Chapter 7
Customer Lifetime Value Applications

Abstract The prior two chapters have covered the technical aspects of Lifetime Value Modeling. But how is LTV applied in the real-world and what types of questions can LTV provide answers that traditional marketing analyses can not? This chapter will provide some answers to these questions. We will discuss how LTV models can be used in the real-world and describe some applications from the literature.

We begin this chapter with the basic analysis of customer acquisition, which has been one of the primary applications of LTV analysis. We will then study reactivation strategies in which the firm selects customers to target for reactivation based on their LTV. Next we will provide some examples of how LTV is used to segment the customer base and is then linked to specific market actions. We will end with the use of LTV models to value firm's customer bases and ultimately the value of the firm.

7.1 Using LTV to Target Customer Acquisition

Probably the earliest application of LTV modeling is for customer acquisition. It offers a very different approach than is commonly used in marketing analysis because it looks at the long-term value of a customer to determine if the firm should make the investment required to acquire the customer. By using an LTV metric for customer acquisition, the firm can determine which acquisition strategies provide the highest payouts or are above the hurdle rate.

Most marketing analysis would concentrate on marketing spending and its return (now called return on marketing investment – ROMI). The typical ROMI model does not separate acquisition from retention marketing spending and cannot see if investing in new customers pays out.

The method begins by determining how much is being spent on customer acquisition. This may be difficult to determine in some industries because it is not be possible to separate new from existing customer expenditures. For example, consumer packaged goods firms often do not know to whom their spending is targeted beyond certain demographic groups. To allocate their spending between acquisition versus retention spending is almost impossible. These firms have very little understanding of customer acquisition costs.

It is easiest to understand the steps if we use an example. Suppose a B2B firm wants to compute the cost of acquiring a customer. In our example, we will assume the firm makes 3 calls per prospect before it decides to cut off future sales calls if no progress is being made. Out of every three prospects, one on average becomes a lead – a prospect that provides some type of "buying signal". Each lead receives approximately 4 sales calls and 1 out of 2 leads become a customer. Thus, it takes on average 17 sales calls to generate one sale.

To cost a sales call the firm keeps records on the direct expenses of a sales person. The average sales person can make 4 calls per week and earns approximately $95,000 per year. Further, there are travel and direct support costs for the sales person which average $35,000 per year. Hence the direct costs associated with the sales person is $130,000 per year. A sales person makes approximately 200 sales calls per year (50×4). Thus the cost of a sales call is $650. If it takes 17 sales calls to close one customer, then the acquisition cost is $17 \times \$650 = \$11{,}050$.

The next step is to determine the lifetime value of the customer once acquired. For this example, the average new customer generates $8,000 incremental profit per year and the retention rate for the customer is 0.75. The firm's discount rate is 15%. Using the LTV formula for this set of assumptions (Chapter 5, Equation 5.2), we have LTV $= m(1+d)/(1+d-r)$, where m = the incremental margin, r = the retention rate and d = the discount rate. This results in the LTV per customer of $23,000. This is greater than the acquisition costs of $11,050.

Suppose the firm could develop a segmentation strategy for its perspective customers. The firm has a segment of customers it estimates will have an incremental margin of $3,000 and the same retention rate as above. This yields LTV = $8,625. The firm cannot afford these customers given the acquisition cost of $11,050.

If the firm can identify customers who have low potential sales, what methods can the firm use to acquire them that generates an LTV greater than the acquisition costs? The firm decided to use telemarketing to generate leads in which low potential firms (identified by firm size) were contacted by outbound telemarketing representatives. The average telemarketing representative makes 10 calls to generate a qualified lead and the estimated cost per call was $3. Through these calls the firm qualified a set of firms. It then sent a sales representative who on average made two direct sales calls and had approximately a 50% chance of closing the account. The acquisition costs for these customers are \$30 (for telemarketing) $+\$650 \times 4 = \$2{,}630$ cost per customer acquired.

Assume the customers acquired through telemarketing contribute incremental profit of \$3,000, a 0.75 retention rate and we use a discount rate of 15%. The LTV of these customers is therefore $\$3,000 \times (1 + 0.15)/(1 + 0.15 - 0.75) = \$8{,}625$, which is greater than the acquisition costs and makes it profitable to use telemarketing techniques to acquire lower-valued customers.

The lessons we have learned from using our LTV model for customer acquisition are: (1) LTV is a better metric for customer acquisition than initial year one profit because the firm is likely to lose money during the acquisition year which would result in not acquiring as many customers (or any customers) as is optimal; (2) comparing LTV to acquisition costs may allow the firm to design alternative acquisition strategies and tactics which can result in finding ways to acquire lower LTV customers.

7.2 Using LTV to Guide Customer Reactivation Strategies

In Chapter 6 we discussed some issues associated with customer reactivation which is a very important application of LTV modeling. We will discuss the generic reactivation problem and briefly mention several relevant articles.

According to Stauss and Friege (1999), regaining lost customers is important for several reasons. First, it helps to secure future sales and profits. Second, the acquisition costs associated with replacing lost customers can be reduced. Lastly, the negative effect of word-of-mouth can be controlled.

There are two metrics commonly used to assess the payback from reactivation: (1) the ex-customer's first-year profitability, and (2) the ex-customer's LTV. Stauss and Friege (1999) use Second Lifetime Value (SLTV) – the lifetime value of a recaptured customer – instead of LTV of the terminated relationship for calculating the value of a regained customer.

Griffin and Lowenstein (2001) build on the framework laid out by Stauss and Friege (1999). They propose a general outline for reacquiring lost customers. According to them, firms should segment lapsed customers on the basis of SLTV, and identify the reason why the customers defected.

Griffin and Lowenstein, using data from publishing firms, show that the SLTV decreases with the duration of a customer's lapse. In addition the way in which the customer was acquired also affects the SLTV. For example, the value of the customer is higher if the customer was recruited through a subscription card for that particular publication as opposed to a secondary subscription source that handles multiple subscriptions.

In a related study, Thomas et al. (2004a) looked at how firms should price customers when reacquiring them and how they should price them when they have been reacquired. The most interesting result shows that pricing affects reactivation tenure. An interesting and somewhat surprising result is that the higher the reactivation price, the longer the tenure of the reactivated customer. This finding can be explained by heterogeneity in price sensitivity with the more price sensitive customers being more likely to change their buying behavior due to price. They also hypothesized, contrary to Griffin and Lowenstein's result described above, that the longer lapse durations are positively related to second durations. Thomas et al. found, however, that the relationship was directionally negative although not statistically significant.

The area of customer reactivation strategies and SLTV is an under researched area. Because of addressability and knowledge of historical behavior, firms should be able to identify segments of lapsed customers who have high profit potential (SLTV). In order to determine lapsed customers profit potential, an understanding and analysis of the reliability of SLTV is needed. To date, no one has conducted such analyses.

7.3 Using SMC's Model to Value Customers

In Chapter 5 we discussed a model developed by Schmittlein et al. (1987) which we call the SMC model. The model provided estimates of how many active customers the firm has, the likelihood that each customer is "alive" and what expected number of purchases the customer will make over a given future time interval. For businesses in which there is no contractual relationship, using models provided by SMC or Fader et al. (2005) is essential for determining the customer's lifetime value. We will discuss a paper by Schmittlein and Peterson (1994) in which they applied the SMC model to an industrial firm to answer some of the questions discussed above. They also provided some empirical validation of the SMC model. We will not repeat the equations provide in Chapter 5 but will discuss the application and how SMC's model was used in an industrial setting.

Schmittlein and Peterson's (SP) application relates to an office products firm. The primary data SP obtained were order data, initial date on the file and amount spent per order. Because customers had no contractual obligation, determining when a customer was active was one of the primary modeling requirements.

SP show how SMC can be applied at the individual level. The key quantities needed are: (1) the probability the customer is alive within any time interval, (2) the number of orders to be made and (3) the average order quantity for the customer. In Table 4 of SP (1994, p. 57), which we summarize in Table 7.1,

Table 7.1 5-year predictions for illustrative customer accounts (From Schmittlein and Peterson 1994)

Customer #	**Entry date**	**Most recent order**	**# of reorders observed**	**P(active)**	**Expected 5-year # reorders**	**Observed average reorder size ($)**	**Expected 5-year $ volume**
1	Pre 3/86	4/89	15	0.99	21.0	70.81	1,521.88
2	Pre 3/86	6/89	2	0.18	0.7	42.50	45.98
3	6/86	2/89	1	0.99	2.8	210.50	424.32
4	8/86	5/87	2	0.38	1.7	59.00	125.48
5	10/86	11/88	6	0.95	10.3	113.00	1,148.68
6	12/86	None	0	0.15	0.3	97.00	24.61
7	12/87	1/89	4	0.91	11.1	54.25	703.49
8	3/87	None	0	0.16	0.3	124.00	29.71
9	5/87	6/87	1	0.05	0.2	79.50	15.52
10	Pre 3/86	10/86	2	0.91	3.9	86.00	350.33

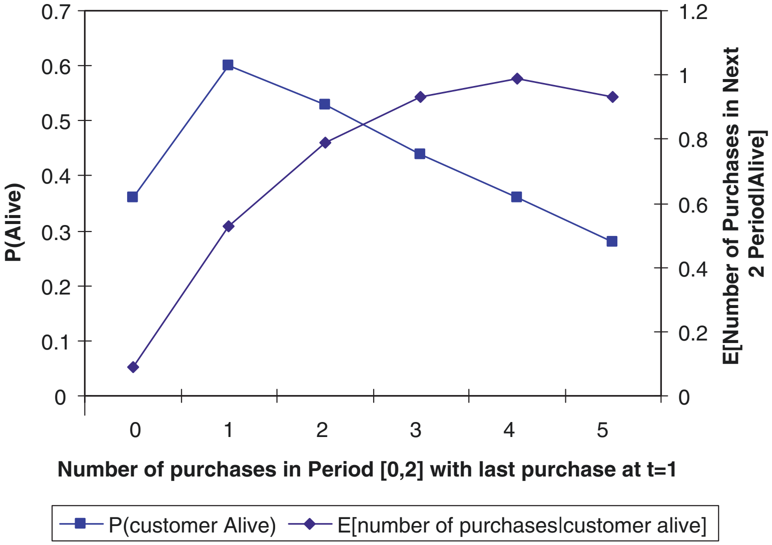

Fig. 7.1 Likelihood the customer is still alive and expected number of future purchases, as a function of number of recent purchases.
Assumptions:
$R = 0.415$
$\alpha = 0.415$
$s = 2$
$\beta = 4$

SP show the computations for a future 5-year period for ten customers. For each customer, they compute the probability the customer is active, the expected number of orders, the average order quantity and the expected dollar volume. Because of the nature of their model, they did not net present value the purchases. What is interesting in Table 7.1 is that customers have either a very high or very low probability of being alive. Only one customer has a probability of 0.38 while the other customer either have a probability of being alive of 0.9 or greater or 0.2 or less. This is an interesting result but may be caused by the choice of the ten customers for whom they displayed the data.

Another interesting insight to be gleaned from the model is the trade-off between the customer remaining active and the expected number of purchases in an upcoming period. Figure 7.1 graphs P(Alive) and E(number of purchases in next 2 periods) as a function of x, the number of purchases between time 0 and $T = 2$, with the last transaction at $t = 1$. Figure 7.1 shows that as the number of purchases increases, the expected number of purchases, *given the customer is alive*, increases. This is because if a customer has purchased a large number of times in [0,2], he or she probably has a high purchase rate, λ, even if the last purchase occurred in period 1. However, if the customer made a large number of purchases in [0,1] and did not purchase in period 2, there is a decent chance that customer has attrited. As a result P(active) follows an inverted U shape as the number of purchases increase. If the number of purchases is small and there have been no purchases in a while ($x = 0$), there is a good chance that customer is no longer alive. If the number of purchases is large but there have not been any purchases in a while, there is also a good chance the customer is no longer alive.

This has important ramifications for customer profitability analysis. The Schmittlein, Columbo, and Morrison method allows one to calculate the probability the customer is still alive, the expected remaining lifetime, and the expected number of purchases in a given time horizon, for each customer. The ideal customer will be one who is likely to be alive, has a long expected remaining lifetime, and has a high purchase rate. Figure 7.1 illustrates that there may be trade-offs among at least two of these quantities.

The other interesting aspect of their application is that they are one of the few authors who modeled the purchase dollar amount per order. Most applications merely assume the dollar amount per order is either the average for the population or for the individual. They used a weighted average estimate for the purchase dollar amount per order based on the information obtained from individual customers and the overall population average amount per order. This is a form of a shrinkage estimator and allows them to develop an individual dollar amount per order.

SP conduct a number of validation tests of the model. They show for example that the model predicts very accurately the number of customers who will make 0, 1, 2, etc., purchases over some future length of time. An important phenomenon the model attempts to capture, however, is the notion of the customer being "alive." To validate this, the authors classified each customer in their database into categories [0–10%], [10–20%], etc. Then they contacted 40 customers in each of these "deciles," and asked them whether they could report "active status." Figure 7.2 shows a clear monotonically increasing relationship between the predicted classification and the percentage who self-reported being alive. There may be some regression toward the mean, e.g., only 55% of those in the [90–100%] group reported being active, while 22% in the [0–10%] group reported being active. However, this could be a function of the way the "active" question was asked. In any case, the strong monotonically increasing relationship is clear and impressive.

SP's study is an excellent application of SCM's model. It shows how the model can be used to estimate individual LTV's and it also provides some model validation results. The other interesting element of their model is that it is applied in a business-to-business setting. Many of the applications of LTV models are better suited for industrial or business-to-business settings because the available databases are much better than for consumer product firms selling through channels of distribution.

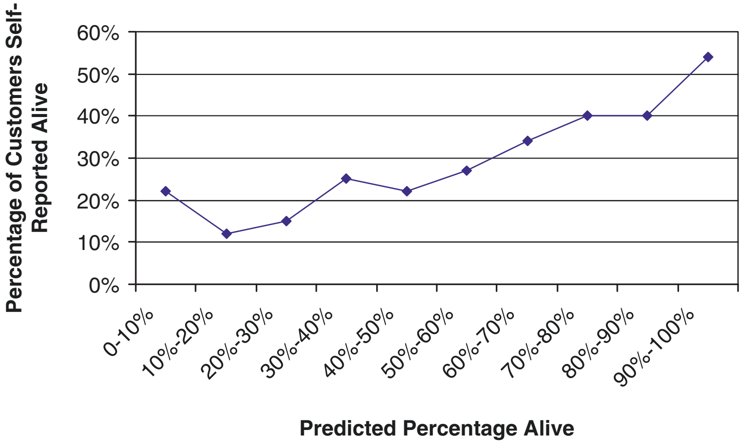

Fig. 7.2 Validation of P[alive] predictions (From Schmittlein and Peterson 1994).

7.4 A Case Example of Applying LTV Modeling

Van Raaij et al. (2003) report the first experience of a business-to-business company with incorporating customer value analysis into their marketing planning. The company, which we will call "DBM," was a multinational firm in the market for professional cleaning products. It sold directly to end-users such as in-flight caterers and professional cleaning services, as well as through distributors. It divided its market into sectors such as healthcare, lodging, or dairy. Sales and profits had been leveling off after years of growth and DBM was worried about new competitors. Further non-product costs (e.g., costs to service customers) had been increasing. The company desired to assign these costs to individual customers and calculate customer profit.

DBM undertook a six-stage process to calculate profit at the customer level and then develop strategies based on the results:

1. Select active customers
2. Design the customer profitability calculation model
3. Calculate customer profit
4. Interpret the results
5. Develop strategies
6. Establish an infrastructure for future applications.

Selecting active customers: The selection of active customers as a first step is quite interesting in view of the models we have reviewed for calculating whether in fact a customer is alive. DBM's approach was quite pragmatic: a customer was active if it had made at least one purchase in the period under consideration. Another issue facing DBM was whether to define the customer as end-users or distributors, or both. DBM decided on end-users because it wanted to develop its marketing efforts from the point of view of the end-user.

This made life difficult because sometimes DBM had to gather information on end-users through distributors, but DBM made the effort and excluded the revenues generated by distributors with whom it could gain agreement for them to supply information to DBM.

Designing the profitability model: The design of the profitability model centered on the assignment of costs. For this DBM used activity-based costing (ABC) (Chapter 6). DBM devised the following list of cost activities (the cost "pool") and cost drivers:

Cost activity	**Cost driver**
Logistics	Costs charged by logistics partner
Order processing	Number of orders placed by customer
Technical service	Service hours spent by mechanics at customer
Customer consultants	Consultant hours spent at customer
Equipment	Cost of equipment placed at customer

Much of the expense involved with each activity was labor costs, which prior to utilizing Activity Based Costing had simply been designated as overhead. Now with the ABC model, DBM could assign labor costs to the particular activities demanded by each customer, and hence calculate customer-specific costs. There were still overhead costs such as product development and non-activity specific sales and marketing costs, and these were allocated to each customer proportional to the customers' gross sales. These should not have been added to the value of the customer since they will distort future marketing costs (Chapter 6). Firm-level overhead such as office housing, etc., were not included in the cost calculations.

Calculating profitability: The calculation of profitability required data on revenues as well as costs, and these were assembled from various data sources within the company.

Interpreting the results: The most obvious initial finding was the great disparity in profitability across customers. While we often talk about the 80/20 rule, where 20% of customers accounts for 80% of revenues or profits, in this case, DBM found that 20% of customers accounted for 95% of profits. The customer "pyramid" is shown in Table 7.2.

Perhaps most interestingly, the top customers provided smaller per customer margins than the large customers. DBM reasoned that this was because

Table 7.2 Customer "pyramid" for DBM example

Customer type	% of customers	% of revenues	% of profits
Top	1	50	49
Large	4	23	25
Medium	15	20	21
Small	80	7	5

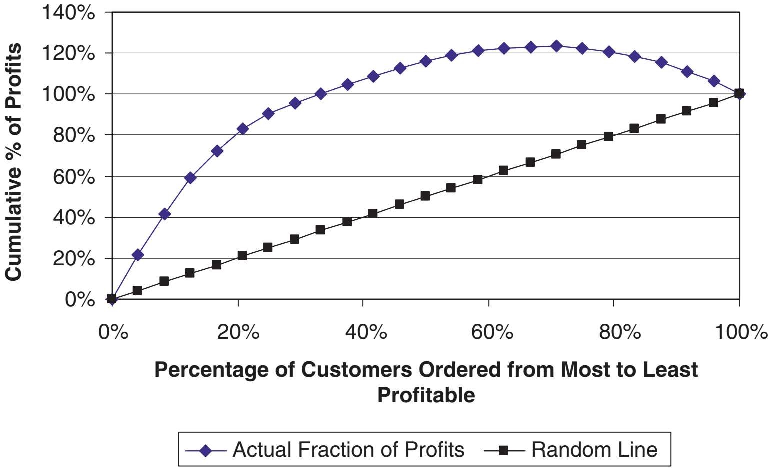

Fig. 7.3 "Stobachoff Curve" showing percentage of profits accounted for by customers ordered by profitability.

top customers had superior bargaining power and hence negotiated lower prices. They also required higher levels of support.

The company used the customer profit curve to plot what they called a "Stobachoff" curve (Storbacka 1997). This is simply the equivalent of a cumulative lift curve. It orders the customers according to profitability, and then plots the cumulative profit accounted for by these customers as one progresses from the highest to lowest profit. Figure 7.3 shows that in this example, 75% of the customers are profitable (the curve increases up to about that point) while 25% are unprofitable. Given that the top 75% of customers accounts for 120% of the profits, the remaining 25% really drag profits down. In this case there are a lot of profitable customers but they are subsidizing a relatively small number (at least a minority) of unprofitable customers. Note that by adding fixed costs through overhead, the firm may be distorting the true profitability of the remaining 25% of the customers. Some of these may be incrementally profitable.

DBM examined these Stobachoff curves for each of its 13 sectors. They classified each sector into one of four groups based on the number of profitable customers being relied on and how much they were subsidizing the unprofitable customers. The result is Fig. 7.4, which depicts four distinct cells. In the low dependence, low subsidizing cell, all customers are profitable and roughly equally so. In the low dependence, high subsidization cell, most customers are profitable but there are a few unprofitable customers who drag down total profits. In the high dependence, low subsidizing cell, there are only a few profitable customers and the rest of customers are unprofitable but not highly so. In perhaps the most dangerous case is the high dependence, high subsidization cell. In this case, there are a few highly profitable customers, and many highly unprofitable customers. This is dangerous because if those

	Low Subsidizing	**High subsidizing**
Low Dependence (Many profitable customers)	High % of profitable customers. Little or no subsidizing 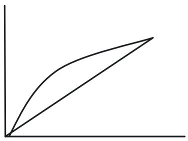	High % of profitable customers. Significant subsidizing of a small number of customers. 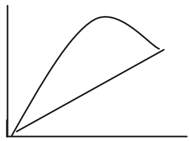
High Dependence (Few Profitable Customers)	Low number of profitable customers. Unprofitable customers are not highly unprofitable. 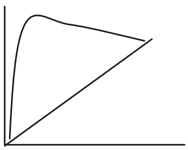	Low number of profitable customers with some highly unprofitable customers. 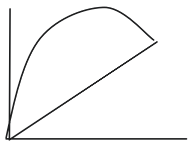

Fig. 7.4 Classification of Storbachoff curves.

few highly profitable customers should defect, the company would suddenly be losing a lot of money.

Develop Strategies: DBM analyzed its Stobachoff curves for each of its 13 sectors. Nine of them showed little subsidization. However, two sectors were in the low dependence, high subsidization cells, suggesting that the few customers dragging down profits should be dealt with individually. The remaining two sectors were in the high dependence, high subsidization cells, meaning the profitable customers had to be nurtured while the others either had to be "fired" or at least made marginally profitable.

In addition to the Stobachoff curve analysis, sector managers were provided with detailed profitability calculations for each customer. This allowed managers to focus on individual customers who were either extremely profitable or unprofitable. Often the unprofitable customers were consuming an inordinate amount of a particular cost activity, e.g., technical services. The manager could work with the customer to reduce these costs.

In summary, the customer profitability analysis (CPA) provided both sector-level and specific customer strategies that vastly improved the way DBM managed its customer base.

Establish infrastructure: DBM decided to perform the CPA every 6 months. The first major hurdle for this company was making sure it could assemble the cost information on a customer basis. Doing so enabled them to calculate individual customer profitability and went a long way in developing its customer strategy. DBM managers realized however that in the long-term, they needed to integrate marketing efforts and incorporate lifetime value and customer potential explicitly in the system.

7.5 Segmentation Methods Using Variants of LTV

7.5.1 Customer Pyramids

Probably the most common analysis of customer profitability data is to rank customers in terms of profit and classify them in a customer "pyramid" (Zeithaml et al. 2001). The pyramid shape occurs because usually a minority of customers account for the majority of profits. Figure 7.5 shows a hypothetical example of a customer pyramid. We saw a similar example in the DBM

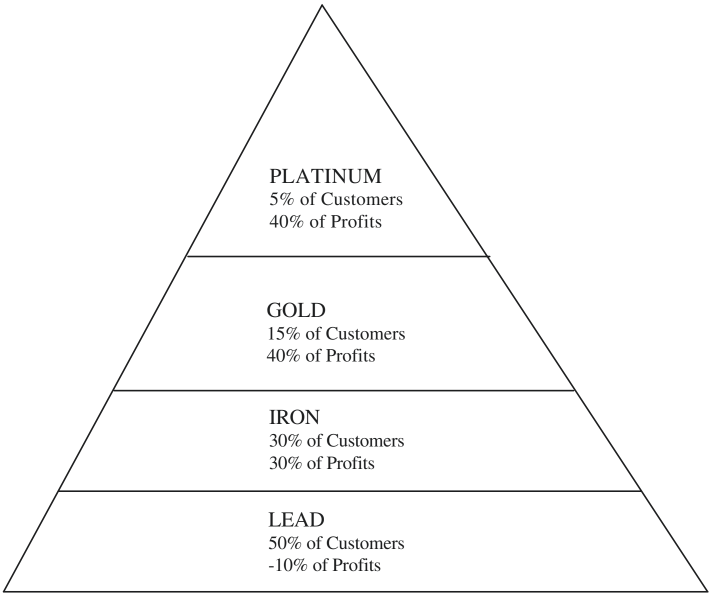

Fig. 7.5 The customer pyramid (From Zeithaml et al. 2001).

analysis above. Customer pyramids form the basis of customer tier programs, which are discussed in Chapter 23.

Figure 7.5 shows that the top 20% of customers account for 80% of profits, although this group can be sub-divided into Platinum and Gold segments, where the Platinum segment is only 5% of customers but accounts for 40% of profits, while the Gold segment is 15% of customers and accounts for 40% of profits. Clearly those in the Platinum segment are key customers. The Iron segment holds its own – it consists of 30% of customers and accounts for 30% of customers. The Lead segment, consisting of half the firm's customers, is unprofitable.

The customer pyramid is simply another way of displaying the Stobachoff curve used in the DBM example. However, the customer pyramid commonly displays the data in four segments and Zeithaml et al. discuss various ways to develop each of these segments. They discuss applications to the marketing research, real estate, and medical industries. Several aspects of creating and managing the customer pyramid are important.

- *Single-measure based*: The single measure can be profit, LTV, sales levels, or potentially even a combination of these. The advantage of a single measure is simplicity; the disadvantage of course is over-simplification. However, in actual applications, firms can look at several measures within each pyramid segment. For example, current period profit might be used to define the pyramid, but the firm could consider LTV as well as customer response to marketing when examining the customers within each pyramid.
- *Firing vs. ignoring vs. developing the customer*: A key decision in customer pyramid management is deciding when to develop a customer, when to leave a customer alone, and when to fire a customer. This corresponds to increasing marketing efforts, leaving marketing efforts as they are, and decreasing marketing efforts. The key questions are for example, "Can an 'iron-tier' customer be converted into gold, or should we avoid investing in this customer, make average profits, and invest in other customers?" These are difficult questions and depend of course on marketing response of the customer as well as competitive and firm resource issues.
- *The pyramid is only as good as the extent to which it differentiates customers*: Zeithaml et al. (2001) emphasize that while by definition the customer pyramid differentiates customers in terms of profitability, they should be differentiated in other key respects as well. For example, customers in different tiers should differ in their preferences for service levels and their willingness to pay for these levels. In addition, customers in each segment should be accessible, i.e., "addressable" in CRM terminology. The customer's potential LTV should be high in order to invest in converting at least some of these customers into higher tiers. Finally the firm should have the resources to be able to invest in potentially high-valued customers. Under these circumstances, customer pyramids are a valuable byproduct of customer value assessment.

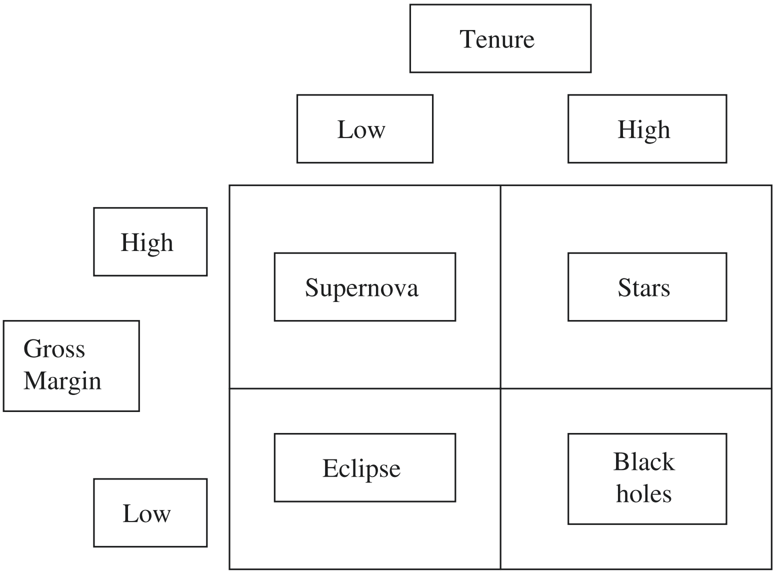

Fig. 7.6 Customer profitability mapping matrix (From Ang and Taylor 2005).

7.5.2 Creating Customer Portfolios Using LTV Measures

Ang and Taylor (2005) provide an interesting application of LTV measures to create a portfolio matrix of customers. They begin with two key measures that affect LTV, tenure and gross margin. Using these metrics they divide customers into four quadrants based on these metrics: low and high tenure and low and high margin. They then name these quadrants. Figure 7.6 below shows the matrix with the names used by Ang and Taylor.

The interesting element of their paper is that they then compute the size and customer profitability for each cell of the matrix. Figure 7.7 shows the results. Ang and Taylor show that 10% out of 14% (5/7) of the customers with low tenure and low profits are unprofitable and 20% out of 42% of the customer with high tenure but low profitability are unprofitable. The firm then has to decide what actions to take.

Ang and Taylor recommend certain actions for each cell. For low-tenure, high-margin customers encourage contracts by offering lower-priced service for entering a 12-month contract. For high-tenure, high-margin customers maintain high level of service and provide avenues for advocacy. For high-tenure, low-margin customers advertise benefits of high priced plans that come from additional features. Finally, for low-tenured, low-margin customers increase prices and reduce service.

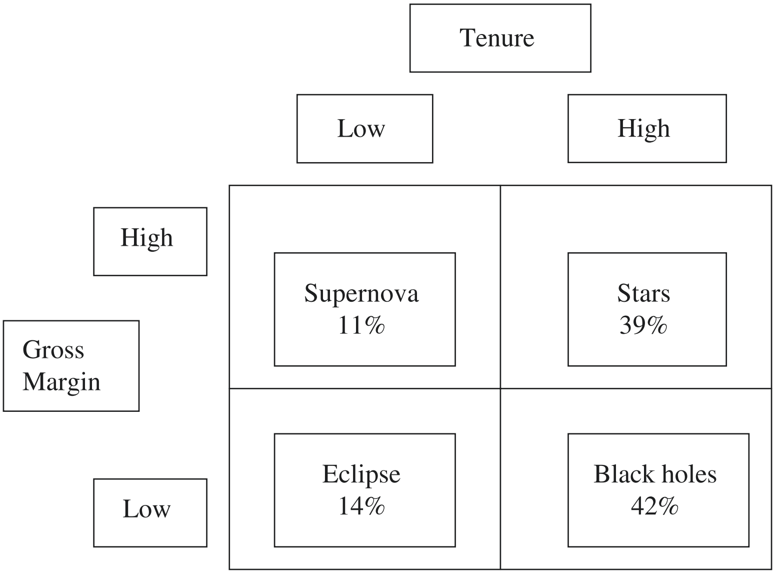

Fig. 7.7 Customer profitability mapping matrix – cell sizes (From Ang and Taylor 2005).

Applying the strategies above, Ang and Taylor described how 18% of the low-tenured, low margin customers were ultimately shifted to high margin customers. Thus, a combination of classifying customers and then applying different strategies, led to higher profits and longer tenure for some customers.

7.6 Drivers of the Components of LTV

In a two papers, Thomas et al. (2004b) and Reinart et al. (2005) discuss the drivers of two components of LTV – lifetime duration and cumulative profits – and the drivers of customer acquisition. Their study used a series of three linked models for duration, profits, and acquisition. Then by applying their model, they showed that many firms are either under or over investing in acquisition or retention marketing.

From the results provide in their papers, we have identified some of the key factors that determine acquisition rates. These are: (1) acquisition expenditure level, (2) demographics and (3) size of market. For duration, the key variables are: (1) retention expenditures, (2) customer usage rates, (3) cross-buying (how many categories purchased), and (4) share of wallet. For profitability, the key variables are: (1) acquisition and retention dollars (2) customer usage rates, (3) lifetime duration, (4) cross-buying, and (5) share of wallet.

Table 7.3 Potential gains from changing market spending (From Thomas et al. 2004b)

Company	How much more or less should be spent	How much profits would increase
B2B	−68.30%	41.52%
Pharmaceutical	31.40%	35.80%
Catalog retailer	−30.70%	28.90%

None of the results for acquisition rates is surprising but their method can be applied to any company's database to make a determination of what factors determine higher or lower acquisition rates. For lifetime duration and profitability, there are some interesting findings. Breadth of purchase and share of wallet are both intuitively appealing but not well-known in the real or academic world.

Thomas et al. (2004b) use their estimated equations for relationship duration, profit contribution, and acquisition likelihood to compute optimal marketing expenditures. They found that firms' spending is very far from "optimal". In Table 7.3 above, we show their key results. It shows that the lowest improvement is 28.9% in increased profitability. The conclusion we draw from their findings is that using models of key LTV components is likely to help firms improve their spending levels. No generalization can be made regarding whether firms over or under spend. However, if Thomas, Reinartz and Kumar's results are replicated across other firms, it appears there is a significant opportunity to increase profits through models of the type proposed in Thomas et al.

7.7 Forcasting Potential LTV

Employing traditional econometric models, Kim et al. (1999) provide a method of forecasting customer-level potential lifetime value for business customers in a telecommunication company. We formally define the potential lifetime value of customer i as:

$$PLV_i = \sum_{t=1}^{\infty} (R_{it} - C_{it})/(1+\gamma)^{t-1} \tag{7.1}$$

where PLV_i is the potential lifetime value of customer i, R_{it} is the revenue generated from customer i at time t, C_{it} is the cost or expenses incurred for customer i at time t, and γ is the discount rate. Depending on the firm's data collection interval, the time interval t above can be monthly, quarterly, or yearly. Similarly, the discount rate γ depends on the data interval t.

The firm's objective is to maximize the potential lifetime values of its customers. The firm should find the optimal C_{it} to maximize PLV_i for each time interval t. Hence, we need to solve a complex optimization problem since the revenue R_{it} is clearly a function of past and current costs/marketing expenditures (C_{it}). We need to specify the functional form of the revenue response curve with respect to the current and previous costs and estimate its parameters with appropriate data. Moreover, the costs consisted of several components such as cost of goods sold, service, and various marketing costs. Different allocation of total cost into these components may change PLV_i. Hence, we make an simplifying assumption that the firm incurs costs proportional to the corresponding revenue, $C_{it} = \beta R_{it}$. We also assume that β is constant over customer and/or time. With this proportionality assumption, Equation 7.1 becomes

$$PLV_i = \sum_{t=1}^{\infty} \frac{(1-\beta)}{(1+\gamma)^{t-1}} R_{it} \tag{7.2}$$

The problem of finding the potential lifetime value of customer i becomes the problem of forecasting his/her future revenue streams. The proportionality assumption allows us to easily calculate the PLV_i from forecasted future revenues. The appropriate specification of forecasting models (for future revenue streams) depends on several factors including data availability, industry characteristics, forecasting horizon, costs, ease of application, and so on Makridakis et al. (1983). So their forecasting model is somewhat customized for forecasting revenues of business customers in a telecommunications company. Compared to residential customers, business customers are better target markets for one-to-one marketing because their average revenues are large and their revenue distribution across business customers is highly skewed. In their study, the top 3% of its business customers account for 60% of total business revenues.

The revenue from customer i at time $t(R_{it})$ are decomposed into two parts and each component is estimated separately. That is, $R_{it} = Q_{it} \cdot CS_{it}$ where Q_{it} is the telecommunication demand of customer i at time t and CS_{it} is the firm's market/customer share. The econometric model for total telecommunication demand includes several independent variables such as the size of business, past telecommunication expenditures and growth rates. On the other hand, market share prediction is mainly based on a customer survey. Upon predicting future revenue streams for each customer, potential lifetime values are derived using the Equation 7.2. They also calculated the realized lifetime values for the corresponding customers that practitioners often use as its proxy. Figure 7.8 shows the relationship between the realized and the potential lifetime values. The correlation is about 0.4 and is statistically significant at $p = 0.05$. However, there are a number of customers whose potential lifetime values are fairly large but their realized lifetime values are small. They may be customers with high growth potentials. There are also some customers who have large realized lifetime values with relatively small potential lifetime values.

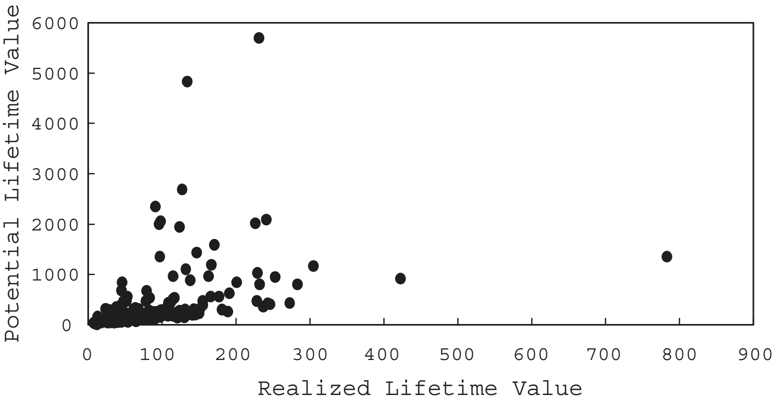

Fig. 7.8 Realized lifetime value versus potential lifetime value (From Kim et al. 1999).

7.8 Valuing a Firm's Customer Base

As mentioned at the beginning of this chapter, one of the important applications of customer valuation is to value the firm's customer base. This value should relate to the firm's stock market value, and hence be of significant interest to senior management. In addition, it helps other firms who might be interested in acquiring the company determine what would be a reasonable price.

Gupta et al. (2004a) applied a simple retention model of lifetime value to assess the value of five firms' customer bases. The analysis is conducted at the customer cohort level, where cohort 0 is the current customer base, cohort 1 is the customers to be acquired next year, cohort 2 is the customers to be acquired the year after, etc. The lifetime value of cohort 0 is:

$$LTV_0 = n_0 \sum_{t=0}^{\infty} m_t \frac{r^t}{(1+\delta)^t} - n_0 c_0 \tag{7.3}$$

where

LTV_0 = Lifetime value of the current customer base.
n_0 = The current number of customers.
m_t = The margin contributed by the average customer in the tth period of their lifetime.
r = Retention rate.
c_0 = Acquisition cost per customer among current customers.

Equation 7.3 is the standard simple retention model with acquisition costs subtracted to produce the net profit contribution of the cohort.

The authors recognize that an important contributor to the long-term value of the company (which supposedly is what the stock market and any acquirer should take into account) is the future set of customers to be acquired by the company.

In general, we have:

$$LTV_k = \frac{n_k}{(1+\delta)^k} \sum_{t=k}^{\infty} m_{t-k} \frac{r^{t-k}}{(1+\delta)^{t-k}} - \frac{n_k c_k}{(1+\delta)^k} \tag{7.4}$$

where LTV_k is the net present value (in $t = 0$) of the lifetime value of the customer's to be acquired k periods from now. Summing Equation 7.4 over all cohorts (k $= 0, 1, \ldots, \infty$) yields the total value of the customer database:

$$\mathit{CustomerValue} = \sum_{k=0}^{\infty} \left\{ \frac{n_k}{(1+\delta)^k} \sum_{t=k}^{\infty} m_{t-k} \frac{r^{t-k}}{(1+\delta)^{t-k}} - \frac{n_k c_k}{(1+\delta)^k} \right\} \tag{7.5}$$

The authors devote considerable time to estimating the key components of their model. To estimate n_k, the number of customers to be acquired k periods from the present, they obtain quarterly data on the size of each firm's customer base and estimate a diffusion-like growth model to project n_k into the future. This of course assumes the pattern of acquiring customers can be projected into the future, but the model is estimated on quarterly data over 5 years, fits well, and incorporates the notion of a peak and subsequent decline, so is realistic. To estimate contribution margin m, the authors use company annual reports and divide by the number of customers. We should note that they are using an average cost, not an incremental cost which will be important in businesses that have high fixed costs such as retailers or distribution businesses. They are able to calculate this number for the cohorts in their historical data, and find them to be fairly stable from cohort to cohort. In addition, they assume that the profit contribution for a given cohort does not change over time, given of course that the customer is still a customer.

The authors estimate acquisition costs by dividing the number of acquired customers by marketing costs. This is a fairly strong assumption, but four of the five firms they investigated were relatively new at the time of their study, so the assumption amounts to the majority of marketing expenditures going toward customer acquisition at least in the early periods. This assumption will break down in the future, but by that time the discount factor would make this not too bad an assumption. The authors estimate retention rate by consulting industry experts and other published data. The retention rates ranged from 70% for Amazon to 95% for Ameritrade. Finally, the authors used a discount rate of 12%.

The key results are shown in Fig. 7.9. The figure shows that customer valuation calculated via Equation 7.5 matches the stock's market value (share price times number of shares) very well for Ameritrade, Capital One, and E* Trade. Interestingly, the customer valuation strongly under-estimates the market value of Amazon.com and eBay.

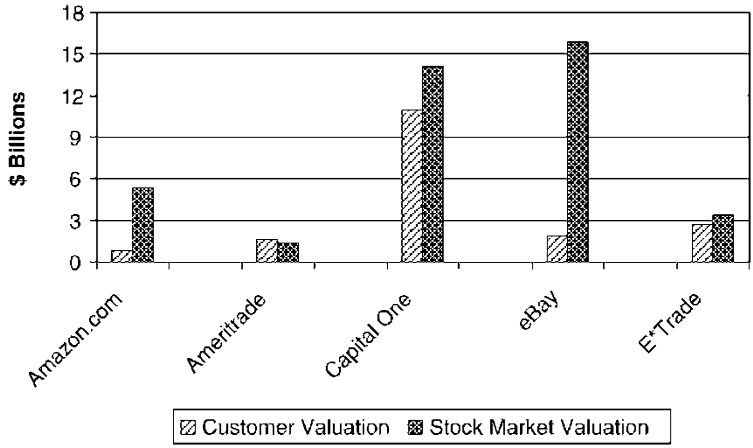

Fig. 7.9 Customer base valuation versus stock market value (From Gupta et al. 2004a).

The authors offer potential explanations, e.g., they cite reports that claim the market might be over-valuing eBay. Also, they discuss that the unique nature of eBay as a representative of buyers and sellers, makes it difficult to calculate the number of "customers."

The authors also raise the possibility that the model is not capturing all the phenomena the market is using the value the customer and hence the firm. For example, the simple retention model does not capture the value of word-of-mouth, which might be significant for Amazon and eBay. The market might also be thinking of extending the Amazon and eBay brand names into other industries. For example, Amazon started as a purveyor of books, and now sells DVD's, electronic equipment, etc. eBay started as a US company, but is now seeking to extend its reach to China. These issues would be factored into the lifetime value calculation by including add-on selling, or allowing margin contributions to increase over time, whereas Gupta et al. assumed they were constant.

There are technical issues in lifetime value calculation that might influence the accuracy of the simple lifetime value formula. First, if the customer base is highly heterogeneous in retention rates, using an average retention rates can underestimate lifetime value. Second, the authors are using a simple retention model on a quarterly basis and a migration model might be more appropriate, where customers could leave and then come back (Chapter 5). This might be accommodated indirectly by the model since the model calculates the number of additional customers each period, not distinguishing whether they are truly new to the firm or customers who are migrating back.

In any case, this is very promising and interesting work. It clearly demonstrates the potential to relate the value of a firm's *long-term* customer base with the market's assessment of the value of the firm. The three cases where the model predicts well show the potential of the approach; the two cases where the model does not predict well opens the door for future research.