Papers

Peter Dorrington

is principal initiative manager at SAS UK — part of the world's largest privately owned software company.

Jason Goodwin

is CRM initiative manager at SAS UK. Both have held posts as consultants in the field of practical CRM application.

The role of lifetime value in customer relationship management

Peter Dorrington and Jason Goodwin

Received (in revised form): 8 April 2002

Abstract

This paper will address a topic common among customer relationship management (CRM) practitioners — the use of lifetime value (LTV) models, especially for customer marketing segmentation. It will draw comparisons between LTV and other common marketing/segmentation metrics and how these are used. It will also discuss some of the challenges facing the analyst who wishes to use LTV and how these might be overcome. Finally, it will illustrate how LTV-based segmentation can be incorporated into a comprehensive model that maximises yield from a finite marketing budget.

Keywords: CRM, lifetime value, LTV, segmentation, profiling, targeting

Introduction

With the introduction of CRM as a discipline, rather than a technology, there has been a resurgence of interest in customer profiling and segmentation — the selection of customers/prospects based upon a defining rule or algorithm. That said, segmentation is nothing new, having existed almost as long as the concept of direct marketing, but it has been evolving.

- First phase — manual selection of criteria (eg postcode).
- Second phase — inclusion of socio-economic/geodemographic data.
- Third phase — data mining/propensity scoring.
- Fourth phase — LTV-based selection.
- Fifth phase (current) — yield-optimised, multi-channel/product segmentation.

Currently a great number of organisations data mine a wide range of data sources, building up profiles of customers/prospects who have an above-average propensity to purchase their products — for example, mortgage lenders actively seek people/families who are likely to move within the near to medium term and therefore require a mortgage. Within some sectors, especially financial services, this modelling process has been extended to try and take into account current and future profitability of clients (LTV) and find more like them — this is where the majority of applications of LTV indicators lie. Of late, a further refinement of the LTV approach has been a focus on 'value' or 'yield' — a topic which will be returned to at the end of this paper.

Peter Dorrington
SAS UK & Ireland
Whittington House
Henley Road
Medmenham
Marlow
Bucks S0L7 26B, UK
Tel: +44 (0)1628 404 204
Fax: +44 (0)1628 404 436
E-mail:
Peter.Dorrington@suk.sas.com

LTV

Definitions of 'value'

Simply put, LTV can be expressed as 'the value of the future revenues from a customer, based on their current and future product holding' — this can be compared to two other important measures.

LTV is a prediction

- CCV (current customer value) — the historic value of the relationship with the customer.
- NPV (net present value) — the current and future value of the relationship with the customer based on current product holding.
- LTV — the current and future value of the relationship with the customer based on current and future product holdings.

The term 'value' in this context refers to the difference between the future revenues generated (eg order price, service agreement premiums, etc) and the costs of generating and servicing these revenues (eg cost of marketing, apportionment of fixed/variable costs, etc). Both NPV and LTV calculations are also based on predictions of customer behaviour.

Therefore, if one were to plot a customer relationship along a timeline, one would end up with a product/value matrix similar to that shown in Figure 1.

In Figure 1, a grossly simplified case, it can be seen that a period of investment (eg marketing) resulted in a sale of the first product, generating revenue (and some cost of service) which can be calculated today as the CCV. If the customer were to retain this product alone, one could calculate a future value of the relationship — the NPV. However, cross-selling a second product to this customer, generating further

Figure 1: Product holding and LTV

revenues, contributes to a significantly higher LTV over the period under consideration.

In Figure 1, a tangible product was used as the basis the revenue stream — if a credit card were to be considered, with a successful campaign to up-sell from a ‘silver’ to a ‘gold’ card, the graph might bear a closer relationship to that shown in Figure 2.

In both of the previous examples it is apparent that an important element of calculating LTV is the ability to predict the future product holding and the ‘value’ of those holdings.

Calculating value

Simply put then, the LTV of any given customer can be expressed as:

$$\text{LTV} = \text{total revenues} - (\text{fixed costs} + \text{variable costs})$$

Revenue is fairly straightforward to measure — one simply adds up the total of all of the orders placed with the organisation. However, more recently there have been a growing number of attempts to model more accurately the *costs* associated with these revenues — practices such as ‘activity-based costing’ and ‘activity-based accounting’. These analyse the time and cost of each step within a complex process and calculate a ‘true cost’ for each sale. This approach is illustrated in Figure 3.

Therefore, in Figure 3, the cost of this particular process would be:

$$\text{Costs} = \pounds_1 + \pounds_2 + \pounds_3 + \pounds_4 + \pounds_5 \dots$$

Costs are more difficult to model than revenue

Figure 2: Product holding and LTV — Credit card example

Figure 3: Activity-based costing/accounting

Portfolio complexity is a major consideration

As one might expect, there are a number of caveats, one of which is that the calculation only holds true for the period under calculation, but a more significant consideration for organisations with a broad portfolio of products and client profiles is *complexity*.

Consider a retail bank — after a successful marketing campaign, it attracts a new client to open a current account. When considering the next product to market to this customer, it may select savings, loans, credit cards, mortgage, insurance and so on. After the second product sale, a third is to be considered and so on. The formula also needs to take into account such things as channel utilisation, fixed and variable cost apportionment and so on.

Risk and bad debt should also be considered

Additionally, in the financial services industry there is always the possibility of ‘bad debt’, where for example a loan or credit card balance remained unpaid. Ideally of course this situation would be avoided through rigorous credit scoring and risk analysis. In practice, however, any risk can only be estimated, and although measures are put in place to mitigate these risks, such as a higher interest rate or a smaller credit limit, bad debts still occur. The fact is that bad debt, or even a predicted bad debt, should be considered a cost when calculating LTV. If bad debts are factored, then take into account the cost of factoring to the calculation. The end result is that all associated costs and revenues are then considered.

Challenges in calculating LTV

Calculations

As mentioned earlier, a significant problem when calculating LTV is complexity — dealing with the bewildering array of choices and possibilities. Figure 4 illustrates one possible timeline.

In Figure 4 it can be seen that there are a number of stages, each with a possible outcome. In order just to obtain an accurate CCV, one would need to calculate what each of these stages costs and the associated revenues. For example, what is the value of each individual marketing campaign? Consider:

- set-up cost
- number of prospects approached
- initial success rate
- processing of ‘non-responders’
- final success rate.

What is the value associated with Product A? Consider:

- this and other campaigns
- revenue profile
- retention rate.

Are some costs transferred between products when one achieves a sale of Product B?

It is no small matter, then, to calculate a ‘true’ value profile, even when using only historic data. In many cases, estimates or ‘proxy’ values are used to represent the costs of stages in a marketing/sales/service process, but these are often little more than estimates based on outdated accounting methods and ‘gut feel’.

But for LTV one needs to take into account the future — and that means modelling and prediction, which come with their own set of problems.

**Proxy values for costs
are often little more
than estimates**

Figure 4: Product marketing/acquisition/retention timeline

There are no 100 per cent certainties in prediction

Greater accuracy comes from either better data or better techniques

Modelling and prediction

One of the most common uses of data-mining technology is to profile customer behaviour in an attempt to model the future. That is, if we understand what drives customer behaviour today, we can reflect this in our sales and marketing efforts — reinforcing the behaviours we desire and reducing the effects of those we do not.

It is undoubtedly true that data-mining-based segmentation of customers can give indications of the propensity to buy for any given market segment (a segment being a group of individuals related by a common set of attributes). However, there is no such thing as a ‘sure bet’ in statistics — only varying levels of probability. The specific challenge for any predictive technique is therefore to make the predictions as accurate as possible.

Consider a simple ‘yes or no’ question: for a given population, we would like to be able to predict how many (and who) will fit the criteria — perfect predictability for ‘yes’ and perfect predictability for ‘no’. A realistic statistical result would give a range of probabilities — from 100 per cent confidence of a ‘no’ answer to 100 per cent confidence of a ‘yes’ answer. If one plots this as a graph (Figure 5) one would see that the majority of the predicted answers are ‘uncertain’ — statistically, somewhere in the middle. If one can improve the accuracy of the prediction, more of the predicted answers will tend towards high confidence values (Figure 6).

Greater accuracy can be achieved in two ways: a better predictive algorithm, or making more *relevant* data available. In the first case we are trying to make better use of what we know, in the second, we are trying to know more.

There are a large number of algorithms available to data miners and these will not be discussed in this paper, as there is a tendency for specific algorithms to be useful only in a given set of circumstances; indeed, the choice of algorithm can be an important exercise in its own right.

Figure 5: Standard model

Figure 6: Increased accuracy model

**Behaviour of
individuals often
appears chaotic**

However, there is a commonly held belief that ‘if only we could get enough data, we would be able to predict the future near perfectly for any given customer’ — this is a fallacy. While large groups of customers may be predictable with moderate levels of confidence, the behaviour of individual customers appears more random (chaotic). Like all chaotic systems, it remains chaotic regardless of how closely observed. An analogy would be the weather — predictable over large areas for moderate time scales, but impossible to predict accurately if it will actually rain on your house tomorrow at 11 o’clock.

For most predictive models, the law of ‘diminishing returns’ applies. That is to say that simply adding more data fields (attributes) does not always guarantee a more accurate prediction. Much depends upon the kind of data collected and their relevance to the modelling process. Figure 7 shows that, initially, adding more relevant data fields dramatically increases the accuracy of the model, but this quite quickly levels out to a point where adding additional data has almost no effect on the accuracy (but almost certainly increases the cost).

**Not all data are
equal in value**

Not all data are equal — some are much more important than others. For example, a customer’s net income is almost always an important factor when considering their propensity to buy — if the prospective customer cannot afford the product, it does not matter how seductive the marketing, they are unlikely to purchase; couples without children are unlikely to become major/regular purchasers of children’s apparel, and so on.

Some of these data can be extracted from direct observation via the interactions we have with our clients, some can be purchased from third parties (eg socio-economic data) and yet more can be imputed. Some data are immutable (eg date of birth); other data may be subjective (eg use of a ‘consumer confidence index’). The right mix of data will change for every organisation, for every application and for every customer segment.

Remember, however, that the collection, storage and use of data can have their own issues, governed by the Data Protection Act. The safest

Figure 7: Effects on accuracy of increasing number of data fields

and most customer-centric approach to dealing with this is to follow the ethics of ‘permission-based marketing’, where the customer’s permission is sought with regards to the use of their data and they are given the opportunity to opt in to receive outbound communications.

The future remains uncertain

As a final point about modelling, concerning the effects of time, like all predictive models, the further into the future we try to model, the greater the margin for error (see Figure 8). Most LTV models are only reliable for the near-to-medium term and can be dramatically affected by unforeseen

Figure 8: Effect of time on model accuracy

circumstances (eg a downturn of the American economy, changes in interest base rates or the 'dot.bomb' effect on the stock market).

Finally, the effects of changing circumstances can dramatically affect the validity of calculations. For example, current low interest rates and a wide variety of competitors have had an enormous impact on the credit card industry. Earnings per pound lent are down and customers are changing cards more frequently as they 'rate chase', but the volume of lending is up (as consumer spending has risen with increased confidence).

If one were to re-examine the previous example of credit card revenues (Figure 9) one might see that more had to be spent on acquiring the customer, earnings were lower than predicted and the customer switched to another lender before break-even was achieved. In this example, the customer had not changed, but the surrounding circumstances had.

In short, calculating LTV is not an exact science and should not be treated as such.

**Calculating LTV is
not an exact science**

Using LTV calculations

Obviously, all organisations want to maximise earnings. Sales and marketing departments are therefore tasked with maximising LTV — that is to say selling more products to the same customers over the lifespan (lifetime) of their relationship with the organisation and finding more customers like these.

In a survey conducted by TBC Research for Accenture, the Chartered Institute of Marketing (CIM) and Siebel Systems,¹ 510 marketing professionals were asked about how and why their organisations were investing in new CRM technologies. Within the findings, two key responses catch one's attention — those shown in Figures 10 and 11.

The reasons these findings were so interesting is because the nature of the questions and responses confirmed that the primary *reason* for investing in new CRM/marketing technology was to increase company revenue, yet the primary marketing *measures* are about lead generation, campaign response rates and 'customer retention analysis'. This begs the

Figure 9: Credit cards — Effect of changing economic climate

Figure 10: Driving reason for new marketing technologies

Figure 11: Marketing effectiveness measures
Source: TBC Research

question, who is measuring the *value* generated by these activities? This phenomenon was specifically mentioned in the executive summary of the report — ‘There appears to be more emphasis on customer acquisition related activities (77%) than customer retention (59%).’

This and anecdotal evidence seem to indicate that LTV calculations are used not only to identify which existing customers are the most profitable, but also to ‘profile’ these customers so more like them can be sought. So ‘value’ may be used to:

- identify those customers who are most profitable and focus retention efforts on them
- find more customers who match the profile of the most profitable customers
- calculate which products or product combinations are contributing most to profit.

Is differentiated service always a good thing?

Using LTV to offer a differentiated service

In an attempt to attract and retain more ‘profitable’ customers, many organisations offer incentives to this segment — in other words, they give

a higher level of service to their good customers than their bad (unprofitable) customers. There is considerable danger in this approach.

- Are your good customers really that good?
- How confident are you about the profitability calculation?
- Is it the customer who is 'bad', or your ability to service them profitably?
- Is a 'non-profitable' customer always going to be so?
- If customers have an unsatisfactory experience of your organisation, will they influence others?

Of late, there have been numerous adverts from a range of organisations promising that they will *not* treat their customers differently. These adverts are not trying to attract unprofitable customers from their competitors, but those 'borderline' cases who could/may be profitable if handled in the right way.

Now many organisations will say that they are not trying to offload their less profitable customers (while secretly desiring to do so) — just that they will handle them using a 'lower-cost' channel, for example by placing a greater emphasis on the Internet rather than the branch network. However, a poor understanding of the true costs of a channel strategy may mean that this is actually more expensive, thereby increasing costs and further reducing profitability. Recently the Routes To Market Association² discussed how adding an additional channel to an established channel strategy actually increased costs — predominantly because of the increased processing complexity and partly because customers continued to use the other channels. Then again, addition of new channels may open up new markets, or drive additional revenue generation in the more established channels.

A poor understanding of channel strategy can increase costs

What next?

There is a fine line between being guided by LTV and being a slave to it. When embarking upon an exercise to calculate LTV we should consider the following:

- without action, measurement is just an academic exercise — we should have clearly understood reasons why we are performing these calculations
- LTV calculations need not be complex, but their inherent inaccuracies need to be acknowledged
- channel strategies should include measures which span channels
- an unthinking adherence to LTV indicators will inevitably lead to missed opportunities and mistakes — constantly monitor performance against predictions
- test your assumptions — where possible, remove proxy costs; where this is not possible, recognise margins for error
- be prepared to think again — an unanticipated change in circumstances may invalidate previous calculations

- do not assume you ‘know your customer’ — facts are worth far more than assumptions.

Yield value delivered today

The successor to LTV — Yield optimisation

Finally, it is no secret that, when calculating which products a customer may have a propensity to buy, our best customers will be candidates for a range of differing, sometimes competing products. Which then should be in the portfolio and in what order do we offer these to the customer?

A further refinement of the LTV approach is to apply priorities to marketing activity based upon the simple tenet ‘what delivers the most return *today* for the least investment?’. Consider the banking example again — how do we choose which product, channel and strategy to use for our next campaign. The answer is simple: whichever will yield the largest revenue for the least outlay, in other words the optimum margin. It is more profitable to promote a product which has a low revenue stream but also a low cost structure than one which gets more money but also costs a lot to sell and service. The only way to make this choice is to get accurate measures of costs across the board and use the very best predictive approaches, constantly tested, to calculate a dependable LTV.

Customer optimisation strategies based on the yield approach where marketing, customer acquisition, customer retention and cross/up-selling are prioritised based on the most profitable product/service combinations can, in theory at least, only lead to greater profitability. Fortunately this yield-optimised approach is now being implemented in a number of enterprise-class marketing suites, when product mix, message variety and channel are all being selected on the basis of higher confidence, cost-aware, short- to medium-term criteria.

References

1. TBC Research (2001) *Managing and Measuring Your Customers In The New Economy*, TBC Research, Richmond.
2. Routes To Market Association (2001) *Achieving Cost Savings Through Effective e-Business*, www.the-rtma.com.