

Paper 138-2012

Applying Analytics to High Performance Customer Profitability Models in Financial Services

Tony Adkins, SAS Institute Inc., Cary, NC

ABSTRACT

Where financial services companies have looked at customers in the past, much of the literature focuses on the “apex” customers. These are the top 10–50 customers by value, whose importance to the business means that they demand an extraordinary service, often at the expense of profitability.

Historically, it was difficult to build models of a scale that produce customer profitability, so models tended to stop at segment level. We have tended to rely on the traditional rather arbitrary groupings used within a business, and this can disguise important trends.

Bringing together analytics and high performance in memory profitability models allow us to discover hidden patterns in behaviors that cross arbitrary boundaries in the financial services industry and refine our interpretation of intermittent risks.

INTRODUCTION

There is a trend for customers to increasingly view suppliers' products and standard service lines as commodities. As a result, what customers now seek from suppliers are special services, ideas, innovation, and thought leadership. Many suppliers have actively shifted their sales and marketing functions from product-centric to customer-centric, through the use of *data mining and business intelligence* tools¹ to understand their customers' behavior – their preferences, purchasing habits, and customer affinity groups. In some companies, the accounting function has supported this shift by reporting customer profitability information (including product gross profit margins) using activity-based costing (ABC) principles. However, is this enough?

THE EARLY DAYS – ACTIVITY-BASED MANAGEMENT

In the early days, activity-based management (ABM) was very much focused on process improvement, and could be seen as part of the whole BPR/six Sigma/TQM movement. After 25 years of these programs, it is probably fair to assume that most companies have reasonably efficient processes. While there might yet be gains to be made in this space, they are unlikely to be dramatic.

Yet ABM literature still largely focuses on the internal efficiency of processes as a whole. It has yet to really address how processes relate to individual customers, and how their varying application impacts profitability. Yet this is where we are now more likely to find the opportunity for dramatic gains in profitability of the firm.

Where we have looked at customers in the past, much of the literature focuses on the “Apex” customers. These are the top 10-50 by value, whose perceived import to the business means that they can frequently demand and receive extraordinary service, often at the expense of profitability. But these few are, by their very definition, atypical. An analysis of their situation does not provide any useful insight in to the factors that impact the profitability of the remaining 99.9% of customers.

Historically, it was very difficult to build models of a scale that could produce individual customer profitability models, so models tended to stop at segment level (for example, all customers from a given standard industry code, geographic area, or other category). We have tended to rely on the traditional rather arbitrary groupings used within a business, and this potentially disguises important information about trends that cross boundaries.

For example, in the following graph we look at a distribution chart (histogram) of an individual customer's average billed value (ABV) for one of these arbitrary groupings. The mean ABV for the group is \$113 (the red flag), but this is biased up by a small group (3%) of high value customers in the group with very large bills. This effectively disguises the fact that the majority (63%) of customers have bills below \$60. In fact, the median (mid-point average, purple flag)

¹ Data mining is the process of extracting patterns from large amounts of stored data by combining methods from statistics and database management systems. It is seen as an increasingly important tool to transform unprecedented quantities of digital data into meaningful information (nicknamed business intelligence) to give organizations an informational advantage. It is used in a wide range of profiling practices, such as marketing, surveillance, fraud detection, and scientific discovery.

Applying Analytics to High Performance Customer Profitability Models – (in Financial Services), continued

ABV was only \$37 and the modal (most common, blue flag) ABV is in the \$0-\$30 category. Given this was a company with a (relatively) costly manual payment reconciliation process, the significance of this problem for certain customers would have been hidden by the averaging effect of grouping all these customers together.




Figure 1. Distribution Chart

THE ESSENTIAL PROBLEM

With the upswing of increased computing power, and particularly transactional-based costing techniques (like TDABM and transaction-based profitability costing), this has changed. Faced with an explosion in the granularity of modern customer profitability models, a model from even a small company could easily include 100,000 customers, which means that the top decile of customers alone includes 10,000 customers. By the time you get up to a fair sized bank, Telco, energy company, or retailer, the numbers of customers in each decile is in the range of 500,000-1,500,000. How do you identify what is behind the variation in profitability then?

Behind the move toward transactional ABM models has been the availability of systematic process data now that most major organizations have introduced ERP systems. This has meant that it is more practical to define activities at a more detailed level, and provide direct cost driver data to support them. This has inevitably led to an increase in the number and sophistication of activities and cost drivers in the model, presenting even more candidates for us to search through to understand what is and is not important.

Finally, when we also factor in more product variations than ever before, and more distribution channels, the complexity of our models is now beyond that at which we can use basic reporting or even OLAP to find our most important insights.

In addition, when we separate customers down to an individual level, we encounter occasional activities associated with a customer (what we might call "risk" incidents) that might have happened in the period under examination, but are likely only to reoccur on an undetermined occasional basis, but have a dramatic impact of the potential profitability of a customer. An example of this is a customer that moves from credit card to credit card just to get a low introductory rate. This could be a significant cost for a bank or credit card company, but one that might be incurred yearly for one customer, but once in a decade for another. For a correct appreciation of the profitability of a customer, we need to not just understand what it costs to process this incident, but also the likely probability of it occurring in any given period.

So it can be seen that the fundamental role of analytics is the following:

1. reduce complexity by identifying what is most important to focus on
2. discover hidden patterns in behavior that cross arbitrary boundaries
3. refine our interpretation of intermittent risks

HIGH VOLUME CUSTOMER PROFITABILITY MODELS

To compound the essential problem, models with large volumes of data, 10s of billions of transactions records, intense cost calculations, and numerous reporting dimensions are beginning to make traditional modeling and analytics difficult to do.

Applying Analytics to High Performance Customer Profitability Models – (in Financial Services), continued

The consequences to this are, poor reporting and response times, long process (run-times) allocations and latencies, and the reporting is not available when we need it.

Traditional modeling tools are optimized for 30-year-old technology (slow disks and CPUs) and have a lot of I/O bottleneck. Essentially, BI and analytics is always an afterthought.

A conventional disk drive accesses data at an average of 5 milliseconds, but in-memory technology accesses data at an average of 5 nanoseconds (1 million times faster than a disk).

According to Gartner (Schlegel, et. al., 2006), “By 2012, 70% of Global 1000 organizations will load detailed data into memory as the primary method to optimize BI application performance.”

In the future, to accomplish the reporting and analytics, tools will leverage in-memory technology that will allow for a much faster, almost instant response time for doing analytics.

BEFORE WE EVEN GET STARTED WITH ANALYTICS WE NEED TO THINK A BIT

The Question Is at the Heart of Everything

When Hiram Bingham set out, he was not exploring the Andes at random. He had heard rumor of a “Lost City of the Incas”, and while he did not know where he would find it, he had a clear idea what it was it he was searching for, and how to look for it. As a consequence, he discovered Machu Picchu in 1911.

Customer analytics is not as glamorous as battling through the Peruvian jungles, but it is just as important that data exploration is not done at random. The objective of the analysis needs to be clearly understood before the exercise begins. When this is understood, the analyst can then determine the following:

- How they need to define their model?
- What data should and should not be included?
- What transformations are needed?
- What statistical tools are appropriate?
- What the outcome of the analysis will look like?
- How it might be interpreted?

Only once all these things are determined can meaningful analysis begin. Remember always, “**Prudens quaestio dimidium scientiae**” – half of science is asking the right questions (Roger Bacon 1214-1294).

The Importance of Ratios

One of the critical first steps when analyzing an ABM model is to understand the data, and how it needs to be transformed to correctly answer your questions. The biggest issue that we face is the impact of volume effects on our analysis, and how they can override any other potentially more important analysis. It is for this reason that often one of the first actions taken is to “normalize the data” through the derivation of key ratios that remove the size effects, and allow deeper insights to be surfaced.

To illustrate this, take the example of a large bank that charges a fixed fee per \$100 paid in as cash over the counter (50 cents). Say there was a known cost of handling this transaction (say \$3), then you could easily create a similar equation to the returns scenario where there are four companies, all of whom pay in all their takings each year over the counter you have:

Company	Total Paid in Over the Counter	Average Amount Paid in OTC	Number of Payments in OTC	Profit from Other Charges	Final Profit after OTC Cash Handling Charges	Profit Margin (as a % of Turnover)
A	\$500,000	\$500	1,000	\$1,000	\$500	0.1%
B	\$500,000	\$2,000	250	\$1,000	\$2,750	0.55%
C	\$5,000,000	\$500	10,000	\$10,000	\$5,000	0.1%
D	\$5,000,000	\$2,000	2,500	\$10,000	\$27,500	0.55%

Example 1: Bank Charges

A simple regression on the number of payments in and final profit will discover a positive correlation, that is, the more payments in you can get from the customer, the more absolute profit you will make, so based on this if you can

Applying Analytics to High Performance Customer Profitability Models – (in Financial Services), continued

encourage people to make more smaller payments you should make more money! A counter intuitive result, but a mistake that a computer would make because it lacks intuition.

However, if you calculate the average amount paid in versus profit margin, then you find the intuitively correct result that the larger the average payment size, the larger the profit margin as a percent.

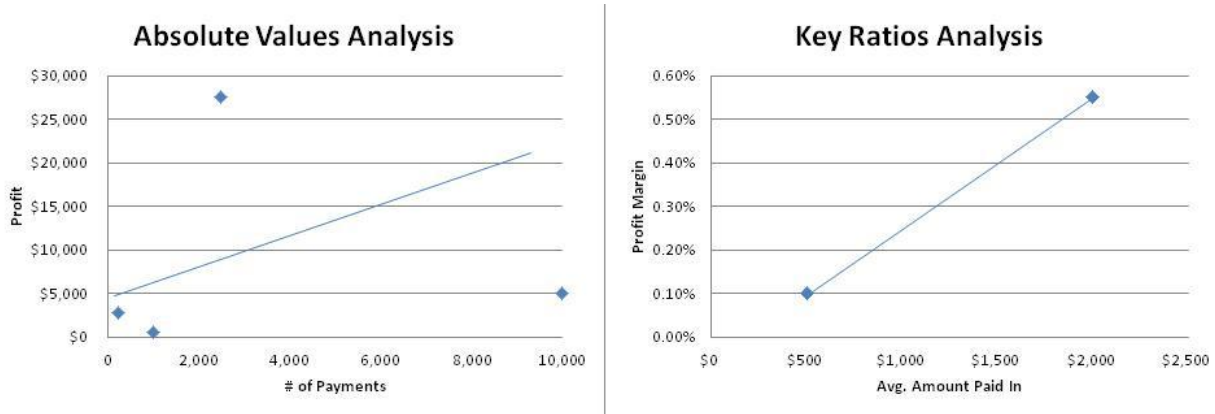


Figure 2. Ratio Analysis

Indeed, the beloved profit cliff (or whale curve) has a tendency to provide a relatively false picture of profitability. Because it orders customers on the basis of absolute profit, it tends to group all the small customers in the center, and place large customers at each end, focusing our attention toward these few significant but anomalous customers, and leading us to miss potentially critical trends and patterns that can be found in the population at large, and which can have a dramatic impact on profitability. By sizing and ordering our customers based on revenue, and showing profit versus revenue, we can transform a relatively benign looking whale curve into a much more dramatic hook curve. The following figure shows that we have a whole set of customers who are in fact more profitable than our largest customer, even though they generate the most absolute profit, and similarly, we have a host of customers who are significantly less profitable than our customer that generates the greatest loss. It is also quite clear that clearing out our unprofitable customers will not have a dramatic impact on our top line, and so we will feel free to attack them without worrying overly about the impact on the share price.




Figure 3. Profit and Revenue Curves

So how can analytics help us with understanding customer profitability?

So we have looked at the importance of transforming the cost drivers in our models to create key ratios, and seen the value that they potentially bring in improving our understanding of the model, but we are still left with the question of how do we find our way through a large model with hundreds of thousands of customers, thousands of cost drivers, and hundreds of products?

In the next section we look at some of the different analytical approaches we can take to improve our understanding of customer profitability using analytics.

Applying Analytics to High Performance Customer Profitability Models – (in Financial Services), continued

TRANSACTIONAL BEHAVIORS

Better Understanding the Impact of Processes

By combining data mining and advanced analytics (in this case a statistical technique called a decision tree) with today's enormous computing power and its access to massive amounts of stored data about customers, one can gain tremendous insight. Decision trees are a simple but powerful form of multiple variable analyses. Produced by algorithms that split data into branch-like partitions, decision trees are developed and presented incrementally as a collection of one-cause, one-effect relationships calculated in a recursive form. The appeal of decision trees lies in their relative power, ease of use, robustness with a variety of data types, and ease with which they can be understood by non-experts.

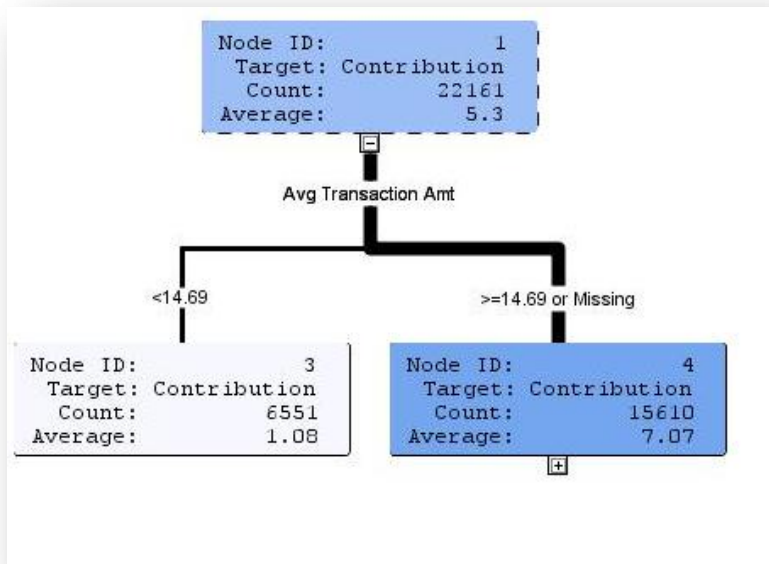


Figure 4. Decision Tree Transaction Amount

Figure 4 displays other potentially useful information:

- It calculates that 5.3 is the average transaction quantity that divides the more and less-profitable customers into two subsets of the whole population of 22,161 customers.
- It calculates that 6,551 customers are the “less profitable” (with their own average transaction quantity of 1.08) – and that 15,610 customers are in the “more profitable subset (with their own average transaction quantity of 7.07).
- It calculates that 14.69 is the dividing amount, with customers above that number being x and those below it being y.

Based on the initial partition, the marketing and sales functions can begin to brainstorm how to alter the behavior of customers in the “less profitable” segment so that they move in the direction of customers in the “more profitable” segment.

But while that brainstorming is occurring, the analysts can delve deeper. After the average transaction quantity is revealed as the most prominent factor, each “more/less profitable” segment will be recursively partitioned.

Applying Analytics to High Performance Customer Profitability Models – (in Financial Services), continued

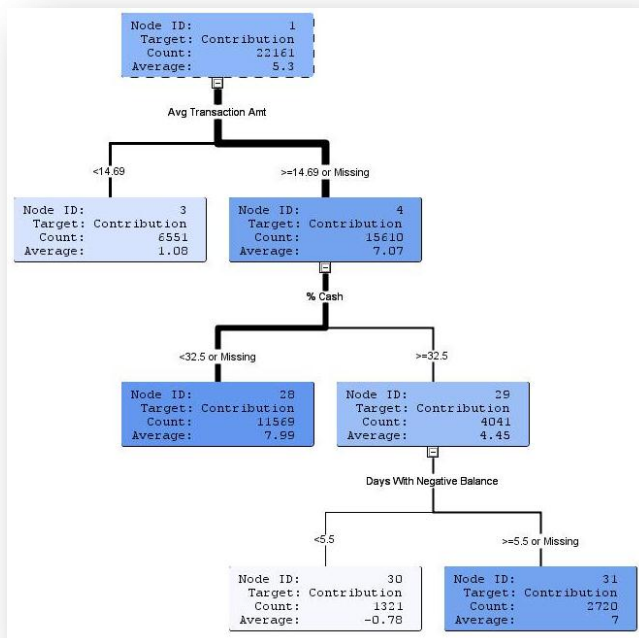


Figure 5. Expanded Decision Tree

Following one branch of the decision tree down, figure 5 reveals that the factor that most differentiates the “more profitable” customers is “% cash”; and subsequently, farther down the tree, a third critical factor – “days with a negative balance” – applies.

At this point, an uncomfortable fact is uncovered. Within the “high average transaction quantity customers” exists a distinct micro-segment. This micro-segment uses a lot of cash and frequently runs overdrafts. Consequently, they are the least profitable customers. Now the marketing and sales functions can focus on this particular micro-segment and brainstorm ideas to change this customer segment’s behavior or their commercial terms, and move them toward profitability. Figure 5 displays the expanding the decision tree diagram.

RISK INCIDENTS

Accounting for Occasional Events in Profitability Analysis

One of the unfortunate side effects of the periodic nature of ABM models is that it captures occasional events against a customer in the period that it happens, registers the impact on profitability in that period, but provides us with little information about whether that event is likely to recur frequently or infrequently.

Examples of this might include things as varied as: issuing of new credit cards, insurance claims, warranty claims, or maybe mobile handset renewals for Telco’s, and home moves for energy companies.

When designing our data exploration model, for these sorts of costs we need to adopt a different strategy for these costs. Specifically, we need to replace the occasional event or behavior itself with a marker indicating the probability of the event occurring. Typically this can be either a direct calculation of the probability of the event occurring for each customer in the period, which can be calculated using a number of techniques, including logistic regressions, neural networks, or indeed a decision tree.

Applying Analytics to High Performance Customer Profitability Models – (in Financial Services), continued

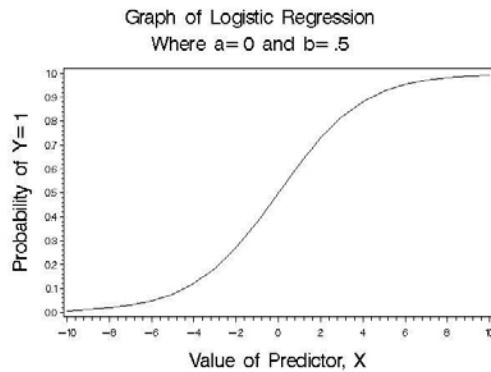


Figure 6. Logistic Regression

Alternatively, we might find a proxy indicator for the risk of the occasional event, these are typically geo-demographic indicators, but could also include products (for example car model effecting the probability of a warranty claim), or channel (returns are more likely for mail order than shop bought products).

Exact choice of approach depends very much upon the following:

- a) the needs of the analysis
- b) the availability of data
- c) the tools available

Customer Strategy

Where ABM has been applied to strategy, it has typically been focused on structural issues: how to organize departments to achieve economies of scale and what markets to continue with. But for a marketer thinking about their customer strategy, customer profitability, and behavioral analysis should be one of the foundation stones. However, their need differs from the person looking to identify and understand the sorts of problem customers that are identified by a decision tree.




Figure 7. Cluster Analysis

With the members of each cluster identified, other information can be overlaid on the analysis to deepen the picture, and appropriate strategies developed.

Applying Analytics to High Performance Customer Profitability Models – (in Financial Services), continued




Figure 8. Further Cluster Analysis

The first group (PG) was happy to pay for a relatively high service level and consequently was very profitable, so a strategy of **Cuddle** was developed. The UNeg group was very similar to the PG group in many ways, but a significant portion of their purchase mix involved products sold at a loss. The strategy here was to **Cure** this problem and return these customers to profit. The core of the business was the PNorm group, low effort with an OK margin whom we needed to **Keep** as customers, and finally there was a class of customers UNorm who asked the earth, but did not want to pay, and for those an active **Cull** strategy was developed where they were offered terms that if they took would make them possible, but if they rejected would leave to them no longer being serviced.

Profitable Customer Acquisition

While much of this paper has been focused on how we identify which behaviors customers exhibit that make them profitable or not, it is not possible to understand how a prospect might behave should we acquire them, and therefore whether they are likely to be worth the effort. Much work has gone on in sophisticated companies to look at expected sales revenues for different demographics. But these analyses can be improved even further by moving from analyzing and segmenting based on expected revenue, to applying the additional insight available from a customer profitability model, which not only has the potential to tell a company what level of profit they might expect from a particular prospect segment, but also how that group is likely to behave, and therefore the potential impact on the company's resources (call centers, loan processing, branches and so on) of these acquisitions.

To achieve this analysis, a decision tree is often the most useful tool, but rather than applying behaviors to segment the customers, we apply demographic indicators, and once a segment is identified, overlay that initial analysis with a behavioral analysis.

Customer Equity Analysis

The final stage in this process is of course to move to some form of lifetime value analysis. Much of the literature assumes that customers advance on an incremental basis, gradually growing over time buying more and new products, right up to the point they leave.

However, with the much greater depth of knowledge we have on customer's behaviors, and which ones are significant, one of the clear findings is that customers are not generally incremental in nature, but rather tend to be relatively static until they go through some form of state change transformation (leaving school, getting married, and losing a job). With our ability to identify how customers typically behave in each of these states, and the propensity of them to move between states, a more realistic approach would be not to look at individual customer lifetime value, but to look at the potential value of a particular segment, and how it will change over time as new customers are acquired either through transfer of acquisition, and are lost either through transfer out, or churn.

CONCLUSION

It's fine to let the computer do the heavy lifting for analytics, but when trying to identify major traits between more or less profitable customers it might not be intuitively obvious. Compound that with the fact that banks are now looking at finding the profit and loss for every single customer in their portfolio. That is, billions of rows of data that need to be calculated and processed before even starting the analytics. It can be like finding a needle in a stack of needles.

The goal is to accelerate the identification of the differentiating drivers so that actions – interventions – can be considered as a way to get high payback profit lift from customers. Even greater benefits and better decisions can come from applying data mining and advanced analytics in high performance customer profitability models.

Applying Analytics to High Performance Customer Profitability Models – (in Financial Services), continued

REFERENCES

Schlegel, Kurt, et. al. 2006. "BI Applications Benefit From In-Memory Technology Improvements." Available at <http://www.gartner.com/id=496876>.

ACKNOWLEDGMENTS

Thank you to Gary Cokins and Dr. Charles Randall who are really the brains and authors behind this SAS Global Forum paper. Please see the recommended readings for Gary Cokins' whitepaper with more information about this topic.

Gary Cokins is an internationally recognized expert, speaker, and author on the subject of advanced cost management and performance management systems. He is a Principal in Global Business Advisory Services with SAS, a leading provider of business intelligence and analytic software headquartered in Cary, NC. Gary received a B.S. in Industrial Engineering/Operations Research from Cornell University and an M.B.A. from Northwestern University's Kellogg School of Management. Gary began his career at FMC Corporation and has served as a management consultant with Deloitte, KPMG Peat Marwick, and Electronic Data Systems (EDS). His latest book is *Performance Management: Integrating Strategy Execution, Methodologies, Risk, and Analytics*. Gary can be reached at gary.cokins@sas.com.

Dr. Charles Randall has built his career in Strategy and Marketing Analytics. Charles' career has encompassed both telecommunications and management consultancy before joining SAS as a Principal Business Consultant. In his current role as Solutions Marketing Manager, his 15 years' experience in these fields provides a depth of expertise and insight which he shares in the numerous articles and papers he writes. Charles received a BSc in Economics and a PhD in Econometrics from the University of Wales. His latest research study is, *Pleased to Meet You: How different customers prefer very different channels*, is a joint project with Professor Hugh Wilson of the Cranfield School of Management. Charles can be reached at charles.randall@suk.sas.com.

RECOMMENDED READING

- Adkins, Tony C. 2006. *Case Studies in Performance Management: A Guide from the Experts* (Wiley and SAS Business Series).
- Cokins, Gary. 2004. *Performance Management: Finding the Missing Pieces (to Close the Intelligence Gap)* (Wiley and SAS Business Series).
- Cokins, Gary. 2009. *Performance Management: Integrating Strategy Execution, Methodologies, Risk, and Analytics*.
- Cokins, Gary and Charles Randall. 2011. SAS Institute white paper. "Increase Customer Profitability Using Data Mining and Advanced Analytics."

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Tony Adkins
SAS Institute Inc.
E-mail: tony.adkins@sas.com
Web: www.sas.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.