

Paper 022-2011

Creating a Marketing Investment Optimization Framework: American Express Travelers Cheques Case Study

Pieter van Ispelen, Director; Landon Starr, Sr. Manager

Global Business Analytics, American Express, SLC, UT

ABSTRACT

Creating a standard, consistent, and effective method of measuring marketing investment effectiveness can be very challenging. Adding a global investment environment, with dramatically varying global consumption patterns, can compound those difficulties. Throw in macro-economic impacts, and the desire to determine the best investment amounts to be allocated to marketing a product, and these challenges manifest themselves into the journey of American Express' Travelers Cheque marketing investment analytics work over the past few years. The purpose of this paper is to share a versatile marketing investment optimization framework, focusing on clustering consumer behavior, leveraging SAS/ETS[®] software for time series sales pattern trends, and utilizing a Monte Carlo type of approach to determine variability and confidence.

INTRODUCTION

Determining one's marketing investment strategy and assessing its effectiveness is a challenge nearly all companies share. Who is our customer? How do I reach them with a relevant message? Did it work? What was the ROI? The answers to these questions can often be a convoluted, iterative process. Throughout the investigation our answers change, customers evolve, even the economy can dramatically change the face of one's investment strategy.

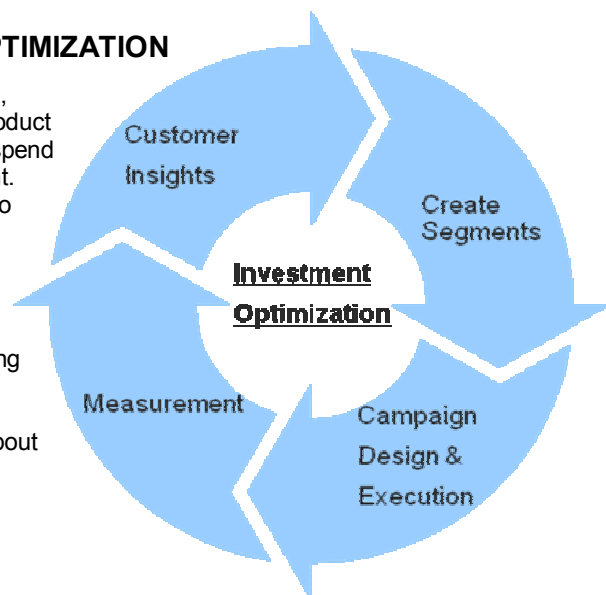
Within the American Express Travelers Cheque (TC) environment, we observe a cash-replacement product generally used on travel for purposes of insurance against lost or stolen currency. The product is sold broadly throughout the world through 3rd party sellers (largely banks). This means that customer information is relatively scant, given the relative anonymity of the customer base.

This paper will illustrate examples and details of the customer insights strategy resulting in strong alignment and communication with marketing for promotion design. Subsequently, the same framework is used in a detailed assessment overview leveraging multivariate clustering provided via the FASTCLUS procedure, and automated time-series forecasting and simulation for calculating sales lift leveraging SAS/ETS[®] software. Ultimately the approach results in the ability to leverage ranked promotion sensitivities to estimate the investment opportunity, and use ROI thresholds to define an investment optimization recommendation.

THE CYCLICAL NATURE OF INVESTMENT OPTIMIZATION

When thinking in terms of marketing investment optimization, consistent with the fluid nature of customer evolution and product competitiveness, investments in promotions and marketing spend requires continuous updates and revisions to remain relevant. Determining customer types, creating actionable segments to differentiate treatments, designing and executing the campaign, measuring the campaign, and starting the process over again are all keys that define a successful investment optimization environment.

The next section outlines the 9 steps that define the marketing investment optimization framework and its application to the Travelers Cheque product. It does this at the hand of a challenge and solution approach to add a level of intuition about each step up front.



Creating a Marketing Investment Optimization Framework, continued

9 STEPS FROM DESIGN TO MEASUREMENT TO OPTIMIZATION

Design

1. Customer Insights
2. Align customer segments with customer touch points
3. Promotion Design Tool

Measurement

4. Campaign execution capture
5. Cluster to optimize test and control
6. Monte Carlo Simulation & confidence assessment
7. Results calculation

Optimization

8. Customer and branch promotion sensitivity
9. Global Investment Optimization

1. ***Challenge: We often do not really know the customer of our products***

Solution: Customer behavioral insights

In an environment where you do not know your customer directly, and therefore cannot target them directly with marketing treatments, one has to be more creative in how to define the customer. The starting point is to identify the behaviors that are known for the customers. Which behaviors can be identified? Do these behaviors have the potential to differentiate your customers? And can we therefore create customer segments which could tell us more about our customers than we had imagined? In identifying differentiating behavior of customers, two principles should be kept in mind. First, the identified behavior needs to be universally available so that the majority of your customers can be classified. Second, it needs to have the ability to be actionable - and therefore reachable - by the variety of marketing actions.

In the case of the American Express TC, a customer is defined for this purpose by two key behaviors: how much did a person purchase and how long did it take that person to encash the cheques in the purchase? Even though we do not know the customer directly, this information can provide a lot of insight into the type of customer we're servicing. Does this customer buy a large amount and encash all of these in a short period of time? That could likely be a customer using the cheques for commercial reasons. Did the customer buy a medium amount but then didn't encash a single check for a year? That would likely be a customer using the cheques as either a savings or back up cash vehicle at home. Following this logic utilizing purchase size and float (the measure of time until encashment), we can define the customers on a continuum between 4 typical groups: commercial use, business travel, leisure travel and cash at home as shown in figure 1.

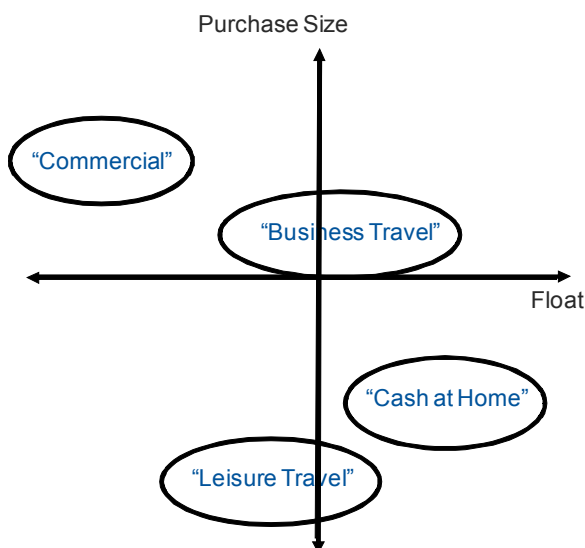


Figure 1: Customer Segmentation

Creating a Marketing Investment Optimization Framework, continued

The observed behavior defines the customer, and provides key insight into the development of the right marketing and servicing treatments. Furthermore, it provides a foundation to develop segments for targeting and measurement.

A direct off-shoot application of this customer segmentation is the ability to identify geographic concentrations of certain customer types as illustrated by Figure 2. This greatly benefits the business organization in developing the right leading strategies from country to country in a first step towards optimizing the return on investment.

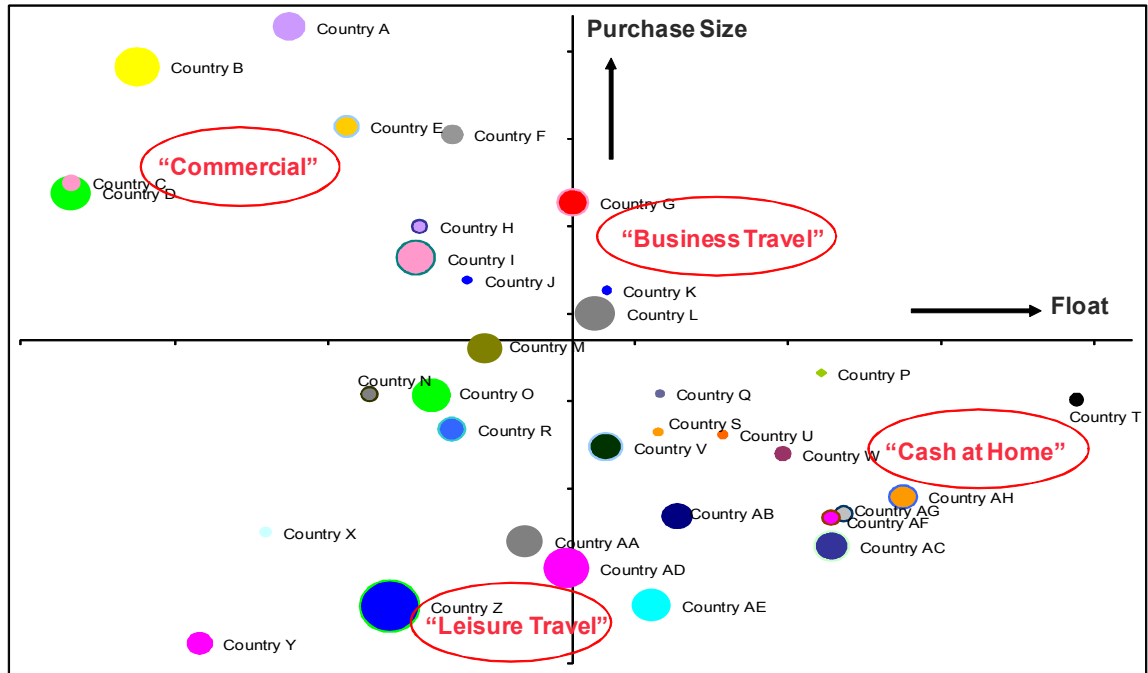


Figure 2: Country Level Segmentation

2. **Challenge:** How do we reach our customer segments?

Solution: Align customer segments with customer touch points

Defining customer segments is the first step towards the ability to develop, target, measure and subsequently optimize. However, in situations where you cannot reach those segments directly, one has an additional step to make. Where is the touch point with the customer, and how can we understand which customers are served by those touch points? This is the step that aims at making the customer insights actionable.

In the case of the TC, the customer touch point is the third party branch location. Customers walk up to the branch and purchase the TCs. Each of these customers is defined by the behavior described above, and the branch can therefore - through aggregation - be defined by the type of customers it serves. For example, certain branches largely serve commercial customers, while others cater primarily to leisure travelers. By defining the type of customers the branch serves, it defines the branch as a target for a particular marketing treatment and therefore part of a segment itself.

Defining each sales branch in terms of the type of customer it serves provides powerful information for marketers. It informs the kind of incentives, the kind of messaging, the kind of offers and the kind of servicing provided at those branches. However, the picture is not quite complete at this point. For marketing investments to pay off, one has to reach enough potential customers to make the investment worthwhile. Providing marketing to a branch that has historically served 3 leisure travelers a month is not likely to pay off similarly as providing marketing to a branch that commonly serves hundreds of customers a week. In segmenting the branches for targeting we therefore need to consider the sales volume of the branch.

In many cases these variables might be sufficient to define branch segments. However, how does one define the time period of the sales volume and what if sales volumes are way up one week, way down the other? That could generate difficulties in targeting the right treatments at the right time. Also it would provide complexity in measuring results as one tries to compare sales lift in a branch with a steady stream of customers with a branch that might be on either a wild up- or downswing in sales. To control for this dynamic, a measure of sales variance is computed for each branch and utilized in the branch segmentation.

3. ***Challenge: How can we ensure proper experimental design setup?***

Solution: Empower marketers with an intuitive Promotion Design Tool

The picture of the branches is now complete, and the next step is to make the picture intuitive to use for marketers. For the most powerful application of this information, a tool is created for marketers that shows every single branch in their region, the segment of customers it serves and its level of sales and variability. Imagine going from just targeting the parent bank partner because of a certain level of sales, to not only the ability to target certain branches, but to target branches that cater to a large number of a certain type of customer! This is pure empowerment for the business, and a key step in optimizing investments. Figure 3 shows an example of the type of information that is provided to marketers. Branches can be highlighted based on a variety of criteria to help select, for example, high sales and high float branches with a high % of the targeted consumer types.

Branch Number	Segment ID	Segment Class*	\$ Sales	Avg. Float	Avg. Purchase Size
111111	17	MHHH	\$62,565	45	1,501
222222	17	MHHH	\$53,391	41	1,700
333333	15	LMLH	\$20,221	24	724
444444	17	MHHH	\$57,294	47	1,470
555555	23	LHHH	\$47,987	10	1,652
666666	7	HLLH	\$12,326	51	582
.... all branches

*Class= Float/Sales/PurchSize/Variability at the following levels: L(Low), M(Medium) and H(High)

Figure 3: Branch Level Customer and Sales Information provided for Planning Purposes

Once this information is available, marketers can start to create their marketing campaign. One of the most important foundations of an investment optimization construct lies in the experimental design of the investments. The promotion design tool is developed to assist marketing managers' in their design of a promotion, which walks the marketing manager through the selection of "test" branches that meet their testing criteria, and ensures an appropriate quantity of accompanying "control" branches. This process ensures the difficulties of statistical significance, proper comparability of the test and control segments, and that a proper experimental design construct can be maintained. From the tool marketers will select the branches they would like to target for the campaign. The tool is critical in this process, as it will guide the manager towards an appropriate experimental design to ensure the ability to measure marketing results accurately post-campaign.

Creating a Marketing Investment Optimization Framework, continued

Figure 4 shows a result of branch selection by segment, including how many more branches remain in a particular segment as 'controls'. Once a minimum of branches in the control is breached, a warning will be shown indicating potential difficulties with assessment accuracy if one proceeds with the current setup. This tool empowers the marketer therefore not only to select the best branches for the promotion, it also provides insight into how many of which kind of branches to leave out of the campaign as controls.

Segment #	Total branches within the country for each segment	Count of Participating Branches by Segment:	Remaining branches for selection in Control	% Test of Total Segment	Float / Sales / Purch / Var
2	63	24	39	38%	HLLL
5	58	32	26	55%	HMMM
6	78	50	28	64%	HHHH
10	90	60	30	67%	MHHH
11	71	41	30	58%	MMMM
13	34	12	22	35%	MLLL
18	31	17	14	55%	LHHH
20	56	29	27	52%	LMMM
22	116	54	62	47%	LLLL
Grand Total	683	372	311	54%	

Figure 4: Segment Level Results of the Selection of Branches to verify Proper Experimental Design

4. **Challenge:** Capture details of the actual campaign executed

Solution: Utilize the Promotion Design Tool to communicate campaign inclusion

This step seems straightforward, but capturing the actual details of an executed campaign is usually quite challenging in the common environment where the analytics function is separated from marketing. Utilizing the promotion design tool not only in the design phase, but also in the execution phase to communicate what was actually done bridges that gap smoothly and ensure proper data capture that is the key input into the subsequent measurement steps.

5. **Challenge:** Translating campaign design into a Test & Control assessment construct

Solution: Cluster distribution alignment between Test & available Control segments

Once the promotion completes, we have a test segment of treated branches as defined by the experimental design setup, and a group of all untreated branches in the region that are candidates for the control segment.

Unexpected sales volatility (such as favorable exchange rates, surges in travel demand, or other macro-economic fluctuations) can add noise to the marketing treatment assessment. To further isolate the impact of the marketing treatment, we aggregate the performance and sales results of the test segment, and derive the appropriate control segment that will not only facilitate cleansing unexpected regional sales fluctuations, but also creates the proxy or baseline of what would have happened if no treatment had been applied.

Depending on the region and newly updated sales data, it's possible that certain variables relationships have shifted. A powerful SAS® software procedure, PROC VARCLUS, is leveraged to confirm the relationship of the 4 clustering variables used in the upfront marketing design: purchase size, float, sales volume, and sales variability.

```
proc varclus data=customerdata proportion=.7 short;
  var  branch_sales_volume branch_sales_variability branch_avg_purchase
       branch_avg_float;
run;
```

Creating a Marketing Investment Optimization Framework, continued

Given that each of the values from the customer defining variables have their own unique scale and standard deviation, a transformation using PROC STANDARD ensures more control over which variables contribute most highly to the definition of the branch cluster/customer profile. Within the standardization step, we take advantage of the opportunity to assign weights to specific variables that may be more important for customer differentiation in a particular region or country by adjusting the "Impact Matrix". By allowing adjustment of the standard deviation of variables planned to be clustered upon, we're effectively changing the importance of a variable in defining the branch clusters, and providing the versatility to customize each assessment when a region has very different usage patterns than another.

When to Increase the Matrix Weight?

	Savings Customers	Commercial Customers	Seasonal Customers	High # New Sellers
Sales Volume				↑
Sales Variance			↑	
Purchase Size		↑		
Float Length	↑			

Variables	Weight
Sales Volume =	4
Sales Variance =	2
Purchase Size =	2
Float Length =	1

The matrix represents the imposed standard deviation on the transformation of each of the variables. This controlled volatility provides a way to "weight" the importance of each variable within the K-Means clustering scheme.

Figure 5: Changing the Weights of the Clustering Variables

```
proc standard data=acct_level2c mean=0 std=4 out=sales_vol_standardized;
var sales;run;
proc standard data=acct_level2c mean=0 std=2 out=sales_var_standardized;
var seller_variability;run;
proc standard data=acct_level2c mean=0 std=2 out=purchase_standardized;
var avg_purch_size;run;
proc standard data=acct_level2c mean=0 std=1 out=avg_float;
var avg_float;run;
```

After standardization, we use PROC FASTCLUS to cluster branches based upon the identified variables. A few subjective PROC FASTCLUS iterations will provide insight to which clusters were collapsed, and which clusters get divided. The more clusters you use in this step, the more customer differentiation you will achieve. However, matching distributions from the test and control segments, which we'll speak of later in the paper, becomes more difficult with the higher number of clusters utilized (A tradeoff to be determined based upon the amount of time you have to assess).

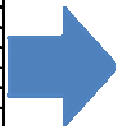
```
proc fastclus maxclusters=50 data=standardized out=clusters50;
var sales seller_variability avg_purch_size avg_float;
run;
```

Once completed with the clustering, we can begin creating our test & control segments. The prevailing concept here, is that we want to create our test & control segments as similar as possible, as the control segment will be representing how the test segment would have performed without being treated.

Given that our branches in similar clusters represent similar customer types & volumes, this becomes a strong method of ensuring similarity between the test & control segment. Since the test segment is what we're trying to measure, we do not alter its cluster distribution. Rather, we take a random 50% or less of the test branches from each cluster. We do not take all the test branches, as the subset will be important in the simulation process.

In creating the control segment, we match the cluster distribution from the test segment clusters within the control segment, creating a very close alignment of sales volume, variability, and customers (purchase size & float similarities), as shown in Figure 6.

All Branches		
Cluster Num	Num Test Branches	Num Control Branches
1	42	87
2	23	35
3	10	2
4	15	39
5	4	112
6	12	15



Sample Size For Simulation		
Cluster Num	Test Input	Control Input
1	21	21
2	12	12
3	5	2
4	7	7
5	2	2
6	6	6

Figure 6: Alignment of Test and Control Cluster Distributions

Depending on the accuracy of the experimental design setup upfront, we may not have enough sellers available in the control segment to properly represent the test segment. In that case we would either accept the under-representation, or borrow branches to back-fill these quantities from the NEAREST cluster. As an example, assuming cluster number 4 is the nearest cluster to cluster 3, we could “borrow” 3 branches from the cluster 4 control segment to backfill its shortcoming in cluster 3.

6. Challenge: How to control for outliers and assess confidence in result?

Solution: Monte-Carlo Simulation + Confidence Assessment

Leveraging the SAS® software MACRO language, 100 random test segments and 100 random control segments are generated from the promo & control pool of branches. For each random test or control segment, distributions across the customer behavior clusters are aligned according to the full test branch distribution to ensure a tight customer type relationship between the test & control segments.



Figure 7: The Creation of 100 Test and Control Samples for Simulation

Upon creation of the 100 test & control segments, a custom time-series forecast is generated for each random segment, resulting in expectations of performance for each subset of branches. SAS/ETS® software is the backbone for this process, and the macro programming language minimizes manual coding for this exercise.




Figure 8: Creating Best Fit Forecasts for each of the Test and Control Samples

```
%forecast(data=Test1, var=sales, id=sales_week, interval=week, stat=arsq,
horizon=60, entry=forecast, out=forecast1, diag=no);
%forecast(data=Test100, var=sales, id=sales_week, interval=week,
stat=arsq, horizon=60, entry=forecast, out=forecast100, diag=no);

%forecast(data=Controll1, var=sales, id=sales_week, interval=week,
stat=arsq, horizon=60, entry=forecast, out=forecast1, diag=no);
%forecast(data=Controll100, var=sales, id=sales_week, interval=week,
stat=arsq, horizon=60, entry=forecast, out=forecast100, diag=no);
```


Creating a Marketing Investment Optimization Framework, continued

An Actual vs. Forecast sales lift is calculated for all 100 of the test segments, and the same is done for each of the control segments, as illustrated in Figure 9. Using the 100 sales lift calculations, we can arrive at a confidence interval based upon the volatility of the sales lift result.

$(\text{Test Actuals1} - \text{Test Forecast1}) / \text{Test Forecast1}$	$(\text{Control Actuals1} - \text{Control Forecast1}) / \text{Control Forecast1}$
...	...
$(\text{Test Actuals100} - \text{Test Forecast100}) / \text{Test Forecast100}$	$(\text{Ctrl Actuals100} - \text{Ctrl Forecast100}) / \text{Ctrl Forecast100}$
Test Sales Lift = Mean	Control Sales Lift = Mean
95% Confidence Intervals Calculated	95% Confidence Intervals Calculated

Figure 9: Calculation of Sales Lift for both the Test and Control

7. **Results Calculation Example**

Figure 10 shows an example of the sales lift calculation using this approach, with Figure 11 showing the simulation results to calculate the confidence interval. The sales lift or sales decline from the control segment represents the unexpected fluctuations found in the region or market.

Results Visualization – Incentive Promotion

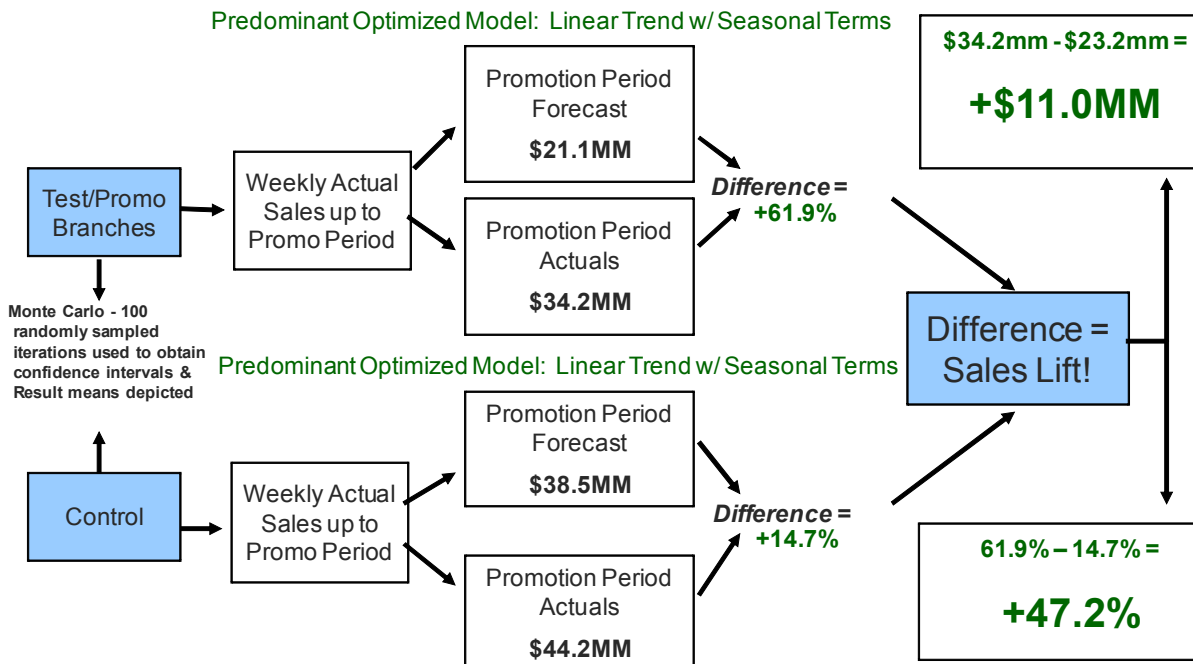


Figure 10: Example of Results Calculation

Promotion (Test) Sales lift results:

61.9% sales lift versus optimized forecast

Monte Carlo Confidence simulation
 (100 Iterations, 4,200 forecasts):
 95% Confidence Upper Bound: 64.4%
 95% Confidence Lower Bound: 59.5%

Control Sample Sales lift results (expected sales performance without promotion's impact):

14.7% sales lift versus optimized forecast

Monte Carlo Confidence simulation
 (100 Iterations, 4,200 forecasts):
 95% Confidence Upper Bound: 16.7%
 95% Confidence Lower Bound: 12.7%

Simulation based Total Sales Lift Interval: 42.8% to 51.7%

Figure 11: Results from the Simulation for both the Promotion (Test) and Control Groups

Creating a Marketing Investment Optimization Framework, continued

8. Challenge: How can I determine the size of the investment opportunity?

Solution: Derive customer and branch promotion sensitivities

The sales lift computed above determines the overall level of success of the promotion. The next natural question would be to determine the roll out potential of the promotion. As part of that question one can also determine which parts of the promotion were more successful than others. This will provide a rank ordering of the opportunity to help optimally invest a finite budget.

The first step in achieving this, is mapping back the results to the underlying customer types. Remember, the very foundation of the methodology was based on the customer types defined by purchase size and float. Using the same general approach as above, one can cluster the customers based on these two variables, and determine which clusters were most impacted, i.e. were most sensitive, to the promotion. Figure 12 shows an example of these sensitivities across different customer clusters.

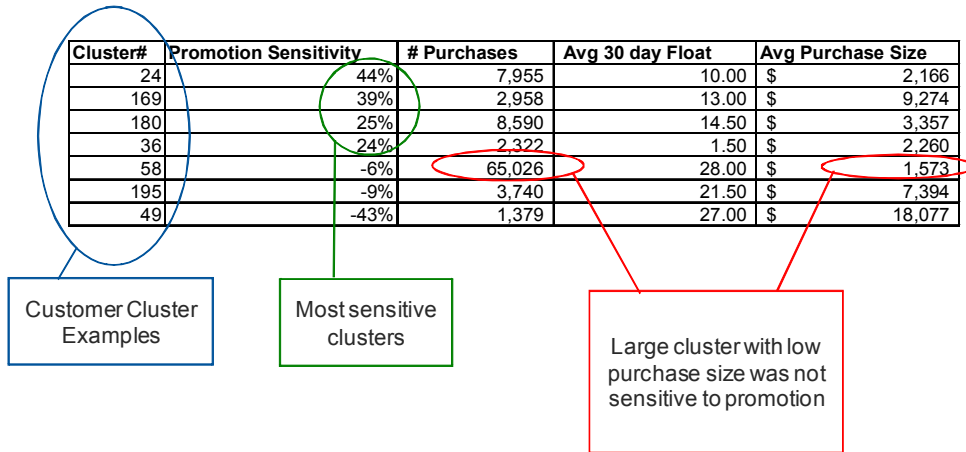


Figure 12: Promotion Sensitivity by Customer Cluster

Since we cannot directly action customer clusters, in the next step we have to translate the customer findings to the customer touch point (branch) level to make them actionable. Simply aggregating the customer sensitivity results at the branch level provides the branch sensitivity based upon the unique mix of customer types they serve. Figure 13 shows an example of branch level sensitivity

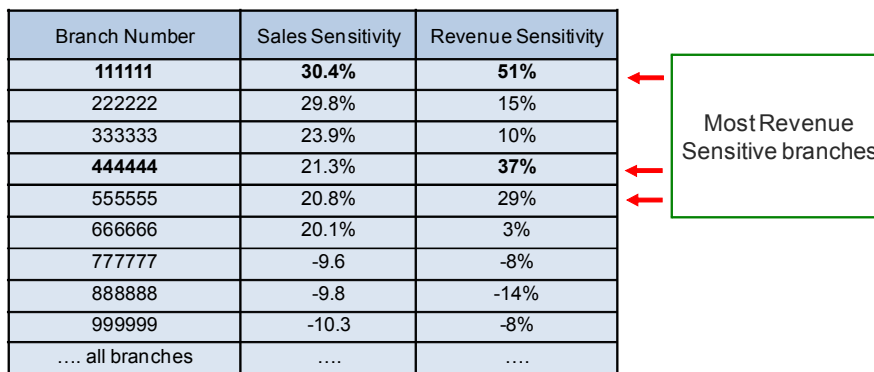


Figure 13: Branch Level Promotion Sensitivities measured by Sales and Revenue

Note that we can define the metric with which to measure the rank ordering. If sales are the objective, the sales sensitivity rank ordering gives you the optimal results. If the objective is more ROI driven, the revenue column is more optimal as this shows the impact of sales to the best – most revenue driving – customers.

Creating a Marketing Investment Optimization Framework, continued

The last step in identifying the opportunity is to apply the customer sensitivities across all available branches, and rank order them based on the desired sensitivity type. Figure 14 shows an example of all branches ranked by sales sensitivity. This provides the business with a rank order of the best targets for this type of promotion. Depending on promotion cost, one can draw the line at a certain ROI cut off and target every branch that meets that threshold – or alternatively if budget is constrained one can target only the very highest ranked branches to get the highest return on the money.

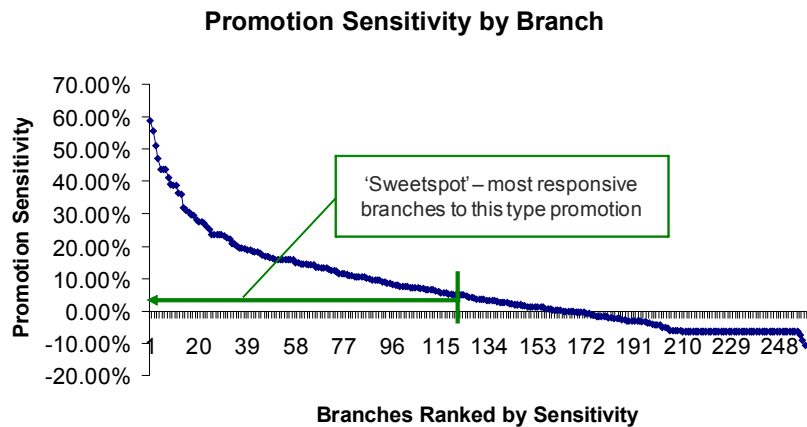


Figure 14: Identification of the Full Roll Out Opportunity based on Branch Sensitivity Rank Ordering

9. ***Challenge:*** *Across a wide variety of promotions, how can I optimize investments globally?*

Solution: *Global Investment Optimization: Rank order branch results and sensitivities across all regions*

Once a wide set of assessments has been completed globally, the optimization component becomes available. Since the methodology is standard across the globe, results in the Far East can be directly compared with results in the West, South and North. Results on branch investment sensitivity, gross opportunity, and ROI can now translate into an optimization model that will direct a finite investment allocation (aka, a budget) into the most impactful segments, regions, and locations!

IMPORTANCE OF INHERENT METHOD FLEXIBILITY

The method described in this paper has worked quite well, helping to book strong results across the world. However, as is often the case in 'the real world' the situation does not always allow for the ideal application of the end to end methodology. In those cases - which one almost always encounters - it is important to have a built in flexibility. In the case of Travelers Cheques, it was not always possible to hold out a control sample. In other cases, testing in a certain region only involved highly variable sellers. In yet another, the total number of branches to choose from were too limited. In all these cases, while not ideal, we were still able to apply a scaled down version of the methodology and keep it part of the overall investment framework.

Creating a Marketing Investment Optimization Framework, continued

Figure 15 illustrates how the methodology is scaled down under certain conditions. Reading from left to right, the methodology goes from the desired full application all the way to the bare minimum on the right. Three main conditions drive the selection of the appropriate method level: quality of the experimental design, number of available branches and sales variability of branches. If sales variability is high, trimming of outliers might not make sense and results will have a wider confidence interval. Once the number of branches get small, a Monte Carlo simulation might not be possible anymore, losing some insight into the variability and stability of the end result. And if the experimental design of the campaign is set up poorly, one might not be able to measure results on the test group against an untreated group, leaving the method basically with comparing forecasted sales with actual sales to compute lift. The latter is definitely not desirable, but in some situations the best method possible.

	Highly Desirable		Less Desirable	
	<i>Method 1 Best</i>	<i>Method 2</i>	<i>Method 3</i>	<i>Method 4</i>
Situation				
Marketing Design	Excellent	Excellent	Good	<u>Poor</u>
# Branches in Data	Large	Large	<u>Small</u>	Small
Sales Variability	Medium	<u>High</u>	Med/Hig	Med/Hig
Techniques				
Optimized Forecasting	√	√	√	√
Test & Control Segment	√	√	√	
Monte Carlo simulation	√	√		
No Outlier Trimming	√			
Allows for				
Accuracy	Optimal	Excellent	Good	Directional
Multiple Treatments	Excellent	Very Good	Good	Poor
Consumer Segmentation	Excellent	Very Good	Reasonable	Poor

Figure 15: Scaling down the methodology based on limitations on design and data

CONCLUSION

True investment optimization at a global scale is a journey of analytics. Applying innovative methods across the full circle of optimization can generate insights that one might have previously thought not to be possible. Even in cases where the customer is largely anonymous, the level of insights, service and customization that can be achieved through innovative analytics is remarkable. Investment optimization is truly a journey. With the ever changing environment and customer preferences, one need to continuously monitor, test, design, measure and optimize. This paper has laid out a methodology that has worked well in the real world. It uses methods such as segmentation, clustering, forecasting and simulation combined with a healthy portion of intuitive thinking. With tweaks here and there, this could potentially be applied across a wide level of industries. No doubt the methodology itself will continue its own journey as well as it adapts and learns and new data becomes historic and future observations current.

CONTACT INFORMATION

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

Pieter van Ispelen
 American Express
 4315 S. 2700 W.
 Salt Lake City, UT 84121
 Work Phone: (801) 945 5129
 E-mail: Pieter.P.vanIspelen@aexp.com
 Personal E-mail: Pieter@VanIspelen.com

Landon Starr
 American Express
 4315 S. 2700 W.
 Salt Lake City, UT 84121
 Work Phone: (801) 945 5131
 E-mail: Landon.V.Starr@aexp.com
 Personal E-mail: Landon_Starr@hotmail.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.