



Which Customers Are Worth Keeping and Which Ones Aren't? Managerial Uses of CLV

Published: July 30, 2003 in [Knowledge@Wharton](http://knowledge.wharton.upenn.edu/article.cfm?articleid=820)

Roy Cardiff runs a mail-order business that tracks sales to each customer. He recently decided to cut costs by curtailing catalogs to those customers who are least likely to buy from him in the future.

His customers break down into three categories: those who made several small purchases throughout the past year; those who made a single purchase but for a much larger amount, and those who have had a long but sporadic relationship with his firm.

Which segment of customers should Smith prune from his mailing list?

According to several Wharton marketing professors who have studied this issue, there is no easy answer, despite new and increasingly sophisticated efforts to measure what is called “Customer Lifetime Value” (CLV) – the present value of the likely future income stream generated by an individual purchaser.

“For many companies, their whole business revolves around trying to understand which customers are worth keeping and which aren’t,” says Wharton marketing professor [Peter Fader](#), who used the mail order example above in a recent co-authored paper entitled, *Biases in Managerial Inferences about Customer Value from Purchase Histories: Intuitive Solutions to the Mailing-List Problem*. “This has led managers from a broad cross section of industries to seek out more refined measures of CLV, using data-intensive procedures to identify top customers in terms of their likely future purchasing patterns.”

The goal is not only to identify customers, but to reach out to them through cross-selling, up-selling, multi-channel marketing and other tactics – all of which are tied to metrics on attrition, retention, churn and a set of statistics known as RFM – recency, frequency and monetary value.

“CLV is a hot area,” notes Wharton marketing professor Xavier Dreze, co-author of a new paper entitled, *A Renewable-Resource Approach to Database Valuation*. Although CLV is by no means new – it has long been used in business markets dealing with large key accounts – the concept has been energized by the increasing sophistication of the Internet “which allows companies to contact people directly and inexpensively.” CLV, Dreze says, “sees customers as a resource [from whom] companies are trying to extract as much value as possible.”

Yet many companies are discovering that CLV – which is one component of Customer Relationship Management (CRM) – remains an elusive metric. First, it is hard to calculate with any degree of



This is a single/personal use copy of Knowledge@Wharton. For multiple copies, custom reprints, e-prints, posters or plaques, please contact PARS International: reprints@parsintl.com P. (212) 221-9595 x407.

certainty; second, it is hard to use.

“The only number a manager can have much confidence in is a customer’s current profitability,” says Wharton marketing professor [George Day](#). “And the basic question becomes, now that you have that data, what are you going to do with it? Some companies use this information to create different programs for different value segments. In the financial services industry, for example, customers get different levels of service depending on how big an account they are. But there is always the risk that by doing this you anger other customers.”

In addition, it’s hard to predict how long a customer will stay with the company or how ‘growable’ he or she is. “In the last analysis,” Day says, “companies don’t really know how profitable customers are.”

Rolling the Dice

CLV is an intuitively appealing concept, but one that for a variety of reasons can be very hard to implement, notes Wharton marketing professor [David Bell](#) in an article entitled, *Seven Barriers to Customer Equity Management*.

CLV, say Bell and others, works best in industries where there is a high cost of acquiring or retaining customers, such as in financial services, airlines and hotels. “It’s also useful in situations where you have a skewed distribution of transactions – i.e. where a small number of people drive most of the business, as in hotels – and where firms can offer rewards and inducements to affect customer behavior,” Bell notes. An example would be airlines companies that can upgrade passengers to first class – a benefit that is considered big to the passenger but whose cost to the company is small.

Collecting data on CLV can offer particular companies a number of benefits, Bell adds. For example, the individual transaction data collected by a hotel helps the company identify its best customers and cross-sell them other products. It also allows company marketers to target that group for customer feedback. Using that feedback, the company can then make smarter decisions about where to most efficiently allocate its marketing resources. Suppose the data shows that a significant percentage of the customers come from upstate New York and are in their 50s; the hotel can use that profile for more accurate outreach, he notes.

Bell points to Harrah’s Casino as a CLV success story. Based on information gleaned from its loyalty program, Harrah’s can now figure out “who is coming into the casino, where they are going once they are inside, how long they sit at different gambling tables and so forth. This allows them to optimize the range, and configuration, of their gambling games.”

Others cite the health care and credit card industries, direct marketers and online email marketers as potential benefactors of CLV data, in part because they are characterized by direct customer contact and easy tracking abilities. For instance sales forces within the pharmaceutical industry, Dreze points out, can use relevant data to decide how often they should visit doctors’ offices to pitch their companies’ drugs.

Basically, says Day, CLV is most applicable “any time you have a database with customer profile and transaction information. But if you are working through channels – using a value-added retailer, for example, or any similar situation where you don’t have a direct relationship with the customer – then it is not as easy to implement.”

Beware Those Angry Customers

Now that marketers can collect better purchase transaction data to help determine a customer’s lifetime value, how should this data be used?

The answer, suggest some researchers, is “cautiously.”

“People are idiosyncratic,” says Bell. “On the individual level, it’s hard to predict customer behavior. It’s easier to predict the behavior of market segments. We can say, for example, that on average the business travel sector will stay “x” number of nights at the Hilton. But if we try to predict how many nights Mr. Jones will stay at the Hilton, it’s a more difficult forecasting problem.”

One of the difficulties with implementing the CLV approach, adds Bell, is that the models forecasters use are very sensitive to assumptions. For example, models frequently make assumptions about how long a customer will continue a relationship with the company, whether that relationship is an active one and how much the customer will spend. Yet some of these assumptions may be inappropriate. “Just because I spent \$100 last year doesn’t mean I will spend \$100 this year,” says Bell. “Or if a customer is inactive, is it because he has temporarily stopped using the product or has switched to a competitor?”

The problem with Internet valuations was that “many companies made inappropriate assumptions about how much customers were worth, how much it cost to acquire them and how long they would stick around,” Bell notes. “The calculation of a dollar value turns out to be very sensitive to these critical assumptions. Any errors that you make can be compounded, which means you can end up with wildly different estimates if just one of the assumptions is off.

“And yet a lot of companies are now using some measure of your lifetime value to determine how they should treat you,” he adds. “If I am an average customer I get put on hold. Otherwise, two rings and I strike right through to a real person. But that assumes that people are fairly static. You have put them in certain buckets and they stay there. Yet perhaps if you had treated me better in the beginning, I would have become a better customer.”

In addition, when firms value their own customers, they are making inferences based on what they know of a person’s history with that firm. “There is missing data. I don’t know what you are doing elsewhere. It could be that you spend \$100 with me every year, but are also spending \$500 with one of my competitors,” says Bell, referring to ‘share of wallet,’ or what a customer spends with your company versus what that customer spends with your competitors. “That is the problem with this methodology. You are trying to assign values to people based on information you have acquired about their transactions with you and nobody else.”

Any model a company uses can provide only one input into the decision process, Bell adds. “Intuition, managerial judgment, have to be there as well.”

Day cites the case of a manufacturer of large scale components who learns he has an unprofitable account. “What do you do? The account may be unprofitable but in these kinds of business markets, that account could be 15% of your sales. It takes a lot to announce that you can’t afford to service them anymore ... Lifetime value is after the fact. The tricky part is forecasting the prospective value; how do you know what this customer will do in the future?” A company’s biggest risk, he adds, is that they “inadvertently turn off customers” who may have become profitable to them in the long term.

Fader suggests that some CLV models ignore the “inherent randomness” of individuals. “These models look at customers’ past behavior and view each one as if he or she were a fixed annuity that pays off at certain stages ... But the pattern of past transactions isn’t the best, or the only, predictor of the future.”

Water Skis and Goggles

While tactics like cross-selling and up-selling have been around for years, these days they are used more frequently and aggressively to try to augment customer lifetime value, says Fader. Their success, he suggests, is mixed.

In cross-selling, a company that has sold you water skis, for example, will try to also sell you goggles. For marketers, the appeal is clear. “It’s easier to sell to somebody that you already know,” says Dreze. “It is trying to maximize the value of the relationship that you already have.” Fader, however, is “somewhat skeptical of the tactic. If someone’s behavior within a category is largely random, then when you take the randomness in one category and cross it with the randomness in another category, it’s often very hard to make any valid connections.”

Up-selling can also be problematic. Consider Amazon, which provides free shipping once a customer spends “x” dollars, or offers a second book for a discounted price once the customer has bought the first book. “In the Amazon example, perhaps a customer would have paid full price for the second book and didn’t need the reduced offer,” says Fader. “Some companies put too much emphasis on up-selling. It’s hard to quantify the true impact of these efforts. Looking at the sales numbers alone doesn’t indicate the amount of incremental profitability that can be directly tied to the marketing effort.”

A sales tactic similar to cross-selling is multi-channel marketing. “It used to be that most companies had only one touch point with the customer,” says Fader, “but now there are many kinds of retail outlets, plus the Internet, direct mail, call centers, etc. It leads to an issue of resource allocation. If one customer uses the Internet and another uses the call center, should we treat them differently? Clearly you might want to push some people to the Internet because it’s cheaper than staffing a call center, but the question is, which customers? What are the behavioral characteristics of people who can be pushed? Should you risk angering loyal call center store customers by trying to move them online or should you focus on less loyal ones even if you can’t get as much value out of them?”

What it gets down to, says Fader, is that “some selling tactics are good, some are bad, but in general it’s hard to sort out the returns on these marketing investments and link them back to ongoing CLV measurement/management. As companies try out many different tactics on their customers, they inadvertently ‘contaminate’ the CLV numbers, making it even harder to figure out which customers to target or ignore in the future.”

Ongoing Research

In a recent paper entitled, *Investigating Recency and Frequency Effects in Customer Base Analysis*, Fader, along with co-authors Bruce Hardie, Chun-Yao Huang and Ka Lok Lee, look at how database marketers assessed the value of different customer groups in relation to their past behavior patterns *before* CLV became so widely-used among managers. “The most popular framework classified prospects based on RFM: the recency, frequency and monetary value of past transactions,” Fader says.

RFM has its roots in direct marketing, one of the most progressive industries in terms of using CLV concepts. Fader and his colleagues wanted to know how the simple RFM measures relate to the more complex CLV estimates, perhaps as “leading indicators” of future purchasing. “If you have a customer who has bought a lot of merchandise but not lately, and a customer who has bought some merchandise lately, which one is better in terms of CLV and is therefore more desirable?” Fader asks, referring back to the opening example. “And how do the tradeoffs between recency and frequency play into this?”

In their paper Fader and his colleagues suggest that simple statistics such as recency and frequency can in fact offer valid estimates of future lifetime values, i.e. “that a limited amount of summarized transaction data, when viewed the right way, can yield CLV forecasts that are just as accurate as those generated from the entire highly-detailed purchase history. The challenge for practitioners is knowing which summary statistics to use, and how to use them correctly. Many common ‘rules of thumb’ don’t lead to very effective managerial policies,” he says.

In *Biases in Managerial Inferences about Customer Value from Purchase Histories: Intuitive Solutions*

to the *Mailing-List Problem*, Fader, David Schweidel and Robert J. Meyer set aside their complex equations in an effort to gain a better understanding of these rules of thumb. Fader recognizes the fact that “in most real world settings, the identification of key customers still has a strong intuitive component.” In other words, despite modeling tools that use purchase transaction data to project future buying patterns, “managers make extensive use of subjective rules for identifying those customers who are likely to be the best (or worst) source of future sales.”

The paper notes that little empirical work has been done examining the ability of managers “to form inferences about customer potential from sales histories...” The researchers address this issue by setting up situations where participants are shown purchase histories for a series of customers and asked to make different assessments about them.

What we found, says Fader, is that managers are inconsistent in their use of summary information such as recency, frequency and monetary value. The ways that managers use these cues vary drastically based on the task they are facing (e.g. figuring out which customers to add to the mailing list and which to drop) as well as the format used in presenting the customer purchase history data to managers. “It is vitally important to understand how managers are affected by these external factors before we encourage them to use any ‘black box’ models ... We need to balance our high-tech model-building efforts with a better understanding of the psychological aspects that underlie managerial decision making.”

In *A Renewable-Resource Approach to Database Valuation*, researchers Dreze and Andre Bonfrer offer a “new way to look at customers. Traditional CLV looks at the net present value of all income generated by one customer. Part of the assumption when marketers compute lifetime value is that at some point the customer will defect,” says Dreze.

But when you make that assumption, he adds, “you severely underestimate the value of the database. If you were trying to optimize your marketing actions based on that formula, you would make the wrong decisions. The reason is because yes, you lose some percentage of your people every year, but you will acquire new ones. You need to take into account the acquisition of new customers when you value the database.” In other words, Dreze says, “it is important to maximize the database value and not the customer value.”

In other research, [Noah Gans](#), Wharton professor of operations and information management, looks at the issue of CLV from an optimization standpoint: If a company has limited resources, which customers should it focus on?

Gans has developed theoretical models looking at how the average time that a customer stays with a service provider is affected by the overall level of service quality. “There can be a strong increase in the expected time a customer will stay with you as you improve the average service quality,” he says. But there are other issues that also must be considered: What is your competitor doing? What does it cost for a customer to switch services? How does the evolution of technology affect the transaction?

At some point a company makes inferences about what kind of customer it is dealing with. “Then it takes an action offering the customer a certain level of service quality. In a call center, for example, this would mean giving the customer priority over other callers. That is an operating control the company is using to manage what the customer gets and the costs of serving that customer.”

Gans says he wants to use marketing models to make better operating decisions. “I am waiting for somebody to hand me a model of how customers behave – how they respond to different levels of services – and then I can describe the costs of providing a certain quality of service.”

He uses the example of cross-selling. “It’s a very simple problem. You decide at the end of a service

whether you should cross-sell. At a call center, for example, cross-selling from an operations point of view adds length to the time of a call and makes other callers wait longer. You need to know how much cross-selling you want to do, when to do it, how much extra capacity it takes, and so forth.

“Any decisions have to take into account the four core marketing factors: price, promotion, product and place of distribution, which all involve marketing but which also have a direct impact on operations.”

Gans addressed some of these issues in a recent paper entitled, *Customer Loyalty and Supplier Quality Competition*. The paper, he says, comes up with mathematical formulas for a service provider's "share of customer" as a function of its and its competitors' overall levels of service.

It then shows that there's a natural service-level "standard" that competing suppliers will converge to. “In real life, you often hear about these things under the rubric of ‘world-class service-level’” he says. “In call centers, for example, answering 80% of the calls in 20 seconds or less is a common standard.” The paper also shows that the more competitors there are in a market, the higher the industry standard, as one would expect.

In terms of maximizing CLV, Gans believes that for companies there is value to tracking the history of what each customer does and deciding, based on that history, what bucket to place the customer in. “Then, based on your inference about the characteristics of that bucket, you can decide how best to treat these customers, whether it's cross-selling, up-selling or whatever. But you have to temper that decision because at any given time a customer comes to visit you, you don't really know what kind of customer he or she is. So your optimal decision has to take into account your uncertainty about how the customer will respond.”

This is a single/personal use copy of Knowledge@Wharton. For multiple copies, custom reprints, e-prints, posters or plaques, please contact PARS International: reprints@parsintl.com P. (212) 221-9595 x407.