A Winning Tool for CPG

The Shapley Value game theory method brings advantages to marketing managers evaluating product line flavor decisions.

By Michael Conklin and Stan Lipovetsky

Executive Summary

The authors introduce the Shapley Value analysis, a method for evaluating product line flavor decisions that has many advantages over the widely used "TURF" approach. The technique can be used in any industry where customers choose to buy one of N offerings. The context of this article refers to consumer food products, but the technique has been used for financial services and consumer durables as well. Shapley Value provides a measure of the overall strength of each possible flavor when considered within all potential product lines. As such, it takes into account scenarios such as out-of-stock situations and competitive flavors. The approach also balances a flavor's ability to appeal to a unique group of consumers with its overall appeal in the population, limiting the appearance of "strange" flavors in the final recommended solution.

Consider the ice cream shop owner who must decide which flavors to sell. The choice can be bewildering, but the storage space is limited. This problem is no less difficult for the ice cream manufacturer who must decide which flavors should be produced, given limited plant resources. Ideally, the product manager would choose the flavors that maximize total sales or profit across the entire product line. The difficulty in solving the problem is that the incremental value of a particular product flavor depends on which other flavors are already in the line. Having 31 flavors may be a great marketing ploy, unless all 31 are variations of strawberry.

This type of problem is one that many marketing managers in the consumer packaged goods industry face on a daily basis. While adding more flavors to product lines such as ice cream, cookies, juices, and carbonated soft drinks will be potentially more appealing to consumers because of the new flavors, such a move creates significant costs. Packaging, advertising, and inventory management costs all rise, as the number of flavors increases. As a result, marketing managers often ask marketing researchers to find an optimal set of flavors that provides the best combination of flavors for a fixed number of products in the line. Typically, total unduplicated reach and frequency (TURF) analysis is used as a tool to recommend solutions to the marketer's dilemma. While TURF answers the specific question that the marketing manager has asked, it doesn't provide information that can help with the strategic decisions that marketing managers must make for their product line. Another alternative, Shapley

Value, serves as a solid solution because it assigns a value for each flavor calculated over all possible combinations of flavors.

TURF ANALYSIS

TURF, a brute-force combinatorial technique, was developed for media-mix models to find the best combination of magazines to place ads to achieve the maximum audience reach and frequency of exposure. When applied to product line decisions, the procedure measures the total number of purchasers (or purchases) for every possible combination of N products taken R at a time. (N is the total number of flavors being considered and R is the number desired in the final product line.) The TURF algorithm simply cycles through each possible combination of R flavors and counts how many consumers would buy at least one of them.

TURF can be conceptualized using Venn diagrams. Exhibit 1 features purchase interest in three flavors (labeled A, B, and C). The area of each circle in the diagram represents the number of people who would purchase the product. If the marketing manager can afford to market only two flavors, then choosing flavors A and C are the best choice. This is because A and C appeal to different sets of consumers. Because A and B overlap substantially, the incremental gain from adding C to A is much greater than any other combination of two products.

The results from a TURF analysis are not always as clear as in this example. In marketing applications, a substantial overlap tends to occur in the appeal of different flavors. This makes sense, because most of the product's appeal results from the base product itself, not the specific flavor. Considering ice cream again: If I want to buy an ice cream cone, and the ice cream shop has only one flavor, I probably still will buy it because I want ice cream. This common overlap in flavor appeal can make TURF analysis very difficult to interpret. Many potential product line combinations can have very similar TURF scores. Exhibit 2 features the results of a TURF analysis from an actual study. The table shows the total reach of all 10 possible three flavor product lines among five flavors that were considered. Three potential product lines are indistinguishable, even if we don't take into account the possibility of sampling error. This type of behavior limits the usefulness of TURF as a decision tool.

By examining the TURF analysis results, you'll see that certain flavors tend to be in the top combinations, while other flavors tend to be at the bottom. This suggests a method for summarizing TURF results. Simply add the number of times a flavor shows up in the set of best combinations. Exhibit 2 shows all possible combinations of five flavors taken three at a time. In more complicated problems, it's not practical to list all of the possible product lines. For example, 38,760 possible six flavor product lines are drawn from 20 possible flavors. With a large number of possible product lines, it becomes difficult to decide how many combinations to look at to generate a summary measure.

SHAPLEY VALUE

We turn to the field of cooperative game theory for a decision tool that provides clearer results in these situations. We can think of the product line flavor decision as a way of building coalitions among players (flavors, in our case) in a game to maximize the total value of the coalition. In the field of cooperative games, a well-known analysis and decision tool for this problem is Shapley Value (SV).

The SV tool was developed to provide an ordering of the worth of players in a multiplayer cooperative game. The key to understanding its utility is that it represents the worth of each player over all possible combinations of players. Extending this to the product line flavor problem, SV assigns a value for each flavor calculated over all possible combinations of flavors.

Conceptually, SV is very simple. To evaluate the strength of any particular flavor, we compare the average strength of all possible product lines that include the flavor to all possible product lines that don't have it. We can evaluate the strength of a product line the same way we do in TURF; that is, we look at how many consumers would find at

Exhibit 1

Hypothetical interest in 3 products

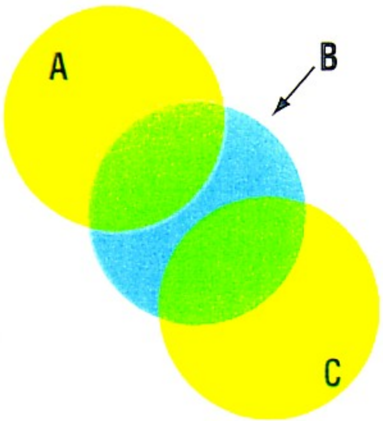

least one product in our line to be desirable. (For more detail on the specifics of calculating Shapley Values, see our reference in the Additional Reading, page 27.)

The SV solution to the problem is closer to the actual marketing situation. In most real-world situations, marketers don't have complete control of all the variables. While a marketer may introduce a product line of a specific size, there's no guarantee that all of the flavors will be available when a consumer makes his or her choice. A specific retail outlet may choose to carry fewer flavors as a result of shelf space limitations, or varying product flavors may be available as a result of out-of-stock situations. SV includes this possibility in its solution set.

The SV solution also is better because it includes comparisons with all flavors, even if they're not in the final product line. This is an advantage of SV. By comparing across all possible combinations of flavors, we include the possibility of competitive action in our analysis. The marketer wants products in his or her product line that provide the maximum incremental value, even when considering competitive offerings. By calculating the SV across all possible combinations of potential product flavors, we assure that we're positioned in the best possible position, even if competitors introduce some of the flavors that we have declined to produce.

TURF analysis provides a way to choose a product line over all possible product lines of a specific size. In contrast to TURF, the Shapley Value analysis provides an ordering of the product flavors that is a "best" solution over all possible combinations and sizes of product lines. Because of this averaging over all possible combinations, Shapley Value provides a theoretically more consistent decision rule.

Theoretical advantages are always nice, but they're even better if they lead to practical improvements. While TURF analysis is a widely used tool for analyzing product flavor decisions, it does not always lead to clear results. Interpretation of the example presented in Exhibit 2 is difficult because the TURF analysis yields no single best choice, and all the solutions are very close to each other.

This can be further demonstrated by bootstrapping the TURF data and determining the frequency with which specific "best" solutions appear. (Bootstrapping is a statistical technique in which repeated samples are drawn with replacement from the data to see how likely we could get similar results with a different sample of consumers.) Exhibit 3 features the results of 500 bootstrap samples of the same real data set. The TURF analysis only actually effectively differentiates the worst possible combinations. Six of the 10 possible combinations are rated best more than 10% of the time. Only the four worst combinations are effectively ruled out. The maximum time that any single combination was judged best was 26.8%, and there were numerous ties within the bootstrap simulations.

Calculating the Shapley Values for each of the flavors provides an ordering of D > A > B > C > E. From this, we conclude that a three-flavor product line should include flavors D, A, and B. In Exhibit 2, this particular combination of flavors

Exhibit 2

TURF results

Product line	# Reached	% Reached
ACD	119	42.8
BCD	119	42.8
BDE	119	42.8
ABD	118	42.4
ADE	118	42.4
BCE	117	42.1
ACE	115	41.4
ABC	114	41.0
ABE	113	40.6
CDE	111	39.9

Exhibit 3

Bootstrap evaluation of TURF

Product line	% of bootstrap iterations where line is rated best by TURF
BCD	26.8
ACD	25.8
BDE	25.8
ABD	15.2
ADE	14.2
BCE	12.4
ACE	3.4
ABC	1.8
ABE	1.0
CDE	0.6

was not one of the three "best" solutions from the TURF analysis, and it also wasn't a leader in the bootstrap TURF analysis summarized in Exhibit 3. This is because TURF only maximizes over a three product line while the SV is calculated over all possible product line sizes.

Furthermore, the ordering of the flavors by Shapley Value is more consistent than the TURF analysis. Again, we demonstrate this by bootstrapping 500 iterations of the SV calculations. All paired comparisons of order are significantly different ($p<.01$). (Results of this bootstrap analysis are summarized in Exhibit 4.)

Shapley Value can be thought of as a weighted sum of terms for each potential size of product line, from 1 to N products. We have discussed why these weighted sums across all possible product line sizes are preferable to TURF results because we cannot guarantee a specific product line size as a result of out-of-stock situations or competitive product introductions. However, it's also unlikely that all potential flavors would be present as a result of shelf space limitations or that only one flavor would be available from manufacturing inefficiencies. To address this potential difficulty, we added some constraints to the Shapley Value calculations.

Our proposed Constrained Shapley Value (CSV) allows the product manager to specify probabilities for specific line sizes. Suppose a marketer is considering introducing a product line with six flavors out of 20 potential flavors. It seems unlikely that stores would carry only one flavor, and perhaps equally unlikely that 20 flavors would be available. On the other hand, the likelihood of exactly six flavors being available in every store is also small. The CSV allows a product manager to gauge the strength of the specific flavors considering only likely product line sizes. The marketer could specify a 30% probability that only five flavors would be available, a 40% probability of six flavors, and a 30% probability of seven flavors. We can then perform the Shapley Value calculations by weighting the results for each line size by the probability. In fact, if we assume a 100% probability of a product line of six flavors, the Shapley Value becomes an excellent summary of a TURF analysis.

The summarization of overall strength of each flavor that is provided by Shapley Value has further advantages over the TURF approach of finding the "best" possible product lines. Because TURF finds the flavor that provides the greatest additional reach from a given base, it tends to add products to the line that have unique appeal to a small group of customers. In essence, TURF tends to add unique and different flavors to your product line. Shapley Value evaluates the overall strength of the flavor across all possible line sizes. This means that strange flavors that appeal to a very small group of consumers still will appear relatively weak in an SV analysis, while they may appear in the best TURF combination. However, as the size of the group that likes the "strange" flavor increases, the Shapley Value of the flavor will increase dramatically. Essentially, the Shapley Value balances between the uniqueness of the flavor's appeal and the overall breadth of appeal. Flavors that uniquely appeal to specific market segments will rise to the top of the Shapley Value analysis only if the size of the segment is big enough.

SV's Implications

The Shapley Value analysis provides a preferred ordering of the product line flavors, which has further implications for business strategy that aren't evident from a TURF analysis. Recall that the product flavor with the highest Shapley Value provides the greatest gain over all possible product combinations. This implies that it's preferable to have that flavor always in the available choice set when consumers are purchasing. The strategy implications of this statement are that the marketer should spend more effort to ensure that the product flavor with the highest Shapley Value has the minimum possible out-of-stock probability. This implies management of inventory in such a way as to give preference to the highest Shapley Value flavor. Following this strategy ensures that no matter which combination of product flavors and competitive flavors are on the shelf at any time, our flavors have the highest Shapley Values, and therefore are the best situation.

These strategy implications also apply directly to our Constrained Shapley Value. The only difference is that the evaluation takes place over all possible product combinations within all likely product line sizes. With large numbers of potential product flavors, this modification is especially important.

Exhibit 4

Bootstrap evaluation of SV

Order comparison	Mean difference	t-value
D>A	0.554	7.22
A>B	0.530	8.23
B>C	0.317	4.19
C>E	0.184	2.81

Finally, Shapley Value will identify flavors that have unique appeal to specific market segments only when those segments are large enough to matter. When TURF will usually provide a "best" solution that has some strange low appeal flavor, the Shapley Value approach only will value such a flavor if the overall appeal is large enough to matter. ■

Additional Reading

Cohen E. (1993), "TURF Analysis," *Quirk's Marketing Research Review*, (June/July), 10-3.

Conklin, M., and Stan Lipovetsky (1998), "Choosing Product Line Variants: A Game Theory Approach," *Computing Science and Statistics*, S. Weisberg, ed. Fairfax Station, VA: Interface Foundation of North America Inc.

Miaoulis G., Free V., and H. Parsons (1990), "TURF: A New Planning Approach for Product Line Extensions," *Marketing Research*, (March), 28-39.

Myerson, R.B. (1991), *Game Theory: Analysis of Conflict*. University Press.

Owen, G. (1982), *Game Theory*. Academic Press.

Roth, A.E. (1977), "The Shapley Value as a Von Neumann-Morganstern Utility," *Econometrica*, 45 (3), 657-64.

Roth, A.E. (ed.) (1988), *The Shapley Value—Essays in Honor of Lloyd S. Shapley*. Cambridge University Press.

Author Information

Michael Conklin is senior vice president of analytic services for Minneapolis-based Custom Research Inc. A graduate of the Wharton School, University of Pennsylvania, he worked at General Electric Co., Stanley Tools Inc., and Black & Decker in marketing research. He specializes in Bayesian methods and the analysis of categorical and ordinal data.

Stan Lipovetsky, a research manager for Custom Research, has an MS in theoretical physics and a PhD in mathematical economics from Moscow University. He worked at Tel Aviv University's Faculty of Management. His primary areas of research include multivariate statistics, multicriteria decision making, and marketing research.