How to Measure Campaign Impact

Written by
Rhonda Carraway Petty
Marketing Insights Data Scientist
rpetty@mathnetix.com

Marketers routinely implement marketing promotions to attract customers and increase sales. Sometimes, these promotions are directed to specific customers (direct campaigns). Other times, promotions are available to all customers in a physical location or geographic market (broad campaigns). Regardless of the type of promotional campaign, Marketers are often left without adequate insight into the performance of their marketing efforts to help them make good decisions for future campaigns.

This paper illustrates, with a simple-to-understand example, how to measure the incremental impact of a marketing promotion. It is intended for the business analyst charged with measuring campaign impact or for the executive who wants to learn more about how measurement works. For the sake of brevity, most illustrations are presented in the context of measuring direct marketing promotions—those targeted toward named individuals.

MEASUREMENT: AN EXERCISE IN PREDICTIVE ANALYTICS

Measuring the impact of a marketing promotion requires knowledge about what would have happened had the promotion not occurred. The promotion's impact is the difference between what actually happened with the promotion and an *estimate* of what would have happened without the promotion. Because measurement involves estimation, it is an exercise in predictive analytics. The better the prediction of what would have happened without the promotion, the more reliable the measurement result. For this reason, measurement needs to be approached with attention to detail.

Figure 1: Conceptual Explanation of Measuring Campaign Impact

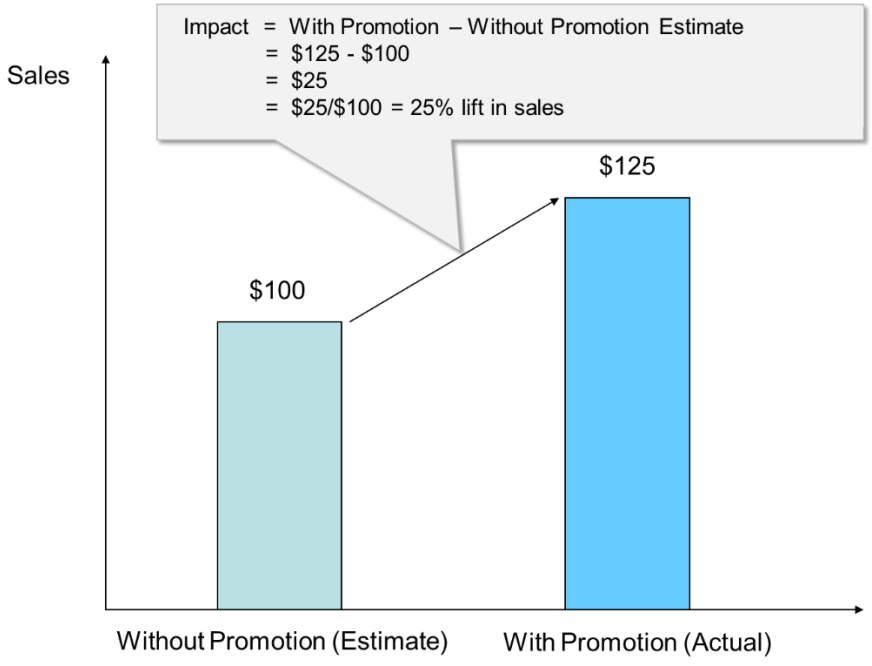

THE NEED FOR CONTROL

Estimating what would have happened without the marketing promotion is often achieved by creating a CONTROL state. The CONTROL state acts as a proxy for what would have happened without the influence of a marketing offer, campaign or message. There are two fundamental ways to establish a CONTROL state:

- **Time-Based Approach (Pre-Post Method).** This approach uses the behavior of the target population in the period *before* the promotion launched as a proxy for a 'no promo' scenario. This Pre-Period behavior is the CONTROL state. This approach assumes the behavior in the Pre-Period would have continued, unchanged, in the Post-Period had there not been a promotion. Holding the target population constant, impact is quantified as the difference in behavior **before** and **after** the promotion launch date (with the Pre-Period and Post-Period established as equal lengths of time). Thus, it is referred to as a 'Pre-Post' method. This approach works well as long as there is no seasonal fluctuation, trend or other event besides the promotion that would cause behavior to shift over time. Because few situations are void of seasonality, trend and unforeseen fluctuations over time, this approach should be used with caution.

Figure 2: Pre-Post Method Holds the Population Constant

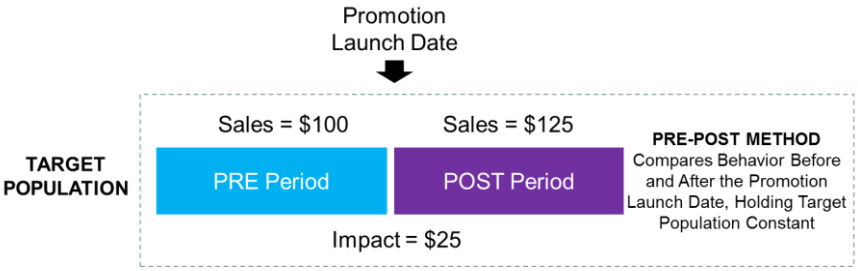

- **Population-Based Approach (Test-Control Method)**. Focusing only on behavior after the promotion launched, this approach compares the behavior of a target "TEST" population to the behavior of an equally-sized lookalike "CONTROL" population. Because the behavior of two populations is compared over the same exact calendar dates, the likelihood of seasonality causing a misread is minimized. But, if the two populations are not clones of one another, sampling bias can occur.

Figure 3: Test-Control Method Holds the Time Period Constant

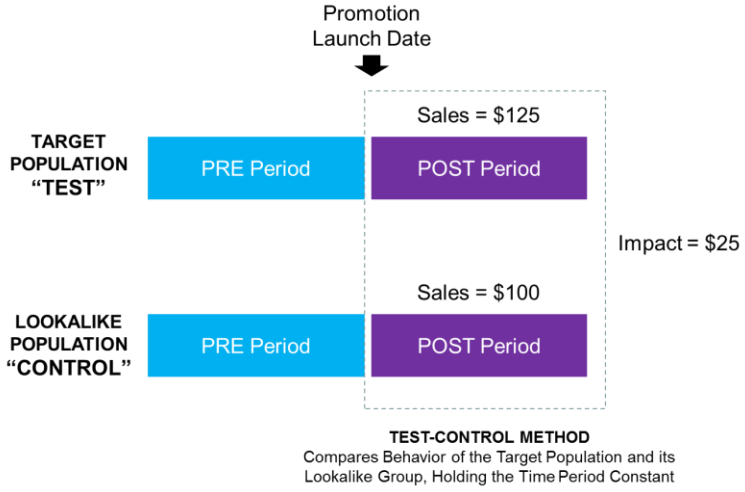

CREATING A "LOOKALIKE" CONTROL GROUP

Because there are few situations where seasonality, trend and undetected outside influences are non-existent, the Population-Based approach to selecting a CONTROL state is often used. With this approach, a Test Group receives the promotion while another "lookalike" Control Group does not receive the promotion. The Control Group is created in advance of the promotion, randomly assembled from the same parent population as the Test Group. In this way, both groups are clones of one another.

To insure the Control Group looks like the Test Group, compare the behavior of the two groups in the period before the promotion starts. If your performance metrics are Sales, Transactions and Unique Customers Buying, make sure the two groups have similar Pre-Period values for Sales, Transactions and Unique Customers Buying. Even better, make sure they also exhibit similar historical trends in these metrics.

Suppose a retailer issued a Recapture offer to its customers who had not purchased in the last six (6) months. In order to estimate the impact of this offer, the retailer *randomly* split the target population of customers into two groups. Half of the customers received the offer (the TEST GROUP) while the other half of the customers did not receive the offer (the CONTROL GROUP). In this way, both groups were sampled from the same population at the same point in time and expected clones of each other.

Figure 4: Equally-Sized Samples Drawn From the Same Population at the Same Time

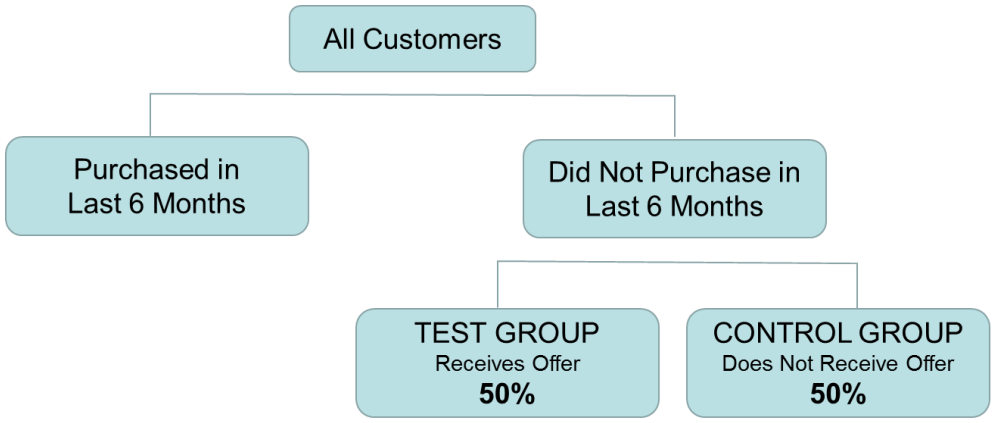

Both the Test and Control groups are sampled from the same population at the same time—just before the promotion starts

Before finalizing the Test and Control Groups, suppress customers that received a recent prior promotion. This will isolate performance to the one promotion that is being measured. If a constant flurry of marketing activity is 'normal' such that it is not possible to isolate performance to the one promotion of interest, be sure results are presented and interpreted as the incremental impact over and beyond routine marketing activity.

Remembering a recent prior promotion, the retailer decided to re-assemble the Test and Control Groups. Customers who were part of a prior recent promotion were suppressed BEFORE selecting the Test and Control Groups. This way, the condition that the groups be sampled from the same population was upheld.

Figure 5: Suppressing Customers Included in Recent Prior Marketing Promotions

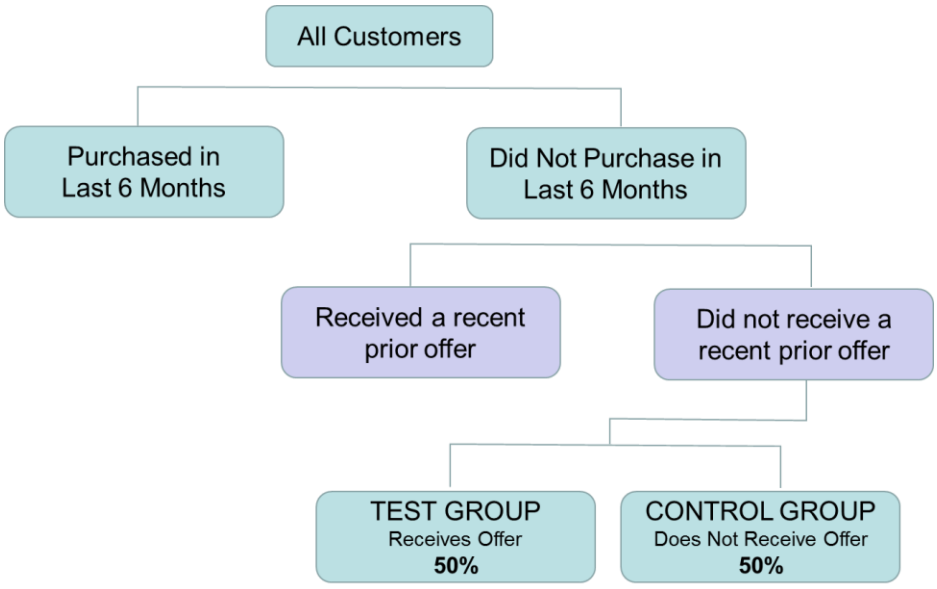

Both the Test and Control groups are sampled from the same population at the same time—just before the promotion starts. Customers who received a recent prior offer are suppressed so that the impact of only the said promotion can be isolated.

A Word on Cell Size

Most Marketers want the Control Group to be as small as possible. When the Control Group is small, more customers can be placed into the Test Group and receive the offer. Placing 90% of the population in the Test Group and 10% of the population in the Control Group is a popular split. If any split other than a 50-50 divide is used to create the Test and Control groups, the impact calculations will need to adjust for unequal cell size.

THE PRE-POST TEST-CONTROL APPROACH

The **Pre-Post Test-Control Approach (PPTC)** is a widely accepted standard because it directly addresses the disadvantages of the other two aforementioned techniques:

- It is less sensitive to seasonality, trend and undetected outside influences—a disadvantage of the Pre-Post Approach.
- It is less sensitive to test and control 'clone' impurities---a disadvantage of the Test-Control approach.
- It has the added advantage of being less sensitive to differences in cell size. The Test Group and the Control group need not consist of the same number of customers.

Before sharing the mechanics of the PPTC approach, let's clarify definitions. The Pre-Period is the period of time just before the promotion launch date. The Post-Period is the period of time just after the promotion launch date. The Promo-Period is the period of time within the broader Post-Period when the promotion is active. Typically, the Pre-Period is established so that it is the same length of time as the Promo-Period.

Figure 6: Semantics Describing Three Analysis Periods

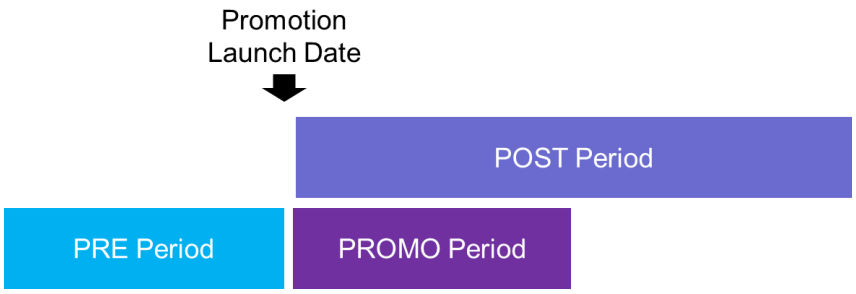

As with the other measurement techniques, the goal of the PPTC Approach is to *estimate* what would have happened in the absence of the marketing promotion. This is accomplished in multiple steps as illustrated with a numeric example.

Step 1: Summarize sales (or other quantifiable performance metric) for the distinct measurement groups. Then, check the Pre-Period alignment of sales between the Test and Control groups, normalizing sales to a 'per-person' basis when the Test and Control Groups are not the same size. Seek an Alignment Index between 0.90 and 1.10.

A retailer launched a promotion that lasted for three weeks. At the end of the promotional period, the retailer wanted to know the impact of the promotion during the period of time when the promotion was active. With this in mind, she summarized sales for four measurement groups and performed a cursory check to confirm Pre-Period behavioral equality between the Test and Control Groups.

Figure 7: Pre-Period Balancing of the Test and Control Groups

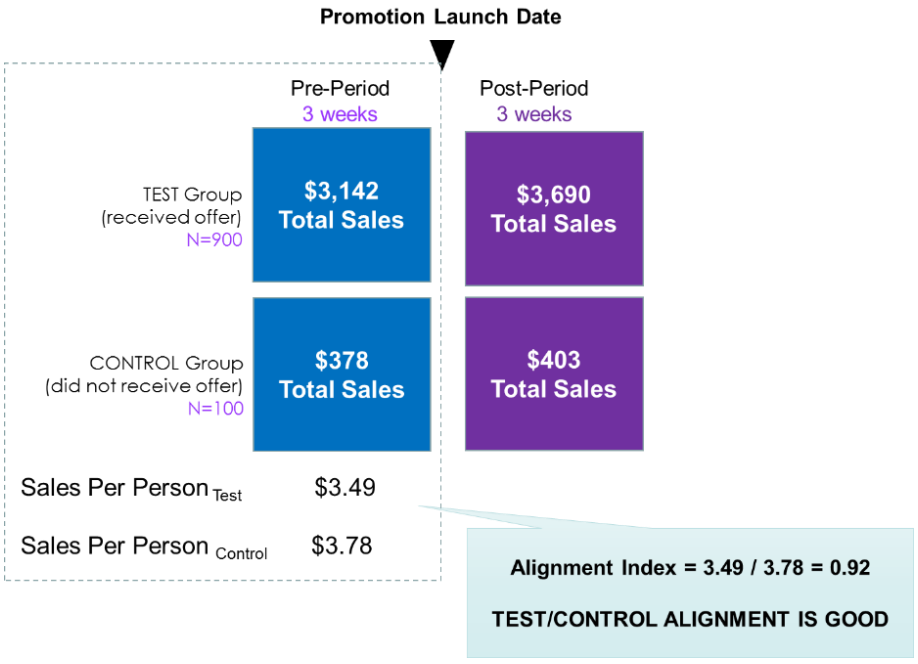

Step 2: Use the Control Group's behavior to calculate a business-as-usual expected shift in performance across time between the Pre-Period and the Post-Period. Apply this expected shift to the Test Group's Pre-Period behavior to estimate what the Test Group would have witnessed in the Post-Period without the promotion.

> Using the Control Group, the Marketing Scientist observed the natural percentage shift in sales between the Pre-Period and the Post-Period which calculated to 6.7%. When applied to the Test Group's Pre-Period sales, an estimate for what the Test Group would have witnessed in the Post-Period, had the promotion not been active, was derived.

Figure 8: Estimate the Test Group's Behavior in a 'No Promo' Scenario

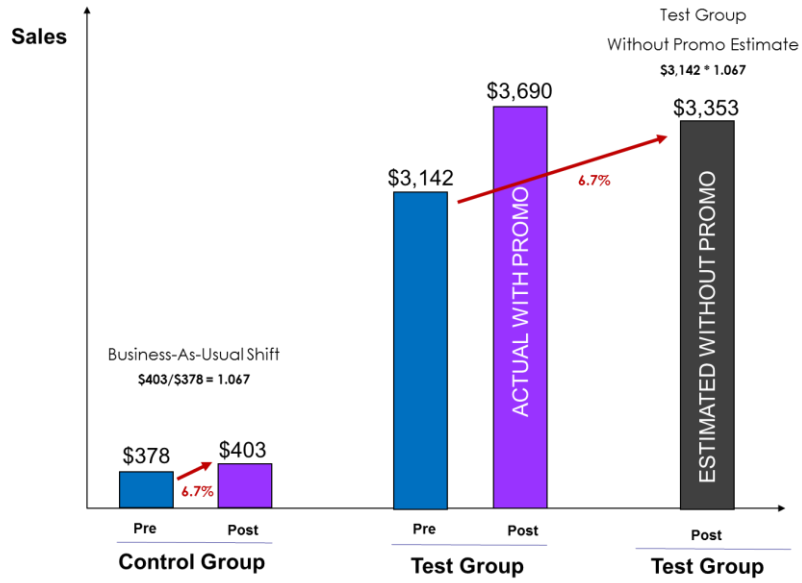

Step 3: Quantify impact as the difference in the Test Group's actual Post-Period sales and its estimated Post-Period sales in a 'no-promo' scenario.

Figure 9: Calculate Impact and Percentage Lift

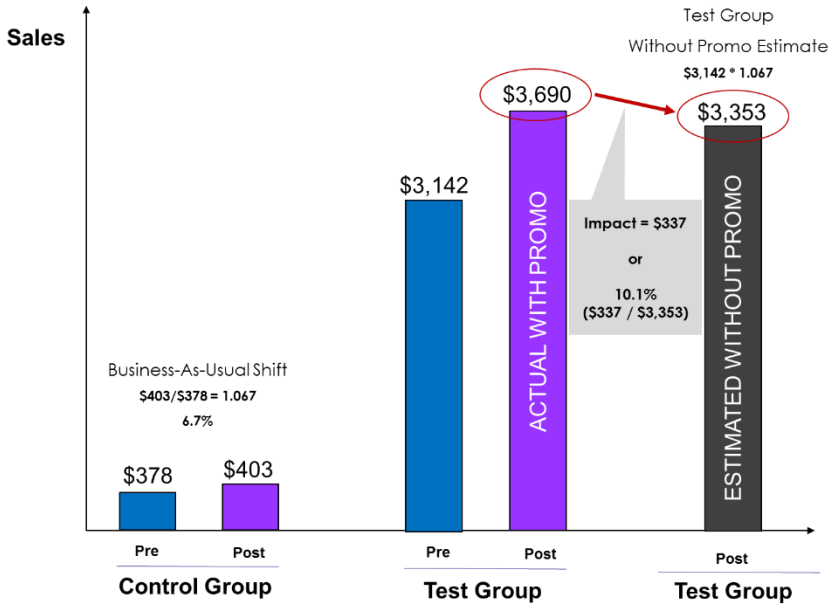

MEASURING SUSTAINED IMPACT

The goal of most marketing programs is to generate incremental business, not only during the promotional window, but also beyond it. Sometimes, a marketing promotion succeeds in shifting the *timing* of business, but not total business. To quantify the sustained success of a promotion while also satisfying management's need to obtain learning from the promotion as soon as possible, at least two promotional 'reads' are recommended---one read when the promotion has completed and a second read at a pre-designated time in the extended Post-Period that follows the end of the actual promotional period.

Figure 10: Anchoring Measurement Results on Two Post-Period Reads

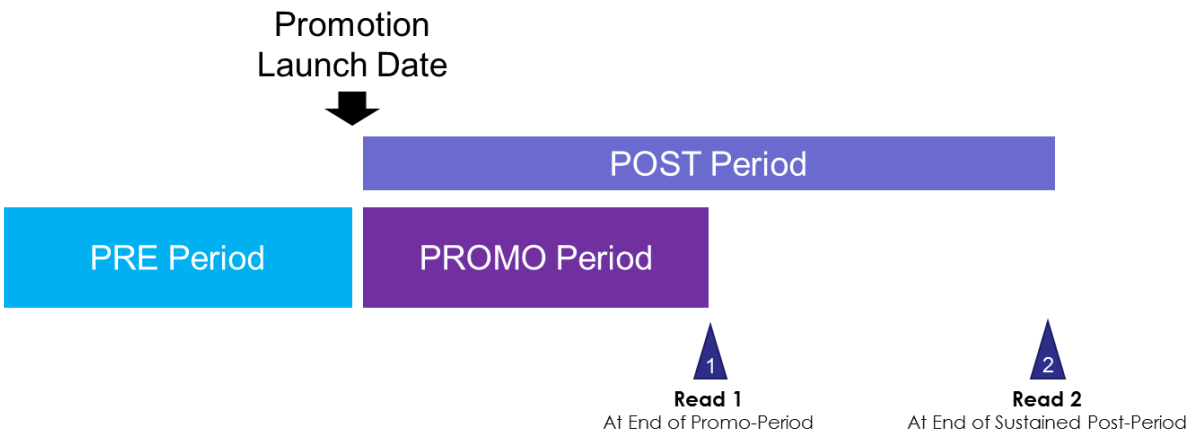

OTHER PERFORMANCE METRICS

Although senior management may only be interested in the bottom line—impact on Total Sales---staff responsible for devising effective marketing strategies will, no doubt, want more insight into the particular behaviors that drive sales. Using the same principles explained above, the impact of the promotion on other key performance indicators can and should be included in measurement output. Consider including the following additional metrics:

- Total Sales
- Total Unique Persons Buying
- Total Positive Transactions
- Sales per Person
- Sales per Transaction
- Transactions per Person
- Total Traffic & Conversion (if collected)

Finally, straight counts of redemptions should also be included if the promotion included an offer to redeem a coupon or other item:

- Redemption Count
- Redemption Rate (Redemption Count / Outgo Quantity)

DETAILED "DEEP-DIVE" PERFORMANCE REPORTS

When a repeat of a promotion is considered, and especially if improvement in return-on-investment from a repeat attempt is sought, a detailed analysis of various subsets of the Test Group will give Marketers focused insight on where to make changes. For example, if one segment of the customer population did not produce a satisfactory ROI, then that segment might be excluded from the repeat launch. Or, if incremental sales of a particular product were not observed, some other incentive may need to be tested to boost sales for that particular category. Deep-dive analysis expands results to include subsets of the Test and Control Groups---by segment, product, channel and/or market. But, as the sub-segments get smaller and smaller, so does their sample size and statistical validity of their results. So, be sure the Test and Control groups which are balanced at aggregate levels of the analysis, remain balanced at the more granular levels also.

Because of their time-consuming nature and tendency to produce an overwhelming abundance of metrics, deep-dive analysis is generally reserved for special campaigns of particular interest. And even then, the deeper analysis is typically directed to only one or two subsets of the analysis sample—by segment or by market, for example. If understanding impact for many slices of the population is desired, it is a better use of time to build a predictive model for scoring customers on their likelihood to respond to a repeat of that campaign. Rather than reporting impact in all possible ways, the response model will produce a sorted file of customers most likely to respond incrementally if the promotion is repeated. And ultimately, this is what the Marketer wants to know---who will respond profitably if the promotion is repeated.

OPTIMIZING THE CONTROL FOR MARKET-WIDE PROMOTIONS

The majority of the illustrations in this document were presented in the context of measuring promotional impact for direct-to-consumer communications where named individuals are the Test Group and receive the promotion. When a promotion is broad in nature, perhaps available to all stores in a given market, the process for selecting a Control Group is a little more complex. If there are ten stores in a market, for example, it might be tempting to set one store in the market aside as a Control Group. But, from a business and customer relations perspective, this will create confusion in the market and disgruntle customers who shop that one particular Control store expecting to receive the advertised offer only to find the store will not honor the promotional offer.

Before launching your market-wide promotion, select one or more other markets (in their entirety) to act as your Control Group. Align the Test Markets and the Control Markets on an array of aggregate Pre-Period behaviors and the composition of the customer base. Use metrics such as Total Sales, Total Transactions, Total Unique Customers, Total New Customers and Total Stores. If possible, also align the markets on the age and type of the stores within the markets. The more alike the Test and Control markets are in the Pre-Period on a variety of different measures and attributes, the more accurate the measurement results.

Mathnetix® has automated the repeatable mechanics of the campaign measurement process in order to increase marketing agility and insight. For more information, please contact Rhonda Petty.

Rhonda Carraway Petty
rpetty@mathnetix.com
601-384-2488