TURF: New Methods for Implementation

TRC
Consumer Insights.
Market Innovations.

By Westley Ritz

White Paper Series

Introduction

TURF is a long-established and quite useful marketing research tool, but not everyone is familiar with how it works, or with the latest developments that can make TURF even more effective. The purposes of this paper are twofold: (1) to explain the technique and (2) to describe the latest methods for implementation

An Illustration

Let's start with a simple product line optimization problem. Suppose the owner of a small ice cream cart wants to optimize sales by offering flavors that will appeal to the largest proportion of customers. The cart has space for 2 flavors, and the owner can make 3 different flavors for equal cost. The 3 flavors are Vanilla, Chocolate, and Pistachio.

Figure 1. Ice cream flavor preferences.

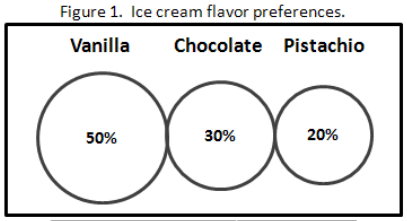

Flavor Combination	% Coverage
Vanilla & Chocolate	**80%**
Vanilla & Pistachio	70%
Chocolate & Pistachio	50%

Figure 1 shows our 3 flavors in a market where each customer prefers only 1 flavor. The obvious solution would be to offer Vanilla and Chocolate, the 2 most popular flavors, to maximize customer coverage at a sum of 50% + 30% = 80%. However, in a more realistic marketplace, there are likely to be overlaps in customer preferences. Let's assume that preferences for Vanilla, Chocolate, and Pistachio overlap as shown in Figure 2.

Figure 2. Ice cream flavor preferences accounting for overlaps.

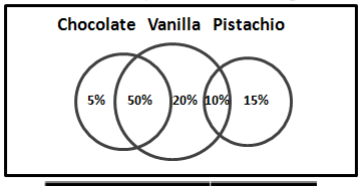

Flavor Combination	% Coverage
Vanilla & Chocolate	85%
Vanilla & Pistachio	**95%**
Chocolate & Pistachio	50%

Now, in a market where customers have multiple preferences, offering Vanilla and Chocolate will reach 85% of customers versus the 80% in a single-preference market. However, Vanilla and Chocolate is no longer the best solution. Instead, offering Vanilla and Pistachio will reach 95% of customers because many Chocolate lovers will buy Vanilla in the absence of Chocolate. It is now obvious that the optimal solution is to offer Vanilla and Pistachio; we only lose 5% market share due to those preferring only Chocolate.

This simple example illustrates how product line optimization can quickly become complex due to customers' preferences for more than only 1 item. In practice, the common approach to predicting sales is to rely on feedback from surveys. However, as in the ice cream example, finding the optimal solution is not as simple as selecting the top scoring products to be included in the product line. This is where an analytical technique like TURF is often implemented.

What is TURF?

TURF stands for Total Unduplicated Reach and Frequency. The analysis runs through every possible combination of products and records for each combination (1) the percentage of those that desire at least 1 product in the given combination (i.e. reach), and (2) the total number of times products are desired in the given combination (i.e. frequency). The definitions for what demonstrates a strong desire for a product (i.e. what constitutes a reach) can be altered to match any a priori notions about matching survey responses to behavior.

Recall the ice cream example in which adding the flavor Pistachio along with Vanilla reaches the most consumers, even though Pistachio by itself is the lowest preferred flavor. Products like the Pistachio ice cream are referred to as niche products; they appeal to a small group that otherwise is indifferent to the more popular products. In general, while overall desire for the product is not high, there may be enough demand that when the majority of the market is satisfied with a couple of the top products being in the product line, the niche product reaches out to a significant portion of the market. No doubt, this is a portion of the market that would have been ignored and revenue lost if decisions were simply based on the products with the highest preference. TURF analyzes all combinations and yields the best combinations as those that have the highest reach value, and secondarily, the highest frequency value. It is common for the top combination to include some product(s) that did not score highly in the desirability ratings. The ability to identify these niche products demonstrates the value of TURF.

Current Practice

Rating scales are a common tool in marketing for measuring preferences of products, or features of products. TURF is also commonly used in conjunction with rating scales for product line

optimization. The first and most important step to applying TURF is in defining what constitutes a reach. A reach occurs for an individual when she desires at least 1 of the products in a given combination of products. Since the interest is in estimating sales, the cut-off point to signify a reach should demonstrate desire for the product. Using a rating scale, say a 10-point scale from 1 to 10, we may define a top-2 box response (a 9 or a 10) as demonstrating desire for the product, and anything less than a 9 shows no desire for the product. While this definition may be stringent and will surely exclude some who may actually purchase the product, it is nonetheless a good estimate for take rate. Now that it is known who desires each product, the TURF analysis may proceed in calculating reach and frequency results for each combination of products.

While rating scales are extensively used in marketing research, their flaws are also quite well-known. Rating scale scores typically lack much variation and hover towards the higher end of the scale (the "halo effect"). The main reason for this is that there is no comparison of products forcing the respondents to focus on both ends of the scale. Because of these issues, other techniques have become increasingly popular among researchers that provide much better data for attribute and product comparisons. Two of these are Maximum Difference Scaling (or Best-Worst Analysis) and Conjoint Analysis. The next section provides details and examples on how TURF may be used in conjunction with each of these techniques.

TURF with Maximum Difference Scaling

The Max-Diff approach to measuring importance groups features or products into small sets and asks respondents to specify their most and least important features from each set. Because the data is focused on both ends of the importance spectrum, we avoid the problematic response tendencies commonly seen in rating scales. Utility scores are then computed for each individual through some statistical procedure, such as Hierarchical Bayes, which depicts the magnitude of relative importance for each feature conveniently scaled from 0 to 100, summing to 100. [Refer to *Identifying Feature Importance: A Comparison of Methods* for more details]

The key step in TURF analyses is in identifying the cut-off point that signifies a strong desire for the feature (i.e. reach). As mentioned previously, the reach for 10-point scales is usually defined as a top-2 or top-3 box response. Alternatively, in Max-Diff the typical approach is to select some quantifiable cut-off, such as a score equal to or greater than 2 or 3 times the expected value. For example, if there are 10 features, the expected value of any one feature is 10, and 2 times that would be a score of 20. Now that the cut-off or reach value has been identified, data are recoded to identify each feature as either a reach or not for each respondent, and the TURF analysis is conducted as normal.

Example

Take as an example a study on health insurance plan components. The client, a large health insurance provider, wished to understand what features were most important for consumers in choosing health plans, and which features should be included in future plans the company may offer. A Max-Diff procedure was performed on 14 features, followed by a TURF analysis. As mentioned before, the key to TURF analyses using Max-Diff results is to identify a suitable cut-off or reach value. The reach for this study was chosen to be twice the expected value of a feature, or (100/14)×2 = 14.29.

Table 1. TURF results summarized: For each best combination of features, the % of max reach and Max-Diff feature ranks are given.

Best Combination of Features	% Reach	Feature Max-Diff Ranks
1	62%	1st
2	82%	1st 2nd
3	87%	1st 2nd 4th
4	91%	1st 2nd 4th 9th
5	93%	1st 2nd 4th 9th 11th

Table 1 summarizes the results of the TURF analysis. The metric '% Reach' measures the percent of respondents who were classified as a reach on at least one of the features in the combination. The best combination of 2 features reaches 82% of the sample, and is comprised of the top 2 ranked features from the Max-Diff scores. This is not surprising. However, when we look at the best combinations of 4 and 5 features, the percentage reached increases by 4 and 2 percentage points while including the 9th and 11th ranked features, respectively. This is a prime example of TURF's ability to identify features that may relate to a niche group of consumers, but is still significant in increasing a product's total consumer base.

TURF with Conjoint Analysis

Conjoint Analysis is a technique that asks respondents to select or rate potential products that are built using a predefined list of attributes and attribute levels. Estimation is typically done at an individual level via some statistical procedure, such as logit or Hierarchical Bayes, to provide utilities for each level of each attribute. These utilities, used in conjunction with a market simulator tool, are used to measure respondents' preferences on any product of interest that can be built from the predetermined attributes and levels. [Refer to *Deriving Value from Research: The Use of Conjoint Analysis for Product Development* for more details]

Using a set of products and their associated preference scores measured on an individual basis, TURF may be applied to determine which products should make up the market offering. Note that this differs from Max-Diff and Rating Scales in which TURF is applied directly to the results. With conjoint, the attribute utilities are first used to build and score specified products, and then these subsequent product preference scores are used in the TURF analysis. Once the product preference scores are obtained, the procedure then is similar to Max-Diff or Rating Scales in which an appropriate reach or cut-off value needs to be identified that signifies a strong desire for the product.

Example

Take another example from the health insurance industry, but focusing on a different health insurance provider that is now interested in determining which insurance plans to offer in the state of New York. Using attributes common to health insurance plans such as deductibles, co-pays, and out-of-pocket costs, a choice-based conjoint analysis was performed and a set of 10 possible plans was created as potential product offerings. Through the use of a market simulator tool, individual preference scores on each of the 10 products were calculated and scaled from 0 to 100. After a close review of the scores, a product was considered a reach if it was the maximum score for the respondent or was within 5 points of the maximum, and was greater than 50 points. So said another way, if a product was among the top products for a given individual and it obtained a preference score greater than 50, then that product was considered desirable to that individual. Now that a reach value has been obtained, each product is identified as either being a reach or not for each respondent, and the TURF analysis proceeds in the usual manner.

In Table 2, the best combinations of 1 to 4 products are described by their percentage of maximum reach and their associated rankings from the product preference scores. The best single product that reaches the largest percent of the sample also had the highest preference score. This result is obvious. However, the pattern quickly changes as the combination of 2 products that have the largest reach (65%) is composed of the 1st and 3rd ranked products. Then, looking at the best combinations of 3 and 4 products, the percentage reached increases by 3 and 1 percentage points while adding the 7th and 6th ranked products, respectively. This outcome is similar to the TURF results when using Max-Diff in that products associated with niche groups of customers outside of the top products add more potential consumers than only choosing the top products.

Table 2. TURF results summarized: For each best combination of products, the % of max reach and product preference ranks are given.

Best Combination of Products	% Reach	Product Preference Ranks
1	51%	1st
2	65%	1st 3rd
3	68%	1st 3rd 7th
4	69%	1st 3rd 7th 6th

Limitations and Conclusion

The goals of this paper were to review TURF analysis and introduce new ways of implementing the method. Max-Diff and conjoint analysis, two well-known survey techniques, were applied in conjunction with TURF to demonstrate additional situations in which TURF may be utilized. TURF is a powerful tool for identifying the products to optimize a product line, and ensuring that the largest number of potential customers is being accounted for. As seen in the previous examples, often the top ranked products alone will not reach the maximum consumer base. TURF identifies those niche products that should be in the product line to appeal to the most people.

Despite the many advantages of TURF, it is not without its shortcomings. A couple of its limitations have been touched upon in this paper, but are worth mentioning directly: (1) the identification of a cut-off value that specifies which products are a reach is completely subjective and may change from study to study, and (2) the percentage of respondents reached may dwindle quickly as the number of products in the combination grows. However, when using sound judgment as to when and how to apply it, TURF remains a very strong and useful product line optimization tool.