Chapter 1 Customer Profile or Customer Model?

Many people think using your customer data is about creating a customer "profile". It's a hot topic. Everybody wants to do it. But what is a customer profile? Here are 2 kinds of customer profiles:

- Customer is married, has children, lives in an upscale neighborhood, and reads Time magazine
- Customer visited the site every day for 2 months, but has not visited the site at all in the past 2 weeks

The first profile is demographic, a set of characteristics. The second profile is behavior-based, involving what the customer is actually doing. It's about customer activity.

Which seems more important to you?

They're both important in their own ways. For someone selling advertising, or deciding on content for a website, the first profile is usually important, because it defines the market for ad sales and provides clues to editorial direction. These are important considerations in attracting customers and generating revenue in the first stages of an online project.

The second profile is about action, behavior, and for anybody concerned about what their customers are **doing**, is more important than the first. Will they visit again? Will they buy again? These are the questions answered by looking at behavior. Customer behavior is a much stronger predictor of your future relationship with a customer than demographic information ever will be. You have to look at the data, the record of their behavior, and it will tell you things. It will tell you "I'm not satisfied". It will tell you "I want to buy more, give me a push." It will tell you "I think your service is awful".

I'd argue the second type of profile is more important longer term, because if the customer stops buying from or visiting the site, you're not going to have much of a chance to serve up the customized pages or ads based on any "profile" given to you. You could customize the heck out of the site based on

demographics or self-reported survey data but customers would never see the results if they never come back. So for the long haul, if you had to choose the more important profile, the profile based on action and behavior would be more critical to you than a demographic one. Customer behavior profiling is critical to a company interested in retaining customers.

Marketers who use data often talk about "customer modeling", instead of customer profiling. Modeling is kind of like profiling, but it is action oriented. Models are not about a static state, like "Customer is 50 years old". Models are about action over time, like "If this customer does not make a purchase in the next 30 days, they are unlikely to come back and make any further purchases".

It sounds so mystical, and it is. To see a mathematical model predict customer behavior is astonishing, to say the least. The model says, "Do this to these people and they will likely do this". The marketer or service provider goes out and does what the model says, and like magic, a good bunch of the customers do exactly what the model said they would. It works like a charm – usually.

Building real models is expensive, because it requires an awesome amount of talent and experience. There are many mathematical techniques used to build models, each with their own pitfalls and gotchas. Success depends a lot on the type of business, the kinds of data available, and the experience of the modeler / analyst in building models for a particular business.

What is a model? Simply, it looks at customers who are engaging in a certain behavior and tries to find a commonality in them. The marketer might say to the modeler, "Here's a list of our very best customers, and here's a list of our former best customers. Is there any behavioral signal a best customer gives before they stop being a customer? What does the data say to you?"

So here's what's in it for you, what this book is about. You can do your own "models", based on the years of experience of what works for Data-Driven marketers and service providers. And while they won't be as good as the "real" models done by Ph.D. analysts, they'll be pretty darn good. Plus, they will help you increase profits while cutting marketing and service costs. This book will show you how to do it, with a spreadsheet. Ph.D. not required.

Then you can use your models to answer some basic marketing and service questions about your customer base. Questions you no doubt have asked many times yourself, such as the following:

- Who do I provide marketing or service programs to? When? How often?
- Should I contact some customers more often than others? (Yes, you definitely should.)
- How much and what kind of incentives should I provide to get a customer to do something I want them to? Can I predict which customers will be responsive to the program? (Yes, you can)
- How can I tell when I'm losing a customer or when service has failed?
- How can I put a value on my different customers and the business as a whole now, and project this value into the future?
- Is my business strong and healthy, or becoming weaker?
- What can I expect in future sales from my existing customers?

And you can also use these behavioral models in combination with demographics and characteristics to produce an even richer picture of the customer. Which of the following seems more useful to you?

- Customer is married, has children, lives in an upscale neighborhood, and reads Time magazine
- Customers who are married, have children, live in upscale neighborhoods, and read Time magazine appear to be disappointed with our site, because a high proportion of them haven't visited the site in the last 30 days

The combination of behavior and characteristics can be very powerful indeed. But without the behavior, demographic characteristics don't tell you much. You will learn how to use both in building your models. First we'll talk about customer behavior, and then add customer characteristics later on.

Chapter 2 Data-Driven Marketing and Service Models: Customer Value Management Basics

I came up with the phrase "Data-Driven" because I needed one name for the process happening in the background of all the marketing and business optimization approaches where customer data is used. As soon as you say "Relationship Marketing" or "Loyalty Marketing" or "1-to-1 Marketing" or "Permission Marketing" or "CRM", all kinds of extra ideas creep in, obscuring what's really going on in the background of all these concepts.

These approaches differ in how they are positioned to the customer, and how they are communicated. But back in the pits where the data analysts are, where customer profiling and modeling take place, they're much the same.

Good marketers and service providers have two objectives with any kind of customer value management, which is what the above approaches are all about:

- 1. Hold on to the most valuable customers
- 2. Try to make less valuable customers more valuable

So whether it's relationship marketing, a loyalty program, permission based, or 1-to-1, you still have to accomplish these goals, and to do it, you have to create marketing or service programs and execute them. This means you have to know the value of your customers and their likelihood to respond to a program, whether the program is customized based on books already purchased, uses loyalty points, or is service-oriented.

The above marketing and service approaches are all "wrappers" around what is really going on — you want the customer to do something, or perhaps not do something. This means you have to reach out to the customer and communicate your marketing and service programs. You need answers to 3 questions — WHO to communicate to, WHEN to communicate to them, and HOW you're going to execute the communication. It doesn't matter what you call your program, what "wrapper" you put it in to present to the customer — you always have to answer these 3 questions (and maybe a few more).

In addition, you probably care about how much you spend on these marketing and service programs. Ideally, instead of blasting out expensive stuff to every customer, you would want to spend money on the customers most likely to do whatever you want them to, and not waste money on those who are not.

You want customers to do something, to take action. You want them to visit your website, make a purchase, sign up for a newsletter, add new services. And once they do it for the first time, you usually want them to do it again, especially since you probably paid big money to get them to do this "something" the first time. You don't want to pay big money the second time. The data can tell you how to accomplish this, no matter what kind of front-end marketing or service program you are running or how you "wrap it up" and present it to the customer. As long as you have the data, you can interpret it for clues as to what steps to take next, and how to save precious marketing dollars in the process.

If you understand Data-Driven Marketing and Business Optimization, you will understand the basic driving forces in all of these customer retention-oriented programs. Here's the basic philosophy of a Data-Driven program operator:

 Data-Driven programs are about allocating resources. All businesses have limited resources, even the dot-coms (eventually). When you spend \$1.00 on a program, you are looking to make back more than \$1.00 in PROFIT (not sales). If you can't make back \$1.00, the dollar is not worth spending. Given multiple places to spend the program dollar, if you can get back \$1.20 in one place and only \$.90 in another, wouldn't you rather spend it where you get \$1.20 back? This approach is called Return on Investment, or ROI, and is the reason why you want to do Data-Driven programs in the first place. Data-Driven programs are among the very few approaches allowing you to accurately measure ROI.

It's about knowing you will make a \$1.20 for every \$1.00 you spend. If you know this for sure, wouldn't it be foolish not to spend every \$1.00 you had in the budget to get \$1.20 back? If you always migrate and reallocate program dollars towards higher ROI efforts, profits will grow even as the program budget stays flat. This idea is at the center of ROI thinking reallocating capital with low return to higher return projects or programs, generating higher profits in the process. ROI is often a difficult concept to understand because there are so many people using ROI in the wrong context and measuring it incorrectly. You will learn the correct way to calculate and use ROI later on in the book.

2. Past and Current customer behavior is the best predictor of Future customer behavior. Think about it. Any entity you can define as a customer – external, internal, distributors, manufacturers, suppliers – they all pursue certain routines, and changes in these routines often indicate an opportunity or challenge is ahead in your relationship with them. When it comes to action-oriented activities like interacting with a web site, this concept really takes on a very important role. You can predict future behavior based on an understanding of past behavior, and use this knowledge to improve marketing or service programs.

We are talking about actual behavior here, not implied behavior. Being a 35-year-old woman is not a behavior; it's a demographic characteristic. Take these two groups of potential buyers who surf around the 'Net:

- People who are a perfect demographic match for your site, but have never made a purchase online
- People who are outside the core demographics for your site, but have purchased repeatedly online

If you sent a 20% off promotion to each group, asking them to visit and make a first purchase, response would be higher from the buyers (second bullet above) than the demographically targeted group (first bullet above). This effect has been demonstrated for years with many different Data-Driven programs. It works because actual behavior is better at predicting future behavior than demographic characteristics are.

3. Customers want to win at the customer game. They like to feel they are in control and smart about choices they make, and they like to feel good about their behavior. Marketers and service providers take advantage of this attitude by offering programs of various kinds to get customers to engage in a certain behavior and feel good about doing it. Customers like to "win" through these programs, whether they are consumer customers taking a discount, B2B customers using a new service offering, distributors selecting your product over other products, or manufacturers working on

supply chain issues. Programs encourage behavior. If you want your customers to do something, you have to do something for them, and if it's something that makes them feel good (like they are winning the customer game) then they're more likely to do it.

This idea has always existed; on the Internet the behavior is much more traceable and obvious than before. Customers now leave evidence of this attitude all over the 'Net, in newsgroups, chat rooms, and so forth.

4. Data-Driven marketing and service programs are all about Action – Reaction – Feedback – Repeat. Marketing and service are conversations, as the ClueTrain Manifesto (http://www.cluetrain.com) and Permission Marketing (http://www.permission.com) have pointed out (if you haven't read these works, you are doing yourself a disfavor). At a high level, service is just another form of marketing – an extremely important one. Marketing and service provision using customer data is a highly evolved and valuable conversation, but it has to be back and forth between the program operator and the customer, and you have to LISTEN to what customers are saying.

That's why I will be talking about the data "speaking to you". The data is, in effect, speaking for the customer, telling you by its very existence (or non-existence) there has been an action, and it's waiting for a reaction. An action or inaction is a raising of the hand by the customer, and the Data-Driven marketer or service provider not only sees the raised hand, but also reacts to it, then looks for the hand to be raised again by the customer.

For example, if a customer visits your site every day and then just stops, something has happened. They are unhappy with the content or service, or they have found an alternative source. Or perhaps they're just plain not interested in you anymore. This inaction on their part is the raising of the hand, the flag telling you something has happened to change the way this customer thinks about your site. You should react to this and then look for feedback from the customer. If you improve the content, e-mail them a notice, and the customer starts visiting again, the feedback has been given. The cycle is complete until the next time the data indicates a change in behavior, and you need to react to the change.

Let's say this same customer then makes a first purchase. This is an enormously important piece of data, because it indicates a very significant change in behavior. You have a new relationship now, a deeper one. You should react and look for feedback. You send a welcome message, thank the customer for the trust they have displayed in your site, and provide a 2^{nd} purchase discount. Then you await feedback from the customer, in the form of a second purchase, or increased visits. Perhaps you get negative feedback, a return of the first purchase. React to this new feedback and repeat the process over again.

The Data-Driven model of marketing / service provision is 2-way, as opposed to the 1-way approach of media advertising or "data-blind" service. It is give and take, an exchange; a communication process. Using a lot of customer communications can be costly in the offline world. But communication costs are generally low on the Internet, so the Data-Driven model is ideally suited for use there.

How is this exchange accomplished? Can the data really "speak"? It can and does, but you need to know its language and learn how to listen. It's not very hard, and I'm going to teach you how to do it. But first, we need more background on customer behavior.

Latency Metric Toolkit

Chapter 3 Trip Wire Marketing

No question about it, the constant drumbeat of the CRM machine over the past several years has confused the heck out of people. I've been doing this stuff for almost 20 years now, and I can tell you it is not as difficult as it is often portrayed. Sure, you can make it very, very complicated if you want to. But if you don't start with the basics, you're going to end up wasting a ton of money. Let's start simple, shall we?

In this chapter I'm going to explain in a more general sense how customer metrics are used, and in particular, address some of the misconceptions people have regarding customer value-based and relationship marketing techniques. Much of CRM is based on these fundamental ideas. Remember, CRM is an approach to managing a business, not a technology. You do not need to live on the bleeding edge of technology to take advantage of a customer-based management philosophy.

Generally, CRM or Relationship Marketing attempts to define customer behavior and then looks for variances in behavior. When you hear people talk about "predictive modeling" or looking for "patterns" using data mining, they are essentially taking a behavioral approach using the latest tools. Once you know how "normal" customers behave, you can do two things with your business approach:

- Formally document normal customer behavior and internalize it systemically, leveraging what you know to improve business functionality and profitability.
- Set up early warning systems, triggering events, or "trip wires" to alert you to customer behavior outside the norm. This variance in behavior generally signals an opportunity to take action with the customer and increase their value online or offline.

What is most important to measure in CRM is change. People spend way too much time worrying about "absolute" numbers, like LifeTime Value. What they should really be looking at is "relative" numbers - change over time. It's not nearly as important to know the absolute value of a customer as it is to know whether this value is rising or falling - called the customer LifeCycle. Knowing and understanding the customer LifeCycle is the most powerful marketing tool you can have. I will show you how to track the customer LifeCycle and use it to increase the ROI of customer marketing later on in the book.

Customers in the aggregate tend to follow similar behavioral patterns, and when any single customer deviates from the norm, this can be a sign of trouble (or opportunity) ahead. For example, if the average new cellular customer calls customer service 60 days after they start, and an individual customer calls customer service 5 days after they start, this customer is exhibiting behavior far outside the norm. Is there a potential problem, or opportunity? Does the customer having difficulty understanding how to use advanced services on the phone? Or is the customer happily inquiring about adding on more services? In either case, there is an opportunity to increase the value of the customer, if you have the ability to recognize the opportunity and react to it in a timely way.

Understand, there is no "average customer", and a business will have many different customer groups, each exhibiting their own kind of "normal" behavior. The tools available to identify and differentiate customer segments using behavioral metrics are discussed at length in this book. For example, the type of media or offer used to attract the customer can have a dramatic effect on long term behavior, and customers who come into the business on the same media and offer will tend to behave in similar ways over time.

In the cell phone case above, the measurement of Latency (number of days until customer service call) serves as the "trip wire", a raising of the hand by the customer, to say to the marketer "I'm different. Pay attention to me." It is then up to the marketing behaviorist to determine the next course of action. Metrics like Latency provide the framework for setting up the capability to recognize the opportunity for increasing customer value.

This raising of the hand by customers, and the reaction by marketers, is the feedback loop at the center of Relationship or LifeCycle-based marketing. It's a repeating Action - Reaction - Feedback cycle. The customer raises the hand, the marketer Reacts. The customer provides Feedback through Action - perhaps they cancel service, or perhaps they add service. The marketer reacts to this

Action, perhaps with a win-back campaign, or with a thank you note. It's a constant (and mostly non-verbal) conversation, an ongoing relationship with the customer requiring interaction to sustain. It is not a relationship in the "buddy-buddy" sense. Customers don't want to be friends with a company, they want the company to be responsive to their needs - even if they never come out and state them openly to the company.

This relationship continues to cycle over and over as long as there is value in the relationship for both the customer and the marketer. If the customer takes an Action and there is no Reaction from the marketer, value begins to disappear for the customer, and they may defect. When value disappears for the marketer (the customer stops taking Action / providing Feedback), marketers should stop spending incremental money on the customer.

Notice I did not say "fire the customer" or any of the related drivel thrown around in some of the CRM venues. All customers deserve (and pay for) a certain level of support. The real question is this: for each **incremental**, or additional dollar spent on marketing to the customer, is there a Return On the Investment? If I have the ability to choose between spending \$1 on a customer returning \$1.10, and \$1 on another customer returning \$3, I would be nuts not to choose the customer returning \$3. I have not "fired" the customer returning only \$1.10; I have just **chosen not to spend incremental money** doing any special marketing to them.

Do you see the difference?

In fact, much of the profitability typical of high ROI Customer Marketing techniques comes from knowing who **not to spend on**. Most of the decreased profitability in any marketing program is a result of over-spending on unsuitable targets with lowered returns. But because marketers tend to look at results in the aggregate, or they are looking at demographically-based segments to measure a behaviorally-based outcome like purchases, they miss important details. For example, certain segments may return \$5 for each \$1 spent while others may lose \$5 for every \$1 spent, even though the campaign as a whole may return \$1.10 for each \$1 spent.

When you are trying to encourage a customer to buy something, you are looking for a behavior to occur. To measure the results of such a marketing campaign using only demographic segmentation without any behavior-based metrics (like Latency) is misleading at best, and lazy otherwise - it's apples and oranges. If you are trying to create behavior, use behavior as your measurement yardstick to define success.

Why is all of this important to understand?

Customers who are in the process of changing their behavior - either accelerating their relationship with you, or terminating their relationship with you - are the highest potential return customers from a marketing perspective. They represent the opportunity to use leverage, to make the highest possible impact with your marketing dollar. You may make money marketing to customers who are just cruising along the LifeCycle, acting like an "average customer". But when you can predict the likelihood of an average customer to turn into a best customer, and you successfully encourage this behavior, or you can reverse a customer defection before it happens, then there are tremendously profitable longer-term implications for the bottom line. You will discover these opportunities by understanding behavior and setting up trip wires to alert you to deviations from normal behavior by a customer.

What about all the rest of the customers, those who are not either accelerating or terminating the relationship? Leave 'em alone. Whatever background marketing you are doing (advertising, branding, service campaigns, etc.) is serving them just fine. High ROI data-driven marketing techniques are best used (and create the highest returns) when they are used to surgically strike at a trend in behavior, not when customers are comfortably plodding along. However, there are not as many comfortable plodders as you think; in fact, from 40% to 60% of your customer base is either in the process of accelerating or terminating their relationship with you right now. The question is this: how do you take advantage of the situation?

Latency, Recency, and all the other metrics described in the Drilling Down book are simply tools for recognizing the opportunity to take an Action in Reaction to the customer raising their hand. If you don't have some kind of system to recognize customers in the process of changing their behavior, you will miss out on most of the highest ROI customer marketing opportunities you have. And don't count on the customer to e-mail you when they're thinking of changing their behavior - we both know that is not going to happen. A more likely scenario: they will just stop taking Action and providing Feedback. And by then, it's too late for you to do anything profitable about it. Set up your trip wires and predict the behavior, folks. It's the only way to sense when an average

customer is ready to become a best customer. And reacting to a customer defection after the fact is a truly sub-optimal way to "manage" a relationship.

Based on a national survey, 50% of marketing managers do not know their customer defection rate, and the other 50% underestimate the true defection rate. After reading this shocking statistic, I figured it was time write the book on customer LifeCycles, which can be used both to track customer defection and define high ROI opportunities to retain customers before they defect. If you understand the customer LifeCycle, you can predict the primary defection points and react to them before customers leave you. This is the highest ROI marketing you can possibly do; it's cheaper than "win-back" (after the customer defects, response is much lower) and preserves the investment and profits you have in the customer.

So we're going to take a little tour though LifeCycle-based marketing land first, and take a look at one of the simplest customer LifeCycle metrics - Latency. Latency is one of the easiest "trip wire" metrics to implement, and you can use it to make more money marketing to customers whether you are using a CRM suite or a spreadsheet to run your business.

At the core of a LifeCycle-based marketing approach is customer behavior. Customers tend to behave in certain ways unique to your business and products, and if you can discover these patterns, you can use them to predict customer behavior. If you can predict customer behavior, you can make a ton of money marketing to your customers, because you can anticipate their behavior and take appropriate steps to try and modify it. Many approaches to customer marketing rely on customer behavior "trip wires". For example, a win-back program is triggered when the customer defects. Have you switched long distance or cellular providers lately? Did you get inundated with win-back calls begging you to reconsider? "Jim, we just wanted you to know we have lowered our rates". Yeah, well, thanks for telling me after over-charging me for the past six months! But could they have known I was about to switch?

Sure. If they had looked at the calling patterns of defected customers like me, they would have seen a common thread in the behavior. These patterns create the "trip wires" for initiating high ROI marketing campaigns before the defection. The proper profit maximizing approach is to wait until I look like I'm going to defect, and then call me and offer a lower rate before I defect. I would humbly submit marketing to the customer after they defect is a sub-optimal approach; the decision has already been made. If you can market to them when they appear likely to defect, you optimize your marketing resources by not applying them too soon or too late in the customer LifeCycle.

An easy to implement and proven powerful LifeCycle trip wire is called Latency. Latency refers to the average time between customer activity events, for example, making a purchase, calling the help desk, or visiting a web site. All you have to do is calculate the average time elapsed (Latency) between the two events, and use this metric as a guide for anti-defection campaigns. Many small business people naturally use Latency in an intuitive way, for example: "Gee, it has been a while since Mary Lou had her hair styled". What the stylist really means is this: Mary Lou is taking longer than the average customer to schedule a "refresh" on her hair. In database marketing terms, her Latency is exceeding the norm. So the stylist calls Mary Lou and finds either a customer who "forgot" and appreciates the reminder, or a customer who has defected to another stylist.

In database marketing, we don't rely on "remembering" the habits of thousands of customers; we measure the behavior and react based on these measurements. When you see a particular customer's behavior diverge from the average customer behavior you have calculated above, you get a trip wire event. Since the calculation of Latency is very simple, and the diverging behavior is easy to spot, this type of anti-defection campaign is an ideal candidate for "lights-out" or automated rules-based customer retention campaigns.

As an example, let's take purchase behavior in a retail scenario. If you examine your customers and find the average time between the second and third purchase is 2 months, you have found "third purchase Latency". Any customer who goes more than 2 months after the second purchase without making a third purchase is diverging from the norm, and a likely defection candidate. It's simple logic. If the average customer makes a third purchase within 2 months of the second purchase, and a particular customer breaks this pattern, they are not acting like the average customer. Something has changed. This particular customer's LifeCycle has become out of synch with the average customer LifeCycle, and this condition is a trip wire for a high ROI Customer Marketing event.

On average, if you divert marketing resources away from customers who have made a 3rd purchase within 2 months after the second purchase, and apply these resources to customers who are "crossing over" the 2 month LifeCycle trip wire without making a third purchase, you will end up spending less money and generating higher profits for any given marketing budget. You are applying

your limited resources right at the time in the customer LifeCycle when they create the most powerful impact - at the point of likely customer defection.

Now, will all these customers respond? No, of course not. But the ones that do become active, loyal customers again, and those that don't are probably not going to be good customers in the future. The behavior of the rest of your customers tells you so. These non-responding customers may not be worth spending money on to "win-back", and in fact, will have much lower response rates to a win-back campaign. They have already demonstrated their lack of interest with their behavior, and you could be better off financially by just letting them go and focusing on more responsive, more profitable customers.

The above example is a relatively crude approach to Latency. As you might suspect, different customer segments will have different Latency characteristics, and the more you fine-tune a Latency campaign, the more profitable it will become. For example, let's say you execute the Latency campaign described above, and succeed in retaining 30% of the defecting customers, making a tidy profit. But you really have two major product lines, software and hardware, each 50% of sales. Could the Latency be different between software and hardware customers? You betcha. Upon further analysis, you find third purchase Latency for software is really one month, and for hardware it's three months. The average 3rd purchase Latency of all customers is 2 months, but the Latency by product line is specific to each line. So you bust the two groups apart, and run separate Latency-based campaigns, one for each product line.

In your original third purchase Latency campaign, you promoted to customers who did not make a third purchase within 2 months of the second purchase. This means you were "late" for software (because the average Latency is really 1 month) and early for hardware (because the average Latency is really 3 months). When you realign the timing based on the line of merchandise, you find instead of retaining 30% of customers, you retain 50% of the customers, because you have synched-up the marketing effort with the true customer LifeCycle.

And that, folks, is what LifeCycle-based marketing is all about - using your own customer's behavior to telegraph to you the most important (and profitable) time to market to them. The customer, through their behavior, raises a hand and asks you to take action. If you synch up your marketing efforts with the natural customer LifeCycle, you can't help but being more successful.

What if you were to look at an entire series of Latencies? The average number of days between the first and second purchases, the average number of days between the second and third purchases, third and fourth, fourth and fifth, etc. You don't have to use purchases, you could use contacts with customer service, visits to a web site, any behavior important to your business. What would that look like, and more importantly, what can it do for you?

It would look like a snapshot of the customer LifeCycle. And what it can do for you is start you on the path to predicting customer behavior and increasing the value of your customer base. Any type of event can be used – purchases, downloads, visits to a web site – but the event must be one that repeats or have an established "sequence" of actions to it.

Let's say you look at average behavior across all customers, and end up with a "Latency Sequence" that looks something like following:

Time between 1st - 2nd event:90 daysTime between 2nd - 3rd event:60 daysTime between 3rd - 4th event:30 daysTime between 4th - 5th event:60 daysTime between 5th - 6th event:90 daysTime between 6th - 7th event:120 daysTime between 7th - 8th event:150 days

What does this pattern say to you? Think about it.

I'll tell you what it says to me. First, as you probably realized, you are now starting to see something that looks like a "cycle", as in LifeCycle of the customer. It's a series of events you can graph with a line and make charts of. If you can measure it, you can try to affect it in a positive way, and determine the results of your efforts. Second, you now have a series of seven "trip wires" you can use as described above to more finely sift and screen behavior looking for deviations from the norm. If the average number of days between events for any single customer starts to exceed the average for all customers, a trip wire call for action is triggered on that customer. And third, somewhere around the 3rd or 4th event, something significant happens to change customer behavior in a very noticeable way. The customer accelerates into the 4th event (the time between events gets shorter and shorter), then begins to decelerate in terms of behavior (the time between events gets longer and longer). Depending on your business, this may be positive or negative.

16

How to use this information?

Regarding the Lifecycle and the trip wires, you could have a series of seven actions ready to take at any point in this LifeCycle where the customer deviates from average behavior. As long as the customer stays on track, save the money and take no action. But as soon as the customer misses or "rolls over" past one of these LifeCycle milestones, you know to pull the trigger on your action. If you follow this model, you will end up maximizing every cent of your budget and driving higher profits, because **you don't spend unless you have to, and when you spend, it creates maximum impact**. This is the recipe for high ROI customer management and marketing. Act only when you have to and always at the point of maximum impact.

Regarding the behavior change, if I was a retailer, this looks negative since the "ramp" in buying behavior reversed and went in the other direction. If I was running a pure service center, this may be a very desirable pattern; perhaps meaning the customer has "learned" the product and no longer needs as much service. It could be negative though, since opportunities to up-sell or cross-sell the customer are decreasing over time. It depends on your business. The important thing to recognize is this: there was a change in behavior, and you should try and determine how you might affect this change in a positive way. Reversals in the direction of a behavior like this are almost always significant turning points in the relationship with the customer.

Human behavior dynamics often take on seemingly "physical" properties. Inertia is one such property - an object in motion tends to remain in motion unless acted on by an outside force. This reversal in the direction of the customer "momentum" after the 4th event indicates there is something about your business - a process (or lack of a process), a product (or lack of a product), something - which causes the average customer to "slow down" and reverse their contact momentum. Your mission (should you decide to accept it) is to find out what it is and try to influence this "something" positively.

If I was a retailer with very limited resources, here is what I'd do. Given the seven possible promotional opportunities listed above, but looking for the **absolutely highest ROI on a single promotional event**, I would send a promotion to the customer immediately after the 4th purchase - and no sooner. I don't want to spend money on a promotion or by reducing my margin if I don't have to, so as long as the customer is accelerating, there is no reason to spend any money. But I would really like the ramp to continue past the 4th purchase,

and any way I can bring that 5th purchase in closer to the 4th is going to affect my bottom line, and perhaps lengthen the ramp into the 5th or 6th purchase and beyond. If I had more money to spend on promotions, I would test each of the seven trip wire opportunities, and pursue only those with the highest ROI, probably using a separate and unique discount approach for each of the seven trip wire opportunities.

If I was a service center, the fact it takes 4 calls to educate the customer might not be acceptable, and I would look for ways to decrease the length of time it takes. If I up-sell and cross-sell, I would look to weight more of this activity early in the process knowing I am not going to get as many chances as time goes on and the customer becomes more likely to defect. Success at either of these actions can create incremental profits with very little expense - you're not necessarily changing what you do, just when you do it, to match more closely with the customer LifeCycle.

Of course, you can begin to subdivide the customer base, just as we did in the hardware / software example above. The Latency Sequence may look quite different for hardware buyers relative to software buyers, and it will certainly be different by the type of campaign you used to attract the customer in the first place. Once you are able to compare and contrast different customer LifeCycles by product, campaign, customer source, or by any other data point meaningful to your business, you will begin to paint a more complete picture of what parameters positively or negatively affect customer behavior. Once you understand the behavior, you can learn to profit from it.

Chapter 4 The Hair Salon Example

There are three main phases to a successful High ROI Customer Marketing program: Measure, Manage, and Maximize. We'll tackle each of these components one at a time in this example.

Two hair salons operate in the same town, Salon A and Salon B. Both are equally competent one-person operations and charge similar prices for similar services and products. And both salons practice CRM.

There is a difference though - Salon A does not use customer data to track and manage the CRM effort, but Salon B does. Salon B's CRM toolset consists of a paper appointment book and a PC with a spreadsheet program. Salon A has only a paper appointment book, and doesn't really track anything.

One day the owner of Salon A is thinking:

Where has Mary Lou been? She's a high value customer who comes in to get the whole job done - hair, nails, massage, the works. Seems to me she hasn't been in the Salon for a while. She's tardy in scheduling her session. I should call her and find out when she is coming in.

The owner of Salon A is practicing CRM. High value customers have been identified, and a change in the behavior of one of these customers has been detected. This situation has been evaluated, and an action to take has been decided on.

But the owner of Salon A is very busy that day, and forgets to call Mary Lou. What's more, the owner has no system for classifying the fact Mary Lou has not been in "for a while". How long is a while? Part of why the owner forgets to call Mary Lou is there is no real urgency; she's just "tardy". But how tardy is tardy? When should the call be made? If there were a rule about "tardy", perhaps there would be more urgency to make the call. But there isn't, so it may seem like a waste of time. The owner thinks later on:

She'll come in sometime soon. I'm too tired to make the call tonight.

As we sit here gazing into Salon A, some other thoughts probably come to mind. How many Mary Lou customers are there? And how "tardy" will they get before the owner calls them? When you are making money cutting hair all day, it's probably hard to face calling Mary Lou customers, right? Time spent on the phone calling customers or sending them postcards is time not spent cutting hair, and the owner of Salon A can't afford to not cut hair. If the owner had only the time or energy to call just three Mary Lou customers, which three would it be?

If the owner has to give up time cutting hair to make calls, these calls better result in more business than was lost by not cutting hair to make calls. This potentially negative outcome is called "opportunity cost". If resources are allocated away from an income producing activity towards another activity, you better make sure these resources create more value than they did before reallocation. If they do not, an opportunity cost has been created. The two fundamental rules of High ROI Customer Marketing are designed to avoid these opportunity costs:

- 1. Don't spend until you have to, and
- 2. When you spend, spend at the point of maximum impact

Over at Salon B, the owner has been thinking along the same lines as the owner of Salon A, about a High Value, tardy customer named Angela. The owner is cleaning up for the night, and thinks:

How many Angela customers do I have? If I keep forgetting to call my Angela customers, I may eventually lose them. But they always come back. Or do they? I'm going to start **Measuring** Angela customers. I'm going to start tracking "tardy" customers and find out exactly what this issue is about. If it's a real issue, I'll worry about it then. If it's not an issue, I can forget about it once and for all, and spend my time cutting hair.

So the owner of Salon B sits down with the paper appointment book, looks through the customer names, and enters all the "High Value" customer names into the spreadsheet, one to a line. The owner reasons the choice to track high value customers in this way:

If there is anything to this "tardy Angela" customer thing, I get hurt the most financially by losing High Value customers. If it's ever going to be worth spending time on this instead of cutting hair, then it will be most worth it to

20

spend the time marketing to high value customers. If it's not worth it for them, it won't be worth it for any customers and I can forget all about the whole thing.

Once the high value customers are entered into a spreadsheet (about 50% of the customers are considered high value), the owner of Salon B then enters all the appointment dates for each high value customer into the columns of the spreadsheet, next to each name. To keep this project manageable, the owner decides to enter only appointments for High Value customers for the past 6 months. The owner also creates columns to subtract the dates from each other for each customer and find the average number of days between visits for each customer. The spreadsheet (nothing special, off the shelf software) is smart enough to know these entries are dates and is able to easily subtract them and convert the result into days, so all these calculations are easy and take less than an hour to create.

The owner of Salon B is then astonished to discover these facts about customers: About 30% of high value customers have not had an appointment in 6 months. Since 50% of all customers are high value, this means 30% of 50% = 15% of all customers are already defected best customers. The average number of days between appointments is very similar across all the high value customers. It is, however, not the 30 days the owner expected, but 40 days.

The owner then assumes a high value, supposedly loyal customer who has not been to the salon in over 6 months is a lost customer - at least for the near future. The owner then calculates the value of the lost business for the 6-month period by multiplying the number of customers lost by the average sale of \$150 per trip at 40 days between trips. Needless to say, the resulting number is a very large, representing many days of total sales for Salon B:

| Total Customers | 200 |
|--|----------|
| Defected Best Customers @ 15 % of Total | 30 |
| Number Trips in 6 months @ 40 days between trips | 4.5 |
| Revenue per Trip | \$150 |
| Lost Revenue: Defected Best Customers | \$20,250 |
| (4.5 x 30 x \$150) | |

The owner of Salon B then thinks:

I must be crazy for not looking at this before. I would make more money by not cutting hair for a couple of hours a week if I could get back even one of these high value customers. I'm going to do something about this right away - before I lose even more high value customers. Now that I have **Measured** this effect and know how much money it is costing me to not address the tardy Angela customers, I need to **Manage** the process somehow. How can I set up some kind of "system" that will help me figure out what to do with this data I have discovered? How can I turn the data into an action plan?

Over at Salon A, the owner knows the names of best customers who "have not been in for a while". But this owner has no system, no way to measure what the dynamics of the situation are. How long is "a while"? But at Salon B, the owner knows the average time between best customer visits is 40 days, and there are customers in this group who have not had an appointment in over 6 months. How can the owner get this business back? The owner thinks:

I'll just mail all these best customers who have not had an appointment in over 6 months a postcard offering them a discount. The postcards will say, "Since you are a best customer, you are entitled to a 15% discount if you come in for a visit within the next two weeks". They will come in and I will start a new relationship with them, and find out why they have not been in. The owner of Salon B prepares the targeted postcards, mails them out, and awaits appointments from these customers

The appointments never come.

A bunch of the postcards come back as "undeliverable", and the owner gets several phone calls from customers saying "I now go to Salon A, take me off your mailing list". Undaunted, the owner of Salon B reasons:

Clearly there is something wrong with this approach. Best customers who have not had an appointment for 6 months must already be "defected" customers. They obviously do not want to come back to me, and feel the relationship is broken already. They have moved on and established new relationships. I will try a new approach with the postcards, and will use the same offer. But this time, I will mail the postcards out as soon as the best customer has not been in for over 40 days. Since the average best customer comes in every 40 days, a best customer who fails to do so is not acting like a best customer. So each week I will use my spreadsheet to identify best customers who have not been in for 40 days, mail the discount postcard out to them, and track the results.

After a month of mailing the postcards to best customers who had not had not had an appointment in over 40 days, the owner of Salon B sat down to analyze the program. Of all the best customers mailed to, 1/3 had made new appointments, and 2/3 had not. But even with the discount, the additional profits from these customers paid for the postcard mailing many times over. High value Customer defection was being **Managed** by the program.

Despite this success, two things bothered the owner of Salon B. The first was what customers who responded said when making their discounted appointments. The second was the 2/3 of best customers who did not respond. The owner thinks:

Half the customers who responded said to me, "I'm so glad you mailed me a discount, I was planning on making an appointment in the next week and would have made one anyway, so it was great to get the discount". So I gave up margin and profits I did not need to give up. And how is it possible that so many of my best customers never responded to my offer? I wonder if there is a way to address these two issues? If I could reduce the number of "would have come in anyway" customers who got a discount, and get more customers to respond overall, I would be really making a ton of money on my best customer retention postcard program. I have **Measured** my best customer defection, and am **Managing** it with this program. I wonder if there is a way to **Maximize**, to make it even more profitable?

Well, fellow Driller, have you got an idea? You know Customer Retention is all about this process: Action - Reaction - Feedback - Repeat. The owner of Salon B has taken an action, and there has been a Reaction. How should the owner go about analyzing the Feedback? The owner of Salon B then has an idea:

What about this group of customers who said "they would have scheduled anyway without the postcard". Are they similar in any way? If there is a common reaction to the postcard among these customers, perhaps there is a commonality in the behavior or backgrounds of the customers. If I can find the key linking these customers together, perhaps I can understand why this is happening with them.

The owner of salon B goes back to the CRM software (a paper appointment book and the customer spreadsheet). The owner has entered "response date" in a spreadsheet column for each customer who responded to the postcard and any comments. The owner sorts the customers by the responders and looks at those customers who said, "would have scheduled anyway without a postcard". For each customer who responded and said this, the owner looks the customer up in the appointment book to find more details.

"Long hair cuts!!!!" the owner exclaims. "They all have long hair cuts!" which the owner immediately realizes is the problem with the discount postcard mailing program. The owner thinks:

Best customers with long hair styles can come in much less often than every 40 days, even through the average of all best customers is a cut every 40 days. So customers with long cuts are getting the postcard too early - they're not really "defected", and schedule a planned appointment with a discount I did not have to offer. They should get a postcard possibly at 60 days, or even 90 days or longer after their last appointment. Since I have a lot of customers with long cuts, most are getting the postcard too early for the cut. This explains the low overall response rate. Best customers with short cuts however, are probably getting the postcard too late. By the time I get them in the mail and they reach the customers with short cuts, it could be too late, they may have already gone elsewhere for their short hair cut.

The owner of Salon B resolves to recalculate the average days between appointments separately for best customers with long cuts and best customers with short cuts. The owner divides the customer base in two - by length of cut, and finds the average time between trips of long cut customers is actually 75 days, and for short cut customers is actually 20 days. Rethinking the retention campaign, the owner resolves to track each group individually, and to do two types of mailings each week - one to long cut customers over 75 days since last visit, and one to short cut customers over 20 days since the last visit.

Using the advanced CRM system (a spreadsheet program with one customer per row), the owner creates a column for acceptable number of days since last visit - 75 days for long cut customers and 20 days for or short cut customers. Using the date of last appointment, the owner creates a simple equation that uses today's date and last appointment date to calculate days since last visit, and to subtract this number from the number in the "acceptable" column. The salon owner thinks:

I have created a "trip wire" system for the best customer retention postcard program. When the number in this column approaches zero or goes negative for a customer, it is time to mail the discount "where have you been" postcard.

Since each customer has an acceptable number of days since last visit based on hair cut length, the timing of the mailings should more closely reflect whether or not the customer has actually defected.

The salon owner tests the new campaign - and it works. Not only does the owner get many fewer customers saying "thanks for the discount, would have been in anyway", the response rate among targeted best customers increases by 30%. The program now is maximized for this level of detail - it makes even more money than it did before, and retains more customers while decreasing the cost of discounts given away. A beautiful thing, the owner thinks. But then another Eureka moment comes to the owner of Salon B:

If I use this system there is another benefit - I should be able to actually forecast what my volume should be months in advance based on customers likely to schedule an appointment. If I see a week coming up where visit volume looks to be low, I can promote to some customers and fill up empty slots, maybe give them a discount for scheduling on a specific day when my traffic is light. That way the customer is happy because they get a special one-time discount, and I am happy because I am maximizing my revenue per day by filling up light traffic days with happy customers!

Just then, the owner of Salon B hears someone walk in the door. A voice calls out, "Can we schedule appointments?" The owner recognizes the voice - it belongs to lost best customer Angela, the one who started this whole project by being tardy in scheduling an appointment. Angela is the reason the owner of Salon B first asked the question, "How many tardy best customers do I have?" But what does she mean "we"?

As the owner of Salon B comes around the corner, Angela smiles and says, "This is my friend Mary Lou. She was going to Salon A, but is dissatisfied with the results she is getting. She would like to try Salon B. And I need a cut too! I tried growing my hair out long, but I decided I like it better short".

The owner of Salon B thinks: I can't predict everything, but my new system is sure better than not predicting anything at all!

Chapter 5 The B2B Software Example

A B2B software company has an appealing pitch to business - their software makes a company more efficient and saves more money for the company than the software costs. The software is modular, with a base application and additional add-ons that are specific to certain business challenges. The selling strategy is to under-price the base application to get market penetration and then make a higher margin on the add-ons. The add-ons drive the profitability of the business, as does the installation and customization of these add-ons.

The company has been quite successful with this selling strategy. But lately the CFO has noticed sales of the base application have risen, but revenue from addons has not risen in the same proportion. In other words, the company is further penetrating the market and gaining new customers but getting less revenue from each customer. The CFO thinks:

I can't understand this. Sales of the base application are rising according to plan but overall company revenue is not growing at the same rate. The only thing I can think of that would create this particular situation is fewer basic application customers are buying add-ons. How can I figure out why this is happening?

The CFO calls the heads of business development and marketing to ask about the situation. They both report they are aware of slowing add-on unit sales per customer, but cannot attribute it this to anything specific. The company is simply penetrating the overall market more deeply they say, and as we penetrate further and further, add-on sales seem to have slowed.

The CFO is not particularly satisfied with this answer, and thinks:

If it shows up in my financial statements, it has to be measurable. I'm just seeing this from too high a view. All the sales of the different base applications and add-ons roll up to total sales, so the data I need to better understand this must exist somewhere. The CFO picks up the phone to call the CIO, and then hesitates. The IT people are going to want to know specifically what I am looking for, the CFO thinks. Do I really know?

What is needed here, fellow Drillers, is quantification, some framework for analyzing the situation. What is the real question to be answered here? The CFO knows IT has limited resources to apply to this kind of ad hoc work - if the request just generates information that leads to another question, then time and resources are wasted.

The CFO could ask for monthly product sales percentage by type over the past year. In a lot of ways, this information would simply confirm what the CFO already knows - sales of add-ons have gone soft. But does it answer the core question of **why** they have gone soft? It does not, and that is the real question at hand. Since customers have different LifeCycles, any monthly sales data will contain customers in various stages of being likely to buy an add-on. So raw monthly financial data - the kind the CFO is used to working with - is not going to answer the "real" question. The CFO thinks:

Customers buy the base package and once they get it integrated and tuned up they start to buy the add-ons. During any one-month period, we have customers who just bought the base package, customers who are in different stages of integration, and customers who are buying add-ons. What I really need to know then is this: what is the average number of weeks between the purchase of addons, this year versus last year? If this number of weeks is rising, that is where the softness in add-on sales is coming from - customers are simply taking longer to make the purchase decision. If this number of weeks is constant or falling, then something else must be going on.

With a definition of the question at hand, the CFO picks up the phone and calls the CIO. The CFO gets the report on the average number of weeks between the purchases of add-ons. The information looks like this:

Average Weeks between Add-On Purchases

| Last Year | This Year |
|-----------|-----------|
| 8.6 weeks | 8.9 weeks |

So it is taking longer for them to purchase, the CFO thinks, and darn it, now I have another question. The IT people are going to have me for breakfast for not thinking this all the way through the first time! I got the information I asked for,

but this information is not actionable, I can't **do** anything with it. There is not enough detail in the information to act.

Fellow Drillers, when you are plumbing the depths of your data, try to think of what you will do with the information you are asking for. Imagine getting back your results, and taking an action based on those results. If you can't imagine the action you would take knowing the information, you are not asking the right question yet. The CFO thinks:

Our add-on modules have different prices and different levels of difficulty involved in their integration. And they are usually installed in a particular sequence. So what I really should have asked for is the average number of weeks between the purchase of add-ons **by add-on** - the time between base purchase and the first add-on, the time between the first add-on and the second, and so forth. Maybe there are problems with installing one of the add-ons due to changes in the next generation of operating systems, for example, and this is slowing the installation of a particular add-on down. If I can get the average number of weeks between add-on purchases by add-on, I can act on it, because I will know which particular add-on is causing the slowdown.

The CFO reluctantly picks up the phone to call the CIO. At least this time, the CFO thinks, I have thought the question out all the way through, and I know what action I can take with the information once I get it. Shortly after a slightly heated exchange involving resource allocation, budgets, and a hiring freeze in IT with the CIO, the CFO gets this report:

Average Weeks between Add-On Purchases by Add-On

| | Last Year | This Year |
|------------------------------|------------|------------|
| Base app to 1st add-on | 12.3 weeks | 12.1 weeks |
| 1st add-on to 2nd add-on | 10.5 weeks | 10.2 weeks |
| 2nd add-on to 3rd add-on | 8.7 weeks | 8.9 weeks |
| 3rd add-on to 4th add-on | 6.1 weeks | 6.7 weeks |
| 4th add-on to 5th add-on | 5.2 weeks | 6.5 weeks |
| Average Time Between Add-Ons | 8.6 weeks | 8.9 weeks |

Fellow Drillers, it would be nice if the pattern were a bit more clear, yes? It appears customers are ordering their first and second add-ons more rapidly than last year, but as they get to the third, forth, and fifth add-ons, they are ordering more slowly than last year. What could this possibly mean? The CFO thinks:

Well, I answered my question, but I've got another. The reason why add-on sales appear soft is a longer purchase cycle for the average add-on, and the reason this is happening is the later add-ons are taking much longer to be purchased than they were last year, even though the first add-ons seem to be cycling much more quickly. What does that mean? I promised the CIO I would be able to act on this information, and I simply do not know how.

Fearing another phone call right away to the CIO, the CFO thinks:

What I have here is change. There has been a significant change in the way this business works for some reason. Change doesn't happen in a vacuum though; something must have caused these changes to happen, a significant event now being reflected by these average weeks between add-on purchase numbers. What could it be?

The CFO remembers the heads of business development and marketing saying the company was "penetrating the overall market more deeply, and as we penetrate further and further, add-on sales seem to have slowed". Was this the change the CFO was looking for? What did it really mean, in terms of how the business may have changed?

Getting the heads of business development and marketing on the phone again, the CFO asks if this market penetration situation had created any changes in the way the company does business. The CFO hears for the first time about a new trade campaign and a new sales person hired to address a particular market segment. This is most assuredly the change the CFO has been looking for!

Gingerly, most humbly, the CFO calls the CIO once again. This time, the CFO wants to see average number of weeks between add-on installs by add-on **by salesperson**. After a promise to review the hiring freeze is extracted from the CFO, the CIO delivers this report:

| | Last Year | This Year | Sales 1 | Sales 2 | Sales 3 | Sales 4 |
|---------------------------|--------------|--------------|------------|------------|------------|------------|
| Base app to 1st add-on | 12.3 | 12.1 | 12.3 | 12.3 | 12.3 | 11.6 |
| 1st add-on to 2nd add-on | 10.5 | 10.2 | 10.5 | 10.5 | 10.5 | 9.4 |
| 2nd add-on to 3rd add-on | 8.7 | 8.9 | 8.7 | 8.7 | 8.7 | 9.3 |
| 3rd add-on to 4th add-on | 6.1 | 6.7 | 6.1 | 6.1 | 6.1 | 8.3 |
| 4th add-on to 5th add-on | 5.2 | 6.5 | 5.2 | 5.2 | 5.2 | 10.2 |
| Avg. Time Between Add-Ons | 8.6 | 8.9 | 8.6 | 8.6 | 8.6 | 9.8 |

Average Weeks between Add-On Purchases by Add-On by Salesperson

And there it is.

Clients of Salesperson # 1, # 2, and # 3 are purchasing add-ons at the same rate they did last year. The clients of the new salesperson # 4 are purchasing in a dramatically different pattern, with much shorter purchase cycles in the beginning and much longer cycle purchases later on. Literally, the LifeCycle of the customers in this market segment are different from the LifeCycles of the average customer from previous years, and dramatically so.

It takes these customers on average 14% longer to purchase any add-on - 9.8 weeks versus 8.6, or 1.2 weeks. Over the entire purchase LifeCycle of the add-ons, this increases the purchase cycle by 4.8 weeks (1.2 x 4). If this new segment is doing a lot of dollar volume compared with the old segments, this could significantly affect sales and make add-on purchases look soft - even though they are in fact getting purchased!

At this moment, the head of business development appears in the door with another person who turns out to be new Salesperson 4. The CFO looks up and the head of biz dev, somewhat sheepishly, introduces the new salesperson.

"Glad to meet you", the CFO says. "By the way, can you tell me something? Do the customers in your new segment purchase and install our add-ons in the order we suggest in our operations manuals?"

"No, they don't" said Salesperson # 4. They install them in a different order, because they are having some difficulty installing a couple of the add-ons, and usually delay those to the end of the purchase cycle when they have more experience with the applications. Is there something we can do about that?"

The CFO just smiles, and thinks:

Looks like I just found the money to pay for unfreezing some hiring in IT.

"I think so", the CFO tells new Salesperson #4, calculating the improvement in cash flow on the fly if these add-ons were installed faster. "I really do think so".

Chapter 6 Turning Latency Data into Profits

Customer LifeCycles are a reality: there is going to be a LifeCycle and you will not be able to stop it. You probably don't know about LifeCycles because you have not measured them. You don't even hear many pundits talking about them. This is most amusing given all the jaw flapping and tongue wagging about LifeTime Value; if you don't understand the customer LifeCycle, how would you ever know when the "LifeTime" was over to measure value? The plain fact is people have it backwards; LifeTime Value is the last thing you want to try to wrestle with when just starting out with customer relationship and value management. You start with the LifeCycle, and only after fully playing out that card, do you move on to the idea of LifeTime Value. You do not have to mess around with calculating absolute customer LifeTime Value to be successful using data-driven marketing. Only after you have nailed down the basics of data-driven marketing do you need to go there; you will learn all about Lifetime Value later on in this book. What you need to understand first is the customer LifeCycle, and how to use knowledge of it to your advantage.

Customers are not just customers one day and then not the next day; there is a process to customer defection, and the smart data-driven marketer creates High ROI Customer Marketing programs by taking advantage of understanding the complete customer defection process.

There are two ways you can increase the value of customers:

- Extend the customer LifeCycle, leaving more time for the customer to increase in value, by increasing the time the customer takes to defect.
- Increase the value of the customer within the existing LifeCycle. The customer still defects pretty much on schedule, but you have done everything you can to increase their value before the defection.

The first approach usually requires some pretty sophisticated tools and can be expensive; loyalty programs are a classic example of extending the LifeCycle. Not for the faint of heart financially and organizationally, loyalty programs also do not work well for every type of business. But they do work and can be extremely profitable if they are designed and executed correctly. If you are

interested in how this type of loyalty program is constructed, visit my website at **http://www.jimnovo.com/download.htm** and download the loyalty case study.

The second approach to increasing customer value above is easier to execute, and for many companies, is the right way to go. It involves what I would call a customer retention or anti-defection program as opposed to a loyalty program, and this is how you go about setting it up.

Recall this table from a Chapter 3:

Time between 1st - 2nd event:90 daysTime between 2nd - 3rd event:60 daysTime between 3rd - 4th event:30 daysTime between 4th - 5th event:60 daysTime between 5th - 6th event:90 daysTime between 6th - 7th event:120 daysTime between 7th - 8th event:150 days

The first place I would look to address the above customer LifeCycle is the fourth event. Why? This event looks to be the one that is "low hanging fruit", since the average customer is accelerating into it, meaning the response rates should be quite high. In other words, we are taking advantage of the natural behavior customers have demonstrated, rather than trying to force them to do something out of the ordinary.

For the average customer, this fourth event happens at 180 days after the first event. How do I know? Just sum the first 3 lines of the table above: 90 days + 60 days + 30 days = 180 days. Any customer who is 180 days old and has not yet made a 4th purchase, a 4th visit to the web site - whatever the event is you are tracking - is acting outside the behavior of the average customer and is a prime candidate for an earlier than normal defection. This is where you focus your efforts. You set up this fourth event as the "trip wire" - if the customer doesn't trip the wire by engaging in the 4th event by day 180, you take action and try to affect this behavior. If you can save just a small percentage of defecting customers, the ROI can be very high, because these customers represent "found profits" which would not have existed without your efforts. And yes, you can measure these found profits - I am going to show you how to do this below.

This may not be the highest short-term ROI promotion we can do, but in terms of reducing customer defection and extending the LifeCycle, it is probably the highest long-term ROI promotion we can do, because we are helping "slow customers" accelerate into that 4th purchase. We have a reasonable expectation, based on looking at average customer behavior, that a certain percentage of customers will do this and continue on into the 5th and 6th events. We are choosing a specific group of customers at a specific time in their LifeCycle to promote to, a group with the highest likelihood of success.

Why concentrate on these defecting customers? The **two fundamental rules of High ROI Customer Marketing:**

- 1. Don't spend until you have to, and
- 2. When you spend, spend at the point of maximum impact

You don't have to spend on customers who make the fourth purchase or visit within 180 days, because they are acting like "average" customers. Why spend on them if everything there is OK and they are behaving normally? You want to concentrate your spending where it will have maximum impact - on the customers who "roll over" the 180-day barrier without engaging in "average" behavior. These customers are the most likely candidates for a complete defection, and by focusing your resources laser-like on these people, you can spend more per customer and really have some impact.

Put another way, let's say you have a customer retention budget of \$20,000 and you have 20,000 customers. You currently spend \$1 per customer each year sending all your customers the same lame retention stuff - statement stuffers that say you care and so forth. But if you could tell which 5,000 customers were the most likely to defect, and only spent on them at the point of maximum impact - when the defection was taking place - you could spend \$4 per customer trying to stop or slow the defection with the same budget, have a much higher success rate, and actually realize the "found profits" I spoke of earlier. Make sense?

How To Execute a Latency-based Promotion

We'll use a retail example because the numbers are easiest to understand and convey. But the same thought process is valid for any kind of business.

1. Determine the timing of your promotion. You normally want to take action as close to the "trip wire" event as is reasonable and practical,

taking into consideration the cost. If you have a ton of customers, there may be enough customers rolling over the "180-day with no 4th purchase" barrier to execute your promotion every week; if not, then gather up enough customers to execute efficiently. Some may be anywhere from 180 - 210 days old with no 4th purchase. That's fine; but don't let them get more than 30 days past the trip wire without taking action.

- 2. Create the offer. In a retailing business, this could be as simple as a discount of some kind. You could sub-divide the 180 day old / no 4th purchase customers into "best" and "other", creating a VIP service offer to best customers and a discount offer to other customers.
- 3. Prepare the list. Select all your 180 day / no 4th purchase customers, and then randomly select 10% of them to not contact. This is called your control group. People will tell you to only use 2% or 3% as control, and statistically they **could** be right about this. But the first time out of the box, I like to go with 10%, for two reasons:
 - a. It's a "no argument" control group size. If your effort works and you can prove it, there won't be chattering from the sidelines about the possibility of a "defective" control group.
 - b. Why spend more than you have to the first time? By taking a large control, you reduce the number of people you are spending on to execute your promotion.

If you created the two groups "best" and "other", you need to take a 10% random sample of each. The other 90% of a group is called the test group; they are the ones who will receive the promotion by direct mail, e-mail, or other means. The creation of proper control groups is absolutely essential to measuring the "found profits" referred to above. If this step has you puzzled, you will read more details on creating control groups and random samples later on in Chapter 19 or see http://www.jimnovo.com/Random-Sample.htm for more.

4. Now you have two lists of people, control and test. Set up your tracking capability, which at minimum is the ability to run a report every 30 days that reveals the sales of each group starting from the

beginning of the promotion, which is when you execute the e-mail, snail mail, or other communication of your offer to the test group. The metric you are interested in here is revenue per customer, so you would take the total sales of each group from the time the promotion is delivered and divide by the number of customers in the group, for both control and test groups.

- 5. Deliver your promotion to the test group.
- 6. Monitor the revenue activity of test and control groups. Run a sales report weekly or every 30 days, and look for divergence in the revenue per customer. The customers in the test group should be registering a higher sales per customer level (you hope). Keep running the report until the increase in revenue between test and control remains stable or begins to fall. When this happens, the LifeCycle of the promotion is over (promotions have LifeCycles too!). Let's say this takes 90 days, so 90 days after the event, you have a revenue per customer number for activity during the promotion, for both the control and test groups.
- 7. Calculate ROI. I'll use some plug numbers as an example. The idea here is to compare the revenue behavior of the test group with the control group, and determine how much additional revenue occurred because of your promotion. Since the control group experienced no promotion, any difference in revenue between test and control can logically be attributed to the promotion. We then take out costs, and see if we added value to the customer LifeCycle in more mercenary terms, did we make money or not?

| | Control | Test |
|--------------------------------------|---------|--------|
| 90 day Revenue per Customer | \$100 | \$110 |
| Gross Margin @ 30% | \$30 | \$33 |
| Additional Margin Due to Promo | | \$3 |
| Per Customer Cost of Promo | | \$.50 |
| Additional Gross Margin per Customer | | \$2.50 |

180 Day / No 4th Purchase Promotion

Here's the key to the above. The people in control generated \$30 in Gross Margin per customer over 90 days; the people in test generated \$33 per customer. So \$3 in additional Gross margin per customer was created because of your promotion, since the two groups are the same in all other ways (if control was truly a random sample).

This \$3 nets down to \$2.50 because the cost of doing the promotion was \$.50 per customer. Note: Nowhere in here are we talking about response rates. Response Rate doesn't matter in the measurement of profitability (it matters a lot in other cases); what matters is actual buying behavior. When you use control groups, you pick up buying behavior you never could have measured by just looking at response rates.

Now, the Per Customer Cost of Event is usually where you get into some arguments. If the event included a discount, the per customer cost of this discount must be included in the calculation:

| Discount | \$5 |
|------------------------------------|---------|
| Number Used | 500 |
| Total Discount | \$2,500 |
| Number of Customers | 5,000 |
| Per Customer Discount | \$.50 |
| Gross Margin / Customer from Above | \$2.50 |
| Gross Margin / Customer - Discount | \$2.00 |

Also, in the strictest sense, there is probably additional overhead attributable to the additional revenue: the cost to take a call and ship the box, the cost of additional salespeople needed to cover the promotion, and so on. These costs would not exist if you had not executed your promotion, so they should be included in the calculation to the extent you can calculate these additional overhead costs.

| Cost of sales people for Promo | \$2,000 |
|--------------------------------------|---------|
| Number of Customers in Promo | 5,000 |
| Per Customer Cost of Salespeople | \$.40 |
| Gross Profit per Customer from Above | \$2.00 |
| Net per Customer Value - Sales Cost | \$1.60 |

This \$1.60 is profit after all expenses have been paid back. You have added \$1.60 in value to the LifeCycle (and LifeTime Value) of the average customer in the promotion. To get to ROI, we need to look at what the promo cost, and compare this to the value we generated; this is the definition of ROI. How much did we invest, and how much did we get back? We know what we got back \$1.60 per customer Net of all costs, so we need to calculate total costs:

(From above)

| Per Customer Cost of Promotion | \$.50 |
|----------------------------------|--------|
| Per Customer Discount | \$.50 |
| Per Customer Cost of Salespeople | \$.40 |
| Per Customer Total Cost | \$1.40 |
| "All Expenses In" 90-Day ROI | 114% |

Note: 1.60 / 1.40 = 114%

You spent \$1.40 and you generated \$1.60 after all costs. It's a 90-Day ROI because the additional revenue generated was measured over a 90-Day period. A 114% return is not something the CFO is going to be against, trust me. In fact, you could make the argument that since ROI in financial circles is usually measured on an annual basis, and this is a 90-day ROI, the real ROI here is 4x the 90-day ROI, or 456% on an "annualized basis".

These are the found profits you have generated from your effort. By comparing the test group with the control group, you have proven these profits would not exist without your 180-day trip wire promotion. A smaller percentage of customers in the test group defected when compared with the control group; at least some portion of test made a purchase, and some kept right on buying for at least 90 days. These are found profits that would not have existed without your effort. You have proven the 180 day / no 4th purchase trip wire promotion added value to the customer LifeCycle, a total of \$1.60 per customer x 5000 customers = \$8000 to be specific, and you did this without costing the company a single dime, since you paid back all your costs with profit from the promotion, and still had \$8000 left over to put in the bank.

I can hear you now. C'mon Jim, looks good on paper, but 485% annualized ROI? An \$8000 profit on a promotion that with every cost imaginable thrown in costs \$7000? How is that remotely possible?

Folks, it's not just "possible", this kind of return is **normal** in LifeCycle-based promotions. Remember the two rules of High ROI Customer Marketing:

- 1. Don't spend until you have to, and
- 2. When you spend, spend at the point of maximum impact

By focusing your resources squarely on the problem, each dollar you spend works much harder. By waiting for the trip wire you narrowed the population you were promoting to, weeding out people you would normally waste money on. And by acting when the wire was tripped, you spent at the point of maximum impact.

Here is why this type of promotion makes so much money. It's anti-defection. You literally kept customers from leaving the company, and the control group proves this. The people you did not promote to in the control continued to slip away, while some portion of folks in the test group were stopped and their behavior reversed. This is where the huge returns come from - it's the relative spending disparity between the groups that creates the "found profits", which would have slipped away had you not done the promotion. It's a "tipping point" kind of idea - if you can be in the right place at the right time with the right catalyst, it doesn't take much change to create a big impact on the scene.

This promotion was not designed to extend the customer LifeCycle, but to add value to the LifeCycle. Did it in fact actually extend the LifeCycle, and how would you measure this effect? All the customers in both the test (received promotion) and control (did not receive promotion) groups were 3x buyers who failed to make a 4th purchase by 180 days after their first purchase. This was the Latency "trip wire" selected to trigger the promotion.

So let's look at tracking these two groups for another 90 days, and look at continuing purchase activity using what I call the Hurdle Rate method.

A Hurdle Rate is simply the percentage of customers in a group who have "at least" a certain amount of activity. You define the behavior hurdle they have to reach, and measure the percentage of customers who have achieved this "threshold" (rate). If you track these percentages over time, you can use them to compare the actual and potential value of customer groups as a whole.

At the point of the promotion, 0% of both groups had made a 4th purchase. Recall we measured the profitability of the promotion over a 90-day period after we sent it to the test and control customer groups.

To track the Hurdle Rates for each group, we ask, "What percent had made at least 1 more purchase at 30 days, at 60 days, and at 90 days after the 90-day promotion was over, in both the test and control groups?" We know some percentage of both groups made a purchase during the promotion, because there were revenues generated in both groups. We made a profit in the first 90 days because the revenues were much higher for the test than control group. So at the beginning of this "post promotion" tracking, we see 1% of control and 3% of test have made 4 or more purchases. For the following 90 days after the promotion was over, data might look like this:

| % 4 or more purchases | Control | Test |
|-----------------------------|---------|------|
| End of 90-day Promotion | 1% | 3% |
| 30 Days After Promotion End | 1% | 5% |
| 60 Days After Promotion End | 2% | 8% |
| 90 Days After Promotion End | 2% | 10% |

Realize this: we have already made money on this promotion, a 114% ROI. We have already added value to the LifeCycle, increasing LifeTime Value - no matter how long a "LifeTime" is (does it really matter, as long as you are making increased profits?)

But as you can see from the chart above, we also extended the LifeCycle itself, because the percentage of customers exceeding the "4 or greater Hurdle" in the test group is far higher than the percentage of customers over the same Hurdle in control, and it appears to be growing over time.

There is a group of customers in the test group who just keep on keeping on - and this percentage (10% at 90 days after Promotion End) is much higher than both the initial group who responded to the promotion and made a 4th purchase (3%) and the test group. What's going on with that?

It's called the Halo Effect. It represents customer activity stimulated by the promotion not occurring within the promotional period. Now, we don't know exactly where it's coming from, and we can't show any measure of profit from it

(we defined our promotion period as 90 days), but it is clearly there, plain as the nose on your face.

Recall when describing the original promotion, I stated, "Response Rate doesn't matter in the measurement of profitability (it matters a lot in other cases). When you use control groups, you pick up buying behavior you never could have measured by just looking at response".

This "buying behavior you never could have measured" is the Halo Effect, working its magic during the promotion. People you have no way to track will respond to the promotion. They want to make a purchase but forget the coupon, for example. So they go ahead and make the purchase anyway - because the promotion "woke them up" to a need for something you sell.

After the promotion is over, the same thing continues. It's the Halo Effect again, working after the promotion. For example, people think about participating in the promotion but wait too long. They've missed it. But they're now in a new state of awareness about your company because of the promotion, and as a result, are more likely to make a purchase given any random positive stimulus. Perhaps some product appears on a TV show. Maybe a competitor promoted a product to them, the customer remembers you sell it also, and prefers your store. It doesn't really matter. Fact is fact, and because of your promotion, you extended the customer LifeCycle. You created a situation where people became more likely to purchase from your company in the future.

Not bad for a beginner. In the first 90 days, your promotion created present value - real bottom line, measurable ROI - adding Value to the customer LifeCycle (LifeTime Value). In the 2nd 90 days, your promotion created future value - accelerated repeat purchase rates - by extending the LifeCycle. CFO sings your praises! At last, somebody who can prove they are making more money than they are spending with marketing!

There is an important lesson here: you will never know how much money promotions really make without using control groups.

Recency Metric Toolkit

Chapter 7 Customer Value-Based Marketing

Over the past five decades, a lot of research and testing has been carried out concerning the profiling of customer behavior based on transactional data. The appearance of computers and "data-mining" has allowed even more extensive studies to be carried out.

The end result? If you had to pick one variable to predict the likelihood of a customer to repeat an action, Recency, or the number of days that have gone by since a customer completed an action (purchase, log-in, download, etc.) is the most powerful predictor of the customer repeating this action.

As each day goes by after the customer completed the action, the customer gets less and less likely to repeat it. Plain and simple. You can run all the fancy data-mining scenarios on "likelihood to buy" or "likelihood to visit" you want to - Recency always comes up as the most important variable in predicting the likelihood of a customer to repeat an action.

Recency is the number one most powerful predictor of future behavior. The more recently a customer has done something, the more likely they are to do it again. Recency can predict the likelihood of purchases, log-ins, game plays, just about any "action-oriented" customer behavior. Recency is why you receive another catalog from the same company shortly after you make your first purchase from them. They know you are most likely to order again immediately after your first order. Recency is the most powerful predictor of future behavior.

Think about this. Latency, or Trip Wire metrics tell you when something bad or good has already happened. Using Recency, you can predict when something bad or good will happen. There is a huge difference between these concepts, my fellow Driller.

It should not surprise you that Recency is also the most powerful predictor of a customer to respond to a promotion - after all, the more likely a customer is to repeat an action, the more likely they are to respond to a promotion asking for

this action (purchase, log-in, download, etc.). If a Recent customer is more likely to repeat an action, and is more responsive to promotions for this action, it follows the more Recent a customer is, the higher their potential value, because Recent customers are the most likely to contribute to profits in the future by responding to your promotions (or simply just coming back by themselves).

Customers who are more Recent have higher potential value than customers who are less Recent, for any given activity. Customers who made a purchase 15 days ago have higher potential value than customers who made a purchase 60 days ago. Customers who logged in last week are much more likely to visit than customers who logged in 30 days ago, and have higher potential value.

Make sense? But how is Recency implemented, how do you actually do anything with this information? Glad you asked.

Chapter 8 The Ad Spending Example

Let's use Recency to compare the potential value of customers coming from two different ads (Ad #1 and Ad # 2) that ran at the same time, for the same duration. The following example uses a spreadsheet, but if you know your way around databases and can query your customer records, have at it your way.

- Identify the groups you want to compare for potential value. In this example, it's the customers who clicked on either of two ads, Ad #1 or Ad #2 (two groups). If you are not keeping the source of customers in your database, start doing it right now it is one of the most important variables you can analyze.
- Decide which activity is most important to you for these groups. If you're a publisher, probably log-ins or page views are most important. If you were selling merchandise, you would use purchases. For this example, we will use purchases. An example using visits (or log-ins, if you don't track visits) is below.
- 3. Import all the purchase records of people who clicked on Ad #1 or Ad #2 into separate spreadsheets. These transactions need to have a date; most interactive activities are date-stamped so this should not be a problem. If an activity you want to profile for potential value has no date stamp, start collecting the dates of activity.
- 4. Pick a time frame to look at Recency. For page views, it might be 1 week; for purchases, maybe 30 days. The exact length is not very critical, because you are interested in comparing the activity between the Ad #1 and Ad #2 groups you want to know which is "better". As long as you use the same time frame for both groups, you are fine. Pick something reasonable based on what you know about your customers. Anywhere from 30 to 90 days would be reasonable for purchases; let's use 30 days.
- 5. Sort the purchase records for Ad #1 from most Recent to least Recent and find out what percentage of the people who clicked on Ad #1 and made a purchase have made at least one more purchase in the past 30

days. Count back 30 days using the transaction dates, total the number of customers making a purchase, and divide by the total people in the spreadsheet. Perhaps it is 20%. Note: The software that comes with the book will automatically aggregate multiple transactions by customer and sort customers by their most Recent transaction for you.

- 6. Run the same analysis for people who clicked on Ad #2 and made a purchase. Let's say only 15% of these people have made at least one purchase in the past 30 days.
- 7. You're done, and you know the answer. A higher percentage of people who clicked on Ad #1 are Recent active and purchasing when compared with Ad #2. This means Ad #1 generates customers with higher potential value. You need to take this into account when analyzing the success of the ads.

Do you understand how powerful this idea is?

If you go through this process for customers grouped by product they bought first, you can determine which products generate new customers with highest potential value. Go through this process for customers grouped by which area of the site they visit most, and you will find which areas generate highest potential value customers. If you go through this process for customers grouped by the demographics or the survey data they provide, you can determine which data points define customers with the highest potential value. All you have to do is create your groups and compare their Recency. The group with the highest percentage engaging in the activity you are measuring over some time period is the group with the highest future value to the company.

This is a simple example of how companies with experience in managing remote shopping customers find ways to maximize sales and minimize expense. The customers, through their actions, tell them which route is the most profitable to take. The most Recent customers for any particular activity are always the ones most likely to repeat that activity, and so have a higher potential value.

You can track multiple activities for the same customer groups. In the first example, you found customers who clicked on Ad #1 and made a purchase are more Recent on purchases, so they have a higher potential value on the activity "purchases". But what about the Recency of people who clicked on the ads for

visits? If they keep coming back, they could be of some future value. Let's see how this Recency study might look.

Visits / log-ins example:

- 1. Import all the visits (or log-ins if you don't track visits) into two separate spreadsheets of people who clicked on Ad #1 or Ad #2 (transactions need a date stamp).
- 2. Pick a Recency cut-off. Again, we are interested in a comparison, so the number isn't critical. Let's use 1 week.
- 3. Sort each spreadsheet from most Recent to least Recent and find out what percentage of the people who clicked on Ad #1 have visited (logged-in) at least once in the past week, as was done above for purchases. You might come up with 10%.
- 4. Run the same analysis for people who clicked on Ad #2. Sort most Recent to least Recent, and do your percentage. You might come up with 30% who have visited / logged-in at least once in the past week.
- 5. You're done, and now you have an interesting situation. It appears the customers who clicked on Ad #1 have a higher potential value on purchases, but people in general who clicked on Ad #2 have a higher potential value on visits. Maybe they're just tire kickers, or maybe they're doing research. We'll take a closer look at finding answers to this situation in a minute.

Note that this method is based on the actual facts of customer behavior - not speculation or "best guess" theories. The behavior of the customer is the most accurate yardstick you will find for assessing potential value. Once you complete studies like these, you can begin to organize all your business practices around the potential value of the customers they generate. If you allocate money away from activities generating low potential value customers, and allocate this money to activities generating higher potential value customers, you will become more profitable over time. It's really as simple as that.

At the beginning of the previous example I specified the ads you were comparing should have "run at the same time, for the same duration". Do you know why? Customer LifeCycles. It's not fair to compare the customer

46

Recency percentage of an ad that ran 90 days ago with an ad that ran 30 days ago, because customers tend to leave you over time. You learned this in the previous Latency model. If one ad has more time to "lose customers", then comparing them would be unfair or biased by the element of time.

This tendency of customers to leave over time has different names depending on the business model - some call it attrition (credit cards), it is an element in churn (cable, long distance, wireless), and in retailing and database marketing it is often called defection. You can compare ads having different start points, as long as they're not so far apart that seasonality comes into play (comparing ads that ran in July with those that ran during November, for example). Simply synch up the LifeCycles and do a Recency analysis at the same point in the customer LifeCycle. This is usually defined as "days from an event", for example, the day they became a new buyer or visitor.

If you want to compare the potential value of an ad running 30 days ago with one running 90 days ago, you have to look at the Recency of the 90 day ago ad 30 days after it ran to take out the LifeCycle effects. Using the previous example, if Ad #2 ran 90 days ago, you would want to find out what percentage of people who clicked on Ad #2 took action 30 days after it first ran. If you always run your analysis referencing the start date of an ad (or any other variable you are measuring), and measure for equal time periods after the start date, you will eliminate most of the LifeCycle effect and can compare the results on an equal basis.

So what about these LifeCycles? Is there a way this information can be tracked and used? Sure; using LifeCycles can solve the little problem we left at the end of the previous example. Thinking about the Ad #1 and #2 example, what if you repeated the Recency query for each ad (made at least one purchase in the past 30 days) every 30 days for 6 months? What would you get? You would have a series of measurements looking at the potential value of the customers generated by the ads over time. You would be able to chart the defection patterns of Ad #1 and Ad #2 customers.

Why is this important? Because if you want to get at the true value of the customers generated by the ads, you have to measure their value over the LifeCycle. You might be surprised. Take a look at the chart below from our previous example:

Recency Percentage Over Time



Customer LifeCycles: Ad #1 and Ad #2

Ad #1 Customers Start with Higher Recency, but have a Shorter Lifecycle than Ad #2 Customers

Top to bottom on the left side of the chart is the percentage of customers making a purchase in the past 30 days; left to right at the bottom of the chart are the months each Recency analysis was performed since the start date of the ad campaigns. Both the Ad #1 (dotted line) and Ad #2 (solid line) start at the percentages we came up with in the Recency of purchases analysis above. If you look at the chart above, you can see that Ad #1 (dotted line) starts at 20% of customers having purchased in the past 30 days, and after6 months, the

percentage drops to less than 5%. Ad #1 also seems to be headed even lower in Recency; these customers are losing even more potential value as time goes on. Ad #2 (solid line) starts at 15% of customers having purchased in the past 30 days and falls into month 3, but then starts rising in later months, ending up higher than it started, and is still rising. Customers from Ad #2 might end up having greater potential value than customers from Ad #1 over the longer term.

48

So it could be that, after looking at the LifeCycle of customers from Ads #1 and #2, you may find even though Ad #1 looks best based on Recency at a point in time, Ad #2 creates higher potential value customers when Recency is looked at over time. In looking at the LifeCycles, we have perhaps come up with a clue to the behavior we saw in the Recency of visits analysis on the previous page. It would appear that the customers from Ad #2 might take a little longer to make a purchasing decision, but become more valuable customers over the long run. This is a very common occurrence in customer behavior mapping and if you are not tracking it, you won't know it is happening, leading to poor decisions about the profitability of your ad campaigns.

This is a picture of the customer LifeCycle at work, and you can conduct this type of study with ads, products, areas of the site, survey data, demographics - any type of customer information you can get Recency data for.

Let's talk about one of the most confusing and misunderstood parts of customer marketing, LifeTime Value. The LifeTime Value of a customer is the net profit the customer generates over their LifeCycle. People tell you not to spend more to get a customer than their LifeTime Value, or you will lose money. This is true on the face of it, but actually figuring out what the LifeTime Value of a customer is can be a difficult task, especially if you don't have the right tools. Besides, what if you are a new company, or have never tracked the data you need to calculate LifeTime Value? Is the concept useless to you?

Not at all. LifeTime value is used to make decisions about allocating marketing to ideas that generate high potential value customers, and away from ideas generating low potential value customers. And to do this, all you need to know is the relative LifeTime Value of the customers generated by each idea. Recall LifeTime Value is the net profit the customer generates over their LifeCycle. So if you know what the LifeCycles look like, you should be able to do a pretty good job of determining who the highest LifeTime Value customers are, relative to each other. If you do your Recency tracking on ads, PPC keywords, newsletter links, and so on, you should be able to compare the relative potential value of the customers generated by each approach and easily decide where your ad budget is most profitably spent.

Let's say you have run 20 campaigns and you know your cost per new customer from each of them. You want to run the top 10 (lowest cost per new customer) but you only have the money for 5 campaigns. With Recency and LifeCycle tracking on the 10 campaigns, all you have to do is choose the top 5 campaigns generating customers with the highest potential value based on Recency. If you allocate your budget to those and away from the bottom five, you are maximizing your budget ROI, regardless of the actual LifeTime Value in dollars of the customers generated. What else could anybody ask for?

Continuing with our Ad #1 and Ad #2 example, based on the LifeCycle chart you just saw, can you make a judgment about which ad generates customers with higher potential value? Looks like Ad #2 to me. Ad #2 appears to generate customers with a longer LifeCycle, so their relative LifeTime Value is higher when compared with Ad #1 customers, given the costs of acquiring and maintaining customers from both ads is roughly the same. Period.

And by the way, with the LifeCycle information in hand, is cost per new customer really the issue? Probably not, because you have to weigh the cost per new customer against the length of the LifeCycle. Customers who are the cheapest to acquire may have the shortest LifeCycles, and customers who are expensive to acquire might have very long LifeCycles. So you really need the potential value and LifeCycle tracking to get the whole picture.

The problem people run into with LifeTime Value is the whole question of determining a LifeTime. There's no easy way to do it, and so the whole idea gets tossed. People get frustrated because there's nothing to grab on to, and no easy way to make comparisons. But when you track the LifeCycle, you know for a fact one group has a longer LifeCycle than the other. Who needs the absolute LifeTime Value number in dollars and cents? As long as you allocate money towards higher potential value customers and away from lower potential value customers, you are maximizing your resources in everything you do. And that is the reason people want to look at LifeTime Value in the first place.

If you really need a hard number, don't be afraid to call an end to the customer LifeTime. They are much shorter than you think. When you are tracking your LifeCycles, and they start to approach 0% of customers making a purchase in the past 30 days (or whatever standard you're using), the LifeTime generated by these particular ads is over. Don't hope customers will magically come back; it usually doesn't work like that. Once you call the end to the LifeTime, subtract your costs (cost of products sold, ad costs, an allocation for service costs) from the revenues for both Ad #1 and Ad #2 customers, and you'll have your LifeTime Value. We'll talk more about Lifetime Value calculations later on.

Chapter 9 Turning Recency Data into Profits

So what do the financials look like on the ad campaign from the previous chapter? ROI, or Return on Investment, is a concept from the financial world frequently applied to database or customer marketing. What people seem to forget is ROI implies the concept of time, because "Return" happens over time. So an ROI calculation really asks, "What was my return over time?" And without looking at time, you can't calculate the "real ROI" of a campaign.

Think about a bank account. The bank says if you put money in, they will pay 2% interest. This is an annual number; if you put in \$100 and don't do anything else, you will have \$102 in your account at the end of a year. Your annual ROI is 2%. At the end of 6 months, it's close to 1%. So ROI depends on what time frame you are using for the calculation. Let's go back to our example of Ad #1 and Ad #2, and look at some ROI numbers.

Ad campaigns #1 and #2 were pretty similar. The ad units were the same and the Cost Per Thousand ads (CPM's) were the same, but they ran on different sites and had different creative. As a result, the response rates (click-throughs resulting in purchase) were different. Here's what they look like at the end of their campaign runs:

| Parameter | Ad #1 | Ad #2 |
|----------------------------|----------|----------|
| a. Cost of Campaign | \$5,000 | \$5,000 |
| b. People Clicking | 2,000 | 1,000 |
| c. Number of Buyers | 185 | 230 |
| d. Average Price | \$90 | \$70 |
| e. Product Sales | \$16,650 | \$16,100 |
| f. Product Margin | 30% | 30% |
| g. Net Margin (f x e) | \$4,995 | \$4,830 |
| h. Campaign profit (g - a) | -\$5 | -\$170 |
| i. Campaign ROI (h / a) | 0% | -3% |

This is short-term ROI, right at the end of the campaign. It's based on what customers first bought during and immediately after the campaign. Ad #1 looks

to be the clear winner, even though it only broke even; Ad #2 has a negative ROI (lost money). You might just leave it at that. But you're smarter now. You know about Recency and customer LifeCycles. Your first Recency tracking is 30 days later, and you find that 20% of the buyers from Ad #1 have made another purchase and 15% of the buyers from Ad #2 have made a purchase (these numbers are from our original example). The rest of the numbers from the LifeCycle charts are included below. Now what does the campaign ROI for Ad #1 and Ad #2 look like?

| Parameter | Ad #1 | Ad #2 |
|-----------------------------|----------|----------|
| 1 month Repeat Buyer % | 20% | 15% |
| 2 month Repeat Buyer % | 15% | 13% |
| 3 month Repeat Buyer % | 12% | 11% |
| 4 month Repeat Buyer % | 9% | 13% |
| 5 month Repeat Buyer % | 6% | 14% |
| 6 month Repeat Buyer % | 4% | 16% |
| j. Total Repeat Buyer % | 66% | 82% |
| k. Original # of Buyers | 185 | 230 |
| 1. New Purchases (j x k) | 122 | 189 |
| m. Average Price (d above) | \$90 | \$70 |
| n. Product Sales (m x l) | \$10,989 | \$13,202 |
| o. Product Margin | 30% | 30% |
| p. New Net Margin (o x n) | \$3,297 | \$3,961 |
| q. Initial Margin (g above) | \$4,995 | \$4,830 |
| r. 6 Month Margin (q + p) | \$8,292 | \$8,791 |
| s. 6 Month Profit (r - a) | \$3,292 | \$3,791 |
| t. 6 Month ROI (s / a) | 66% | 76% |

Surprised? These ad campaigns at 6 months have ROI numbers many times higher than they did at 30 days. And, Ad #2 has emerged as the winner because at 6 months, it has passed Ad #1 in profits and ROI. If you recall the LifeCycle chart, at the 6-month point, Ad #1 is going downhill fast while Ad #2 is still climbing. Looks like the spread in profits and ROI is going to get even wider still over time.

Now, is this the best and most accurate way to determine the true ROI? No, it's a "down and dirty" approximation of campaign ROI you can make on the spot, a "back of the napkin" kind of idea. You're looking for trends and comparisons in order to get a feel for what's working. But it is a heck of a lot better than just looking at conversion to first purchase.

For this example, we don't know how many of the repeat customers may have made multiple purchases in a month - we only know they made "at least one". And we're using "average price" from the original campaign; this has probably changed. The point of this example is making comparisons between the potential values of customers generated by different ad campaigns, and using these comparisons to make more money with your advertising. Once you see important trends emerging, you can decide whether it's worth spending the time and resources to take ROI down to the last penny.

Could the customer LifeCycle for Ad #2 suddenly turn down and undercut the customer LifeCycle for Ad #1? Sure, it's possible, but not likely. This late in the LifeCycle, when a group of customers is moving in a certain direction, they tend to keep moving in the same direction. Good customers tend to remain that way (until they leave you) and poor quality customers tend to remain as they are. That's one reason there is so much money wasted in customer marketing; marketers are not targeting using the LifeCycle, and as a result they're making untargeted offers at the wrong times to most customers. Early in the LifeCycle, it can be difficult to tell if a customer will become profitable or not. That's why it is so critical to track trends like this from the beginning; later on, it becomes less important as the customers tend to remain either profitable or unprofitable. There's no reason to guess though, is there? You'll be tracking these LifeCycles, because that is what smart marketers do.

Recency in Promotions

You will generally see response rates to a promotion asking for a specific action (purchase, visit, click a link) fall as the number of weeks or months since the customer last engaged in the activity you are trying to encourage rises – in other words, as the customer becomes "less Recent". This relationship is a smooth curve and quite predictable once you establish the "slope" of it for your business. Response rate by Recency might look like this:

Customer inactive for 31-60 days, Response rate = 20%Customer inactive for 61-90 days, Response rate = 10%Customer inactive for 91-120 days, Response rate = 4%Customer inactive for 120+ days, Response rate = 1%

The absolute response rates will be different depending on the business, media used, and offer, but the relative response rates will follow a decelerating curve as shown above, that is, the less Recent the customer, the more dramatic a drop in response rate you will get to your request for an action. In terms of using this information for promotions, you will find some point along the curve where you will hit "breakeven", meaning the cost of the campaign will equal the profits or benefit generated. For example, let's say you offer a discount, gift, or other incentive in your retention / lapsed customer campaign and need a response rate of at least 4% to pay back the campaign cost. This is your breakeven point. The implication for this 4% breakeven campaign contained in the Recency information above is this: don't bother to promote to any customer who hasn't engaged in the activity you are trying to encourage for over 3 months, because you're wasting your money. Response will be too low to pay back the cost of the campaign with any customer who has been inactive for over 3 months.

This Recency effect is very stable over time, allowing you to predict in advance what response to a campaign will be, once you do an "establishing" campaign to see what your response rate is for any particular offer. Recency will predict average response rate for any specific combination of offer and media used. You can save a tremendous amount of money by forecasting your response by using Recency, and not promoting to customers unlikely to be profitable.

Let's set up and execute a Recency test. Classify customers in 30-day Recency segments by the last date of the activity you want to profile for Recency. If you want to profile purchases, customers could be segmented by date of last purchase, for example:

31 – 60 days ago 61 – 90 days ago 91 – 120 days ago 120+ days ago

Take a 10% random sample of customers from each segment (every 10th person in the segment), and send all of them a promotion with the same offer, say 20% off any purchase in the next 30 days. Look at the response rate by these 30-day

segments. You will find response falls off significantly as you look at Recency segments further back in time. If you repeat the test using the same offer to a different sample of each 30-day segment, the response rate by segment will be very close to the response rate by segment in the first test. This kind of stability allows accurate predictions of marketing ROI before promotions are even sent out to customers.

The response rate in any one of the 30-day segments above will be influenced by the value of your offer, and both response rate and cost of the offer have significant impact on the profitability of your campaign to any segment. As offer value increases, so does response rate, and so do costs. Ideally, you want to find the ideal mix of response rate and offer value creating the highest profitability for each segment you promote to.

You can use Recency to "ladder" the promotional discount, gift, or incentive value offered in a promotion, boosting overall response while cutting expenses by minimizing discount or other incentive costs.

Let's use purchases as an example, and say you usually e-mail all your customers a 10% discount when you do a promotion. If you were using a Recency ladder approach for this purchase incentive, you might apply your discount strategy this way:

Customer inactive for 31-60 days, Response rate = 20%, discount = 5%Customer inactive for 61-90 days, Response rate = 10%, discount = 10%Customer inactive for 91-120 days, Response rate = 4%, discount = 15%Customer inactive for 120+ days, Response rate = 1%, discount = 20%

Using this approach, you are allocating the most "bang for the buck" discountwise where you need it most - the least Recent, lowest response customers, and pulling back on some discounting where you don't need it as much - the most Recent, highest response customers.

Since your most Recent customers are most likely to respond, you can back off on their discount and you reduce the cost of giving discounts to customers who "may have bought anyway without a discount". You then reallocate this discount money to where it is needed most – boosting the response rates of those much less likely to respond - the less Recent customers. Your response rates will vary depending on the offer, media used, and your business. You have to test these ladders with different combinations of offer and media to find the optimum profitability for each Recency segment. The interesting and quite useful benefit of this approach is the "automatic" overall customer retention effect discount ladders have.

Using a ladder of this type means your promotional discount budget is automatically working harder and harder to keep a customer active with you as they drift further and further away from you. The less Recent a customer is, the less likely they are to buy or visit again, and by using a discount ladder you are counteracting the customer LifeCycle (the tendency of customers to leave you over time) with stronger discounts as the defecting customer behavior plays out. If a most Recent customer does not respond to the 5% offer, as they get less Recent, they automatically get offers rising in value, and at some point, many will take advantage of an offer. The customers who run through this system without taking any offers were likely lost to you as a customer already, and not worth the extra expense to try and keep promoting to them. Let's set up and execute a discount ladder test.

Pick any one of the segments from your Recency test above and now test discount level for the segment. Let's say you used a 20% discount in the first test. Pick a segment (say 91 - 120 days), and create a 20% random sample of the segment (every 5th customer) divided into 4 equal test groups. Send each test group a different discount - say 5%, 10%, 15% and 20%. Look at your response rates and calculate the profitability for the 91 - 120 day segment at each discount level. You will find your result looks similar to the following:

| Customer Sample | 1000 | 1000 | 1000 | 1000 |
|---------------------------------|---------|---------|-------------|---------|
| Discount Offer | 5% | 10% | 15% | 20% |
| Response Rate | 2% | 4% | 6% | 8% |
| Responders | 20 | 40 | 60 | 80 |
| Average Price | \$80 | \$80 | \$80 | \$80 |
| Totals Sales | \$1,600 | \$3,200 | \$4,800 | \$6,400 |
| Gross Margin | 30% | 30% | 30% | 30% |
| Gross Profit | \$480 | \$960 | \$1,440 | \$1,920 |
| Discount Cost | \$80 | \$320 | \$720 | \$1,280 |
| Net Profit before Media Expense | \$400 | \$640 | \$720 | \$640 |

As you can see, the most profitable offer to the 91 - 120 day Recency segment is 15% off. If you offer 20%, you get a higher response rate but lower profits;

any offer under 15% significantly diminishes response rate. Repeat this test for each Recency segment, and you will find the most profitable discount rising as the customer becomes less Recent, creating your discount "ladder".

When you implement your promotions based on a Recency / Discount ladder, as customers become less Recent and therefore less likely to respond to a promotion, they will be automatically offered a higher discount – one that maximizes profit for each Recency segment the customer passes through. Discount ladders create in effect a "lights-out" customer retention program suitable for automation.

There is a subtle but important side benefit to using a Recency / Discount ladder approach to manage e-mail efforts. Instead of blasting out indiscriminate offers to the whole customer base, taking a ladder approach more closely matches the offer value to the "attitude" or point in the LifeCycle a customer has reached. Following the mantra of Permission Marketing, this is called being "relevant", and will tend to increase open rate and response as customers begin to put a higher value on your e-mail relative to other offers they may get.

In addition, as e-mail clutter and execution expense increase, response will fall and profits will decrease, as customers get tired of receiving multiple promotions. Over time, you will find it is simply more profitable to e-mail customers less often, because you know for a fact the most profitable offer to make and when to make it, based on the Recency / Discount Ladder. Using this approach will generally help you rise above the clutter by sending fewer, higher impact promotions. The Recency / Discount ladder approach to creating a customer retention program is clean, simple, and easy to implement. And if you don't have any formal customer retention program in place, much better than what you're using now!

"Now hold on just a minute, Jim," you say. "Recency is a very cool concept, but I can think of some specific instances where it can't possibly work. A person who just filed a tax return 30 days ago is not more likely to file one than a person who filed one 60 days ago, and the same thing is true for people who bought a new car. Explain yourself!"

There are two issues to consider when using Recency – external forces and time frame. If there are powerful external forces shaping behavior – like the April 15th tax deadline – these forces may overcome the Recency effect. An accountant trying to manage customer relationships would probably look more

to Latency and set a trip wire: I will call best customers who don't schedule an appointment by March 15, for example. The tax deadline is simply too powerful a force and overcomes normal human behavior.

One also needs to consider Recency in light of the cycle of normal behavior. It is unrealistic to think of Recency in new car buying in terms of 30 and 60-day periods, when the normal purchase cycle may be 3 or 4 years long. It's not a rational use of the Recency metric. However, for the dealer selling the original car to the customer, as this purchase gets to be 3 or 4 years old, the longer it has been since the purchase, the less likely the customer is to make the car purchase from this same dealer.

Recency is a very powerful metric, but there are times when it simply is not appropriate to use without some adjustments. If there are powerful cycles acting on behavior, Recency often takes a back seat to Latency. Often the two concepts can be used together – there is first a Latency trip wire and then Recency kicks in. For example, up until the April 15th deadline, the accountant is really operating in the world of Latency. If customers don't call by a certain day, they are unlikely to be using the accountant for their tax return. Once the April 15 deadline passes though, the accountant is in the Land of Recency – the longer it has been since the last tax filling, the less likely it is the customer will be using the accountant next year. The accountant needs to get on the phone with these high value customers and find out what happened right away if the customer is to be recaptured.

The new car dealer is in a similar situation. Let's say the average customer trades in every four years. Up to four years after the new car purchase, the dealer is in the Latency world – there is a trip wire at 4 years, and any customer who has not purchased again at the 4-year point is in danger of being lost. After the 4-year point passes, the more time passing, the less likely it is the customer will come back – the Recency effect. As time goes by after the trip wire triggers, it becomes more and more urgent the customer be contacted and made an offer. And don't forget this: the more time that passes, the higher the offer will have to be to get the customer to come back – the Discount Ladder effect.