Academia Revista Latinoamericana de Administración

Why individualized contact policies are critical in the mass affluent market

Gabriel Bitran Susana Mondschein

Why individualized contact policies are critical in the mass affluent market

Por qué las políticas de contacto individualizado son importantes con los clientes de alto poder adquisitivo

Received 24 April 2014
Revised 14 March 2015
Accepted 16 March 2015

Gabriel Bitran
Sloan School of Business, MIT, Cambridge, Massachusetts, USA, and
Susana Mondschein
Facultad de Ingeniería y Ciencias, Universidad Adolfo Ibañez, Santiago, Chile

Abstract

Purpose – The purpose of this paper is to study the optimal contact policies for customers that belong to the mass affluent market.

Design/methodology/approach – The authors formulate a stochastic dynamic programming model to determine the optimal frequency of contacts in order to maximize the expected return of the company.

Findings – The authors show that personalized marketing strategies provide a competitive advantage to companies that contact their customers directly through, for example, phone calls or meetings. The authors show that a threshold policy is only optimal for customers with increasing sensitivity to contact. In all other cases, optimal policies might have a less intuitive structure. The authors also study the importance of the size of the customer database and determine the optimal maximum recency when maintenance costs are present.

Practical implications – Contact policies should be tailored for each company/industry individually, due to their sensitivity to customers' purchasing behavior.

Keywords Marketing and advertising, Stochastic dynamic programming
Paper type Research paper

Resumen

Propósito – En este artículo se estudian políticas óptimas de contacto de clientes que pertenecen al segmento de mercado con alto poder adquisitivo.

Diseño/metodología – Para ello formulamos un modelo de programación dinámica estocástica para determinar la frecuencia óptima de contactos durante un período de tiempo, de modo de maximizar el beneficio esperado de la empresa.

Resultados – Los resultados muestran que una estrategia de marketing personalizada entrega una herramienta competitiva a aquellas empresas que contactan directamente a sus clientes, a través, por ejemplo, de llamadas telefónicas o reuniones periódicas. También se muestra que las políticas con un

JEL Classification — M19

Academia Revista Latinoamericana de Administración
Vol. 28 No. 2, 2015
pp. 251-272

1012-8255
DOI 10.1108/ARLA-03-2014-0037

umbral a partir del cual es óptimo contactar a todos los clientes son óptimas solo en el caso de clientes con una creciente sensibilidad al contacto. En todos los otros casos, las políticas óptimas tienen una estructura menos intuitiva. Finalmente, se estudia el tamaño óptimo de la base de datos de clientes y hasta que antigüedad es recomendable mantenerlos en el sistema.
Implicaciones prácticas – Las políticas de contacto a clientes deberían ser confeccionadas a la medida de las características individuales de la empresa, puesto que ellas son altamente sensibles al comportamiento de compra de los clientes.
Palabras clave Marketing y publicidad, Programación dinámica estocástica
Tipo de papel Trabajo de investigación

1. Introduction

Different definitions of the mass affluent market are found in the literature, depending on the country, year of the article, or simply on the personal preferences of the author. Some definitions include "households with $500,000 or more in net worth, not including primary residences" (Sill-Levy and Hagan, 2008), "people with $100,000 to $1 million in assets to invest" (Business Week, 2004), "households with a net worth between $250,000 and $2 million" (Paikert, 2008). To reconcile these differences, we use a qualitative definition for this market where mass affluent customers are those who fall between the middle class and the wealthiest consumers. The traditional division in marketing between mass and high-end market is no longer true due to the rapid and steady growth of the affluent mass market, where customers are willing to pay a premium for brand and quality. For example, the number of consumers with an annual income of US$80,000 or greater has been steadily increasing from 1996 to 2005 and more than double in those ten years. A similar trend can be observed for consumers with an annual income of US$100,000 or more, where the number of households falling in this category increased from 8.2 percent in 1996 to 17.2 percent in 2005 (US Census Bureau). The shape of income distribution in the USA also has changed dramatically since 1970. Nowadays, it is possible to observe a ski slope instead of a bell curve. Currently, the income distribution starts with a large mass at the left-hand side, which corresponds to lower-income households, and then the curve slopes downward until quickly flattens out, forming a very thin tail of higher-income households Although, the income distribution still shows a significant number of low-income households, it does not show a huge middle-class mass market. Instead, it shows a mass-market that has shifted upscale over time (Nunes and Johnson, 2004). This change in the income distribution shows clearly the growth of the so-called mass affluent market that is the focus of our study. For a complete discussion of the qualitative and quantitative features of this new market refer to Nunes and Johnson (2004) and Sam (2007).

As it is documented in the literature, companies are either using mass marketing strategies to target this new segment or sales associates are implementing contact policies based, in most cases, on their own personal experience. Nunes and Johnson (2004), provide an excellent example of an ineffective mass marketing contact policy of an affluent customer: "[…] the experience of one watch owner demonstrates that some firms have a long way to go in effectively communicating with the moneyed masses. Though this customer has purchased three watches from a certain top luxury brand, he has never been personally contacted or invited back to the store. Instead, the company regularly mails him a slick (and clearly expensive) in-house marketing publication built around hip topics like the modern art scene, something he cares nothing for. Even making an allowance for its auto-personalized letter, these mass mailings are so badly matched to his expectations that he is unsure he will ever be as interested in the company's products the same way again." Other authors have also addressed the importance of a contact policy in terms of the adequacy of the sales environment. For example, Oechsli (2005) states a "Proactive relationship-building principle: From the initial contact forward, your emphasis must be to take the initiative and proactively build a long-term, professional relationship: everyone in your organization who sells and supports must understand and fully engage in the process of building long-term professional relationships." Traditionally, for the mass-market segment, monetary incentives have been a successful tool to boost sales in the form of coupons, rebates, and limited-time offers. However, these are not necessarily good incentives for the more affluent customers. According to Nunes and Johnson (2004), marketers should employ promotions that promise knowledge and connoisseurship. One effective channel to implement these more intangible attribute is through personal contact with customers.

There is a vast literature that addresses the effect of marketing mix variables on sales, namely the response models for marketing management. These models are used to study the market behavior and, ultimately, the impact of different marketing efforts on sales. The common methodology behind the response models is econometric and time series analysis. A complete analysis of these models can be found in Hanssens *et al.* (2001). In our paper we study the effect of one particular kind of marketing effort: individualized contacts policies on the new and growing Mass Affluent Market. We differentiate from previous approaches in the use of an optimization model (stochastic dynamic programming model) to determine the optimal timing when these contacts have to be made, as a function of the attributes of each individual customer (e.g. recency, frequency and monetary order). Ma *et al.* (2015) propose a model to answer the question of when it is the best time to reactive an inactive customer. We argue in this paper that, to maximize the lifetime value of a customer, it might be even worth to contact highly active customers with the purpose of increasing their purchasing response rate.

The Supernova project at Merrill Lynch (Oliva and Bitran, 2003) provides an example of a successful application of an individualized contact policy for the mass affluent market. Before implementing Supernova, some financial advisers (FAs) routinely contacted clients to check on them, offer advice on existing investments with the company, or solicit additional business. However, the majority of the FAs rarely contacted the smaller clients in their portfolio, except when they were called to initiate a trade, or report a problem. FAs were extremely independent in the way they conducted business as long as they kept within the letter of the law. FAs managed their own "book" of clients (broker's list of clients having an account with the firm), and usually felt safe with a big book. As a result, service, customer retention, and profitability did not matter much. According to one Merrill FA "all those names really make you feel secure, and that's important in a business like this where you can rely on yourself and you have to keep producing if you want to eat." The Supernova project came as an initiative to improve significantly the quality of the service provided by the FAs to better serve the more affluent customers. As the "father" of the project put it "The objective of Supernova is to create the ultimate client experience." Thus, the executives at Merrill believe that three aspects of a relationship were critical to a client satisfaction: the frequency and quality of contact, rapid response to problems, and attention to details. They strongly believe that in order to accomplish these goals, each FA had to have a book that did not exceed a total of 200 clients, instead of the average of 550 clients per FA. This reduction would make them serve well those customers that make significant business with them, for example those clients having at least \$1 million in annuitized assets at Merrill. And therefore, FAs had to focus on the affluent mass market, with the strong believe that serving fewer but more profitable customers would increase the total expected revenue for them and the firm. With the new scheme, FAs dedicate between six to eight hours each day to these meetings. Of course, a lot of practical issues arise when implementing such a policy. For example, what to do with the less profitable customers? How to convince FAs to move from the old scheme to the Supernova project? and, how to help FAs to make this transition? These critical issues of implementation are discussed in Section 3.

We show in this paper the crucial role of individualized contact policies in successfully managing and retaining customers in the mass affluent market. We study the dependency of these policies on the status of the customer in terms of her/his purchasing pattern and on the type of products or services being offered. We show, for example, that the shape of the learning curve for the use of a new product or service determines the structure of the optimal contact policy, i.e., when a customer should be actively contacted. Thus, a contact policy that is optimal in the clothing industry might not be suitable for the cell phone industry. Furthermore, we show that some optimal contact policies challenge the common managerial intuition. For example, in some cases, it is advisable to leave "loyal customers" alone, and only spend resources on recovering "bad customers." We also show that the customers' database maintenance costs play a central role when deciding who should be kept in the sales associate costumers list. Thus, the common perception in the industry that the larger the customers' database, the more profitable it would be proved not to be necessarily true when maintenance costs are significant. We explore the structure of efficient and effective (reaching the right customers at the lowest cost) strategies for contact policies. Finally, a well-designed contact policy should also consider acquiring new mass affluent customers. Therefore, prospecting should be a planned activity for the sales associate: how to contact, when, and by what means.

The graph at the top of Figure 1 shows the net profit of the never and always contact policies as a function of the maximum net profit for different contact costs. The maximum net profit is the profit obtained when the optimal contact policy is implemented. The never contact policy is defined as a policy where customers are never contacted independently of their purchasing history. Similarly, the always contact policy is such that all customers in the sales associate customers' list are contacted periodically independently of their purchasing behavior. As expected, the never contact policy improves its performance as the contact cost increases, and the always contact policy worsens its performance as the contact cost increases. However, it is interesting to notice that the structure of the optimal contact policy changes in a "non-continuous" way as the contact cost increases. Thus, at the bottom of Figure 1, we observe that when the contact cost is low, the optimal contact policy coincides with the always contact policy. As the contact cost increases, the optimal policy becomes a "threshold" policy where it is optimal to contact only those customers whose last recency[1] is smaller than a certain "threshold" value. However, as the contact cost further increases, the optimal policy has a mixed structure, where it is optimal to contact those customers with intermediate recencies and not to contact customers that have bought recently or a long time ago. Finally, for a contact cost high enough, the optimal contact policy coincides with the never contact policy. We observe with this example that although net profits behave as expected when contact costs increase, the structure of the optimal contact policy behaves in a less intuitive way, and therefore, to

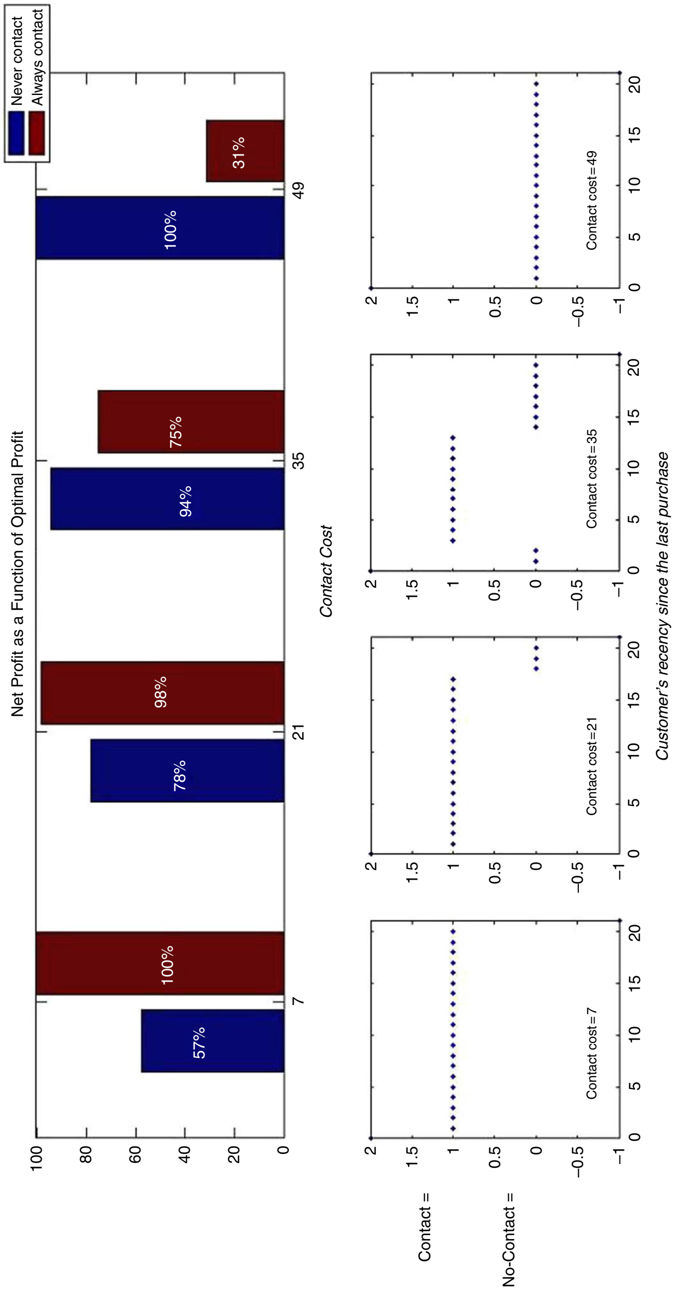

Figure 1.
Performance of never and always contact policies as a percentage of the expected profit with the optimal policy

"guess" its shape using personal experience or successful applications in other industries might lead to important loses in profits. In Sections 2 and 3, we analyze in detail which are the main factors that determine the structure of the optimal contact policies depending on customers' behaviors and type of industries.

There are two main questions that need to be addressed when defining the optimal contact policies: when and who to contact in each period during the planning horizon. In Section 2, we study what are the factors that determine the answers to these questions. When refers to the recencies or states for which it is optimal to contact a customer as her/him status evolves dynamically as a function of time and purchasing behavior. Who refers to the set of customers that is profitable to keep in the sales associate customers' list. We show that there is a tradeoff between managerial costs of keeping customers in the database and future expected profits associated with them.

2. When and who to contact

In this section we use a qualitative classification of customers as a function of their recency. We call active customers those who have a "low" recency, and therefore have bought recently from the company. Mature customers are those with a "medium" recency. Finally, passive customers are those with a "high" recency that still belong to the sales associate's customer list but have not bought from the company for a long time. We assume that the purchasing probability is a non-increasing function of the recency, i.e., a passive customer is less likely to buy in the current season compared to an active or mature customer, and that a contact performed by the sales person increases the purchasing probability in a given period of time.

2.1 When: customers' classification in terms of sensitivity to contact

As we mention in "Introduction," the question of when to contact a customer depends largely on her/his sensitivity to respond to that contact. We define the sensitivity to contact as the increase in the purchasing probability when contacting a customer vs not contacting him/her. There are two main factors that contribute to the sensitivity to a contact: the type of industry and the recency of the customer within the company.

We classify sensitivity to contact depending on the industry in four categories: increasing, decreasing, quasi-concave, and quasi-convex sensitivity. We measure sensitivity to a contact as the difference between the purchasing probability with and without a contact for a given recency.

In Section 3, we will show that depending on the type of sensitivity to contact or equivalently on the type of industry, optimal contact policies might have different structures, and therefore, what is optimal in one industry (or company) might not be necessarily optimal for another. The following classification is important to understand why contact policies are complex and intuition and experience might not always work appropriately. In what follows we describe the different type of sensitivities where industries might be classified. We notice that this is a qualitative classification, and actual classification must be performed using empirical data:

- Increasing sensitivity: in this category we include industries where the impact of contact on customers increases with recency. Thus, the increase in the purchasing probability is higher for those customers with a higher recency. We observe in this case, that customers that have not bought from the company for a longer period of time, "passive customers," are more responsive to a contact.

One example in this category is the plastic surgery industry. Contacting a customer that just had a surgery may be a bad idea because he/she might still be sore or the positive effects of surgery might not be evident yet. However, as time goes by, the customer might be more receptive to "new improvements," so the impact of a contact is higher for those customers that have recovered and forgot about the unpleasant procedure. In general, increasing sensitivity to contact can be observed in businesses where the rewards of the product or service are delayed in time compared to the moment it was acquired. Other examples include landscaping and preventive medicine[2].

- Decreasing sensitivity: in this category we include industries where impact of contact on customers decreases with recency. In these industries, we observe that customers that have not purchased from the company for a longer period of time are less sensitive to a contact for a variety of reasons including switching to the competition or moving to another income niche. Therefore, the short term benefits of contacting an active customer are higher compared to those when contacting mature or passive customers. It is important to mention that higher short term benefits do not necessarily imply that only active customers are worth to contact when maximizing total profit over the planning horizon. We discuss this in detail in Section 3. One industry that fits in this category is luxury hands bags (Louis Vuitton, Hêrmes). In this case, if a customer has become a passive customer (she has not bought from the company in the last seasons), most likely she has switched to the competition because she likes another brand better or she cannot afford this brand anymore. And therefore, the impact of a contact on the purchasing probabilities is higher when she/he is still an active or mature customer of the company. In general, decreasing sensitivity is observed in industries where customers have immediate rewards from the products and services and consumption is periodic. Other industries include clothing, cosmetics, and gourmet food.
- Quasi-concave sensitivity: in this category we include industries where impact of contact on customers reaches a maximum for mature customers. Thus, for low or high recencies, the impact of a contact is relatively low compared to the effect on customers with medium recencies. An example in this category is the cellular phone industry. In this case, new users might need time to learn the new technology before acquiring additional services or upgrading the existing ones. Thus, the acquisition of new or complementary products or services might not be attractive when recency is low, i.e., new customers. It might be more sensible to wait until customers are familiar with the service before offering new products through contacting the user. On the other hand, when recency gets too high, it might be an indication that the customer is less inclined to acquire new products. She/he might not be a "high-tech" person that will update the product on a regular basis. In general, industries with a clear learning curve for the use of their services and products can fit well within this category. Other industries include cable TV, computer services, and financial services. Another category of industries that follow this patter for customers' behavior is the durable goods industry, where customers replace their products only after several periods of time. In this category, we have cars, furniture, and appliances. According to the CEO of an exclusive furniture store in Connecticut, customers come back to the store every two to three years on average. Therefore, contacting a customer

right after she/he has made a purchase might not have the same impact as waiting for several periods before offering new products.

- Quasi-convex sensitivity: in this category we include industries where impact of contact on customers reaches a minimum for mature customers. Thus, for low or high recencies, the impact of a contact is higher compared to the effect on customers with medium recencies. An example within this category is the SPA industry, where customers might have a high response to a contact in the months immediately after receiving the service because of the instant gratification of the experience. Given the luxury nature of the service, as time goes by, customers might not be highly sensitive to a contact from the company, mainly because customers tend to forget the intangible benefits of the service and just remember the costs. On the other hand, for customers with high recency that have not spend money on this service for a certain time, might feel that they deserve to be pampered again. Other industries in this category include luxury or exotic trip and cruises (Asia, Alaska, Patagonia, and the Galápagos). However, clothing, in some cases, can also fit in this category. For example, after using a brand for years a customer can get bored and switch to the competition. After a while she/he might be more sensitive to a switch back again.

We notice that the purpose of the description above is to provide a framework to analyze the optimal contact policies and their differences according to the different sensitivity classifications. We remark that some companies within an industry could have a better fit in a different sensitivity category. Ultimately, the classification depends on the empirical estimation of the purchasing probability functions. In what follows, we use a Venn diagram in Figure 2 to show a qualitative classification of some industries according to the customers' sensitivity to contact. We notice that some industries can fit in more than one segment depending on the customers' behavior. Thus, this classification is only intended as a qualitative illustration. Figure 3 shows

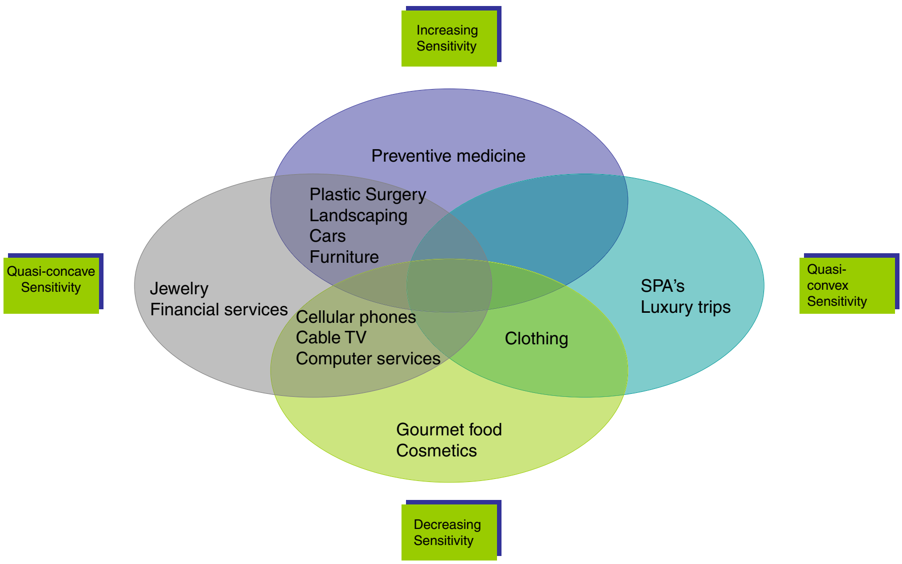

Figure 2.
Classification of industries according to their responses' sensitivity to contact policies

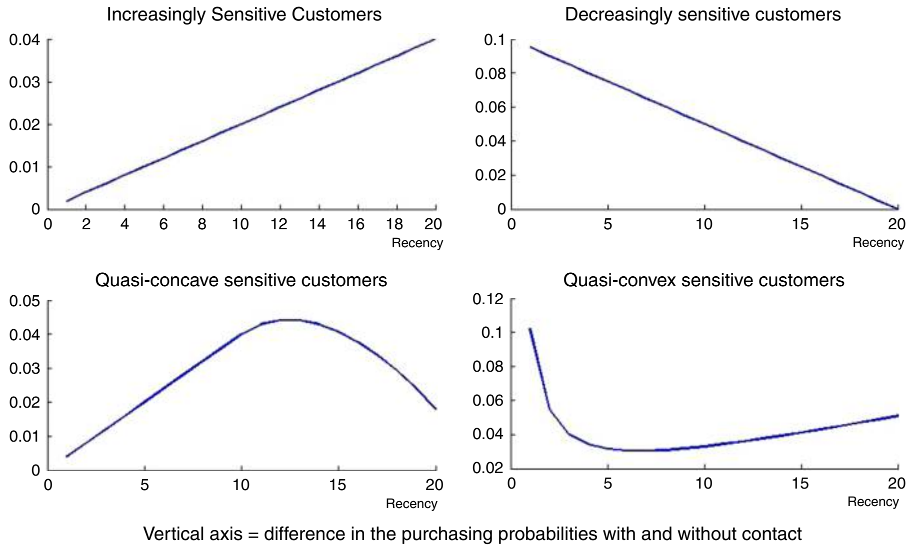

Figure 3. Absolute difference between the purchasing probabilities with and without a contact as a function of the customer's recency for different customers' sensitivities

examples of the dependency of purchasing probabilities as a function of the type of sensitivity and recency.

2.2 Who: choosing the optimal size for the sales associate customers list

The popular belief that a larger database of customers is a synonymous of a more valuable asset is not necessarily true. In many cases, it might reflect the fact that the database contains some irrelevant information that only costs money to the company but contributes marginally to profits. Who are the "value customers" whose names and personal information are worth to keep in the company's database? This is the fundamental question that we analyze quantitatively in Section 3.2.

In the academic literature, the managerial cost of building and maintaining a customers list is usually neglected. For example, see Bitran and Mondschein (1996) and Gönül and Shi (1998). However, in practice this might be an important barrier to efficiently manage a customers list. As stated by the owner and CEO of an upscale furniture company in Connecticut[3], one of the big obstacles in using the purchasing information contained in the customers list is the cost in time and money involved in making this information accessible for marketing purposes. Even for the cases where the information is already "computerized," the hardware and/or the software might not be appropriate for manipulating the information on a periodic basis. Furthermore, there is an important periodic monetary cost associated to the maintenance of the customers' database: adding and removing names and updating customers' information (addresses, e-mails, phone' numbers, socio-economic data, and purchasing history). Thus, the larger the size of the associate's customers list, the larger the total periodic maintenance cost. And therefore, there is a tradeoff between the maintenance costs of more customers in the associate list and the expected purchasing profits associated to them.

Additionally to the direct monetary cost, it is also critical to analyze the indirect cost associated to a potential deterioration in the quality of service due to an unnecessary large database. For example, in the Supernova project at Merrill Lynch

(Oliva and Bitran, 2003), this was a crucial factor when reducing the customers list from 550 to 200 clients. As it was stated by a FAs: "[…] if a client called, we spoke to them. We didn't have time to make calls to clients because we were busy dealing with clients calling us – wanting us to fix problems they were having, hold their hand when the markets declined, or do trades. Extra time was spent prospecting – and we had to do a lot of that given the number of clients who quit. It was hard to respond to problems quickly. We used to get overwhelmed […]".

For example, if we consider that customers are classified uniquely in terms of their recency, then the size of the associate's customers list is defined by the maximum recency allowed for a customer, i.e., if a customer reaches this maximum recency, she is removed from the associate list. In Figure 4 we show the effect of the value of the maximum recency allowed in the customers list on the total expected profit of a customer whose current recency is equal to one (she purchased in the last period). The horizontal axis corresponds to different sizes allowed for the customers' recency to remain in the list. For the case where there is no maintenance cost, the expected profit is a monotonically increasing function of the size for the maximum recency allowed, and therefore, the larger the size of the list the higher the profits. The intuition behind this result is simple; as long as there is a positive probability of a purchase and there is no cost of maintaining a customer in the associate list, it is always beneficial to keep his/her name as a potential customer. However, when we increase the maintenance cost per customer to \$3 per period, we do not observe this asymptotic behavior anymore. In fact, the net profit has a quasi-concave shape reaching a maximum when the maximum allowed value for the recency is equal to 10. If the unit maintenance costs increase to \$5 per period, then the optimal maximum recency for the associate list decreases to 5. As expected, we will keep in the customers' list only those customers that are very profitable.

In Section 3, we discuss a mathematical model to quantitatively study the questions introduced in Section 2. Using a dynamic programming formulation, we show that the optimal contact policies depend on the customers' current and future stochastic purchasing patterns, on the purchasing probabilities with and without a contact, on the monetary cost of contact, and on the size of the maximum recency allowed in the

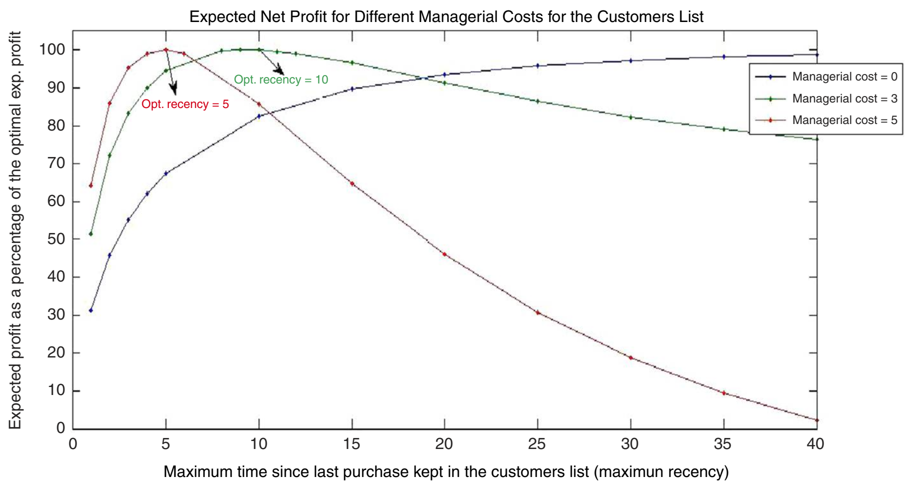

Figure 4.
Optimal expected profits for different managerial costs as a function of the maximum recency allowed in the customers list

customer database. Therefore, we will show that the fact that a customer is in a state where she is highly receptive to a contact does not necessarily mean that a contact must be made. This decision must consider all future dynamic for the customer behavior. The opposite is also true; if a customer is particularly insensitive to a contact, it is not always true that the optimal decision is not to contact that customer. On the contrary, it might be the case where a contact is valuable and optimal, because of the chances that the customer will be "activated" and move to a more profitable state.

3. Model

In what follows we introduce a simple model to describe the customers' behavior as a function of the contact policy. We assume that customers are segmented according to their recency, i.e., last period (season) since they bought from the company. Thus, if a customer bought in the last period her recency is equal to 1, and in general, it is equal to i if the last time the customer bought from the company was i periods ago.

3.1 Single client model with an infinite planning horizon

The following model maximizes the lifetime value of a customer whose current recency is equal to i. We consider that for every period there is a probability that she makes a purchase, which increases when the customer is contacted by the company. Additionally, we consider that there is a unit cost of contacting a customer (time and/or opportunity cost of the sales person and any other additional expenses). In this model, we assume that there is no maintenance cost for the customer database and consider a maximum recency that is allowed for the customers; when this is exceeded the customer's information is deleted from the database. Every period, the company has to decide whether or not to contact each customer in order to maximize the expected profit associated to her.

Notation.

- $V(i)$ is the lifetime value of a customer that has a recency equal to i at the current period.
- R the average monetary purchase. We assume that customers in different states of recency have the same average monetary purchase. However, the model can be easily extended to the case when the monetary order depends on the customer's recency.
- c the unit contact cost.
- u the binary decision variable that takes value 1 is the customer is contacted in the current period or 0 otherwise.
- $p_u(i)$ the purchasing probability for a customer with recency equal to i and contact policy u is implemented.
- α the discount rate.
- I is the largest recency kept in the customers' list.

$$V(i) = Max_{u \in \{0,1\}} \left\{ g_{i,u} + \alpha p_u(i) V(1) + \alpha(1-p_u(i)) V(i+1) \right\}, \quad \forall i = 1, \cdots, I \qquad (1)$$
$$V(I+1) = 0.$$

where $g_{i,u} = Rp_u(i) - uc$.

We consider that the probability of purchasing is higher (or equal) when the customer is contacted by the company, compared to the case when she/he is not. Therefore, $p_1(i) \geqslant p_0(i)$ $\forall i$. We notice this is a realistic and fundamental assumption in this paper, otherwise contact policies would not be object of research:

P1. For the case of increasing sensitivity to contact the optimal contact policy is a threshold policy such that if it is optimal to contact a customer with recency i then it is optimal to contact her at any higher recency.

Proof. See Appendix.

We emphasize that *P1* does not use any assumption regarding the shape of the purchasing probability functions for the cases with and without a contact, other than assuming that the purchasing probability for a given recency is positively affected by a contact.

For the cases of decreasing sensitivity, quasi-concave sensitivity, and quasi-convex sensitivity to contact, the optimal contact policies are not necessarily a threshold policy. Thus, depending on the functional form of the purchasing probability functions, we can find different forms for the optimal contact policies. For example, for the case of decreasing sensitivity to contact, there are cases where the optimal policy is to contact customers with low recencies and not to contact those with high recencies. However, there are other instances where the optimal contact policy behaves exactly in the opposite way, i.e., it is optimal to contact those customers with higher recencies instead of those with lower recencies. It is important to point out that the latter case challenges our intuition. At a first glance, we would be more inclined to think that, because customers are less receptive to any contact when they have not bought from the company for several periods, then they probably switch to the competition or simply change the consumption habits. And therefore, it would not be beneficial to contact them. For illustration purposes, we consider the case of a purchasing probability function without/with contact equal to $p_0(i) = 0.7/i$, $p_1(i) = p_0(i) - 0.005i + 0.1$ (which corresponds to the case of a decreasing sensitivity to contact), an average monetary purchase of $R = 1$, a largest recency $I = 10$, and a discount rate of $\alpha = 0.9$. For the case of a unit contact cost $c = 0.20$, the optimal contact policy is not to contact the most active customers with recency equal to 1 and contact the rest ($r = 2$-10). However, when the unit contact cost increases to $c = 0.25$ the optimal policy switches to not to contact costumers with the highest recencies from 1 to 3, then contact those with recencies from 4 to 9, and then switch to not to contact the largest recency equal to 3.

Table I shows the different structure of the optimal policies as a function of the customers' sensitivity to contact. If a cell has a Yes, it implies that the pair (contact

	Increasing sensitivity	Decreasing sensitivity	Quasi-concave sensitivity	Quasi-convex sensitivity
No contact Contact	No	Yes	Yes	Yes
No contact Contact	No	Yes	Yes	Yes
No contact Contact	No	Yes	Yes	Yes
No contact Contact	Yes	Yes	Yes	Yes

Table I.
Optimal contact policies as a function of customers' sensitivity

policy, type of customer sensitivity to contact) is feasible. If it has a No, then the corresponding pair is infeasible.

3.2 Choosing the optimal database

There is a common perception in direct marketing that the larger the number of customers in the database, the better. This might be a true statement in the cases where maintenance costs are negligible. However, in practice, often these costs are relevant and play a significant role when deciding the composition of the sales associate customers' list:

- In this subsection, we discuss a model to determine the optimal maximum recency (time since the last purchase) to be kept in the company's or sales associate database. We define m as the maintenance cost per period and per customer to keep an updated list and $\tilde{V}_I(i)$ as the lifetime value of a customer that has a recency equal to i at the current period when the maximum allowed recency is equal to I. Thus, the model to maximize the lifetime expected value of a customer is given by:

$$\tilde{V}_I(i) = Max_{u \in \{0,1\}} \left\{ g_{i,u} - m + \alpha p_u(i) \tilde{V}_I(1) + \alpha(1 - p_u(i)) \tilde{V}_I(i+1) \right\} \quad \forall i = 1, \cdots, I$$
$$\tilde{V}_I(I+1) = 0. \tag{2}$$

And, the optimal maximum recency in the database I^* is given by:

$$I^* = Arg \max_{I \in Z^+} \tilde{V}_I(1) \tag{3}$$

P2. For the case where the maintenance cost is zero ($m = 0$) the optimal maximum recency in the database goes to infinity.

Thus, if $m = 0$ then $I^* \to +\infty$.

Proof. The proof is straightforward. Consider that the optimal contact policy for the case where the associate list allows a maximum recency equal to I is u, where $u(i)$ is the optimal contact decision when customer is in state i, $\forall i = 1, \ldots, I$. Then, the same policy u is feasible for the case where we increase the maximum recency allowed in the associate list to $I+1$ by adding $u(I+1) = 0$ (we do not contact customers in state is $I+1$). Clearly, the expected return does not decrease in the latter case; adding this additional state cannot lead to negative benefits because even if we do not contact customers in state $I+1$, they can still purchase with a positive probability and then add to the total expected profit (note that the maintenance cost is zero). There might be cases where the purchasing probability is zero for this additional state, and therefore the expected profit remains unchanged. Finally, we note that we are using a feasible contact policy for the case when $I+1$ states are admissible. And therefore, the optimal contact policy for this case could lead to an even larger expected profit. ■

P2 states the intuitive result that when there is no maintenance cost for the database, customers should not be eliminated from the customers associate list, even in those cases where she/he has not bought from the company for a long time. This way, as long as there is a positive purchasing probability, the customer is still potentially profitable. Therefore, it is optimal to keep all customers in the database independent of their purchasing history. We note that this result is valid independently of the shape and the quality of the empirical estimations of the purchasing probability functions

However, in practice there is a maintenance cost for customers in the database; this cost consists of: updating customers' information (address, telephone number, e-mail, etc.); keeping track of customers' purchasing behavior; eliminating customers due to "bad" credit; and software and hardware costs. Therefore, optimization problem (3) must be solved to find the optimal number of recencies for which customers are profitable to be kept in the sales associate customers list. We remark that the optimal contact policy depends on the maintenance cost per period (m) for the associate customer list as well as on the other costs and benefits of the business. For example, a contact policy that is optimal when the maintenance cost is zero might be sub-optimal when the maintenance cost increases; it might be optimal to shorten the expected time in the associate list to avoid the higher maintenance costs. This can be accomplished by using a less proactive contact policy.

The expected time that a customer stays in the database when the maximum recency allowed is equal to I, starting at any period of time when her/his recency is equal to 1, can be computed as follows:

$$\omega = \frac{1+p_{12}+p_{12}p_{23}+p_{12}p_{23}p_{34}+\cdots+p_{12}p_{23}p_{34}p_{45}\cdots p_{I-1I}}{1-\left(p_{11}+p_{12}p_{21}+p_{12}p_{23}p_{31}+p_{12}p_{23}p_{34}p_{41}+\cdots+p_{12}p_{23}p_{34}p_{45}\cdots p_{I-1,I}p_{I1}\right)}$$

where $p_{i,i+1}=(1-p_{u^*}(i))$ and $p_{ii}=p_{u^*}(i)$ represent the purchasing probabilities associated to the optimal contact policy μ^*. This calculation uses the Markov Chain representation of the stochastic evolution of a customer over the planning horizon and the fact that the last state is an absorbing state.

3.3 Implementation issues

Implementing a contact policy can pose several practical problems. In this subsection we analyze the importance of these issues as well as some approaches to solve them.

Transition. If the firm is already established, then an important transition takes place from the traditional marketing approach to the new scheme of contacting clients as a strategy of boosting purchasing responses. How to persuade a sales associate to spend more time in the telephone making calls instead of spending that time on the sales floor? For example, in the Supernova project at Merrill Lynch described in Section 1 (Oliva and Bitran, 2003), more than 80 percent of the FAs resisted the change and were not willing to adopt the new contact policy immediately[4]. For this purpose, road shows were put in place around the country, and for those FAs who did not attend these road shows, one-to-one meetings were set up with their managers who made a compelling argument in favor of the program. On the other hand, a number of clients had to be removed from the sales associate list (those who were not worth to be kept in the sales associate customers list). A concern to both the FAs and the company was the potential loss of these customers. From the company point of view, these customers can be profitable if served by an appropriate channel like a call center. In the Supernova project, the less profitable customers were transferred to Merrill's centralized facility for smaller accounts. This center proactively called the clients at least four times a year to ensure that their needs were being met. Many clients did not object to this new style of service. In fact the retention rate for the center was higher than that of the average FA's.

Incentives. An important question that executives have to address is how to align the incentives of the sales associates and the firm. In our formulation we consider an infinite planning horizon where the total expected revenue from each customer is maximized. However, it is not unusual in many industries, for sales associates to have a short or medium term planning horizon. Therefore, they may resist to implement the optimal contact policies because they might result in losses for them in the short term. For example, sales associates might believe that spending all of their time on the selling floor is more profitable for them, than spending considerable time on the phone contacting customers. When implementing the Supernova project at Merrill Lynch, one of the FAs pointed out that although the first months were difficult, when they were settle into the new routine, the phone did not ring very often, they met with the people that paid them, and got rid of the rest, and earned, in many cases, considerably more. When hiring new FAs, managers were looking for people who were service-oriented, as opposed to transaction-oriented. In the retail industry, this transition can be addressed in a similar fashion. Additionally, the implementation of a career type of commitment might help the sales associates to consider a medium or long-term horizon to make their selling decisions.

Parameter estimation. The purchasing probability functions are critical parameters in the model and need to be estimated using the past purchasing history of customers. We assume that customers are segmented according to their recencies, and therefore, for each of these states the purchasing probability can be estimated with and without contact. For example, using the maximum likelihood method, an estimate for these probabilities would be given by:

$$\left(\hat{p}_u(i) = \frac{k_u(i)}{n_u(i)},\right) \quad \forall i,\ u \in \{0,1\}$$

where $k_u(i)$ is the number of observations of customers in state i that have bought when contact policy u was applied and $n_u(i)$ is the total number of customers in state i when policy u was applied (independently if they bought or not). The mean and variance are equal to $p_u(i)$ and $p_u(i)(1-p_u(i))/n_u(i)$, respectively. Therefore, for a given error on the estimation of the purchasing probabilities, we can determine the minimum number of observations required in each segment i.

However, when analyzing the mathematical structure of the model, we observe that a crucial parameter to determine the optimal contact policy is the difference between the purchasing probabilities with and without contact. Thus, when deciding for the optimal policy (whether or not to contact a customer with recency equal to i), we compare the objective functions $g_{0,i} + \alpha p_0(i) + \alpha(1-p_0(i))V(i+1)$ and $g_{1,i} + \alpha p_1(i) + \alpha(1-p_1(i))V(i+1)$.

For example, at state i, when making the decision whether or not to contact a client with a recency i, the purchasing probability difference has to satisfy the following inequality in order to be optimal to contact the customer:

$$\Delta p(i) = p_1(i) - p_0(i) \geq \frac{c}{\hat{R} - \alpha V(i+1)}.$$

where $\hat{R} = R + \alpha V(1)$ and

$$V(i+1) = \hat{R}\left[\begin{array}{l} p_0(i+1) + \alpha(1-p_0(i+1))p_0(i+2) + \alpha^2(1-p_0(i+1))(1-p_0(i+2))p_0(i+3) + \cdots + \\ \alpha^{I-i}(1-p_0(i+1))(1-p_0(i+2))(1-p_0(i+3)) \times \cdots \times (1-p_0(I-1))p_0(I) \end{array}\right]$$

The expression above for the optimal expected profit for a customer starting at state $i+1$, $V(i+1)$, is valid when the optimal decisions for states $i+1, \ldots, I$ are not to contact the customer. Slightly different expressions for the right hand side of $V(i+1)$ are found for any other combinations of optimal policies for states $i+1$ to I. Given this dependency on the difference between the purchasing probability with and without a contact, it is critical to spend resources in making a good estimation of this difference and not only on the absolute values of these probability functions Thus, a similar analysis has to be done for the estimate of $p_1(i)-p_0(i)$ which mean and variance are equal to $(p_1(i)-p_0(i))$ and $\left((p_1(i)(1-p_1(i)))/(n_1(i)) + (p_0(i)(1-p_0(i))/(n_0(i))\right.$.

Figure 5 shows an example where the actual purchasing probability functions where either overestimated or underestimated by an amount of e of their true value. In this example, the value of e corresponds to 5 percent of the purchasing probabilities without contact. We computed the total expected profit for the different probability estimations as a percentage of the optimal expected profit when using the true purchasing probabilities, for different contact costs. The first case (dark blue) corresponds to the case where the purchasing probability function without contact is overestimated by an amount e and the purchasing probability function with contact is underestimated by the same amount e, and therefore, the absolute difference in the probability functions is equal to $2e$. A similar situation is observed for the last case (dark red), but in this case we underestimate the probability function without contact and overestimate the probability function with contact. For the two cases in the center (light blue and yellow) we under or overestimate the probability functions, but the difference in the probability functions is equal to the true difference. We observe that the performance for the same error in the purchasing probabilities is better for the cases where the difference in these probabilities is

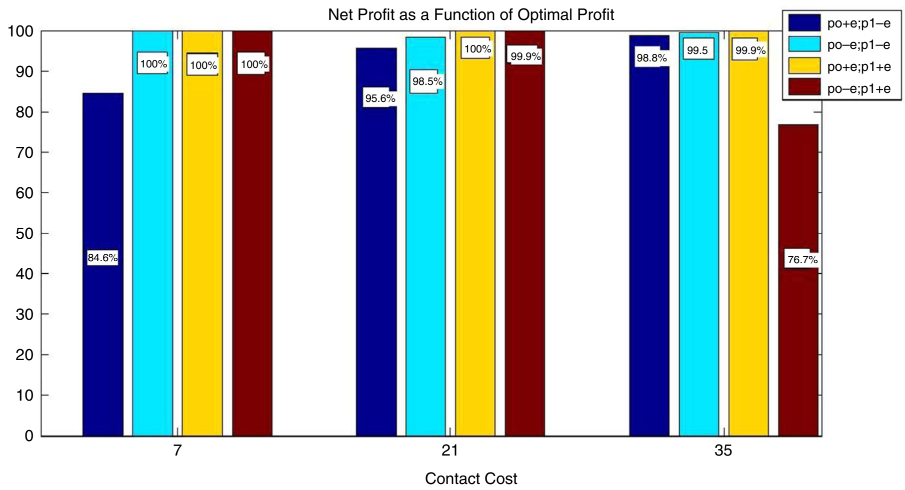

Figure 5.
Impact of estimation errors in purchasing probability functions on expected profit

equal to the true value. Furthermore, in all the cases where the difference is well estimated, the loss in the expected profit is <1.5 percent.

In Figure 6, we compare the expected profits when the difference in the purchasing probabilities and their true value is equal to **e** (implementing the sub-optimal policies given by the imprecise estimations of the purchasing probability functions) against the cases where the never contact and always contact policies are implemented. We notice, however, that the difference between the purchasing probability functions with and without contact remains the same, i.e.:

$$(p_1(i)-p_0(i)) = ((p_1(i)-e)-(p_0(i)-e)) = ((p_1(i)+e)-(p_0(i)+e)) \text{ for all } i.$$

We use the same parameters of the previous example. We observe, that in this case, the contact policies obtained when there are errors in the estimation of the probability functions lead to higher expected profits than in the cases where the never or always contact policies where implemented. Therefore, despite the estimation errors of the purchasing probabilities when implementing the optimal contact policies, they still lead to significantly better results compared to the cases where either we do not contact anybody (never contact policy) or we contact everybody (always contact policy).

We use Figures 5 and 6 to illustrate the impact of parameter estimation errors on the total expected profit when implementing contact policies for a given purchasing probability function. In what follows, we extend these results and consider 800 purchasing probability functions stochastically generated with the purpose of showing the impact of estimation in the values of the optimal objective functions. For each of these 800 hundred functions, we generate a number (n) of random purchasing outcomes and we use these observations to estimate the true purchasing probability functions with and without contact. Using these estimates we solve the mathematical model and determine the (sub-optimal) contact policies. Finally, we compute the difference in the objective functions when using the sub-optimal contact policy with the true purchasing

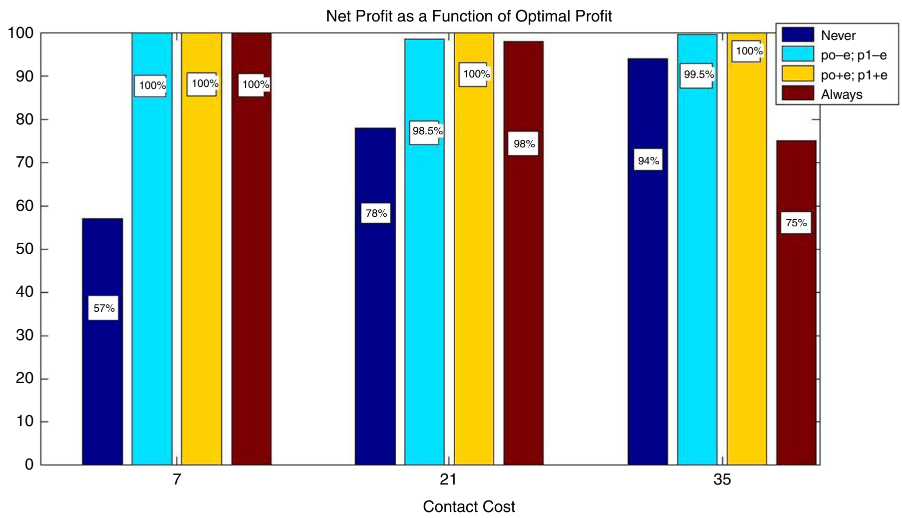

Figure 6.
Impact of estimation errors in purchasing probability functions on expected profit

probability functions with those obtained using the optimal contact policy (obtained when using the true purchasing probability functions).

In Figure 7 we show three histograms for the cases where 100, 500 and 1,000 purchasing observations were randomly generated using the true purchasing probability functions. We observe that when 100 observations were used to estimate the purchasing probabilities, in 56 percent of the cases the losses in the expected profit are < 1 percent, in 68 percent of the cases the losses are < 2 percent and in 75 percent of the cases the losses are < 3 percent. As expected, these results significantly improve when more observations are used to estimate the purchasing probability functions. Thus, when 1,000 observations are used, in 92 percent of the cases the expected losses are < 1 percent, and in 99 percent of the cases are < 2 percent.

In Figure 8 we use the same 800 purchasing probability functions to determine the number of cases where the never and always contact policies lead to losses in the corresponding percentage range. We observe that the losses are significantly higher compared to the cases where we use the optimal contact policies with the estimates of the purchasing probability functions. In the never contact policy, only in 6 percent of the cases the losses were < 1 percent and in the always contact policy this number increases to 30 percent of the cases, but still further away from the 92 percent of the cases were 1,000 observations for the estimations where used. Similar results were obtained in a number of additional simulations performed using randomly generated purchasing probability functions.

4. Conclusions

In this paper we study optimal contact policies for customers in the mass affluent market. We formulate a stochastic dynamic programming model that determines, in every period of time, whether or not to contact a customer in the company's database

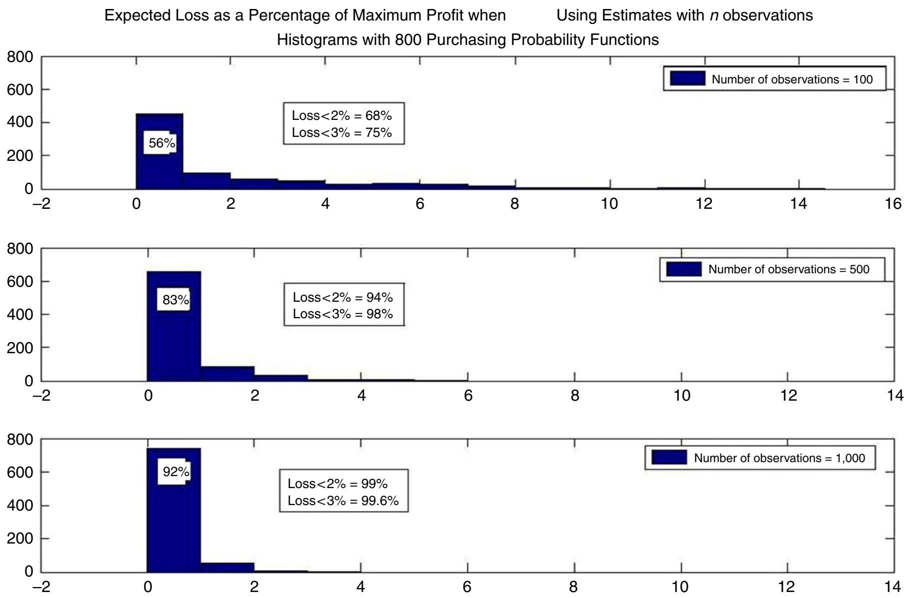

Figure 7. Number of cases (out of 800) with losses within each percentage range due to the use of the estimate of the purchasing probability functions using n observations

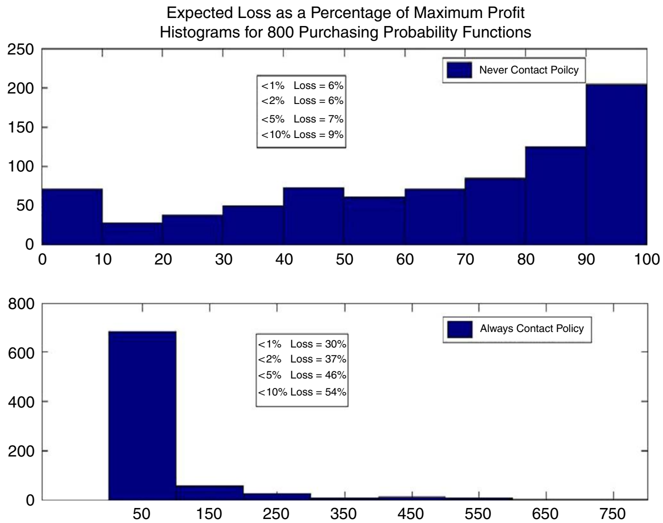

Figure 8. Number of cases (out of 800) with losses within each percentage range due to the use of the never and always contact policies

in order to maximize the expected profit. This decision strongly depends on the state of the customer regarding his/her purchasing behavior.

We characterize customers/industries according to their sensitivity to contact, i.e., the impact of a contact on the purchasing probability functions. Given this classification, we obtain optimal threshold policies only in the cases where customers have an increasing sensitivity to contact (the impact of a contact increases with recency). However, there are a number of other cases where the structure of the optimal contact policy behaves in a less intuitive way, and therefore, to "guess" its shape using the manager's or other industries' experiences might lead to sub-optimal solutions.

We also study the impact of the managerial cost of keeping an updated customers database. We show that when this cost is not negligible, then there is an optimal maximum recency for the customers to be kept in the database; customers whose lifetime falls below a certain value are not profitable when compared to costs of keeping their information in the database.

Finally, we analyze the importance of the parameter estimates when solving the model in practice. The most critical parameters correspond to the purchasing probability functions with and without contact. We show that when making the optimal decisions, the difference between these probability functions is more relevant than the absolute values of each of these curves. We illustrate these results using computational experiments that show that the optimal policies using estimates are close to the "true" optimal policies when using an estimate such that the difference between the purchasing probability functions with and without contact is small (although the actual estimates might have a considerable error). We also show that, even when using estimates instead of the true values for the purchasing probability functions, the expected profits obtained are higher compared to those obtained when simpler contact policies are implemented, such as the never or always contact policies.

The main goal of this paper is to understand the impact of contact policies in the affluent mass market. Therefore, we identify and analyze the main factors that influence the optimal contact policies. Thus, we consider that the optimal policies can be implemented and there are no relevant constraints regarding the amount of time available for contact, interaction among customers, and interaction among sales people. As a future research, a more general model could be developed to include these constraints.

Notes

1. We define recency as the last period (season) when the customer bought from the company. Therefore, a customer with recency equal to one is a customer that acquired a product or service in the last period. In general, if the last time a customer bought from the company is n periods ago, then her/his recency is equal to n. Recency is a measure of the customers' activity in terms of purchasing behavior.
2. Some hospitals or private clinics offer physicals including blood tests, heart tests, and abdominal Cat Scan's on a preventive basis for early detection of medical problems. Although they are controversial, they became a lucrative business in some countries.
3. Triffin, K. 2008, Owner and CEO of Fairhaven Furniture. Personal conversation.
4. FAs had to get rid of approximately 70 percent of their list of customers and contact the remaining clients in a 12-4-2 fashion: minimum of 12 monthly contacts, of which four were portfolio reviews, and two were face-to-face meetings.

References

Bitran, G. and Mondschein, S. (1996), "Mailing decisions in the catalog sales industry", *Management Science*, Vol. 42 No. 9, pp. 1364-1381.

Business Week (2004), "Schwab: courting the 'mass affluent' ", *Business Week*, March 7, available at: www.businessweek.com/magazine/content/04_10/b3873084_mz020.html

Gönül, F. and Shi, M. (1998), "Optimal mailing of catalogs: a new methodology using estimable structural dynamic programming models", *Management Science* Vol. 44 No. 9, pp. 1249-1262.

Hanssens, D., Parsons, L. and Schultz, R. (2001), "Marketing response models", *International Series in Quantitative Marketing*, Vol. 12, Chapter 1.

Ma, S., Tan, H. and Shu, F. (2015), "When is the best time to reactivate your inactive customers?", *Marketing Letters*, Vol. 26 No. 1, pp. 81-98.

Nunes, P. and Johnson B. (2004), *Mass Affluence: Seven New Rules of Marketing to Today's Consumer*, Harvard Business School Press, Boston, MA.

Oechsli, M. (2005), *The Art of Selling to the Affluent: How to Attract Service and Retain Wealthy Customers & Clients for Life*, John Wiley & Sons Inc.

Oliva, R. and Bitran, G. (2003), "Merrill Lynch: supernova", Case Study No. N9-604-053, Harvard Business School.

Paikert, C. (2008), "Investment news", Mass Affluent Market up for Grabs, available at: www.investmentnews.com/apps/pbcs.dll/article?AID = /20080612/REG/626470777

Sam, S.L. (2007), "A control mechanism for sales associates in high-end retail", Doctoral thesis, Massachusetts Institute of Technology, Boston, MA.

Schweidel, D. and Knox, G. (2013), "Incorporating direct marketing activity into latent attrition models", *Marketing Science*, Vol. 32 No. 3, pp. 471-487.

Sill-Levy, E. and Hagan, J. (2008), "What do the affluent want from their providers?", June, available at: www.onwallstreet.com/asset/article/609851/do-affluent-want-their-providers.html

Appendix

Proof of *P1*. For the case of increasing sensitivity we know that the following inequality holds: $p_1(i)-p_0(i) \leqslant p_1(i+1)-p_0(i+1)$. For the purpose of contradiction, we assume that:

- It is optimal to contact a customer with recency equal to i

$$V(i) = Rp_1(i)-c+\alpha p_1(i)V(1)+\alpha(1-p_1(1))V(i+1)$$

- It is optimal NOT to contact a customer with recency equal to i+1

$$V(i+1) = Rp_0(i+1)+\alpha p_0(i+1)V(1)+\alpha(1-p_0(i+1))V(i+2)$$

We know that:

$$\begin{aligned} V(i) &= Rp_1(i)-c+\alpha p_1(i)V(1)+\alpha(1-p_1(i))V(i+1) \\ &\geqslant Rp_0(i)+\alpha p_0(i)V(1)+\alpha(1-p_0(i))V(i+1) \end{aligned}$$

And therefore:

$$V(i+1) \geqslant \frac{R}{\alpha}+V(1)-\frac{c}{\alpha(p_1(i)-p_0(i))} \tag{A1}$$

Equivalently:

$$\begin{aligned} V(i+1) &= Rp_0(i+1)+\alpha p_0(i+1)V(1)+\alpha(1-p_0(i+1))V(i+2) \\ &\geqslant Rp_1(i+1)-c+\alpha p_1(i+1)V(1)+\alpha(1-p_1(i+1))V(i+2) \end{aligned}$$

Thus:

$$V(i+2) \geqslant \frac{R}{\alpha}+V(1)-\frac{c}{\alpha(p_1(i+1)-p_0(i+1))} \tag{A2}$$

Using (A1), we have that:

$$\begin{aligned} V(i+1) &= Rp_0(i+1)+\alpha p_0(i+1)V(1)+\alpha(1-p_0(i+1))V(i+2) \\ &\leqslant \frac{R}{\alpha}+V(1)-\frac{c}{\alpha(p_1(i)-p_0(i))} \end{aligned}$$

And therefore:

$$V(i+2) \leqslant \left[\frac{R}{\alpha}(1-\alpha p_0(i+1))+V(1)\left(1-p_0(i+1)-\frac{c}{\alpha(p_1(i)-p_0(i))}\right)\right]\Bigg/\alpha(1-p_0(i+1)) \tag{A3}$$

Using (A2) and (A3), we obtain:

$$\frac{R}{\alpha}+V(1)-\frac{c}{\alpha(p_1(i+1)-p_0(i+1))} \leqslant \left[\left(\frac{R}{\alpha}+V(1)\right)\left(1-\alpha p_0(i+1)-\frac{c}{\alpha(p_1(i)-p_0(i))}\right)\right]\Bigg/\alpha(1-p_0(i+1))$$

Considering that $\frac{1-\alpha p_0(i+1)}{\alpha(1-p_0(i+1))} \geqslant 1$, we have:

$$-\frac{c}{p_1(i+1)-p_0(i+1)} \leqslant \frac{-c}{\alpha(p_1(i)-p_0(i))(1-p_0(i+1))}$$

But, we know that $p_1(i)-p_0(i) \leqslant p_1(i+1)-p_0(i+1)$, and thus we get a contradiction. ■

About the authors

Gabriel Bitran is a Professor of Management Science and former of MIT Sloan School of Management. His research interests lie in the field of operations management in manufacturing and the service industry. He has published and consulted widely in operations management.

Dr Susana Mondschein is an Associate Professor of Operations Management at the School of Engineering and Science of the University Adolfo Ibañez. Her research interests include the study of optimal decision making under uncertainty in manufacturing and service management problems. She has worked on applications to retailing, catalog sales, hotel reservations, pollution control, and healthcare among others. Dr Susana Mondschein is the corresponding author and can be contacted at: susana.mondschein@uai.cl