Paper SAS6581-2016

SAS® Analytics for Multi-armed Bandit Market Analysis

Thomas Lehman, SAS Institute Inc.

ABSTRACT

This presentation describes the new SAS Customer Intelligence 360 solution and its automated analytic test functions. The new SAS solution provides three automated testing methods; Traditional A/B testing, Multi-Armed Bandit Testing and Multi-Variant Testing. This paper will focus primarily on multi-armed bandit (MAB) analysis of controlled experiments. Multi-armed bandit analysis has been in existence since the 1950's and is sometimes referred to as the *K*-or *N*-armed bandit problem where a gambler at a row of slot machines (sometimes known as "one-armed bandits") has to decide which machines to play, as well as the frequency and order to play each machine with the objective of maximizing returns. In practice, multi-armed bandits have been used to model the problem of managing research projects in a large organization. However, the same technology can be applied within the marketing space where, for example, variants of campaign creatives or variants of customer journey sub-paths are compared to better understand customer behavior through collecting their respective response rates. Distribution weights of the variants are adjusted as time passes and conversion events are observed to reflect what customers are responding to thus optimizing the performance of the campaign. The SAS Customer Intelligence 360 analytic engine involves the collection of data at regularly scheduled time intervals and processes the data through a simulation engine to return an updated weighting schema. The process is automated in the digital space and led through the new solution's content serving and execution engine for interactive communication channels. Marketing gains are studied during the process to illustrate the business value of the multi-arm bandit test. SAS has provided capability to switch from traditional A/B testing to MAB testing when conditions warrant by feeding the required inputs and results to date to feed the multi-armed bandit engine which will be discussed.

INTRODUCTION

SAS has a strong set of Customer Intelligence tools which continuously perform in outstanding fashion in all Analyst reviews, always landing in the magic quadrants. Tools like Marketing Optimization, Marketing Automation, Real-time Decision Manager, Intelligent Advertising and several others provide a marketing organization or department the intelligence required to be effective when interacting with customers. These tools study behavior at the micro and macro level to understand trends to provide intelligence which can drive the bottom line by providing more conversion activity.

Over the last several months, SAS has undergone a major effort to integrate these tools together into a solution known as CI360. CI360 has one interface largely driven through SAS Visual Analytics. SAS Visual Analytics has been designed for expertise of all levels. A beginner data scientist can immediately create reports while a seasoned Data Scientist can explore the data in great detail, creating inferential models and other more advanced functions. SAS is not new to the CI space. The difference SAS brings is 40+ years of Analytic knowledge.

To use the CI360 Analytics application the user sets up parameters and goals in a very intuitive interface and the individual tool usage is transparent. Analytic knowledge gained includes segmentation structures and performance gain opportunities, tests of performance and inferential statistical based drivers of successes. CI360 focuses on the customer journey, integrating past actions with event stream data to make optimal decisions. SAS CI360 is a hosted solution run in various Amazon Web Services instances across the world dictated by customer proximity.

This paper will focus primarily on one of the marketing effectiveness test routines; Multi Armed Bandit. The interface to run this test is very intuitive and the underlying analytics very powerful making the setup and performance of the test ideal for a beginning or seasoned marketer or data scientist.

Figure 2 shows the CI360 interface.

Figure 1. CI360 Interface

METHODOLOGY

The periodicity of how often MAB results are updated is typically dictated by data volumes. SAS by default has chosen 12 hour intervals. Most process every 24 hours. The benefit of more frequent updates, if the data volumes justify, is that the weights and therefore the content delivery distributions are slightly more optimal which produces greater gains throughout the course of the test and can allow it to end slightly quicker. SAS runs MAB as a hosted service where all data is managed in Amazon Web Services in a multitenant design in what we refer to as Data Hub. Data Hub supports all CI360 analytics and other functions.

MAB is a self-learning method. Once it is set up the user can let run hands-off. It features a Thompson Sampling Monte Carlo Simulation Engine. The way this works is for each variant, or creative, being tested in a MAB test a dataset is created and each variant gets a column. Create many rows, usually 100,000 is plenty. Populate the data cells with values from a Beta distribution, using appropriate parameters for the distribution based on previous performance for each variant respectively. The wining variant of each row is recorded and a frequency count of the winners is performed and those frequencies become the new content delivery weights. Though the distribution is Binomial – convert or non-convert, the transformation from Binomial to Beta is required to feed the simulation engine.

MAB does support adding and removing variants during the course of a test and can help aid in those decisions. SAS provides automatic conclusion of test using 5% and 95% confidence limits by default. In the graphic shown below the MAB has four variants which each started with equal distributions weights of 25%. As the test ran Variant C outperformed the other three and after a month emerged as the winner once its probability of being the optimal variant crossed the 95% confidence limit. The other three naturally feel beneath the 5% limit.

Figure 22 shows MAB output.

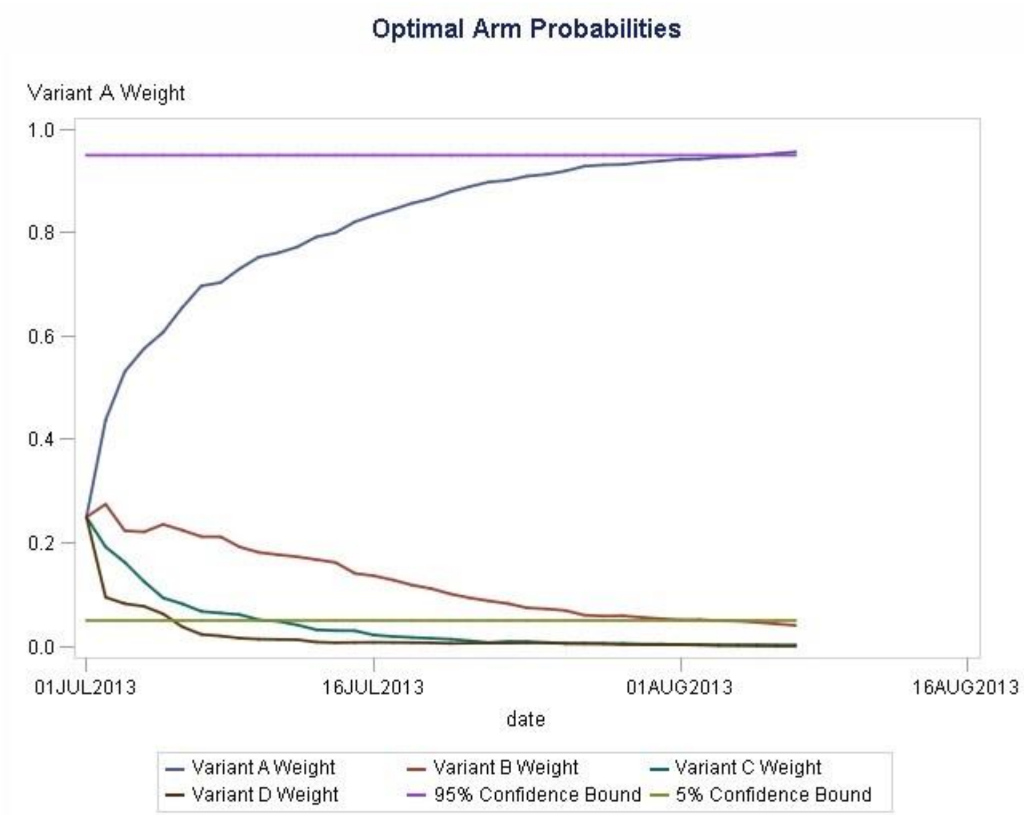

Figure 2. MAB Weights Adjusted over Time Based on Response Rates

CONVERSIONS GAINED METRIC

The Conversions Gained macro code was written as an enhancement for Multi-Armed Bandit (MAB) testing that can be applied in other areas of CI TNG Analytics. The basic premise behind the code is to quantify the delta in conversion activity as a result of running the test. The calculation takes into account variant response rates, variant content serving rates and population sizes of unique visitors both current and historic. The calculation is meant to provide an estimate of conversions for the current time frame had the initial test weightings been held constant versus the variable weightings from MAB provided to drive additional response activity. Once this calculation is derived across the variant population of the test it is compared with actual conversion activity for the current time period and the delta between the two provided a time based view of the value the test is providing. The resulting time series data provided integer based data that can be plotted together with time based variant weights and confidence values (probability of being optimal) to visualize the data. For example, consider the line plot below. It shows the performance of five variants over a nine day period. This plot will drive the logic discussion below. Note that in this test there are five variants shown, confidence limits at 5% and 95% and conversions gained. The confidence limits and variant weights all anchor on the left vertical axis where conversions gained anchors on the right vertical axis. In this test all variants started with a uniform weight of 20%. Since the process makes all comparisons to the first time period, naturally the value of conversions gained for the first time period will be set to zero since the test is comparing that period to itself.

Figure 23 shows the Addition of the Conversions Gained Metric.

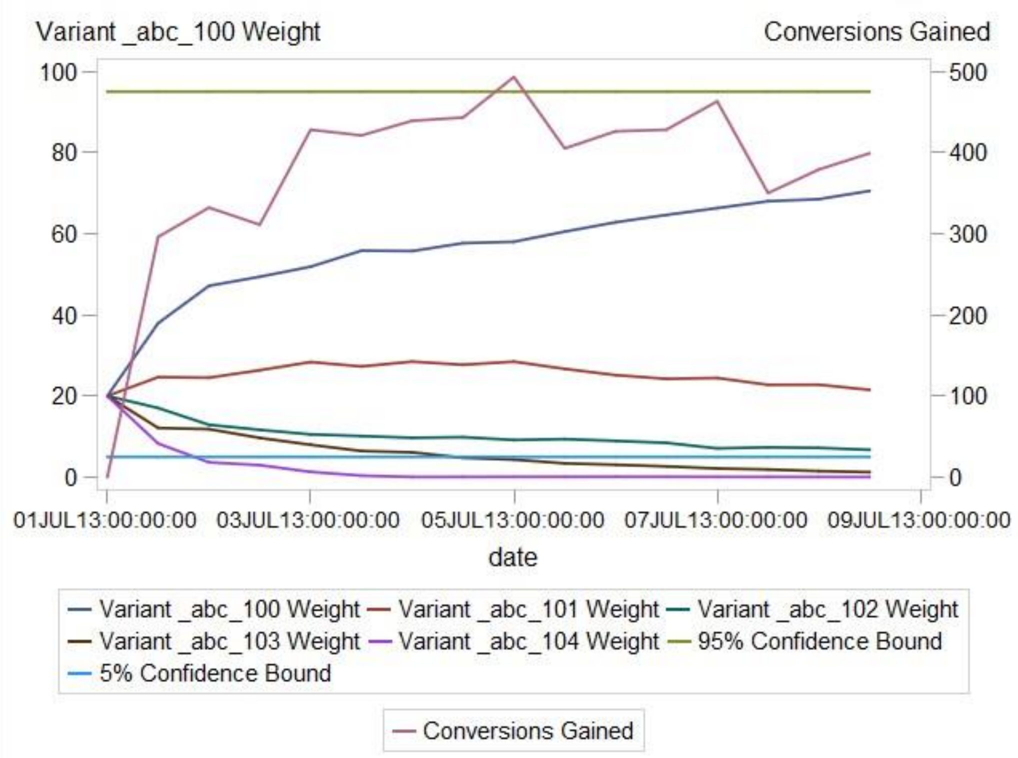

Figure 3. MAB Conversions Gained Over Time Based on Delivery Rules

MAB WORKFLOW

MAB Tests are designed to run in automated fashion. A user interface is the conduit between the marketer and the data at the onset of the test. Data is captured for the variants (Arms) to include in the test and their respective initial distribution weights, test run periodicity if other than 12 hours and a confidence stop rule threshold can be set. Next the Data Hub is queried to collect the data specific to the defined task for the last *N* hours where *N* is defined as the weight update periodicity. That data is then passed to an analytics layer where the data is summarized to get the number of impression counts, conversion and failure counts per arm and the Beta distribution parameters for each arm are calculated. Those parameters are then fed through the Bayesian Bandit simulation engine to gather the updated weights (probability of being optimal) and the gain and cost of the experiment are studied. Next, the resulting data is sent to a persistent holding area where the data is accumulated for each run to serve the reporting layer. Figure 24 shows the Multi-Armed Bandit workflow.

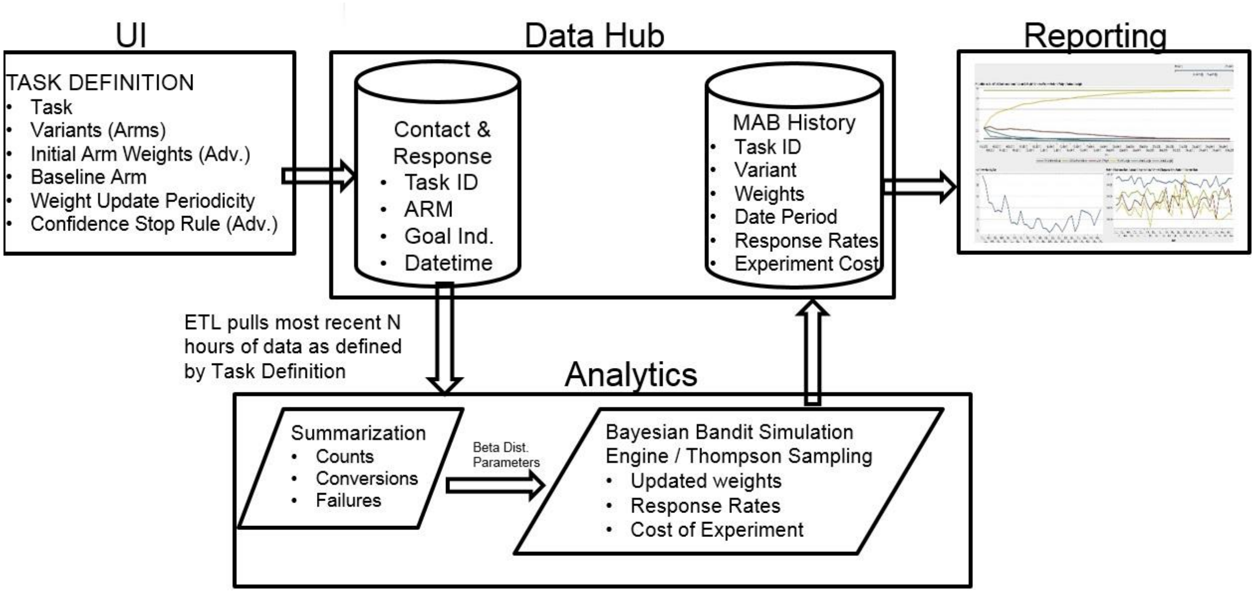

Display 4. Workflow for Multi-Armed Bandit Analysis

CONCLUSION

Multi-Arm Bandit analysis can be a very useful tool for marketers to gauge the effectiveness of different versions of an advertisement within a campaign. Its construct often allows for quicker execution than a traditional A/B test. SAS has an analytic construct which supports more than two variants in the A/B test construct. In the event that an A/B test seems to be running excessively long the user does have the ability to switch from the A/B construct to an MAB construct while the test is in flight.

MAB analysis can be very sensitive at the onset of the test, especially when response rates are very low. One must also be careful to watch for ties for more than one variant in an MAB test. That is if two variants are performing equally well then over time both would converge towards 50% probability of being optimal which would prevent any stop rules from occurring. Under these circumstances it may be advised to consider a Multi-Variant test approach.