Predictive Modeling: Innovative Ways to Target Customers and Improve Marketing Processes

Executive Summary

Marketing has traditionally operated with little or no quantitative data outside of market research. With the advent of customer databases and external market data, predictive models can be developed to drive better marketing outcomes. The impact can be far-reaching—resulting in a better understanding of cross-selling programs, purchasing propensity, and sales dynamics. These data models can also result in better metrics and measurement. By understanding key drivers of customer behavior, one can develop better leading indicators to support executive marketing dashboards and ultimately marketing strategy.

Here are case examples of how leading companies are using predictive modeling to their advantage:

What is Predictive Modeling?

A rigorous approach which uses existing patterns in current data to statistically predict the behavior of future customers.

Predictive Model services include:

- Database sampling and aggregation
- Data mining and Model Building
- Model testing
- Model deployment
- Ongoing challenger vs. champion updates

1. Cross-Selling Programs

The Problem

A financial services company wanted to identify which customers were likely to upgrade to Gold status in the next 12 months. Gold status customers received access to special services and provided the company with greater income.

The executives called a meeting that included two Vice Presidents and resulted in three opinions on the subject. The prevailing opinions were:

1. Customers with the most products, regardless of what the products were, seemed a good bet to reach Gold status.
2. Customers with the highest balance, even if it was in only one account, were obviously prime candidates.
3. Customers with an increase in transaction volume over recent months were clearly headed for the Gold circle.

After a heated discussion and ensuing gridlock, no one was backing down.

Data Mining Solution

Ultimately, data mining techniques that included merging a customer demographic database and a customer transaction database revealed a profile of the customer most likely to reach Gold:

.

- The best and most promising customers were those who started out with a mortgage and then moved on to other services and products.
- The worst type of customer started out with one credit card, or a checking account with a balance of under $100, and moved on from there. This customer rarely reached the Gold circle.

Data mining also revealed that gender and age were not drivers for Gold status candidacy. A high income was. Also, the account balance didn't need to be extremely high, only above a certain threshold. And finally, transaction volume was not a driver.

Strategic Implications

With a clear picture of the most valuable customers in the long-term, the company could reassess its marketing strategy and better allocate resources. For example, instead of offering freebies and perks on new checking accounts, or mass-mailing credit card invitations, the company needed to sell the service that most often led to Gold status: the mortgage. While it's more difficult to chase this customer, it's where the company focus should be and where the highest payoff lies.

Data Mining Models Drive Marketing Strategy & Dashboards

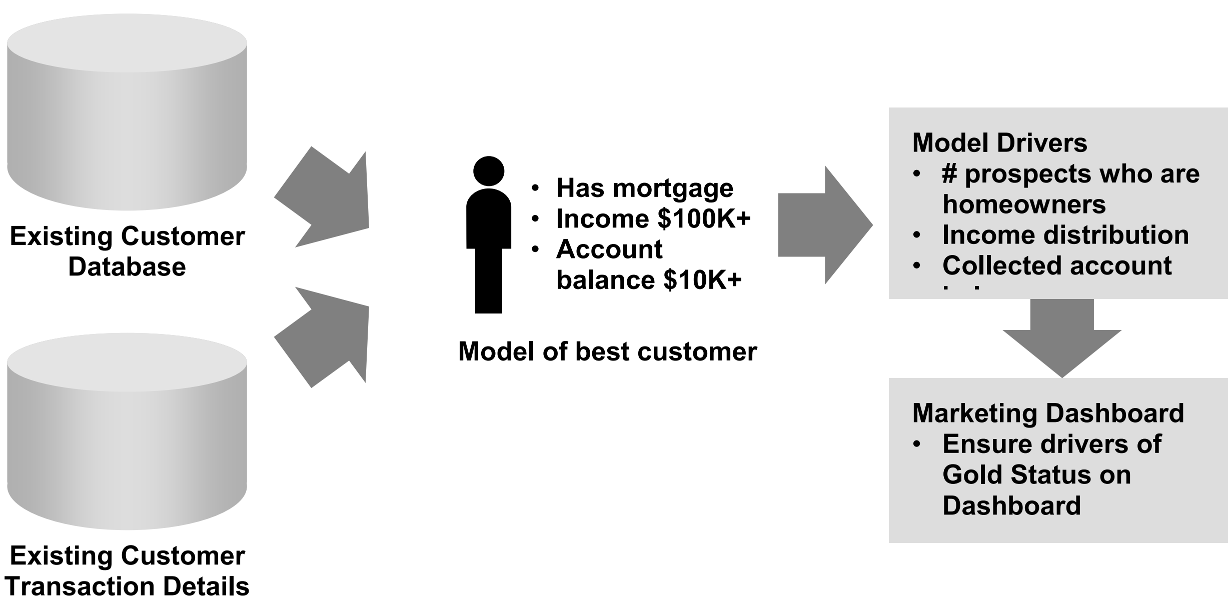

The Problem

Another example is a software company that had purchased and compiled a list of sales leads—companies it hadn't yet approached or sold to, but clearly wanted to. The software company also had an internal customer database, containing the size, geographic location and other information for existing customers. The company needed to know how to allocate marketing budget for maximum return on the list of sales leads.

Purchasing Propensity Models

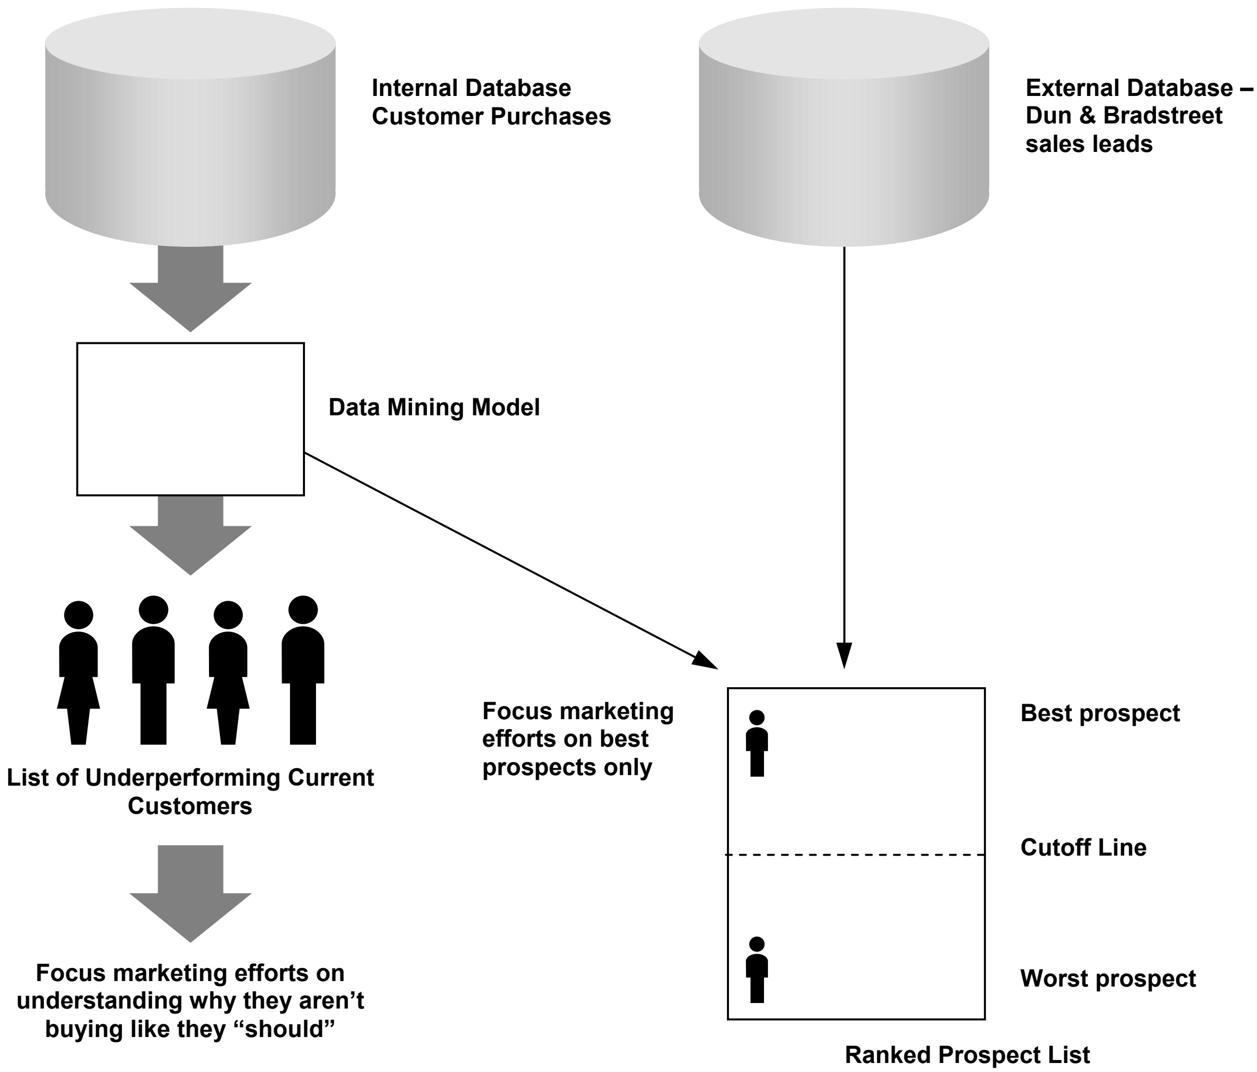

Data Mining Solution

Data mining allowed the software company to analyze its existing customer database and profile top customers. This analysis allowed the company to Identify existing customers who were under-purchasing, in terms of size, location and other factors, and score the new sales leads on a scale from 0 to 100, where 100 represented a sales prospect very likely to buy.

Strategic Implications

The company could immediately utilize the data in two ways:

- For existing customers who were not buying as much as the model predicted they "should," the software company could ask "Why?" It could glean answers from the data and from interviews, focus groups, and casual conversations with the client, and work to remedy the problem.
- For new sales leads, the company could give the sales force a list that included rankings for all leads. The sales team also received a suggested cutoff point, such as "Please concentrate on prospects with a score of 70 or higher out of 100." The sales team could then focus efforts based on statistically significant results, instead of working their way down the list alphabetically or by gut feel.

3. Customer Feedback

Data Mining Solution

Many companies have learned the value of customer feedback and are chasing it through phone- and online-surveys, forums and frequent, informal communications. But how do you know if what your customer says matches what the customer does?

In one case, an industrial products company wanted to know why some of its established customers were buying less and less. When the executive team approached the customers for honest feedback, they were told, "The salespeople are often late, don't know the product and dress poorly." The company's previous knee-jerk response to such criticism would be to wrangle the sales team for an intense product-training seminar, hire a "dress for success" consultant, and institute a corporate "lateness policy."

Data Mining Solution

This time, the company relied on data mining to uncover drivers of customer purchase behavior. When each customer's purchasing pattern was correlated with the related salesperson, areas like promptness, product knowledge and dress weren't critical factors. In fact, these items weren't even on the radar. Other factors, such as the salesperson's number of years of experience in the industry, or whether the salesperson clearly asked for the sale, were also in this case not statistically significant drivers of customer purchasing behavior.

It was all about customer loyalty. Customers bought from salespeople whom they felt were trustworthy and communicated well. In fact, the combination of high trust and good communication skills could overcome lack of product knowledge, lateness, or substandard dress. A rep who said, "I don't know the answer to that, but I'll find out for you and have it for you by next Tuesday," and then followed through on that commitment every single time, often sold more than a rep who knew the answer offhand or who attempted to cover up lack of knowledge.

Strategic Implications

Instead of an immediate over-emphasis on product training, the company could better focus efforts and budget where it would have most impact—training Human Resources to break from a recent hiring focus on product familiarity and job tenure. Instead, HR needed to recruit the "Rainmaker" personality, one who communicates well, is a true "people person," and innately builds client trust and confidence.

Once the company brings the right personality on board, they can easily train that salesperson on products, rather than trying to teach a product expert the more elusive art of being a great communicator.

How, Exactly, Does It Work?
Solving Challenges with Database Sampling and Model Calibration

Some who work in large companies may be wondering: "How can I build a predictive model if our customer information is in multiple, distributed databases? "

This can be easier than many people think if you have good database administrators and IT personnel on hand. After the first time, it can be mostly automated and run quarterly to keep updating your model with the most recent data available.

- First, we decide upon the target variable, which is what we want to predict, such as revenue per customer.
- Then, with help from the I.T. personnel, we identify which databases may contain useful information -- in this example, a customer demographic database may contain the ZIP code and customer company size; a transaction database may contain the revenue information, and a customer service database may contain a record of any customer issues and their resolution
- Next, we ask the I.T. personnel to create a download of the relevant information for the same customers from each database, making sure to include a unique identifier such as Customer ID.
- Finally, the Data Mining team merges the disparate downloads into one large statistical master file, which is used for model building and testing.

What if our customer database is extremely large?
Won't building a model from this be like trying to drink from a fire hose?

Just as market research conducts a statistical sampling of a population to represent the entire population, Database Sampling is a statistical technique that enables you to collect data from multiple customer and external databases and aggregate it into a single model.

- Sometimes Database Sampling can be as simple as taking every 100th record from existing databases and then checking to ensure it's a representative sample.
- Other times, when the model wants to focus on a relatively rare target such as a very high net revenue customer, we will deliberately oversample the very high net revenue customers while deliberately undersampling everybody else. This will allow us to build a model which is more finely tuned to the very high net revenue customers, while still able to handle the rest of the population.
- When there is a huge amount of data, aggregation can work to reduce the size and improve overall data quality. For example, instead of trying to analyze data for each of 50 states, we may choose to aggregate into three geographic regions: East, Central, and West; often formerly invisible patterns come into focus when one takes a higher-level view.
- Some companies have a time series component to their databases. This often occurs in the IT industry, when customers may complete a large IT purchase every few years. Just before the purchase, the customer is a great prospect, while immediately after the purchase, the customer is not. In this case, the database sampling may include data from fewer customers over a longer time frame, which will allow the model to track purchasing patterns over time.

Database Sampling and Aggregation

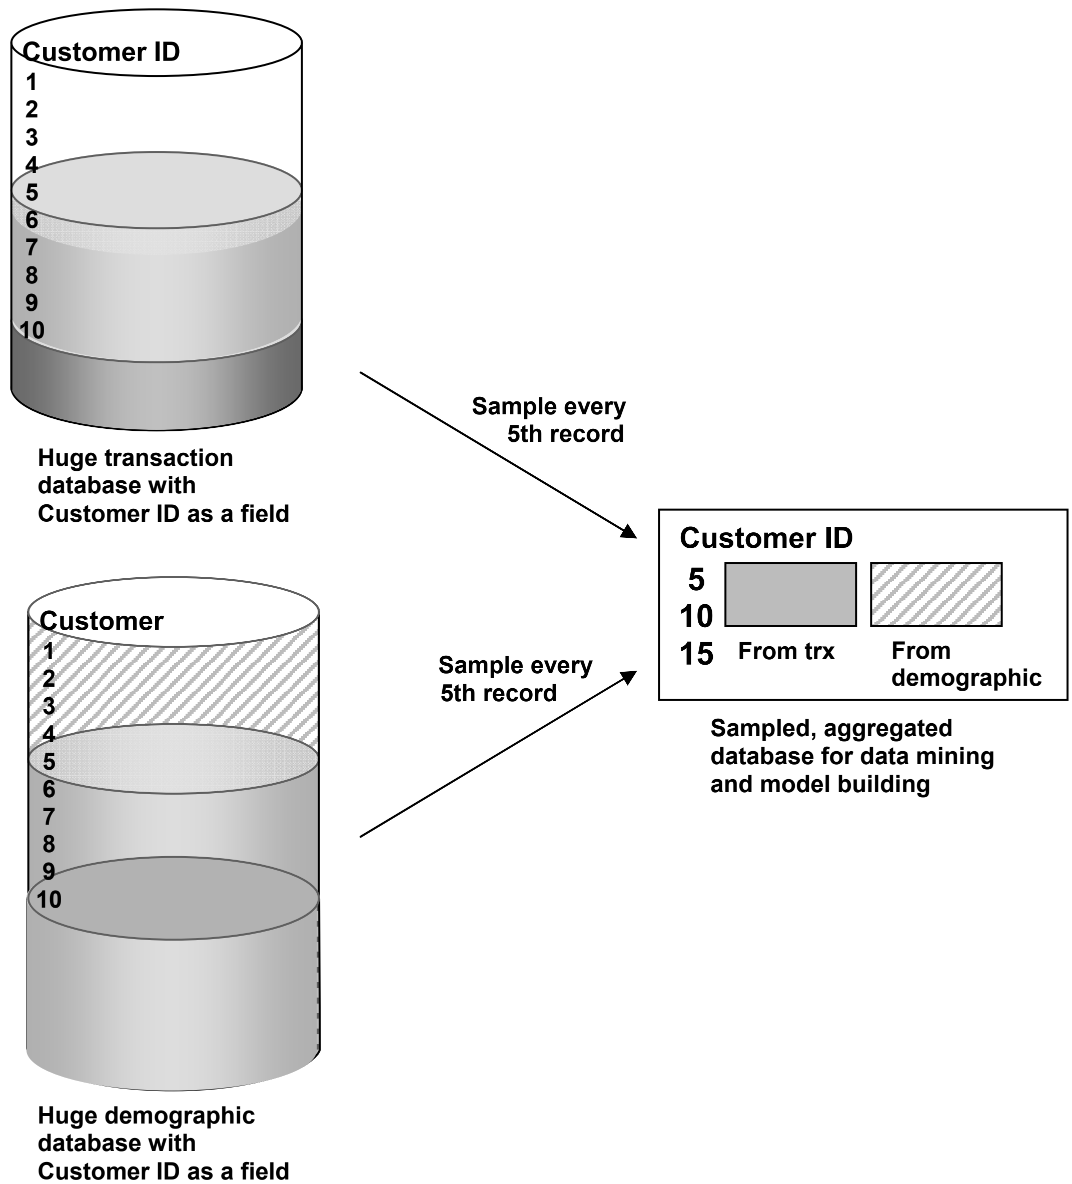

How will we know if the Data Mining Model is any good?

When conducting market research studies, we develop a 95% confidence interval that the data will be within +/- 5% of our findings. When building Data Mining Models, a different methodology is often used.

- We generally gather at least twice as much data as is needed to build the model.
- We build the model using half of the data (the “build” set), and then test the model on the other half (the “test” set).
- We can look at the test set results, and see what would have happened with our actual customers and our existing business if we had used a particular model. It is a little bit like looking at last year’s investment portfolio results, calculating what the results would have been under a series of different investment strategies,

and then picking the best one for moving forward. This gives us a relatively low-risk way of evaluating several competing models and selecting the best model for our data and our company.

About Exponential Edge, Inc.

Exponential Edge helps global clients to build new roads to revenue through strategic opportunity assessment, alliances and marketing and sales operational effectiveness. Exponential Edge can be reached at info@exponentialedge.com or www.exponentialedge.com

.

Disclaimer: *Identifying details have been changed to protect client confidentiality.*