

Paper 286-2012

Improving Retail Decisions with Customer Analytics: Leveraging Actionable Customer Insights across the Retail Enterprise to Build Sales and Profits

Wanda Shive and Dwight Mouton, SAS Institute, Inc., Cary, NC

ABSTRACT

Customer-centric decision making is a key differentiator for retailers. Collecting customer data is a common practice. Major challenges include making customer offers more personal and delivering actionable customer insights to merchants at the point of decision for assortment, pricing, and promotions.

This paper highlights how retailers are increasing sales and profitability by using optimization to improve redemption rates of customer-specific offers and by adding the customer dimension to key merchandising decisions.

INTRODUCTION

Retailers around the world collect millions of customer data points per day. However, most retailers are not leveraging the vast insights buried within their customer databases. The terabytes of data can become overwhelming. Retailers quickly drown in customer data and gasp for actionable insights. Customer analytics identifies profitable growth opportunities that otherwise go undetected.

Three key ways to integrate customer insights into retail decisions are to focus on customer economics, stealth marketing, and customer-centric merchandising.

CUSTOMER ECONOMICS

The first step toward integrating customer insights across the enterprise requires customer data exploration and analyses. Understanding your customers' behavior over time is critical. Two major outcomes from these initial customer data analyses include customer segments and customer metrics. These outcomes provide the entire organization with both a framework for measuring results and a common customer language that can be understood and leveraged across the organization.

CUSTOMER SEGMENTATION

Customers can be grouped or segmented based on many different measures. Fit-for-purpose segmentations enable retailers to incorporate customer metrics into a variety of business decisions.

Characteristics of proven customer segmentation approaches include:

- identifiable segments that merchants understand
- focus segments that reflect corporate strategies
- actionable segments that are effective for enterprise-wide decisions

Common customer segmentation methods used by retailers include:

- Recency, Frequency, and Monetary (RFM)
- lifestyle and life stage
- price consciousness

RFM

A core financial segmentation method deployed by many retailers is known as the Recency, Frequency, and Monetary, or RFM, segmentation method. Historical purchase behavior found in the retailer's customer database is used to group customers by measuring how recently they shopped, how often they shop, how profitable they are to the retailer, and how much they spend.

RFM segmentation quickly identifies the small portion of customers who contribute the majority of the sales and profits. For most retailers, the top 20-30% of customers account for 50-70% of sales and profits.

Customer-centric retailers use this core segmentation method in tandem with typical monthly financial reporting to

Improving Retail Decisions with Customer Analytics, continued

measure performance across the enterprise.

Lifestyle and Life Stage

For this segmentation method, retailers group customers by behavioral characteristics. Many of these attributes are gleaned from the customer's shopping history, which is found in the retailer's customer database. In addition to shopping behavior within the retailer, external demographic and economic data are often incorporated to enrich this view of the customer. These external data elements include household income, education, household size, and presence of children. This segmentation method paints an identifiable picture of each customer group. Segment names like "Shoppers on a Budget", "Affluent Professionals", "Convenience", and "Busy Suburban Moms" are common among lifestyle and life stage segmentations.

Price Consciousness

Some retailers segment or group their customers by their sensitivity to price and promotions. The customer's purchase history within the retailer's customer database is analyzed to group customers by answering questions like:

- What percentage of items that were purchased was on promotion?
- What percentage of the customer's sales dollars was discounted?
- What percentage of the customer's baskets contained discounted items?

Retailers typically apply these customer groups as sub-segments within the RFM or lifestyle and life stage segments. For example, a retailer might strategically tailor customer-specific promotions to price-sensitive "Busy Suburban Moms."

CUSTOMER METRICS

After customer segments are in place, the next step toward using customer data across the enterprise is to establish key customer metrics. The measurements that answer key questions about the retailer's customers include:

- How many customers shop with us? (Number of customers or households)
- How much do they spend each time they shop? (sales per week per customer [SPW])
- How much do they buy? (items per week per customer [IPW])
- How often do customers shop? (visits per week per customer [VPW])
- How profitable are my customers? (margin per week per customer [MPW])
- How price sensitive are they? (% of promoted items purchased / total items purchased)

A retailer's primary objectives are to attract and retain customers. These key customer metrics are tracked over time and can identify trends such as when a retailer's best customers' purchasing habits are changing. Armed with customer insights, the retailer can take action to quickly reach out and reengage with those customers.

One way in which retailers use the customer segmentation and metrics outlined above are in customer-specific promotion offers. Some retailers have millions of identified customers and thousands of promotional offers per week. They strive to communicate customer-specific offers that are most relevant to each customer. The sheer number of potential customer-specific offer combinations and the complexity of their business rules exceed the capabilities of most existing systems. The goal is to communicate promotions that are most relevant to each customer based on insights derived from the retailer's customer database, while considering the profitability to the retailer.

STEALTH MARKETING: COUPON OPTIMIZATION

Whether online, through mail, or at the register, there is a science to making sure that a coupon is relevant to each customer. Many retailers or their marketing service providers (MSPs) use segmentation, modeling, and prioritization to select coupons. Only a few use mathematical optimization, but they are getting huge returns in both efficiency gains and profitability.

Interestingly, there are two very different goals depending on the organization. Retailers want to maximize redemption to bring in customers. Coupon MSPs want to maximize the number of coupons. On top of these goals are contracts with the manufactures, budget constraints, and customer preferences.

In this section, we will look at two different cases where optimization is used to solve very different goals and to account for conflicting business rules.

WHAT IS COUPON OPTIMIZATION?

Even before we define Coupon Optimization, it would be prudent to define the type of optimization that we will discuss in this section. By “optimization” we mean using mathematical techniques to find the best possible solution, given certain business rules or constraints. The mathematics in this case is under the umbrella of Operations Research. The mathematics goes beyond using sorting techniques and selection criteria, although those have their own place and can be effective, depending on the circumstances.

In coupon optimization, the mathematics is used to maximize a goal like revenue, redemption rate, or number of printed coupons. Each of these goals leads to a different outcome, and the choice of which one to use depends on whether the company is optimizing as a retailer or as a marketing service provider. Regardless of the goal, the business rules or constraints that are used in coupon optimization are very similar across the different use cases. In the next two sections, we will detail the typical scenarios that we find when marketers want to optimize their coupons.

RETAILER COUPONS

Over the years, several coupon optimization scenarios have emerged, but generally they fall into two categories for retailers. The first is direct marketing, in which the retailer selects coupons in a sheet or booklet format and sends them to customers through direct channels, primarily printed mail. The second is coupons printed at the point of sale or at the cash register after customers have paid for their purchases. This second category is similar to how manufactures and MSPs operate, so we will focus on the direct marketing type in this section.

The goal for the retailer is to get the customer to redeem the coupon and hopefully buy other products while they are in the store. From an optimization standpoint, this means that the objective is to maximize response, or, if the right data exists, to maximize the expected revenue or profit that results from the redemption of the coupon (including other purchases they might make). These values are represented in the optimization as model scores. The optimization software uses them, in conjunction with the costs of the offer and other constraints, to select the best set of coupons for each customer.

Most of the optimization rules or constraints are fairly generic across retailers. Retailers have budgets that they have to comply with and contracts with manufactures. These contracts with manufactures are what can often make an optimization scenario difficult to solve because they are minimum constraints, such as “a minimum of 2,000 Brand X soda coupons per week”. If there are only 3,000 customers eligible for those coupons, then the constraint becomes tight and therefore harder to solve. The next level of constraints contains those within and between brands. For example, a printed coupon sheet can have only four coupons from the soft drinks category, but they cannot be from directly competing brands like Coke and Pepsi. See Table 1 for examples.

Business Rule	Optimization Parameter	Expression
Coke must get at least 4000 coupons	Cell Size Constraint	Cell Size \geq 4000 AND Communication = “Coke”
If a customer gets Coke, they cannot get Pepsi, and vice versa	Blocking Contact Policy	Communication = “Coke” BLOCKS Communication = “Pepsi”
Each customer has exactly 32 coupon spaces	Customer Level Constraint	Min “Coupon Space” = 32 AND Max “Coupon Space” = 32 Using SUM of “number_of_spaces” variable from Communications data

Table 1. Example of Parameters for Coupon Optimization

The last set of constraints is related to the number of coupons on a sheet. In our example, each sheet or booklet has only 32 spaces, but some coupons take up more than one space, making it even more complex.

Loyalty Marketing Example

One SAS customer uses SAS® Marketing Optimization to select coupons for a booklet that it distributes periodically to loyalty card members. Each customer can be eligible for as many as 220 potential offers. The retailer uses business rules and advanced analytics to determine the eligibility of the customer for the offers. A simple rule might be for example the identification of vegetarian households so that those members are not eligible for any coupons for meat products. By defining the eligibility, the offer table becomes less dense and is easier to process. The process of defining eligibility takes place before the data is loaded into SAS Marketing Optimization.

Improving Retail Decisions with Customer Analytics, continued

The first set of criteria for the optimization is to define the number of slots available for coupons. Because some coupons can take more than one slot, the problem needs to be structured a little differently. As mentioned in the previous section, these pamphlets contain 32 slots for the coupons, but because some of the coupons take up more than one slot the optimization analyst has to adjust the scenario to handle that fluctuation (see Figure 1). The analyst does this by adding a column in the communication table (a table that lists all the coupons) that contains the number of spaces that a coupon will use.



Figure 2. Example of Printed Coupon Sheet

Another constraint for this retailer is category limits. The retailer has a limit of no more than four coupons from any one product group (for example, soft drinks) per customer. These are fairly straightforward constraints, and the analyst simply creates one constraint for each category. What could make this situation more complex is if the retailer wanted to ensure that certain categories were forced into the coupon booklet. Just having the maximum constraint of four coupons would not guarantee that soft drinks were represented. If they wanted to ensure that there was at least one soft drink coupon for every customer, then another constraint would need to be added for that category, such as a minimum of one soft drink coupon. Note that tightening the constraints in this way makes the optimization problem more difficult (the scenario runs longer) because the solvable region is smaller. In other words, the optimization solver has fewer options to make decisions that meet all the criteria.

The last set of constraints that this retailer must comply with is the minimum and maximum limits from the

Improving Retail Decisions with Customer Analytics, continued

manufacturers. The manufacturers want to sell more of their products and hopefully gain new customers, but the retailer wants the customers to buy more than just what is on the coupon. The tradeoff is in the manufacturer's constraints. These constraints are typically at the brand level and often tie to media promotions, so they can change over time. For example, Coke might want 10,000 coupons printed in the week leading up to the Super Bowl, but drop that number down to 4,000 the day after the game. Although these kinds of constraints are typically stated in wide ranges, sometimes they are expressed as exact numbers: "Print exactly 10,000 coupons." As mentioned in conjunction with the category constraints, these reduce the size of the feasible region, which makes it harder, in general, to find a solution.

All optimization problems must have an objective, and in this case it is to maximize the probability to purchase. This retailer uses a statistical model to feed into the optimization to determine this probability. The simplest measure is the probability of redeeming any coupon, but that does not give the optimization solver as much to work with as would having a model for each coupon or at least for each category. An even more complex objective would be to maximize the value of the customer's visit based on the probability of redeeming each coupon. This is, of course, a much more difficult task, and most retailers are not able to arrive at these kinds of measures, either due to a lack of skills or the availability of the data.

MARKETING SERVICE PROVIDER COUPONS

While individual retailers are tackling this problem, there are some marketing service providers using SAS and providing coupon optimization at the point of sale (POS) as a service (see Figure 2).



There are a few key differences between optimizing at the POS and direct marketing, but first "real time" needs to be defined. So far, no retailers are running "real-time coupon optimization" in the way that we have described optimization. Some might call their process optimization, and that is fine, but the process is not using mathematical optimization. It's not that the technology does not exist; it's just that the focus has not been on the analytics and the data required to pull it off. So, for this example, we will be discussing batch optimization, or preselecting the coupons that the customer will receive when they check out, rather than optimizing between the time the cashier hits Total and the coupons are printed for the customer.

The first distinction between POS and direct marketing is the number of coupons involved. Rather than trying to fill a booklet, retailers are looking for only a few coupons per customer, often a maximum of five. Another distinction between the retailers and the MSPs is the goal or objective of the optimization. Although the retailers make more revenue when the customer buys more than just the coupon item, that has no effect on the MSP's revenue. MSPs are often paid by the manufacturers, so their revenue is dependent on the number of coupons that are printed. In that case, the objective might be to maximize the number of coupons.

Register-Printed Coupon Example

One MSP (a SAS client) had a project for one of their grocery customers for which there were 64 potential coupon offers representing 20 different product groups or categories. The grocer had 2 million loyalty card members who could potentially receive the coupons at check-out.

This project was set up quite differently from the previous retailer example. Rather than needing to print exactly 32 coupons (or spaces for the coupons), the MSP wanted the shopper to get between three and five coupons, or none at all, and the size of the coupon did not matter. Their category constraint was that each category had a limit on the total number of coupons that could be printed. Moreover, each shopper could get only one offer per category. The objective of the optimization was to maximize the total number of printed coupons while minimizing the number of

Improving Retail Decisions with Customer Analytics, continued

customers who would not receive any coupons. In this case, the scores were at the category level. Because the constraints were not very tight, the optimization completed in only a few minutes. This gives analysts the ability to run more projects per day, thereby increasing the MSP's revenue capacity. In addition, the results (number of coupons printed and number of customers receiving a set of coupons) were substantially better than those obtained with the old manual prioritization method.

Some retailers' marketing departments leverage customer analytics to generate marketing offers similar to the ones described above. However, only a few retailers are leveraging customer analytics within their merchandising decisions.

CUSTOMER-CENTRIC MERCHANDISING

For many retailers, customer data is still primarily managed by the marketing department. Customer insights are shared with the merchandising team on an infrequent, fragmented, ad hoc basis via manual reports. This situation leads to inconsistent use of customer analytics for merchandising decisions.

Two areas in which retailers are beginning to show progress toward integrating customer insights into merchandising decisions are assortment and promotions.

ASSORTMENT

Creating localized assortments is the mantra of many retailers. Integrating customer insights into the process of determining which products to offer by location, channel, or both provides a new dimension to those decisions. Understanding the importance of a product to key customer segments is critical information to consider when making drop or keep decisions about products within an assortment.

Retailers are beginning to create loyalty scores that contain customer metrics like loyalty to the retailer, category, and brand. Measures of how exclusive a product or brand is to certain key customer segments provide guidance in determining whether demand for one product can be transferred to another product. A weighted loyalty score is used in assortment optimization to ensure that lower-selling products that are critical to key customer segments are not discontinued or dropped.

PROMOTIONS

Mass promotions, advertised via newspaper circulars, are still a primary promotional vehicle for many retailers. The process for determining which products to feature and promote on the front page of those circulars is generally cyclical. Many of those decisions are based on the prior year's promotions.

A few retailers are beginning to incorporate customer insights derived from the retailer's customer database into their ad planning process. They use key customer metrics like household penetration, revenue per household, and units per household to estimate the potential sales and profits of products that they are considering for their front page. These retailers estimate the buying behavior of primary customer segments so that overall profitability from the circular is maximized.

CONCLUSION

Leading retailers are incorporating customer insights into key decisions across their corporations. They understand customer economics. They regularly measure the behavior of each customer segment and understand the related effects on overall business results. They demand insights from customer analytics when determining strategies. Some retailers intentionally differentiate customer-specific marketing offers and use optimization to ensure that those offers are the most relevant to each customer. And a few retailers are beginning to infuse customer insights into their merchandising decisions. Customer analytics can be a substantial competitive advantage for retailers. We have only scratched the surface of the many ways in which using customer insights can improve customer satisfaction and grow retailer profitability.

RECOMMENDED READING

Davenport, Thomas H., and Jeanne Harris. 2007. *Competing on Analytics: The New Science of Winning*. Boston, MA: Harvard Business School Publishing Corporation.

ACKNOWLEDGMENTS

We appreciate the help from our colleague Michelle Opp.

Improving Retail Decisions with Customer Analytics, continued

CONTACT INFORMATION

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

Wanda Shive
SAS Institute, Inc.
SAS Campus Drive
Cary, NC 27513
919-531-531-2564
E-mail: Wanda.Shive@sas.com

Dwight Mouton
SAS Institute, Inc.
SAS Campus Drive
Cary, NC 27513
919-531-531-1912
E-mail: Dwight.Mouton@sas.com

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

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