The Shapley Value in Marketing Research: 15 Years and Counting

Michael Conklin & Stan Lipovetsky

We review the application of the Shapley Value to marketing research over the past 15 years. We attempt to provide a comprehensive understanding of how it can give insight to customers. We outline assumptions underlying the interpretations so that attendees will be better equipped to answer objections to the application of the Shapley Value as an insight tool.

Imagine it is 1998. My colleague Stan Lipovetsky, is working on a TURF analysis (Total Unduplicated Reach and Frequency) for product line optimization. Stan, being new to marketing research, asked the obvious question – "what are we trying to do with the TURF analysis?"

TURF[1] is a technique that was first used in the media business to understand which magazines to place an advertisement in. The goal was to find a set of magazines that would maximize the number of people who would see your ad (unduplicated reach) as well as the maximizing the frequency of exposure among those who were reached. This was adapted for marketing research for use in product line optimization. Here, the idea was to find a set of products to offer in the marketplace such that you would maximize the number of people who would buy at least one of those products. The general procedure at the time was to ask consumers to give a purchase interest scale response for each potential flavor in a product line. Then the TURF algorithm is run to find the pair of flavors that maximizes reach (the number of people who will definitely buy at least one product of the two), the triplet that maximizes reach, the quad that maximizes reach and so on. TURF itself is an *np*-hard problem. To be sure you have found the set of n products that maximizes reach you must calculate the reach for all possible sets of n.

Stan looked at the calculations we were doing for the TURF analysis and said "This reminds me of something I know from game theory, the Shapley Value." "So, what is the Shapley Value?" I asked. And so began a 15 year odyssey into the realm of game theory and a single tool that has turned out to be very useful in a variety of situations.

The Shapley Value

Shapley first described the Shapley Value in his seminal paper in 1953.[2] The Shapley Value applies to cooperative games, where players can band together to form coalitions, and each coalition creates a value by playing the game. The Shapley Value, allocates that total value of

---

[1] Wikipedia
[2] (Shapley, 1953)

the game to each player. By evaluating over all possible coalitions that a player can join in, a value for each specific player can be derived.

Formally the Shapley Value for player *i* is defined as:

$$\phi_i = \sum_{S\text{ - all subsets}} \gamma_{n(s)} \left[\upsilon(S) - \upsilon(S - \{i\}) \right]$$

Where: $$\gamma_{n(s)} = \frac{(s-1)!(n-s)!}{n!}$$

So, summing across all possible subsets of players *S*, the value of player *i* is the value of the game for a subset containing player *i* minus the value of that same subset of players without player *i*. In other words, it is the marginal value of adding the player to any possible set of other players. The summation is weighted by a factor that reflects the number of subsets of a particular size (*s*) that are possible given the total number of players (*n*).

When we apply the concept to the TURF game we have a situation where we create all possible sets of products, and calculate the "value" of each set by determining its "reach", or the percent of consumers in the study who would buy at least one item in the set. By applying the Shapley Value calculation to this data we can allocate the overall reach of all of the items to the individual items. This gives us a relative "value" of each individual product. The values of these products add up to the total value of the game, or the reach, of all of the products.

The fact that we can apply this calculation to the TURF game doesn't necessarily mean that it is useful. And, it certainly appears that the Shapley Value is an *np*-hard problem as well. We need to calculate the overall reach or value of every possible subset of products to even calculate the Shapley Value for each product.

Fortunately, the TURF game corresponds to what is known in game theory as a simple game. A simple game has a number of properties. In a simple game, the value of a game is either a 1 or a 0. All players in a coalition or team that produce a 1 value have a Shapley Value of *1/r* where *r* is the number of players in the team that can produce a win. In the TURF context, a consumer is reached by a subset of products. Those products all get a Shapley Value of *1/r* where *r* is the number of products that are in that subset. All other products get a Shapley Value of 0.

Another property of simple games is that they can be combined. In our TURF data, we treat each consumer as being a simple game. To combine the simple games represented by the consumers in our study, we calculate the Shapley Value for each product for each consumer and then average across consumers.

We solve the problem of how to calculate the Shapley Value for TURF problems by considering the TURF game as a simple game. But we still are not sure what this "value" represents. For this we need to look at the problem from a marketing perspective.

A Simple Model of Consumer Behavior

Consider this simple model of consumer behavior:

1. A consumer plans to buy in the category and enters the store
2. She reviews the products available and identifies a small subset (relevant set) that have the possibility of meeting her needs
3. She randomly chooses a product from that subset.

Now clearly most of us are not explicitly using some random number generator in our heads to choose which product to buy when we visit the store. Instead we evaluate the products available and choose the one that maximizes our personal utility, that is, we choose the product we prefer…at that moment. The product that will maximize our utility depends upon several factors. One factor is the benefits that the particular product delivers. A second factor is the benefits delivered by other competing products that are available in the store. Benefits delivered are evaluated in the context of needs. If one has no need for a benefit then its utility is non-existent. If one has a great need for a particular benefit then a product delivering that benefit will have a high utility and a good chance of being the utility maximizing choice.

When we observe consumer purchases, for example by looking at data from a purchase panel, one can see that the specific products available, and their benefits, stay relatively constant, but none the less, consumers seem to buy different products on different trips to the store. This would seem to indicate that the driver of choice is the degree to which a person's needs change from trip to trip. Hypothetically, we can map an individual's needs to specific products that maximize utility when that need is present. This means that if we can observe the different products that a person purchases over some time period, then we can infer that those purchases are a result of the distribution of need states that exist for that consumer.

If the distribution of need states for a specific consumer were such that the probability of choosing each product in the relevant set was equal then the purchase shares of each product would be the equivalent of the Shapley Value of each product. Therefore, we can think of the Shapley Value calculation as a simple choice model, where the probability of choosing a particular product is 0 for all products *not* in the relevant set and *1/r* for all *r* products in the relevant set.

An alternative to the Shapley Value calculation would be to estimate the specific probabilities of choosing each product using a multinomial logit discrete choice model. If, we can estimate the probabilities of purchase for each product for each consumer, then this should be a superior estimate of purchase shares since the probabilities estimated in this manner would not be arbitrarily equal for relevant products and would not be uniformly zero for non-relevant products. But, is it feasible, in the context of a consumer interview, to obtain enough choice data to accurately estimate those probabilities of purchase, especially if the product space is large? In addition, it is not possible in the course of a 20 minute interview to ask consumers to realistically make choices across multiple need states.

Application of the Shapley Value to Consumer Behavior

If we weight the consumers in our study by the relative frequency of category purchase and units per purchase occasion then the Shapley Value becomes directly, a measure of share of units purchased. This moves the Shapley Value from being an interesting research technique to being a very useful business management tool.

Anecdotally, we understand that category managers at retailers obtain a ranked sales report for their category and consider the items that make up the bottom 20% of volume to be candidates for delisting or being replaced in the store. Since the Shapley Value provides an estimate of the sales rate for each product (in any combination), we can create a more viable recommendation for a product line. Instead of choosing products that maximize "reach", we can use a dual rule of maximizing reach subject to the restriction that no products in the line fall into the bottom 20% of volume overall.

To effectively do this analysis, one needs to collect data a little differently from TURF. In a typical TURF study one asks respondents to give some purchase interest measure to each of the prospective products that would go in the product line. A consumer is counted as "reached" if she provides a top-box response to the purchase interest question. The problem with this approach is two-fold. First, the questioning procedure is very tedious, especially as the number of products in your product line increases. For that very reason, competitive brands are not typically included. But, competitive brands are critical. Those are the products you want to replace on the retailer's shelves. The Shapley Value analysis can show you which of the competitor's products your proposed line should displace, but it can only do so if you have included the competitive products in your study.

Our suggestion is to ask respondents which products, from the category, they have purchased in some limited time period. (The time period should be dependent on the general category frequency of purchase). This data can be used to calculate Shapley Values and optimize a product line if all we are considering are existing products in the marketplace.

When considering new product concepts the problem is how to reliably determine if a new product would become part of a consumer's relevant set. This is especially problematic since consumers are well known to overstate their interest in new product concepts. A method we have found effective is to ask the typical purchase intent question for the new product and supplement it by asking consumers to rank order the new concept amongst the other products they currently buy (i.e. the ones selected in the previous task). We count a new product as entering an individual consumer's relevant set if, and only if, they rated it top box in purchase intent *and* they ranked it ahead of all currently bought products. In our experience, this procedure appears to produce reasonable estimates from the Shapley Value. (Since there is no actual sales data in these cases a true validation has not been possible).

Going beyond TURF – Other applications of the Shapley Value

Recall that the Shapley Value is a way of allocating the total value of a game to the participants in a fair manner. There are plenty of situations where we only know the total value of something but we want to understand how that value can be allocated to the components that create that value. One clear example is linear regression analysis. Here we want to understand the value that each predictor has in producing the overall value of the model. The overall value of the model is usually measured by the R^2 value. Frequently we wish to allocate that overall R^2 value to the predictors to determine their relative importance.

In 2000, my colleague Stan was working with one method of evaluating the importance of predictors, the net effects. Net effects are a decomposition of the R^2 defined as:

$$\beta' R \beta = NE$$

Where the betas are vectors of standardized regression coefficients and R is the correlation matrix of the predictor variables. The *NE* vector, when summed equals the R^2 of the model. This particular decomposition of R^2 is problematic when there is a high degree of multicollinearity amongst the predictors. In those cases there can often be a sign reversal in the beta coefficients which can cause the net effect for that predictor to be negative. This makes the interpretation of the net effects as an allocation of the total predictive power of the model illogical.

My experience with the Shapley Value caused me to wonder if the Shapley Value might be a solution to this problem. The Shapley Value is an allocation of a total value. The individual Shapley Values will therefore sum to that total value, and they will all be positive. We can easily (although less easily than the line optimization case) calculate the incremental value of each predictor across all combinations of predictors.

In the Shapley Value equation we substitute for the value term the R^2 of each model:

$$\phi_i = \sum_{S\text{ - all subsets}} \gamma_{n(s)} \left[R_S^2 - R_{S-\{i\}}^2 \right]$$

This is no longer a simple game in the parlance of game theory so it becomes an *np-hard* problem again. But, for sets of predictors that are smaller than 30 it is a reasonable calculation on modern computers.

I was convinced that this was an excellent idea. As is often the case with excellent ideas, it turned out that there were many others doing research in other fields who had also come up with essentially the same idea.[345] Many other related techniques also appear in the literature.

We did, however, take the approach one step further. Going back to the net effects decomposition discussed earlier we realized that both of these techniques, net effects and Shapley Value were trying to do the same thing, allocate the overall model R^2 to the individual predictors. So, if we assume that the Shapley Values are approximations of the Net Effects then we can "reverse" the decomposition and calculate new beta coefficients so that they are as consistent as possible with the Shapley Values.

$$\beta' R\beta \cong SV$$

This requires a non-linear solver but we can estimate a new set of beta coefficients that result in Net Effects that are very close to the Shapley Values. These new coefficients can then be used in a predictive model.

Gromping and Landau have criticized this approach[6]. We show in a rejoinder[7]that in conditions of high multicollinearity, the model with the adjusted beta coefficients as described above does a better job of predicting new data than the standard OLS model. We do recommend only utilizing the adjusted coefficients in those extreme conditions.

Of course, there are other decompositions of R^2 in the literature besides the Net Effects decomposition. One decomposition, which was first described by Gibson[8], and later rediscovered by Johnson[9] decomposes the R^2 as follows.

$$R^2 = \beta' R^{1/2} R^{1/2} \beta$$

This produces two identical vectors of weights ω that when squared, sum to the R^2 of the model. These can be interpreted as importance weights and are very close approximations to the Shapley Values. The advantage of using this approximation of the Shapley Values for importance is that this particular decomposition is not an *np-hard* problem like the Shapley Value calculation and therefore is much easier to compute with large numbers of predictors.

Moving on from Linear Regression – Other Allocation problems

3 (Kruskal, 1987)
4 (Budescu, 1993)
5 (Lindeman, Merenda, & Gold, 1980)
6 (Gromping & Landau, 2009)
7 (Lipovetsky & Conklin, 2010)
8 (Gibson, 1962)
9 (Johnson, 1966)

One of the nice things about the Shapley Value is that the "value function" is abstract. You can define value in any way that you want, turn the Shapley Value crank and output an allocation of that value to the component parts.

Consider the customer satisfaction problem. The Kano theory of customer satisfaction[10] suggests that different product benefits have different types of relationships to overall satisfaction.

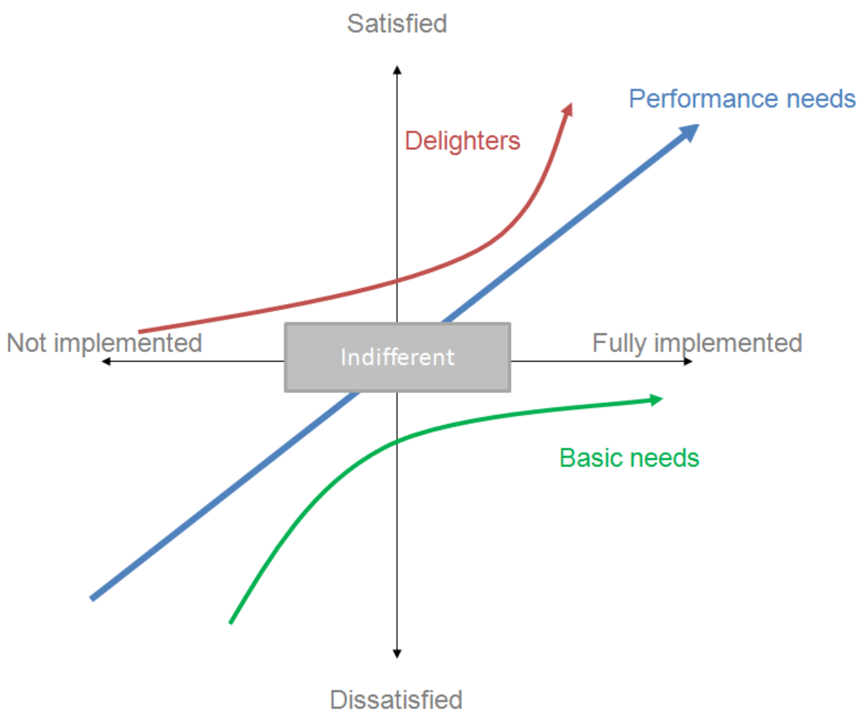

Graphic by David Brown - Wikipedia

Identifying attributes that are "basic needs" or "must-be" attributes is critical in customer satisfaction research. These are the items that cause overall dissatisfaction if, and only if, you fail to deliver. The interesting thing about these attributes is that they are non-compensatory, that is, if you fail to deliver on any one of these attributes you will have overall dissatisfaction, no matter how well you perform on other attributes.

Standard linear regression driver model approaches clearly don't work here. There are two issues, first a linear regression model is inherently compensatory, and second, the vast majority of the data is located in the upper right quadrant of the graph above.

As a result, we construct a model like this:

First – let $\bar{A}$, $\bar{B}$, $\bar{C}$ ... $\bar{K}$ represent customers dissatisfied with A,B,C….K respectively

also let $\bar{Y}$ represent customers dissatisfied overall.

[10] (Kano, Seraku, Takahashi, & Tsuji, 1984)

We want to find a set of items such that

$$\{\overline{A|B|C}\} => \bar{Y}$$

in other words, dissatisfaction with A or B or C implies dissatisfaction overall

One way of evaluating this is by calculating the reach into $\bar{Y}$. In other words, the percent of dissatisfied people $\bar{Y}$ that are dissatisfied with any item in the set. But, this cannot be the end of the calculation because we need to subtract from this the percent of people who are satisfied overall Y but are dissatisfied with one of the items in the set. In other words we need to subtract the false positive rate. This statistic is known as Youden's J[11] and we can use it to evaluate any dissatisfaction model of the form noted above.

In our case, we treat Youden's J statistic as the "value" of the set of items. We can search for the set of items that maximizes Youden's J and then use the Shapley Value calculation to allocate that value to the individual items[12]. This provides a priority for improvement.

Summary

Since we started using the Shapley Value in marketing research problems a decade and one half ago we have found it to be a very useful technique whenever we need to allocate a total value to component parts.

In the case of line optimization it immediately generalizes to a reasonable model of consumer behavior making it an extremely useful business management tool. Other applications have also proved to be quite useful. Business management, after all, seems to be primarily about prioritization and the Shapley Value procedure provides a convenient way to prioritize the components of many business decisions when direct measures of value of those components are not available.

Budescu, D. (1993). Dominance Analysis: a new approach tot he problem of relative importance in multiple regression. *Psychological Bulletin*, 114:542-551.

Conklin, M., Powaga, K., & Lipovetsky, S. (2004). Customer Satisfaction Analysis: identification of key drivers. *European Journal of Operational Research*, 154: 819-827.

Gibson, W. A. (1962). On the least-squares orthogonalization of an oblique transformation. *Psychometrika*, 11:32-34.

Gromping, U., & Landau, S. (2009). Do not adjust coefficients in Shapley value regression. *Applied Stochastic Models in Business and Industry*.

[11] (Youden, 1950)
[12] (Conklin, Powaga, & Lipovetsky, 2004)

Johnson, R. M. (1966). The Minimal Transformation to Orthonormality. *Psychometrika*, 61-66.

Kano, N., Seraku, N., Takahashi, F., & Tsuji, S. (1984). Attractive quality and must-be quality. *Journal of the Japanese Society for Quality Control*, 39-48.

Kruskal, W. (1987). Relative Importance by Averaging over Orderings. *The American Statistician*, 41:6-10.

Lindeman, R. H., Merenda, P. F., & Gold, R. Z. (1980). *Introduction to Bivariate and Multivariate Analysis.* Glenview, Il: Scott, Foresman.

Lipovetsky, S., & Conklin, M. (2010). Reply to the paper "Do not adjust coefficients in Shapley value regression”. *Applied Stochastic Models in Business and Industry*, 26: 203-204.

Shapley, L. S. (1953). A Value for n-Person Games. In e. A. H.W. Kuhn & A.W. Tucker, *Contributions to the Theory of Games, Vol II.* (pp. 307-17). Princeton, NJ: Princeton University Press.

Youden, W. (1950). Index for rating diagnostic tests. *Cancer*, 3: 32-35.