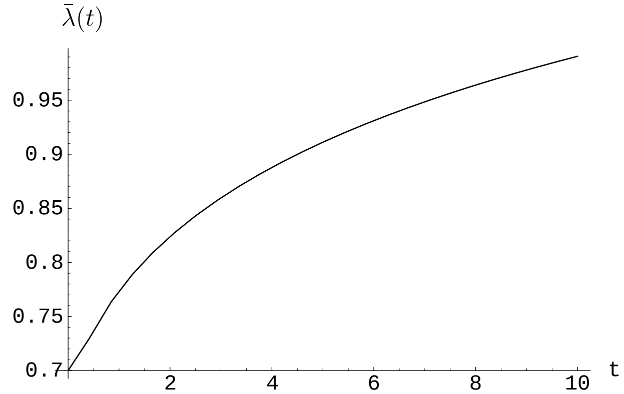

The proportion of A types among active customers over time;
$v0.1, c_l = 0, c_h = 1, r = 0.05, \sigma = 1.5$, and $\lambda_0 = 0.7$.

3.5 Implications for Aggregate Profits

Without customer information, the firm's knowledge about customers remains the same, and the firm uses the same decision rule for all customers at all times. As a result, the expected profits will remain the same over time. However, with customer information, the firm can learn and act on customer information through service customization or customer selection, and thus, its expected profits increase over time. We show that when $\lambda_0 < \lambda^*$, the firm implementing CIM loses money in earlier periods. Although the firm eventually makes up for the loss of profit in later periods, it needs to know the intertemporal pattern of profits and losses for two main reasons. First, the firm can determine the duration of financial losses before profits turn around. Second, the firm can identify its maximum loss, as well the accumulative amount of loss it will incur and make financial plans accordingly.

Because learning is independent across customers, the population density is the same as the probability density. In the customization case, the aggregate profit across customers for all t is given by

$$\pi(t) = \int_i \pi_i(\lambda_{it})di = E[v - c(\lambda_{it})] = v - c(E[0.5 + |\lambda_{it} - 0.5|]) = v - c(0.5 + E[|\lambda_{it} - 0.5|]).$$

Notice that λ_{it} is a martingale, and $v - c(0.5 + E[|\lambda_{it} - 0.5|])$ is a convex function of λ_{it}.

Consequently, $v - c(0.5 + E[|\lambda_{it} - 0.5|])$ is a submartingale, and $\pi(t)$ is an increasing function of t. Using the probability density function of λ_{it} obtained in the Appendix, we have

$$
\begin{aligned}
E[|\lambda_{it} - 0.5|] &= 0.5 - \lambda_0 + \int_0^\infty \lambda_0(1 - \exp(-y))n\left(\frac{y - y_0 - \sigma_s^2 t/2}{\sigma_s\sqrt{t}}\right) d\left(\frac{y}{\sigma_s\sqrt{t}}\right) \\
&= 0.5 - \lambda_0 N\left(\frac{y_0 + \sigma_s^2 t/2}{\sigma_s\sqrt{t}}\right) - (1 - \lambda_0)N\left(\frac{y_0 - \sigma_s^2 t/2}{\sigma_s\sqrt{t}}\right). \qquad (9)
\end{aligned}
$$

where

$$
y_0 = \ln\left(\frac{\lambda_0}{1 - \lambda_0}\right).
$$

The cumulative aggregate profit at time t is given by,

$$
\begin{aligned}
\Pi(t) &= \int_0^t e^{-r\tau}\pi(\tau)d\tau \\
&= \frac{-e^{-rt}}{r}\pi(t) + \frac{\pi(0)}{r} - \int_0^t e^{-r\tau}\pi'(\tau)/r \\
&= \frac{-e^{-rt}}{r}\pi(t) + \frac{\pi(0)}{r} - \frac{c_h - c_l}{r}\int_0^t e^{-r\tau}\frac{\lambda_0 y_0}{\sigma_s\sqrt{t^3}}n(\frac{y_0 + \sigma_s^2 t/2}{\sigma_s\sqrt{t}})d\tau \\
&= \frac{-e^{-rt}}{r}\pi(t) + \frac{\pi(0)}{r} - \frac{c_h - c_l}{r} \\
&\quad \times \left[N\left(\frac{-y_0 - \sqrt{\sigma_s^4 + 8r\sigma_s^2}t/2}{\sigma_s\sqrt{t}}\right) + \exp\left(-y_0\sqrt{1 + \frac{8r}{\sigma_s^2}}\right)N\left(\frac{-y_0 + \sqrt{\sigma_s^4 + 8r\sigma_s^2}t/2}{\sigma_s\sqrt{t}}\right)\right]. \qquad (10)
\end{aligned}
$$

Determining $\pi(t)$ and $\Pi(t)$ allows the firm to learn when it will make positive profit flows and the maximum amount of profit loss it will suffer.

In Figures 9 and 10, we plot the profit flows and cumulative profits for the customization scenario. Over time, the firm is more likely to match each customer with the right service, and as a result, the profit flow increases gradually. As t increases, $\pi(t)$ will cross 0 at a point $t = t_0$. From time t_0 on, the firm will have positive cash flows. Consequently, $\Pi(t)$

Figure 9: Profit Flow under Customization

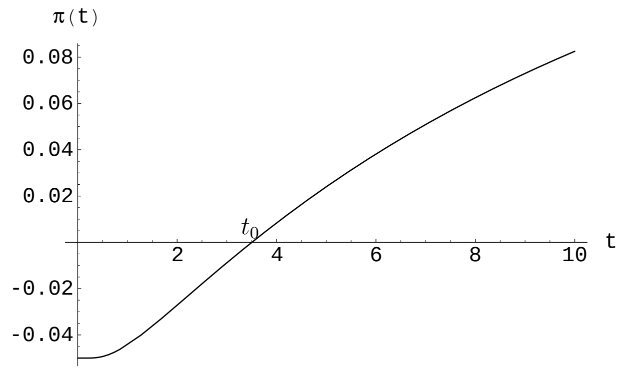

Profit flow over time; $v = 0.2, c_l = 0, c_h = 1, r = 0.05$, and $\sigma = 1.5, \lambda_0 = 0.75$.

first decreases over time and later increases for $t > t_0$. At time point t_0, $\Pi(t_0)$ reaches the maximum loss possible and increases thereafter. Finally, there exists a time point $t_1 > t_0$, such that the firm fully recovers its losses and Π_{t_1}=0. $\Pi(t)$ then becomes positive for $t > t_1$, where $t_1 > t_0$.

Next we consider the customer selection case and again focus on $\lambda_0 > 0.5$. Consider the case in which $\hat{\lambda} > 0.5$. Remember that $H(t)$ is the fraction of customers who have hit the absorbing barrier $\hat{\lambda}$ at time t, and $\bar{\lambda}(t)$ denotes the population average of λ_{it} among the remaining customers. The profit flow at time t is

$$\pi(t) \equiv E \int_i (v - c(\lambda_{it}), D_{it}) di = (1 - H(t))(v - c(\bar{\lambda}_t)) = v - c(\lambda_0) - H(t)[v - c(\hat{\lambda})]. \quad (11)$$

Equation (11) offers some insights regarding the point at which the firm should cut off unprofitable customers. At time t, the firm improves its profit flows because it no longer serves customers with $\lambda_{it} = \hat{\lambda}$. Recall that $v - c(\lambda^*) = 0$. Thus, $\hat{\lambda}$ must be less than λ^*, and customers with $\lambda_{it} = \hat{\lambda}$ lose money for the firm. If $\hat{\lambda}$ is very close to λ^*, the firm does not gain much from discontinuing services for unprofitable consumers, because they only marginally lose money for the firm. In other words, when $\hat{\lambda}$ is too close to λ^*, the firm makes too much type II error by stopping services for many customers that it should continue to serve. If $\hat{\lambda}$ is very low, the firm waits too long to remove customers with very negative profits and thus

Figure 10: Cumulative Profits under Customization

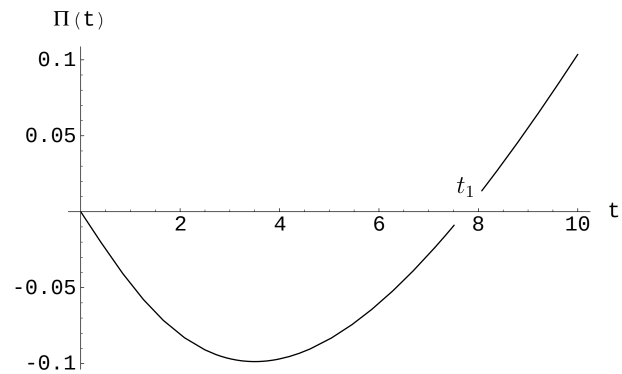

Cumulative discounted profit over time; $v = 0.2, c_l = 0, c_h = 1, r = 0.05, \sigma = 1.5$, and $\lambda_0 = 0.75$.

makes too much type I error by serving customers that it should have dropped. The best choice of $\hat{\lambda}$ trades off these two types of errors. Mathematically, when $\hat{\lambda}$ is low, $v - c(\hat{\lambda})$ will be more negative, but it will take longer to reach $\hat{\lambda}$ and thus, $H(t)$ will be smaller. When $\hat{\lambda}$ is high, $v - c(\hat{\lambda})$ will be less negative, yet customers will be quicker to reach, $\hat{\lambda}$ and thus, $H(t)$ will be larger. Therefore, the optimal $\hat{\lambda}$ provides the best trade-off between type I and type II errors, such that the cumulative discounted savings $-H(t)[v - c(\hat{\lambda})]$ is maximized.

Given the expression of expected flow, we next determine the cumulative discounted profits. We define the cumulative discounted profit function up to time t as Π_t and derived

$$\Pi(t) = \int_0^t e^{-r\tau}\pi(\tau)d\tau = \frac{-e^{-rt}}{r}\pi(t) + \frac{\pi(0)}{r} - \int_0^t e^{-r\tau}h(\tau)[v - c(\hat{\lambda})]/r.$$

We obtain the following expression:

$$\begin{aligned}\Pi(t) \quad = \quad & \frac{-e^{-rt}}{r}\pi(t) + \frac{\pi(0)}{r} - \exp(\frac{\sigma_s^2 - \sqrt{\sigma_s^4 + 8r\sigma_s^2}}{2\sigma_s^2}(\hat{y} - y_0))\frac{\lambda_0(v - c(\hat{\lambda}))}{r\hat{\lambda}} \\ & \times[N(\frac{(\hat{y} - y_0) - \sqrt{\sigma_s^4 + 8r\sigma_s^2}t/2}{\sigma_s\sqrt{t}}) + \exp\left((\hat{y} - y_0)\sqrt{1 + \frac{8r}{\sigma_s^2}}\right)N(\frac{(\hat{y} - y_0) + \sqrt{\sigma_s^4 + 8r\sigma_s^2}t/2}{\sigma_s\sqrt{t}})].\end{aligned} \tag{12}$$

Figure 11: Profit Flow under Customer Selection

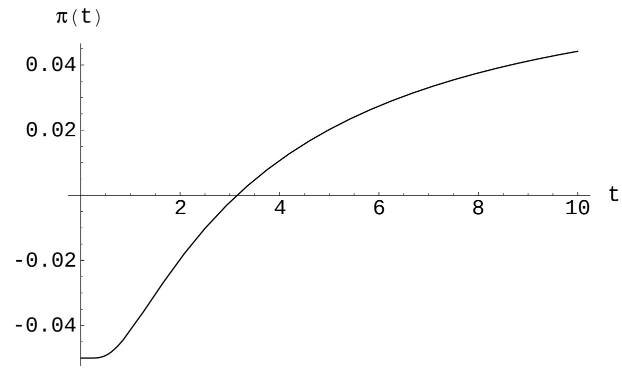

Profit flow over time. We have $v = 0.1, c_l = 0, c_h = 1, r = 0.05, \sigma = 1.5, \lambda_0 = 0.85$.

In Figures 11 and 12, we plot the firm's expected profit and cumulative profit flows as a function of t. For $\lambda_0 < \lambda^*$, the initial profit flows is negative and only becomes positive in later periods. The cumulative profits first declines and then turns around and becomes positive for large t.

Proposition 6 *Suppose that $0.5 \leq \lambda_0 < \lambda^*$. The expected profits are initially negative but become positive subsequently, and the cumulative positive profits dominate the earlier losses eventually. There exist time points such that $\pi(t_0) = 0$ and $\Pi(t_1) = 0$. At time point t_0, the firm will earn positive profits, and the firm's maximum cumulative loss occurs at time point t_0. At time point t_1, the firm will have recovered all its losses.*

3.6 Comparison with a Myopic Strategy

We assume firms are forward looking, but what happens when a myopic firm learns about customers but still discontinues customer service whenever its expected profit is negative? We show that such a strategy can make learning worthless in some situations.

Proposition 7 *When $\lambda^* \geq 0.5$, the value of learning is 0 under the myopic strategy. However, when $\lambda^* < 0.5$, the myopic strategy coincides with the optimal strategy.*

Figure 12: Cumulative Profits under Customer Selection

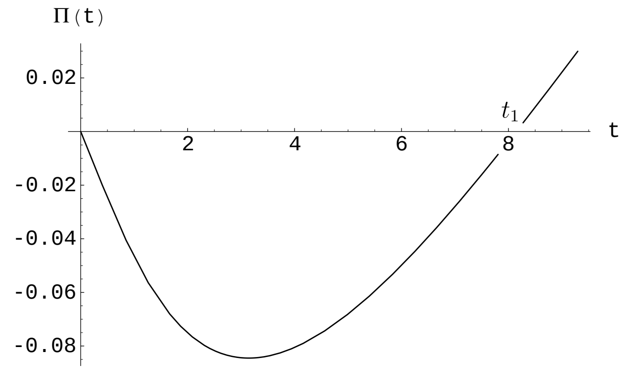

We have $v = 0.1$, $c_l = 0$, $c_h = 1$, $r = 0.05$, $\sigma = 1.5$, and $\lambda_0 = 0.85$.

Proposition 7 describes situations in which the myopic or *ad hoc* strategy makes learning optimal or worthless. When $\lambda^* < 0.5$, we have $\hat{\lambda} < \lambda^* < 0.5$, and the firm follows the optimal strategy, as in the customization scenario. Both the optimal strategy and the ad hoc strategy make the firm switch services whenever λ_{it} crosses 0.5, and thus, the ad hoc strategy is optimal.

When $\lambda^* \geq 0.5$, two possibilities emerge. First, consider $\hat{\lambda} < 0.5 \leq \lambda^*$. Starting from $\lambda_0 \geq 0.5$, the myopic firm will stop at λ^*, such that information gains are 0. However, the optimal strategy is to continue serving customers at all times with the lower cost service. Second, when $0.5 \leq \hat{\lambda} < \lambda^*$, the myopic strategy and the optimal strategy recommend that the firm stop serving customers at points λ^* and $\hat{\lambda}$, respectively. However, in the case of the myopic strategy, it discontinues services for those whose expected profits are 0. As a result, the expected profits from the remaining customers stay the same, and learning creates no value. In contrast, the optimal strategy removes unprofitable customers who reach $\hat{\lambda}$ from the population, so the expected profits increase. Moreover, $\hat{\lambda}$ gets chosen such that type I and type II errors balanced each other to maximize life time value from customers through learning.

Although in the case of $\lambda^* < 0.5$, the myopic strategy happens to be optimal, our analysis indicates that when $\lambda^* \geq 0.5$, learning without optimization could make all efforts to learn fruitless.[4] As a result, it is important to integrate learning into forward-looking decision making to obtain the optimal strategy.

3.7 Firms with Endogenous Speed of Learning

Thus far, we have assumed that the speed of learning is unrelated to the firm's decision. In reality, there are many ways a firm can speed up its identification of customers. In this section, we extend the model to a case with multiple units of sales in which pricing affects the quantity of purchase, which in turn affects the speed of learning. In other words, assuming units of sales affect the speed of learning, we discuss how learning affects the firm's promotion strategy when the firm can internalize the speed of learning.

We assume that the consumer has a downward sloping demand curve of

$$p_{it} = \alpha - \beta q_{it}.$$

We further assume that the precision of information learned is proportional to the units sold

$$dc_i(q_{it}, S_{it}, t) = A(q_{it})[c(T_i, S_{it})dt + \frac{\sigma}{\sqrt{q_{it}}} W_{it}],$$

where $A(q_{it})$ reflects the higher cost due to more service units. We assume that $A(1) = 1$ (as a normalization) and $A(q_{it})$ is a twice-differentiable, increasing, and convex function of q_{it}. In addition, the precision of information increases with units sold q_{it}. We assume that $q_{it} \geq 0$, and thus, the expected profit flow at time t is given by

$$\pi_{it} = p_{it} q_{it} - A(q_{it}) c(\lambda_{it}).$$

where $c(\lambda_{it}$ is given by (5) Because there is a one-one correspondence between p_{it} and q_{it}, the optimal strategy of the firm is to determine S_{it} and q_{it}. The life time profit of customer i is

$$\Pi(\{q_{i\tau}, S_{i\tau}, \tau \geq t\}, \lambda_{it}) = E \int_t^{\infty} e^{-r\tau} \pi_{i\tau} d\tau,$$

[4] If we normalize the gains from the optimal strategy with adaptive learning to 100%, the gains from the myopic strategy are 0% in the customer selection case and 100% in the service customization case. More generally, with more complex heterogeneity, the gains from the myopic strategy fall somewhere between 0% and 100%.

and the value function, or maximized cumulative profit on the optimal path, is given by

$$V(\lambda_{it}) \equiv \max_{\{q_{i\tau}, S_{i\tau}, \tau \geq t\}} \Pi(\{q_{i\tau}, S_{i\tau}, \tau \geq t\}, \lambda_{it}). \tag{13}$$

Due to the symmetry of cost functions for the two services A and B, the optimal service strategy is as described by Equation (4); that is, provide service A if $\lambda_{it} \geq 0.5$ and B otherwise. We next determine the optimal quantity choice q_{it}. Following the standard approach for a Bellman equation, we get

$$0 = \max_{q_{it}}[p_{it}q_{it} - A(q_{it})c(\lambda_{it})]dt - rV(\lambda_{it})dt + \frac{q_{it}\lambda_{it}^2(1-\lambda_{it})^2(c_h - c_l)^2}{2\sigma^2}V''(\lambda_{it})dt. \tag{14}$$

For an interior solution $q_{it} > 0$, we divide both sides of Equation (14) by dt and take the first-order condition of the right-hand side with respect to q_{it}; thus

$$0 = \alpha - 2\beta q_{it} - A'(q_{it})c(\lambda_{it}) + \frac{\lambda_{it}^2(1-\lambda_{it})^2(c_h - c_l)^2}{2\sigma^2}V''(\lambda_{it}). \tag{15}$$

Let

$$M(q_{it}) = 2\beta q_{it} + A'(q_{it})c(\lambda_{it}), \tag{16}$$

then the optimal solution for $\hat{q}_{it}$ is

$$\hat{q}_{it} = M^{-1}\left(\alpha + \frac{\lambda_{it}^2(1-\lambda_{it})^2(c_h - c_l)^2}{2\sigma^2}V''(\lambda_{it})\right).$$

Suppose that the firm decides not to collect any more information from time t on. Then the term

$$\frac{\lambda_{it}^2(1-\lambda_{it})^2(c_h - c_l)^2}{2\sigma^2}V''(\lambda_{it})$$

drops out of the solution. Without further learning at time t for customer i, the optimal solution of q_{it} becomes

$$q_{it}^* = M^{-1}(\alpha),$$

which corresponds to the myopic solution of maximizing immediate profit. When it implements learning, the firm increases its sales volume to interact with the customer more often and collect more information.

Rearranging the terms in Equation (14), we obtain the following ordinary differential equation for the value function:

$$V(\lambda_{it}) = \frac{(\alpha - \beta q_{it})q_{it} - A(q_{it})c(\lambda_{it})}{r} + \frac{\lambda_{it}^2(1-\lambda_{it})^2(c_h - c_l)^2 q_{it}}{2r\sigma^2} V''(\lambda_{it}).$$

Notice that $A(q_{it})$ is convex, and thus, M is increasing. Moreover, information always adds value, so V'' must be positive. As a result, $\hat{q}_{it} > q_{it}^*$ and the firm charges a lower price with learning. Let $d_{it} = p_{it}^* - \hat{p}_{it}$ denote the price discount due to CIM. The firm needs a lower price to achieve higher sales volume, and thus, the discount must be positive.

As the firm learns more about the customer, λ_{it} changes over time, and we must determine how price changes with λ_{it}. In the case without endogenous learning, price always decreases with λ_{it} as the marginal cost decreases for higher λ_{it}. However, the case with learning is more complicated. When λ_{it} increases, though marginal cost decreases, the need for experimentation also falls. As a result, the firm may want to reduce the amount of experimentation and increase the price. How price will be affected by λ_{it} is thus not clear.

We have the following results regarding price discounts and the comparative statics of price with respect to λ_{it} in the presence of CIM:

Proposition 8 *Suppose that $\lambda_{it} \geq 0.5$. The firm will have a promotion for customers with uncertainty when the firm implements customer information management. The promotion moves toward 0 as t approaches infinity. Mathematically, $d_{it} \geq 0$ and $d_{it} \to 0$ as $t \to \infty$. When $rV' - (c_h - c_l)(A(q_{it}) - q_{it}A'(q_{it}))$ is positive (negative), the price decreases (increases) with λ_{it}.*

Proof: Promotion follows due to monotonicity of function M. As t approaches infinity, the firm's uncertainty regarding the remaining customers moves toward 0. Consequently, λ_{it} moves toward either 0 or 1. As a result, $q_{it}^* - \hat{q}_{it}$ and $d_{it} = \hat{p}_{it} - p_{it}^*$ both move closer to 0. The remaining proof appears in the Appendix.

In the case without information, as λ_{it} increases, the price always decreases because the marginal cost decreases. However, with information, this trend may not be true because of

Figure 13: Price Discount

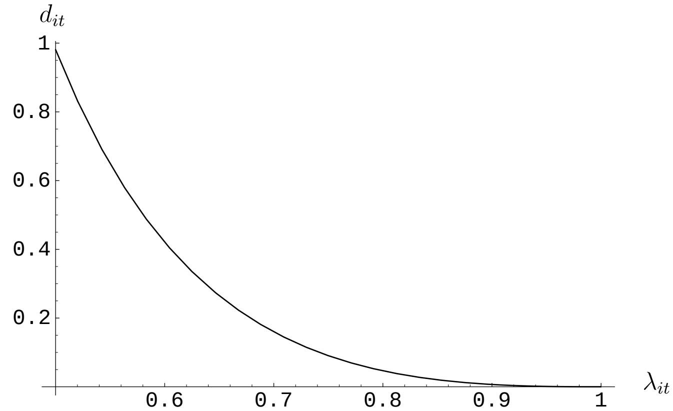

Price discount vs λ_{it}; $\alpha = 1, \beta = 0.5, c_l = 0.3, c_h = 0.7, r = 0.05$, and $\sigma = 2.5$.

two effects. The first effect is that marginal cost decreases with λ_0, and when marginal cost decreases, the firm charges a lower price. The second effect is that when λ_{it} increases, there is less uncertainty and thus less promotion, which causes the price to increase. Whether price decreases or increases with λ_{it} depends on which effect dominates.

In Figure 13, we plot the price discount as a function of λ_{it}. Interestingly, d_{it} is a decreasing function of λ_{it} for $\lambda_{it} \geq 0.5$ in the numerical example. Intuitively, for smaller λ_{it} in the interval $[0.5, 1]$, there is more uncertainty, and the firm can learn more about customers through experimentation. As a result, the firm is willing to give a bigger discount to customers with lower λ_{it}.

4 Managerial Implications

The recent trend in CIM contains three components: analysis of information, forward looking, and optimal use of information. Our analytical results shed light on the value of learning and acting on information, the time to discontinue service to a customer, the duration and amount of short-term financial losses before learning pays off, and how the firm can modify its decisions to speed up the learning process, as we discuss in more detail next.

Value of Customer Information

Marketing managers must realize that they have incomplete information about individual customers and they also need to understand that the limits of static segmentation approaches that dominate current business practice. Incomplete information is subject to type I (the firm provides services to the wrong customer) and type II (the firm wrongly turns away profitable customers) errors, which lead to suboptimal decisions. Because the ongoing information revolution has exponentially increased the amount of information about each individual customer and created highly interactive marketing communication environments, firms should recognize the importance of tracking the interactive history of each consumer, especially when consumers change their demand and preferences during the course of a long-term relationship with the firm. Learning from the feedback continuously collected from recent decisions enables the firm to mitigate type I and type II errors, improve the accuracy of customer identification, and track consumer development. More important, firms should realize that acting on information is the ultimate means to realize the value of information. Our results indicate that learning and acting on information can improve marketing effectiveness and marketing efficiency, which can be transformed into improved long-term profits.

Forward-Looking into Long-term Profit

In some situations, a firm that only cares about myopic profit can not benefit from implementing learning. Rather, a firm must be forward looking for its learning efforts to be fruitful. This finding supports the idea of customer lifetime value management advocated by Reinartz and Kumar (2000, 2003) and Berger et al. (2002). Furthermore, firms should understand that learning might incur a short-term loss of profit. However, the firm can benefit from more accurate customer information in the long term because of its improved ability to identify customer preferences and customize its marketing decisions accordingly. As a result, customers are better served (effectiveness), and firms better control their costs (efficiency), which translates into greater long-term profits. Therefore, a firm implementing CIM should take a forward-looking view and be prepared to tolerate short-term financial losses during the early stages.

Optimal Adaptation of Marketing Strategy

Not only does the result of learning improve a firm's marketing decisions, but the process of learning also affects a firm's day-to-day marketing strategies before the firm reaches a steady state. Learning enables firms to take more risk and enter markets with more customers who may seem unprofitable. Making targeting decisions based on static segmentation approaches may be suboptimal. The proper point at which to discontinue service depends on the optimal trade-off of type I and type II errors. The tolerance level depends on the importance of future profits and information precision. Firms should understand the pattern of inter-temporal profit flow to estimate the amount and duration of short-term financial loss they will incur by learning. Accordingly, they can evaluate their investments and the payoff of implementing a CIM system.

Customer-Centric Decisions

Learning and acting on information gives the firm the ability to update its customer knowledge and improve its marketing strategies in a continuous fashion. Improved customer knowledge lets the firm act on information and provide products or services tailored to the preferences of each individual customer (or more practically, each segment). In addition, a forward-looking firm can take into account the future marketing consequences of its current marketing decisions. Consequently, it will sacrifice its short-term profit, proactively act on customer information, and provide the customer better service to maximize customer lifetime value. That is, its decision is customized and proactive, which is akin to the idea of customer-centric marketing.

Endogenized Speed of Learning

Firms should view the speed of learning itself as an endogenous variable. Learning is part of the decision process, and decision making can facilitate data collection. For example, when the firm can learn faster by inducing the customer to buy more units of services, it should provide discounts to unfamiliar customers to explore the potential that they may be profitable customers. This offers an alternative explanation to our observation that some firms, especially those with some monopoly power, such as the cable industry, offer promotional prices to new customers but keep the price high for their existing customers.

In summary, learning and acting on customer information, as formulated in our framework, integrates data mining and analytical decision support systems and thus help firms improve their knowledge of customers and adapt their decisions to maximize long-term profit. With the emergence of software and Web-based automated decision support systems, we anticipate that more and more firms will employ technologies that allows immediate access to customer databases, learn about customers' intrinsic preference, solve dynamic programming problems or their simplified heuristics to obtain optimal marketing intervention decisions.

5 Conclusion

Recent technological developments offer massive possibilities for tracking the purchase history of and collecting detailed information about each individual customer. However, customer information hidden in a database cannot become firm knowledge unless it is appropriated analyzed. Furthermore, a firm's knowledge about an individual customer cannot be transformed into profit unless the firm adapts its decisions to match its knowledge about individual customers. Realizing that customer information can create potential comparative advantages, as promised by customer-centric marketing, firms have started to develop learning and interactive marketing strategies to explore the opportunities enabled by data collection, data mining, and decision-making technology. Therefore, research must investigate the value of learning and how learning interacts with marketing strategies.

We consider a market with two kinds of customers and two kinds of products. When a firm matches the right product with the right customer, it can reduce its service costs. We formulate the company's matching, targeting, and pricing decisions as solutions to a stochastic dynamic programming problem with uncertainty, in which the firm needs to learn about heterogeneous customer types, take into account the dynamic effect of its current decisions, and optimally balance the short-term cost of learning with the long-term profit gain it obtains from its improved knowledge about its individual customers. We analytically derive the closed-form solution for the firm's matching, selection, and pricing strategies and compare them for firms with and without learning. We investigate the properties of the dynamic learning process, the interaction between learning and decision making, and the evolution of profit over time. We also examine the effects on decision making and profit implications of parameters such as discount rate and the precision of information. By extending the model to a setting in which the precision of the signal is proportional to the purchase quantity, we study how a firm's pricing strategies can be modified to speed up the learning process.

Learning and acting on information helps a firm improve the accuracy of its knowledge about each customer and adapt its individualized marketing decisions on the basis of the most updated information. As a result, customers are better serviced and/or the service costs decrease, which then leads to improved profit. In addition, a firm with learning behaves more strategically than a firm without learning. Without learning, it is optimal for the firm to keep serving customers with positive expected profits from the very beginning. In contrast, learning allows a firm to gradually identify and discontinue its service to unprofitable customers and match the right products to the right customers according to their preferences. Instead of immediately discontinuing service to a customer with negative expected profit, a learning firm chooses to serve the unprofitable customer to avoid type II errors. The firm also must maximize long term profits, because the myopic use of information can make learning worthless. Furthermore, learning might make the firm incur some financial losses during the early periods, which can be paid off in the long term. Finally, the learning process can be sped up by modifying marketing decisions, such as providing price promotions to new customers.

By demonstrating whether and how CIM helps firms improve their long-term profit and how learning modifies their marketing strategies, our research clarifies the value of investing in CIM and informs firms how their day-to-day marketing decisions may be influenced by learning. In turn, we delineate the short-term costs and long-term benefits of learning and acting on information and describe how marketing decisions can facilitate learning. A better understanding of these issues remains imperative, especially considering the increasing demand for and supply of real-time solutions for integrating database and analytical marketing decisions to achieve more customer-centric marketing. We focus on learning about heterogeneity in service costs in a stylized example, and this same approach can be applied to other forms of consumer heterogeneity, such as consumer demand and preference.

This paper is subject to several limitations that suggest avenues for further research. First, for simplicity, we assume customers differ in terms of service cost; additional research might consider other forms of customer heterogeneity, such as willingness to pay, channel preference (Ansari et al. 2005), price sensitivity, and so forth. Modeling heterogeneity in customer demand allows us to investigate customer acquisition and retention decisions. Our model also might be extended to include the arrival of new customers or changes in customer preferences. For example, if customers' preferences change over time, the firm needs to use the correct model to update its beliefs about customers. Second, it would be interesting to study how competition affects learning. Customers may switch back and forth between the firm and its competitor, so customers' future switching behavior should be modeled in a forward-looking analysis. In addition, learning could intensify with greater competition, and firms might consider information sharing in such settings. Third, it would be useful to develop feasible statistical algorithms that can measure customer insights and develop optimization routines for decision-support systems that automate the implementations of marketing decisions. Fourth, we have only analyzed two types of products. Further research should extend the model to markets with more complex customer heterogeneity and product lines. Moreover, product design could be endogenized to accommodate learning.

6 Appendix

Proof of Theorem 1: [xxx where is proposition 6?]

On the optimal path where $D^*_{it} = 1$ over the interval $[t, t+dt)$, we have

$$V(\lambda_{it}) = [v - c(\lambda_{it})]dt + (1 - rdt)E_t[V(\lambda_{i(t+dt)})]. \tag{17}$$

By Ito's lemma,

$$V(\lambda_{i(t+dt)}) \quad = \quad V(\lambda_{it}) + V'(\lambda_{it})d\lambda_{it} + \frac{\sigma^2(\lambda_{it})}{2}V''(\lambda_{it})dt. \tag{18}$$

Taking expectations on both sides of Equation(18), we have

$$E_tV(\lambda_{i(t+dt)}) = V(\lambda_{it}) + \frac{\sigma^2(\lambda_{it})}{2}V''(\lambda_{it})dt. \tag{19}$$

Applying the expression of $E_tV(\lambda_{i(t+dt)})$ to Equation (17) and ignoring higher-order terms of dt, we get

$$\begin{aligned} V(\lambda_{it}) &= (v - c(\lambda_{it}))dt + (1 - rdt)E_tV(\lambda_{i(t+dt)}) \\ &= (v - c(\lambda_{it}))dt + V(\lambda_{it}) - rdtV(\lambda_{it}) + \frac{\sigma^2(\lambda_{it})}{2}V''(\lambda_{it})dt. \end{aligned} \tag{20}$$

When we subtract both sides of Equation (20) by $V(\lambda_{it}) - rV(\lambda_{it})dt$ and divide both sides by rdt to eliminate the dt term, then divide both sides by r, we get the ordinary differential equation in the text. Q.E.D.

We next determine the general solutions for the Bellman equation. For such an ordinary differential equation, the general solutions are given by:[5][xxx, find this reference.]

$$V(\lambda_{it}) = \frac{v - c(S_{it}, \lambda_{it})}{r} + b_1\lambda_{it}^{(\gamma+1)/2}(1 - \lambda_{it})^{-(\gamma-1)/2} + b_2\lambda_{it}^{-(\gamma-1)/2}(1 - \lambda_{it})^{(\gamma+1)/2},$$

[5]See Kartsatos (1997).

where

$$\gamma \equiv \sqrt{1 + \frac{8r\sigma^2}{(c_h - c_l)^2}} > 1,$$

and b_1, b_2 are coefficients that must be determined later.

In the setup, the parameters with respect to services A and B are symmetric, so $V(\lambda_{it}) = V(1-\lambda_{it})$. Let b_1, b_2 denote the coefficients for the case $\lambda_{it} \geq 0.5$ and b_1', b_2' the coefficients for the case that $\lambda_{it} < 0.5$; then, $b_1 = b_2'$ and $b_2 = b_1'$. Without loss of generality, we focus on the case for $\lambda_{it} \geq 0.5$, for which the firm must provide service A to customer i if the service is continued. However, because $\gamma > 1$, as $\lambda_{it} \to 1$, $b_1 \lambda_{it}^{(\gamma+1)/2}(1-\lambda_{it})^{-(\gamma-1)/2}$ approaches infinity for $b_1 > 0$. Consequently, we must have the following condition:

$$b_1 = 0, \text{for } \lambda_{it} \geq 0.5.$$

Let $b \equiv b_1$; we get the results stated in the proposition. Q.E.D.

Derivation of the Probability Density of λ_{it}

There are λ_0 proportion of the customers who are A type customers. For these customers, their posterior belief λ follows the following stochastic differential equation:

$$d\lambda_{it} = -\frac{\lambda_{it}(1-\lambda_{it})(c_h - c_l)}{\sigma^2}[c_l - c(\lambda_{it})]dt - \frac{\lambda_{it}(1-\lambda_{it})(c_h - c_l)}{\sigma}dW_{it}.$$

Let

$$y = \ln(\frac{\lambda}{1-\lambda}, \sigma_s = \frac{c_h - c_l}{\sigma}),$$

then we have

$$dy = \frac{1}{2}\sigma_s^2 dt + \sigma_s dW.$$

The probability density of y is

$$g_1(y,t) = \sqrt{\frac{1}{2\pi}}\frac{1}{\sigma_s}\exp(-\frac{y - y_0 - \sigma_s^2 t)^2}{2\sigma_s^2}).$$

The probability density for λ at time t is

$$h_1(\lambda, t) = g(y,t)dy/d\lambda = g(y,t)/[\lambda(1-\lambda)].$$

Furthermore, we have

$$\lambda_{it} = \frac{(\lambda_0/(1-\lambda_0))\exp((c_h-c_l)^2 t/(2\sigma^2) + [(c_h-c_l)/\sigma]W(t))}{1+(\lambda_0/(1-\lambda_0))\exp((c_h-c_l)^2 t/(2\sigma^2) + [(c_h-c_l)/\sigma]W(t)).}$$

There are $1-\lambda_0$ proportion of customers who are B type customers. For these customers, their posterior belief λ follows the following stochastic differential equation:

$$d\lambda_{it} = -\frac{\lambda_{it}(1-\lambda_{it})(c_h-c_l)}{\sigma^2}[c_h - c(\lambda_{it})]dt - \frac{\lambda_{it}(1-\lambda_{it})(c_h-c_l)}{\sigma}dW_{it}.$$

The probability density of y is

$$g_2(y,t) = \sqrt{\frac{1}{2\pi}}\frac{1}{\sigma_s}\exp(-\frac{y-y_0+\sigma_s^2 t)^2}{2\sigma_s^2}).$$

The probability density for λ at time t is

$$h_2(\lambda, t) = g(y,t)dy/d\lambda = g(y,t)/[\lambda(1-\lambda),]$$

and

$$\lambda_{it} = \frac{(\lambda_0/(1-\lambda_0))\exp(-(c_h-c_l)^2 t/(2\sigma^2) + [(c_h-c_l)/\sigma]W(t))}{1+(\lambda_0/(1-\lambda_0))\exp(-(c_h-c_l)^2 t/(2\sigma^2) + [(c_h-c_l)/\sigma]W(t))}$$

The probability density is

$$h(\lambda) = \lambda_0 h_1 + (1-\lambda_0)h_2.$$

Q.E.D.

Proof of Proposition 3: It is easy to verify that $d\hat{\lambda}/d\gamma > 0$ and $d\gamma/d\sigma > 0, d\gamma/r > 0$. Therefore, we have $d\hat{\lambda}/d\sigma > 0, d\hat{\lambda}/r > 0$.

Proof of Proposition 4: For $\lambda_{it} \le \lambda^*$, we have $VI(\lambda_{it}) = V(\lambda_{it})$ and $dVI(\lambda_{it})/d\lambda_{it} = dV(\lambda_{it})/\lambda_{it} > 0$. When $\lambda_{it} > \lambda^*$, $VI(\lambda_{it}) = b\lambda_{it}^{\frac{-(\gamma-1}{2}}(1-lambda_{it})^{\frac{\gamma+1}{2}}$. Moreover, $VI''(\lambda_{it}) =$

$V''(\lambda_{it}) > 0$, $VI'(\lambda_{it})|_{\lambda_{it}=1} = 0$. Thus, $VI'(\lambda_{it}) \leq 0$. In addition, we have

$$dVI/dr = -\frac{VI}{r} - \frac{c_h - c_l}{2r\gamma}\lambda_{it}^{\frac{-(\gamma-1}{2}}(1-\lambda_{it})^{\frac{\gamma+1}{2}}(\ln(\frac{\lambda_{it}}{1-\lambda_{it}})(\frac{d\gamma}{dr}) < 0$$

$$dVI/d\sigma = -\frac{c_h - c_l}{2r\gamma}\lambda_{it}^{\frac{-(\gamma-1}{2}}(1-\lambda_{it})^{\frac{\gamma+1}{2}}(\ln(\frac{\lambda_{it}}{1-\lambda_{it}})(\frac{d\gamma}{d\sigma}) < 0.$$

Proof of Proposition 5: Convergence of market size follows the convergence of the hitting probability $H(t)$. The market size decreases with σ and r because $\hat{\lambda}$ increases with σ and r. $\bar{\lambda}(t)$ moves toward 1 because $H(t)$ goes to $(1-\lambda_0)/(1-\hat{\lambda})$.

Proof of Proposition 7: When $\lambda^* < 0.5$, we have $\hat{\lambda} < \lambda^* < 0.5$, so the service customization strategy is the optimal strategy, which coincides with the myopic strategy. When $\lambda^* \geq 0.5$, the profit flow under the ad hoc strategy is

$$\int_i (v - c(\lambda_{it})D_{it}di = v - c(\lambda_0) - H^*(t)(v - c(\lambda^*)),$$

where $H^*(t)$ is the proportion of customers hitting the ad hoc stopping point λ^* in the ad hoc strategy. However, $v - c(\lambda^*) = 0$. The profit flow of the firm does not improve over the case without learning, and the value of information is 0 under the ad hoc strategy.

Proof of Proposition 8:

We can derive the Bellman equation on the optimal path.

$$V(\lambda_{it}) = \max_{q_{it}}[p_{it}q_{it} - A(q_{it})c(S_{it}, \lambda_{it})]dt + (1 - rdt)EV(\lambda_{i(t+dt)}). \tag{21}$$

From the Ito's lemma, we have

$$V(\lambda_{i(t+dt)}) = V(\lambda_{it}) + V'(\lambda_{it})d\lambda_{it} + \frac{1}{2}V''(\lambda_{it})(d\lambda_{it})^2$$

. When we expectations, we derive

$$E_t[V(\lambda_{i(t+dt)})] = V(\lambda_{it}) + \frac{q_{it}\lambda^2(1-\lambda)^2(c_h - c_l)^2}{2\sigma^2}V''(\lambda_{it}).$$

Plugging the expression of $E_t[V(\lambda_{i(t+dt)})]$ into Equation (21), we get

$$0 = \max_{q_{it}}[p_{it}q_{it} - A(q_{it})c(S_{it}, \lambda_{it})]dt - rV(\lambda_{it})dt + \frac{q_{it}\lambda^2(1-\lambda)^2(c_h - c_l)^2}{2\sigma^2}V''(\lambda_{it})dt. \quad (22)$$

The first-order condition is

$$0 = \alpha - 2\beta q_{it} - A'(q_{it})c(\lambda_{it}) + \frac{\lambda_{it}^2(1-\lambda_{it})^2(c_h - c_l)^2}{2\sigma^2}V''(\lambda_{it}), \quad (23)$$

and we let

$$F(q_{it}, \lambda_{it}) \equiv \alpha - 2\beta q_{it} - A'(q_{it})c(\lambda_{it}) + \frac{\lambda_{it}^2(1-\lambda_{it})^2(c_h - c_l)^2}{2\sigma^2}V''(\lambda_{it}).$$

Thus, $\partial F(q_{it}, \lambda_{it})/\partial q_{it} = -2\beta - A''(q_{it}) < 0$. By the implicit function theorem, the sign of $dq_{it}/d\lambda_{it}$ is the same as the sign of $q_{it}\partial F(q_{it}, \lambda_{it})/\partial \lambda_{it}$. In turn,

$$\begin{aligned} q_{it}\partial F(q_{it}, \lambda_{it})/\partial \lambda_{it} &= +(c_h - c_l)q_{it}A'(q_{it} + \left(\frac{\lambda_{it}^2(1-\lambda_{it})^2(c_h - c_l)^2}{2\sigma^2}V''(\lambda_{it})\right)' \\ &= rV' - (c_h - c_l)[A(q_{it}) - q_{it}A'(q_{it})]. \quad (24) \end{aligned}$$

The derivative of q_{it} with resepct to λ_{it} has the same sign as $rV' - (c_h - c_l)[A(q_{it}) - q_{it}A'(q_{it})]$, which implies that $dp_{it}/d\lambda_{it}$ has the opposite sign.

Q.E.D.

References

[1] Anderson, Eugene W. (2006), ???Linking Service and Finance,??? Marketing Science, 25(6), 587-89.

[2] Ansari, Asim, Carl Mela, and Scott A. Neslin (2005), "Customer Channel Migration," Working Paper, Tuck School of Business, Dartmouth College.

[3] Bergemann and Valimaki (1996), "Learning and Strategic Pricing," Econometrica, 64, No. 5, 1125-1149. Sep., 1996.

[4] Berger, Paul D., Ruth N. Bolton, Douglas Bowman, Elten Briggs, V. Kumar, A. Parasuraman, Creed Terry (2002), "Marketing Actions and The Value of Customer Assets: A Framework for Customer Asset Management." Journal of Service Research, 5, 39-54.

[5] Bitran, GR and S. Monschein (1997), "Periodic pricing of seasonal product in retailing." Management Sci. 43 427-443.

[6] Blattberg and Deighton (1991), "Interactive Marketing: Exploit the Age of Addressability." Sloan Management Review, 33, 5-14

[7] Bolton, Ruth N., Katherine N. Lemon, and Peter C. Verhoef (2004), "The Theoretical Underpinnings of Customer Asset Management: A Framework and Propositions for Future Research." Journal of the Academy of Marketing Science, 32(3), 1-20.

[8] Bult, JR, Wansbeek, T (1995), "Optimal selection for direct mail," Marketing Science, Vol. 14 No.4, pp.378-44.

[9] Chen, Yuxin, Chakravarthi Narasimhan, Z. John Zhang (2001), "Individual marketing with imperfect targetability," Marketing Science 20(1) 23-41.

[10] Chintagunta, Pradeep, Tulin Erdem, Peter Rossi and Michel Wedel (2006), "Structural Modeling In Marketing: Review and Assessment," Marketing Science. 25(6), 604-616.

[11] Gonul, F. and Shi, MZ (1998), "Optimal Mailing of Catalogs: A New Methodology Using Estimable Structural Dynamic Programming Models," Management Science, 44, 1249-1262.

[12] Erdem, Tulin, Michael P. Keane (1996). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. Marketing Science 15 (1), 1-20.

[13] Erdem, Tulin, Susumu Imai, Michael P. Keane (2003). Brand and quantity choice dynamics under price uncertainty. Quantitative Marketing and Economics, 1, 5-64.

[14] Gonul, Fusun, Kannan Srinivasan (1996). Impact of consumer expectations of coupons on purchase behavior. Marketing Science 15 (3), 262-279.

[15] Haeckel, SH (1998), "About the Nature and Future of Interactive Marketing," Journal of Interactive Marketing, 12 (1) 63-71.

[16] Harrington, JE, (1995), "Experimentation and Learning in a Differentiated-Products Duopoly," 66, (1), 275-288.

[17] Hitsch, GJ, (2006), "An Empirical Model of Optimal Dynamic Product Launch and Exit Under Demand Uncertainty," Marketing Science, 25 (1), 25-50.

[18] Judd, K. L. and Riordan, M. (1994), “Price and quality in a new product monopoly,” Review of Economic Studies, 61,773-789.

[19] Kamakura, A. Wagner and Russell, Gary J, ”A Probabilistic Choice Model for Market Segmentation,” Journal of Marketing Research. Chicago: Nov 1989. Vol. 26, Iss. 4; p. 379-90.

[20] Kamakura Wagner et al (2006), “Choice Models and Customer Relationship Managements - Summary Paper for the Sixth Choice Symposium,” forthcoming at Marketing Letters.

[21] Krishna, Aradhna (1992). “The normative impact of consumer price expectations for multiple brands on consumer purchase behavior.” Marketing Science 11 (3), 266-287.

[22] Lewis, Michael (2005), “A Dynamic Programming Approach to Customer Relationship Pricing.” Management Science, 51(6), 981-994.

[23] Lipster, R.S. and Shiryaev, A.N. (1977), “Statistics of Random Processes, I: General Theory.” Springer-Verlag, New York.

[24] McLennan (1984), “Price dispersion and incomplete learning in the long run, Journal of Economic Dynamics and Control. 7 (3): 331-347.

[25] Narayanan, Sridhar and Puneet Manchanda (2006), ???Heterogeneity Learning and the Targeting of Marketing Communication for New Products.??? Working paper, University of Chicago.

[26] Netzer, Oded, James Lattin, V. Srinivasan (2005). “A hidden Markov model of customer relationship dynamics.” Working paper, Columbia University.

[27] Peppers, Rogers, and Dorf (1999), “Is your Company Ready or One-to-One Marketing.” Harvard Business Review, 77, 151-160.

[28] Reinartz and Kumar (2000), “On the Profitability of Long-Life Customers in a non contractual Setting: An Empirical Investigation and Implications for Marketing.” Journal of Marketing, 64, 17-35.

[29] Reinartz and Kumar (2003), “The Impact of Customer Relationship Characteristics on Profitable Lifetime Duration.” Journal of Marketing, 67, 77-99.

[30] Rossi, Peter E., Robert E. McCulloch, and Greg M. Allenby (1996), ???The Value of Purchase History Data in Target Marketing,??? Marketing Science, 15(4), 321-40.

[31] Rust, Roland T and Tuck Siong Chung (2005), “Marketing Models of Service and Relationships.” Forthcoming, Marketing Science.

[32] Rust, Roland T., Katherine N. Lemon and Valarie A. Zeithaml (2004), “Return on Marketing: Using Customer Equity to Focus Marketing Strategy.” Journal of Marketing, 68 (1), 109-127.

[33] Rust, Roland T., Peter C. Verhoef (2005), “Optimizing the marketing interventions mix in CRM.” Marketing Science.

[34] Sanford J. Grossman, Richard E. Kihlstrom, and Leonard J. Mirman (1977), “A Bayesian Approach to the Production of Information and Learning by Doing,” The Review of Economic Studies, 1977, 44, (3) 533-547

[35] Sawhney, Mohanbir, Sridhar Balasubramanian and Vish Krishnan (2004), “Creating Growth With Services”, MIT Sloan Management Review, 45(2), Winter.
bibitem Schmittlein, D.C., Peterson, R.A., 1994. Customer base analysis: An industrial purchase process application. Marketing Science 13, 41-67.

[36] Shankar, Venkatesh, Amy Smith, and Arvind Rangaswamy (2003), “The Relationship Between Customer Satisfaction and Loyalty in Online and Offline Environments,” International Journal of Research in Marketing, Vol. 20, No. 2 (June), p. 153-175.

[37] Shugan, Steven. 2004. The impact of advancing technology on marketing and academic research. Marketing Science 23 (4), 469-475.

[38] Sun, Baohong and Shibo Li (2005), “Learning and Acting Upon Customer Information - With an Empirical Application to the Service Allocations with Off-Shore Centers,” working paper, Carnegie Mellon University.

[39] Sun, Baohong, Shibo Li, and Catherine Zhou (2006), “Adaptive” Learning and ”Proactive” Customer Relationship Management. forthcoming at Journal of Interactive Marketing.

[40] Trefler, Daniel (1993), “Optimal Learning With Endogenous Information,” International Economic Review 34, 565-581.

[41] Venkatesan, Rajikumar, V. Kumar (2004), “A customer lifetime value framework for customer selection and resource allocation strategy,” Journal of Marketing 68 (4), 106-125.

[42] Wind and Rangaswamy (2001), “Customization: The Second Revolution in Mass Customization,” Journal of Interactive Marketing, 15(1):13-32.

[43] Winer, Russell (2001), “A Framework for Customer Relationship Management,” California Management Review, 2001, 43. (Summer), 89-105.