CUSTOMER VALUE MANAGEMENT

Mark Stirling, Director, Paragon Business Solutions,

3 Greenock Rd, London, W3 8DU

Mark Stirling has over 15 years experience in the credit scoring industry as an IT analyst and software developer for Infolink and, in the last 10 years as a director of Paragon Business Solutions (a specialist company providing credit risk management solutions for financial institutions). His recent experiences have been focussed on decision engine-based systems.

Abstract

Both ‘customer management’ and ‘customer value’ are phrases in common usage in the consumer credit industry, yet their scope and potential is not always fully understood nor exploited. This paper aims to explain these concepts and show how they can become good commercial practice by using customer value as the core decision-making tool throughout the whole credit life cycle. Consistent decision-making is fundamental to maintaining and enhancing customer relationships, and an integrated methodology for achieving this is essential.

Introduction

In consumer credit, successfully managing customers involves a cradle-to-grave approach. At the outset is the marketing/sales function, which needs to be able to offer the customer the right product at the right time and at the right price. This applies equally to attracting new customers as it does to offering incentives to retain existing customers.

During the life of the relationship, many contacts will be made with the customer, initiated by either party for a variety of reasons. A proactive lender will maximise these opportunities by attempting to cross-sell other products - preferably the most appropriate product which fulfils the customer’s immediate needs.

Towards the end of the life cycle, the collections and recoveries people need to be able to apply the most appropriate strategies for recovering monies rather than looking after customers. Early in the collections cycle, the objective should be customer rehabilitation (with dignity) but there will come a time when that customer relationship has irretrievably broken down, and the aim is recoveries of debts in the most effective manner. Knowing the customer means being able to apply the most appropriate approach at every stage of all these processes.

These operational groups all need to provide a cohesive, non-contradictory approach to the customer. It is perhaps more important that the customer sees that the approach is co-ordinated and consistent and that the decisions handed out by each group are not conflicting in any way.

Historic situation

To put the current situation in perspective, it is worth looking briefly at how the industry has changed over the last 20 years.

In the early 1980s, there were four main banks, a plethora of smaller building societies, and two credit card issuers. Each of these had large numbers of customers, virtually all of whom exhibited a life-long, monogamous relationship. Customer loyalty was assumed and standard. Credit or risk management departments were the exception not the rule. Scorecards were primarily used for application assessment and only evaluated future risk. Operational systems were largely independent of each other. Applications processing systems were not linked to the accounts system and specialist systems for collections and recoveries were the exception.

On the other hand, consumers largely did not understand the concept of the ‘cost of money’. They were driven by how much they could afford in monthly repayments.

Current perspective

Today there is extensive competition (particularly in the credit card market) and building societies have merged to become banks of equal standing with the traditional banks. As a result customer loyalty has waned. Far more customers are moving their accounts when they perceive a better offer is available elsewhere and they are spreading their financial requirements amongst multiple lenders.

Risk management departments are the norm, and scorecards are used to predict every conceivable outcome. Operational systems are far more integrated providing a smoother transition between the different stages in the credit life cycle.

Changes are also occurring in the wider world. Consumers increasingly understand the cost of money. This is particularly evident in the growth and success of balance transfer offers from the credit card companies. Mergers and acquisitions between banks and building societies in particular, mean that these organisations have inherited customer databases that are incompatible with their own systems making customer management more difficult in the short term. The government is increasingly imposing constraints on our use of data. (The most recent example of this is the voters roll and any restrictions in this regard will have major adverse consequences across the industry.)

In parallel with these changes, there has been an exponential growth in computing and analytic power. Not only have PCs arrived but also they have become as powerful as mainframes were 20 years ago. New programming languages have provided new capabilities and both are sitting on our desktops under individual control. Neural networks and artificial intelligence have emerged from the confines of academia to become a much more widely used tool for predictive modelling.

All these changes mean that it is increasingly difficult to get to know customers in the first instance, and over time, to keep up with their changing needs. To achieve this a lender requires systems that can:

- apply consistent decisions,
- make multi-dimensional decisions,
- use multiple decision-making techniques,
- accommodate multiple data sources,
- interface with existing operational systems, and
- provide dynamic flexible monitoring.

Consistency of approach

Consistency (in the sense that the lender is congruent or compatible with the customer) means that when the customer has a particular need, the lender is able to fulfil this need, and preferably able to anticipate it before the customer asks for it. Consistency is vital because it engenders customer credibility and confidence, which in turn means customer loyalty. At a minimum it represents professional behaviour on the part of the lender.

Multi-dimensional decisions (or customer value)

In today's complex environment the most effective way to achieve this consistency is by the use of decision engines.

These are fully integrated systems that fulfil all the functional needs listed above. For any given situation, decision engines will take inputs from multiple sources (both internal and external) and assess all the associated data to produce a multi-dimensional decision, i.e. decisions that are based on more than one criterion. If for example, the required decision is whether or not to mail a target on a mailing list, the traditional approach would be to assess the likelihood of obtaining a response. Now however, this can be complemented by also assessing the chances of being accepted, of using the product and becoming a future profitable customer. The probabilities associated with each of these components can calculated (using statistical and/or artificial intelligence techniques or both) and summarised into a single figure of customer value. This value is then fed into the applicable strategy and the decision engine will recommend or take the most appropriate action as a result.

What is customer value management?

Customer value is sometimes thought of as net present value (NPV) and the ideas are similar. Customer value management can be defined as an 'integrated approach managing the customer relationship by modelling the value provided by that customer, by adopting a multi-dimensional approach to decision-making. This is achieved by adopting a true risk and reward balance with a real measure of the anticipated value of the applicant.' (It should be noted here that, although this paper is discussing *customer* management, this is only operating at account level, and assessing the worth of the customer *within* that account relationship only.)

In simple terms, this means making decisions on customers not on the basis of a single factor (which was traditionally risk) but on the multi-dimensional basis of risk, attrition, expected income and whatever else is appropriate. The path taken to arrive at that decision will vary for each applicant and will be the one that combines the most appropriate factors in each case.

This concept can be applied to decisions that need to be taken at any stage of the credit life cycle in full awareness of the total extent of the customer's relationship with the lender. For example, if a customer asks his branch for an increase in his overdraft limit, then the lender would be aware of his commitments elsewhere in the organisation, e.g. that he was overdue on his credit card or his mortgage was in arrears.

Customer value based strategies

Having used a decision engine to calculate a customer value, it is important to understand that different decisions and/or actions might arise compared with the traditional risk-based approaches. To illustrate the new possibilities, two examples based on credit card products will be considered. These are all simplified to demonstrate the principles whereas in reality many more criteria would be assessed in making each decision. (In the following graphs, the customer value is shown as a single-digit number – to illustrate its relative nature – on the vertical axis, and the horizontal axis is time, split into months.)

New applications

In a new application scenario that uses customer value, acceptance strategies have been defined for three distinct groups:

- Low-risk applicants who want the card, are likely to use it and become full payers, and can, but may not pay,
- Middle group who need the card, will definitely use it, take extended credit and are more likely to default, and
- High-risk applicants who need the card; will use it extensively but can't pay.

Using a decision engine each of these factors can be quantified, and aggregated up into a single customer value for each case. (This methodology is discussed later in this paper.) Figure 1 gives four possible strategies.

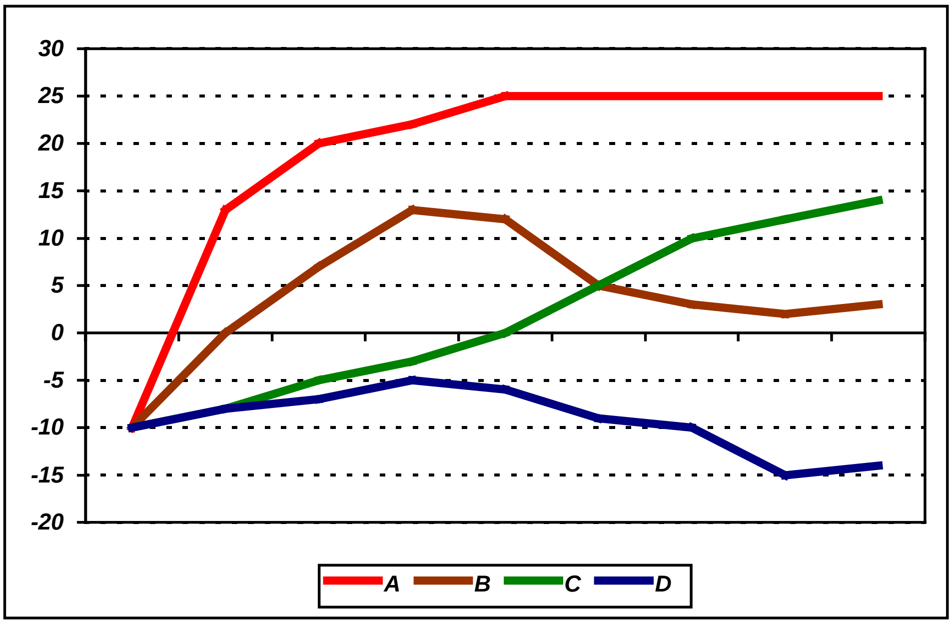

Figure 1: Acceptance strategies

The first point to note is that the curves all start at the same negative value. At this point they all represent *potential* customers and their value is directly linked to the costs of the recruitment process.

Customer A (top line) looks very promising in that he will become profitable very early on and remain profitable. This applicant can be accepted and given a high level of facilities. At the other end of the spectrum, Customer D (bottom line) is never going to become profitable, and he will be declined at the outset.

However, the two middle lines illustrate a paradox. The short-term profit of Customer B (second line) is just that, short-term only, and Customer C (third line) is the opposite - slow to grow into profitability but yet he will probably remain so thereafter. But the difficult question is how should these customers be treated? Firstly are they accepted or declined? For instance should the short-term gain be traded-off against for the longer-term loss, or vice-versa? Or should they be declined – or asked for further information on which to make a more informed decision? In practice, the answers to these questions will depend on the lender's policies.

In a traditional context when risk was the only factor under consideration, Customer A would probably have been given a gold card. However, now that multiple criteria are being assessed it is no longer valid to say that high profitability equals low risk, and a different strategy should be adopted.

Collections example

In this situation the required decision is to define the most appropriate action. There are three customer groups all of whom have missed the payment date on their credit card and are five days in arrears. The three criteria that are being considered when calculating the customer value are:

- is he likely to roll onto higher levels of delinquency
- is he a 'can't pay' or a 'won't pay', and
- is he an extended credit taker or a full repayer?

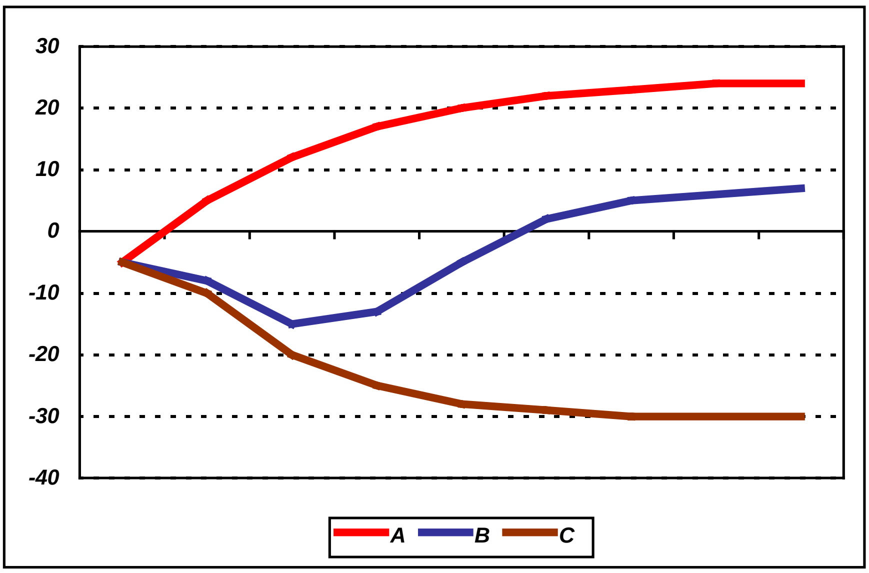

Figure 2: Collections strategies

The first point to note in Figure 2 is that the curves all start with a negative customer value. This is shown here as -5 for the sake of simplicity but in reality this would vary in each case depending on the past history.

Customer A (top line) is a previous good customer of long standing who has seldom defaulted. He is very unlikely to roll on, and he is both able and willing to pay. He is typical of the customer who has missed his payment simply because he is away on holidays; his future customer value quickly returns to the positive. As a result, no action is taken.

Customer B (middle line) is more problematic. He is likely to roll onto at least 60 days down, but he does have both the ability and willingness to pay, only not immediately. So his customer value drops off before it returns to the positive where it continues at a much lower level than customer A. This scenario might be typical of someone who has had a recent financial hiccup like redundancy. In this case the action might be a statement message, followed by a letter at 30 days.

Customer C (third line) is the worst case. He is likely to roll onto recoveries, he probably can pay but he has no intention of doing so. So his customer value drops to the lowest point possible where it flattens out as it runs along the bottom. In this instance it is probably prudent to phone the customer even at this early stage to ascertain the situation with the objective of minimising the potential losses.
Strategies based on this approach can be applied at any point that there is customer contact. Credit limit increases and cross-selling opportunities can be maximised for the good customers, whereas multiple collections and recoveries strategies can be defined for delinquent cases.

Calculating Customer Value

For the purpose of explaining the process of calculating customer value, this paper will only consider new applications for a credit card product. As we have already seen, it is possible to apply the process at any stage of the life cycle.

Credit risk will remain an integral part of the assessment as will initial recruitment costs and on-going administration costs. Whereas all of these factors have established methods of evaluation (scorecards, fixed cost allocation etc), a new factor of ‘time since account opened’ will be added. To explain in terms of credit risk, traditionally risk has only been assessed at a single, fixed point in the future, often 12 months ahead. However if risk is also considered at several future points, say, quarterly in the first 12 months then a more precise assessment can be achieved. Other items likely to be evaluated are early closure, type of activity (i.e. full payer or extended credit taker), levels of outstanding balances and potential dormancy.

Calculation

To consider all these elements in a multi-dimensional way is a complex process. The problem is how to bring together all the dimensions so they are comprehensible and useable. The approach followed here is to calculate NPV for each customer based on predictions of the NPV components. An apparently straightforward approach would be to perform NPV calculations using spreadsheets. However, the difficulties with spreadsheets are that they are cumbersome, and inherently contain embedded assumptions and estimates. This serves to make operational implementation more problematic. Taking a tranche of new applications at various ages, a typical NPV spreadsheet might look like the example in Table 1.

Age of tranche of apps	**1 Month**	**2 Months**	**3 Months**	**4 Months**
No new apps in tranche	800	800	800	Etc
% full payers	25%	25%	25%	
Balance	600	600	600	
% dormant	10%	10%	10%	
Balance	540	540	540	
Etc				

Table 1: Initial Spreadsheet

This approach does nothing to differentiate between cases because two major assumptions (the shaded lines) have been used. In this example, the lender knows that his book contains 25% full payers and 10% dormancy, but he doesn’t know is what those levels actually are in the early months.

If these figures could be calculated for each month, and these used instead, then a more realistic model will result. Using statistically derived scorecards, it is possible to build models to predict the levels of these activities. If the resulting estimates are then used in the NPV calculations, the resultant spreadsheet may resemble the example in Table 2.

Age of tranche of apps	**1 Month**	**2 Months**	**3 Months**	**4 Months**
No new apps in tranche	800	800	800	Etc
% full payers	**45%**	**40%**	**25%**	
Balance	440	480	600	
% dormant	**2%**	**5%**	**8%**	
Balance	431	456	552	
Etc				

Table 2: Refined spreadsheet

In Table 2 the percentage of full payers starts off at a high level but quickly reduces, whereas the dormancy shows a reverse trend. These 'actual' figures also serve to illustrate the extent of the corresponding generalisations used in Table 1. Specific figures would be repeated for every other item in the spreadsheet.

Processing Applications

In the above example differentiation has been applied on the basis of the age of a tranche of accounts. However, this approach does not permit any degree of deeper analysis and the ideal situation would assess the all the criteria for each applicant individually. To achieve this a decision engine is required. This is designed around multiple, interlinking processing modules (PM) each of which combines a single decision tree together with its complementary scorecard (see Figure 3).

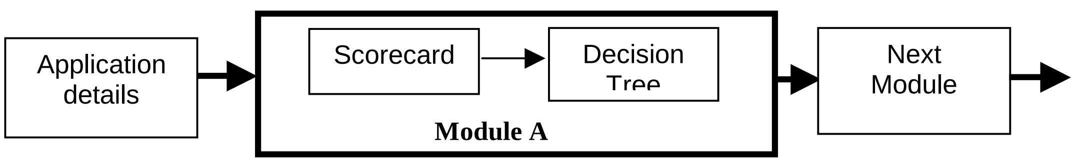

Figure 3: Processing Module for each variable

Any single module will use the scorecard to predict a driver-variable (drivers are variables which determine the direction and path of the process) and the decision tree will then use the outcome from the scorecard to determine which is the next most appropriate module for each individual application.

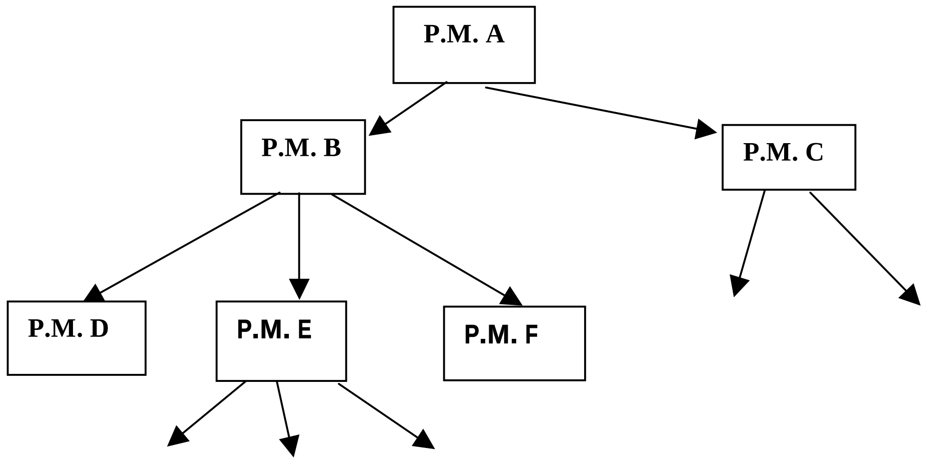

Figure 4: Decision Shrub

With conventional decision trees, the end points (or terminal nodes) generate decisions about specific actions (e.g. accept or decline and/or offer the appropriate credit limit). However, because this decision engine is a collection of modules, each with its own decision tree, the terminal node of each tree is not an action, but a decision about which next module to connect to, i.e. these trees have a conjoined module at each of their terminal nodes. The path taken through the first tree will determine the next most appropriate module to link to and work through, and this process will be repeated until a satisfactory decision has been arrived at. This interlinking network of modules has been styled a 'decision shrub'.

In this example the first driver-variable is the potential for the applicant to become a full payer as shown in the spreadsheet in the above tables. (In reality the first driver-variable for assessment will have been determined in the development process, but for illustrative purposes we are assuming it is still full payer potential.) The scorecard in the first module (shown as PM A in Figure 4) will model this outcome and the results will be fed into its complementary decision tree. The next step in the path (in this example, either PM B or PM C) will be determined by the specific result from the scorecard.

In the spreadsheet, the next variable for evaluation for every customer regardless was likely dormancy. However, using a decision shrub, PM A may determine that, for Customer X, residential stability is the next most critical variable to be assessed followed by employment status, but in the case of Customer Y, it may be that lifestyle followed by potential risk, for example, is more valuable. Thus the specific path for each customer will be different.

This example has been simplified to explain the underlying principles. In reality the decision engine would be many layers deep and many modules wide, thus creating a huge number of potential paths through the shrub. At the end of all these branches and layers, the terminal nodes will deliver a measure of the likely customer value of the individual applicant. This is then translated into the most appropriate action for each individual customer in the manner illustrated in Figure 4. An important point to note in this representation is that the processing modules are themselves decision making units which may contain scorecards, decision trees, etc.

Other data sources

Each operational function has traditionally used limited data sources that were considered appropriate for its needs. For instance, marketing would use lifestyle data, but it was unlikely that this would also be used to make decisions about credit limit increases or collections activities. By using a decision shrub approach, all possible data sources can be used to make any decision at any point. As well as enhancing the assessment of any particular situation, this approach also adds to the consistency of the decision-making process that was discussed earlier.

Monitoring

Introducing sophisticated decision engines and designing complex strategies based on customer value is a major undertaking and its effectiveness must be quantified and compared with the original objectives. Undertaking strategy testing before implementation is part of this process but the ongoing requirement is in-depth monitoring.

Traditionally monitoring largely assessed scorecard-based dynamics, but during recent years there has been a move towards the use of strategy-based assessment for both scorecards and portfolio performance. It is important to monitor strategies to:

- Ensure that any assumptions made during the development process remain valid when the strategies are applied to the live environment,
- Be certain that the outcomes predicted in test strategies actually materialize when the strategy is implemented,
- Assess the performance and trends of the developing portfolio and determine a need for any appropriate action.

Monitoring the strategies should result in the circular process of aiding the development of new and more effective strategies. But above all strategy monitoring will enable the organisation to increase its awareness of its customers relationships and needs.

Conclusion

It is now possible to generate an objective measure of individual customer value. By combining traditional scorecards with decision trees to create sophisticated and complex decision engines the ability to treat each customer on his own merits is realisable.

Today's lending environment is likely to become even more complex, aggressive and competitive in the foreseeable future. To remain a serious contender in the consumer credit industry, customer loyalty will become paramount and more sophisticated means of achieving this will become commonplace. Consistency of decision-making will need to be maintained, and above all, the customer will perceive that a personal, tailored relationship has been established and long-term customer loyalty will be restored.