Appeared in the *Proceedings of the 1997 Sawtooth Software Conference*.

Obtaining Product-Market Maps from Preference Data

Terry Elrod
University of Alberta

Abstract

This paper introduces the reader to product-market maps and sho ws how they can be used to explain and predict brand preference and, ultimately, brand choice. It then considers why it might be desirable to estimate product-market maps from consumer preferences for existing brands. A model for accomplishing this is described and its utility explored using the data analyzed by Wittenschlaeger and Fiedler, whose paper also appears in this volume. A product-market map is fit to pairwise preferences for existing brands obtained from users of air traffic management systems. An additional analysis of brand perceptions assists in interpretation and verification of the map.

An Introduction to Product-Market Maps

A product-market map uses a picture to characterize both products (i.e. brands) and market (i.e. customers) in terms of the benefits that drive consumer brand preference and choice. Product-market maps are best explained using a simplified example.

Suppose that choice of toothpaste is driven by how the brands are perceived in terms of two fundamental benefits: health and social. And suppose that there are only three brands perceived by consumers as shown in Figure 1. Because this picture portrays only products, it is a product map. You may think of this picture as simply a plot of the three brands in terms of their average ratings on these two benefits using a seven-point ratings scale. In this hypothetical example, Crest enjoys a strong perception in terms of the health benefit, but is weak on the social benefit. Ultra-Brite and Close-Up are more similar to each other than either is to Crest—both are relatively strong in terms of the social benefit but weaker in terms of the health benefit. The position of a brand in a product map signifies how much of each benefit it delivers as perceived by consumers.

Figure 1. A Hypothetical Product Map

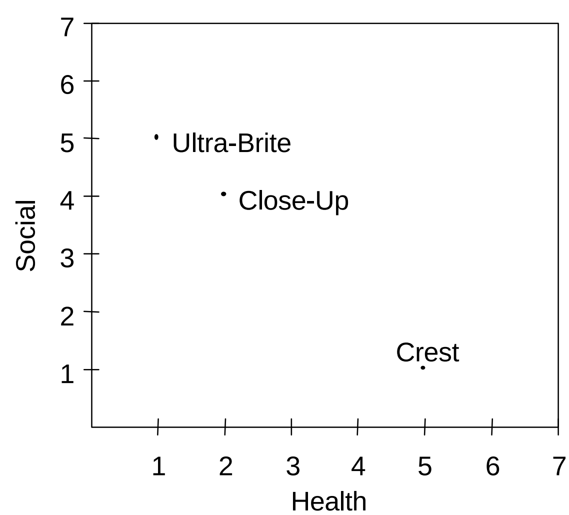

A product-market map adds consumers to a product map. Two such segments are added to the product map of Figure 1 to yield the hypothetical product-market map of Figure 2. The location of the two segments in the map reflects how each segment uses each benefit to determine its preferences for the brands.

Determining Brand Preferences from a Product-Market Map

There are two common models for relating product-market maps to brand preferences: the ideal-point model and the vector model. The ideal-point model assumes that the location of a consumer segment in the map represents the consumer's "ideal brand" in terms of the benefits underlying the product category. Brand preference is inversely related to

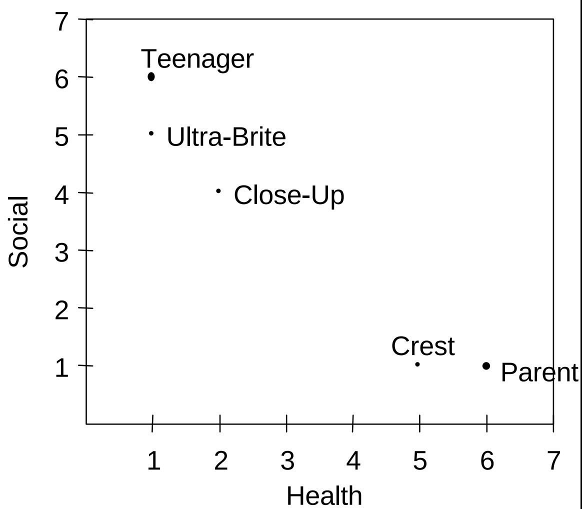

Figure 2. A Hypothetical Product-Market Map

each brand's distance from the consumer's ideal point. The distances underlying brand preference for the teenager segment of our hypothetical example are shown in Figure 3.

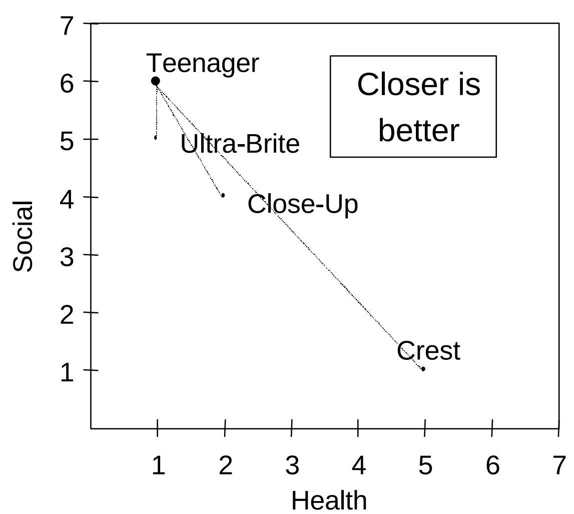

Figure 3. Ideal-Point Model Illustrated for the Teenager Segment

The vector model assumes that more of a product benefit is always better, although consumer segments still differ in terms of how much importance they attach to each benefit. The vector model for the teenager segment is illustrated in Figure 4. Because the vector model underlies the product-market map described later in this paper, we will examine its properties more fully.

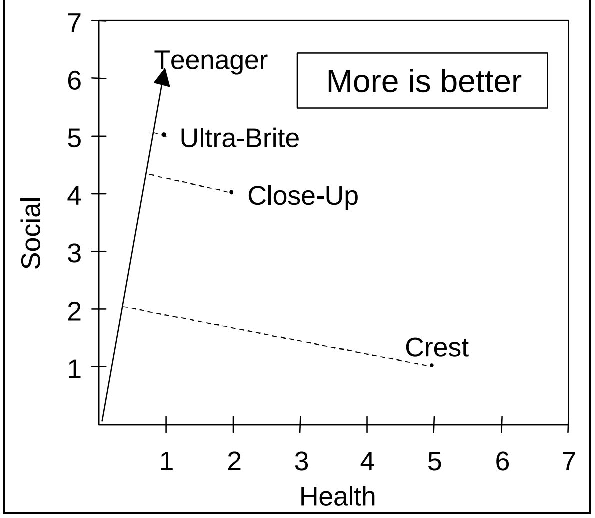

Figure 4. Vector Model Illustrated for the Teenager Segment

In a vector model, each consumer segment can be represented by an arrow originating at the origin of the map and ending at the location of the segment in the map. The arrow emphasizes the relative importance of the two benefits to the segment. In a vector model, brands that lie farthest in the direction indicated are most preferred, while distance of a brand *from* the arrow is irrelevant. The brand preferences are proportional to the projection (at a right angle) of the brands onto the arrow. These projections for the teenager segment are also indicated in Figure 4. They indicate that Ultra-Brite is most preferred, but Close-Up is a close second. Crest, with its poor perceived performance in terms of the social benefit, is a distant last preference for this segment.

Using the Product-Market Map

Figure 5 portrays the hypothetical product-market map for both segments, as well as the implied preferences for all brands and segments using the vector model. This product-market map, simple as it is, can be used to illustrate much about the importance of mar-

Figure 5. The Brand Preferences for All Brands and Segments Implied by the Vector Model

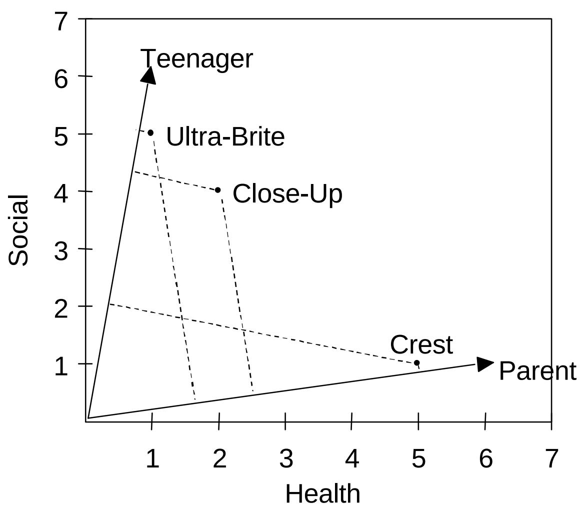

ket segmentation, product differentiation, and the intimate connection between them.

Crest is much preferred by the parent segment, which means that Crest enjoys a near-monopolistic position vis-à-vis this group. This is due to Crest being the only brand that is tailored to the greater importance the parent segment places upon the health benefit. Its poor perceived social benefit is not important to this segment. However it is also apparent that Crest has no prospect of attracting significant sales from the teenager segment. Its marketing should therefore be directed to the parent segment. It may be priced at a premium and still be preferred by this segment.

While Ultra-Brite is the preferred brand for the teenager segment, Close-Up is a close second. Aggressive pricing and advertising by Close-Up may suffice to attract significant sales from the teenager segment, and this prospect prevents Ultra-Brite from enjoying large profit margins. Neither brand can hope to attract appreciable sales from the parent segment.

Table 1. Tabular Representation of the Toothpaste Product-Market Map

	Product benefits		Brand preferences	
	Health	Social	Parent	Teenager
Crest	5	1	31	11
Close-Up	2	4	16	26
Ultra-Brite	1	5	11	31
Parent	6	1	↗	
Teenager	1	6		↗

A Tabular Representation of the Toothpaste Example

All of the information in the product-market map of Figure 5 can also be shown in tabular form as in Table 1. The first two columns of numbers in the table show the locations of the brands and segments in the map. The last two columns show the brand preferences for both segments as implied by this vector map. The preference value of the parent segment for Crest is obtained as the sum of the benefits of Crest weighted by the importances the parent segment attaches to these benefits: i.e. $5 \times 6 + 1 \times 1 = 31$. This is a straightforward application of the multi-attribute utility model familiar to marketers.

Estimating Product-Market Maps from Preference Data

So far we have considered product-market maps and how they may be related to brand preferences without considering how such maps may be obtained. Obtaining meaningful product-market maps is a nontrivial exercise because they represent brand *perceptions* in terms of the product *benefits* that underlie brand preference and choice. Neither consumer perceptions nor the product benefits underlying preference are directly observable.

There are three primary methodologies for obtaining product-market maps. The oldest is simply to have consumers provide a product-market map directly by rating the brands and themselves in terms of benefits specified by the researcher in advance. Because fundamental benefits are intangible and rating scales somewhat artificial, this approach tends not to yield product-market maps that predict choices well.

Two other methods exploit the redundancy among the three types of data shown in Table 1: brand perceptions, consumer importances for benefits, and brand preferences. Knowing any two of these types of data allows calculation of the third by application of the multi-attribute utility model. For example, suppose that you have obtained the product part of the product-market map and that you have also collected information from consumers about their preferences for the brands. What is known and is not known to you under this scenario is shown in Table 2.

Table 2. Unknown Consumer Importances for Toothpaste Benefits

	Product benefits		Brand preferences	
	Health	Social	Parent	Teenager
Crest	5	1	31	11
Close-Up	2	4	16	26
Ultra-Brite	1	5	11	31
Parent	?	?	➚	
Teenager	?	?		➚

Because of the redundancy of information in Table 2 it is possible to estimate the importance each segment attaches to each benefit. This may be estimated separately for each segment using regression analysis. The dependent variable would be, for example, parent preferences for the three brands, and the two independent variables would be the values each of the three brands have on the two benefits.

With only three brands there are only three observations, so the regression estimates will not be very stable. I will discuss a method for stabilizing these estimates later, but clearly having more than three brands in a market would help to obtain more reliable estimates of the consumer locations in product-market map.

The article by Wittenschlaeger and Fiedler that also appears in this volume provides an excellent demonstration of how a product map might be obtained as the first step towards developing a product-market map. Their product map could serve to determine both the benefits that underlie brand preference as well as the locations of the brands in terms of these benefits. The regression analysis I just described could then be performed as a second procedure. Obtaining a product-market map by analyzing brand preferences using a given product map is known as an "external" analysis of preferences. The word external refers to the fact that the product map was obtained in advance using other information.

In this paper I illustrate what is known as an "internal" analysis of preference data. The goal of the analysis is to obtain both product and consumer segment locations in a product-market map simultaneously, and only by using brand preference information. This is a more ambitious task. A glance at Table 3 shows how much we seek to estimate from the brand preference data alone.

Table 3. Obtaining Both Brand and Consumer Locations from Brand Preferences

	Product benefits		Brand preferences	
	Health	Social	Parent	Teenager
Crest	?	?	31	11
Close-Up	?	?	16	26
Ultra-Brite	?	?	11	31
Parent	?	?	➚	
Teenager	?	?		➚

A Model for Obtaining a Product-Market Map from APM Preferences

This paper will employ the vector model to estimate a product-market map from brand preferences. Ideal-point product-market maps are often very difficult to estimate from brand preferences. Ideal-point models are more general than vector models. However if a vector model accounts adequately for brand preferences, as is often the case, then the data contain little information to allow estimation of the additional generality of the ideal-point model.

Flexibility of the Vector Model

Note that the vector model assumes more is better when it comes to *benefits* revealed by the model, but not necessarily for the brand *attributes* that characterize the brands. This distinction is important and can be illustrated using our toothpaste example.

I recall an issue of *Consumers Report* some years ago that contained a review article for toothpastes which stated that the primary determinant of the tooth-whitening ability of a toothpaste is its abrasive content. Abrasives help to remove stains from teeth (as well as plaque), but too much abrasive content can accelerate the wearing away of tooth enamel.

If consumers believe that a toothpaste can have too much tooth-whitening power, then this attribute would be nonlinearly related to the health and social benefits. Figure 6 illustrates this by showing four points of a 7-point scale that relates a brand's perceived tooth-whitening ability to its location in the map. A toothpaste with a rating of 1 has virtually no tooth-whitening ability. I postulate that such a brand would have a low rating on the social benefit of toothpaste (assuming that some tooth-whitening ability is deemed essential to this benefit), while I place the brand (arbitrarily) at slightly above the midpoint of the scale for the health benefit.

Figure 6. Hypothesized Relationship Between Tooth-Whitening Ability and the Benefits of Toothpastes

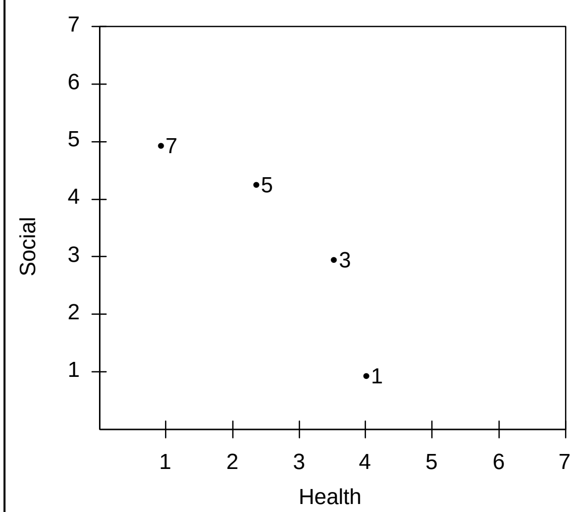

Increasing this brand's rating to 3 on tooth-whitening ability helps its performance on the social benefit with little harm to its perceive health benefit. Increasing its perceived tooth-whitening rating further, however, yields diminishing improvement on the social benefit while harming the brand's perceived health benefit.

Consumer segments which attach approximately equal importance to the health and social benefits of toothpastes (and who would be represented in Figure 6 by a vector pointing to the upper-right corner) would prefer toothpastes with tooth-whitening ratings in the range of 3 to 5 on the 7-point scale. This prediction is consistent with an ideal-point model relating brand preference to brand attributes even though the model used to estimate the product-market map from preferences is vector-based. Extreme segments, such as the parent and teenager segments of the example, might still prefer extreme ends of the tooth-whitening scale.

This discussion of tooth-whitening ability and how it may be nonlinearly related to the benefits underlying product-market maps serves to indicate the flexibility possessed by such maps based on the vector model when the maps are estimated from brand preferences. Given the flexibility of such models, it is not surprising that the additional generality of ideal-point models of product-market maps often makes them hard to estimate reliably from preference data.

Particulars of the Model Employed

The method developed for this paper to derive a product-market map from preference data is a variant of factor analysis. Factor analysis is familiar to many marketers. The output of a factor analysis invariably shows how the variables included in the analysis are related to the two or so factors estimated by the analysis. When analyzing a data matrix of brand preferences, with as many rows as consumers and one column for each brand, the "variables" are the brands and the "factors" are the benefits. Hence the result of a factor analysis applied to such data is a product map. However factor analysis also estimates "factor scores," one for each respondent. Factor scores characterize the respondents in terms of the same factors. They are the coefficients for the respondents that, together with the product map, best reproduce the pattern in the preference data. Thus the factor scores are also the importance weights for the individual respondents.

For the purpose of analyzing consumer preference or choice data to obtain a product-market map, I have adapted factor analysis in four respects to better accommodate this particular marketing application. The remainder of this section provides a brief description of these differences.

(1) A factor analysis usually begins by standardizing the data. The raw data are rescaled so that the mean for each column is zero and the standard deviation is one. It makes sense to do this when some of the variables included in the factor analysis differ from other variables in their units of measurement. However in this setting every variable is a measure of preference and the different columns simply refer to different brands. Differences in average preference across brands is vital information that is retained and accounted for by the analysis.

(2) Factor analysis assumes that the different variables may be measured with different amounts of error. This also makes sense when the variables are measured on different scales or ask fundamentally different questions. However here each variable is an expression of brand preference on the same scale, and the only difference between questions is the brand being rated. Therefore I have not allowed the error variances to differ arbitrarily from one brand to the next.

(3) Factor analysis is often estimated by maximum likelihood assuming that the factor scores have a multivariate normal distribution across respondents. Using a distribution such as the multivariate normal to characterize consumer heterogeneity is a good idea for two reasons. First, the respondents invariably represent a sample from the population of consumers. When a *sample* has been taken in order to learn about the *population* from which the sample has been taken, then it is appropriate that the analysis explicitly recognize this fact. This principle is widely overlooked in marketing research. Estimating factor scores at the individual-level for each respondent (or conjoint part worths, for that matter) is simply incorrect. The analogous error in an analysis of variance would be to model random effects as if they were fixed effects.

A second reason to use a statistical distribution such as the multivariate normal to characterize customer heterogeneity is that it ameliorates the problem of trying to estimate too many coefficients from too few data. To illustrate the economy that results, assume for the moment that we have a two-dimensional map and 300 respondents. Estimating importance vectors separately for each respondent requires the estimation of 600 parameters. Estimating the mean and variance-covariance matrix for a bivariate normal distribution, on the other hand, requires estimation of only 5 parameters.

Nevertheless, the assumption of multivariate normality for consumer importances is a strong assumption. I have relaxed this assumption in two ways. First, the multivariate normal distribution has been replaced by the more robust and general multivariate t distribution. The t distribution has "longer tails" than the normal, so it is more robust to outlying respondents. The degrees of freedom is estimated along with the other unknowns of the product-market map. Second, rather than assume that the distribution applies to all respondents, I have only assumed that it applies *within* each segment. Thus consumer heterogeneity within segments is explicitly accounted for. Because of the high degree of indeterminacy in estimating product-market maps from brand preferences, it is possible without loss of generality to scale the maps so that consumers within each segment have independent standard t-distributions about the segment mean. This is convenient because it allows us to represent each segment by its mean alone, without also having to portray consumer heterogeneity within each of the segments.

(4) A final extension to factor analysis is particular to the type of preference data obtained by Sawtooth Software's APM. APM does not provide data on brand preferences as I have described them. Rather, APM obtains *pairwise* preferences using a 100-point "probability of purchase" scale. Thus we do not observe brand preferences directly, but a measure of the difference in preference for pairs of brands. Obtaining product-market maps from pairwise preferences involves additional programming but is not conceptually difficult.

A natural method for analyzing probabilities is to transform them into logits. However stated probabilities can include the endpoints of the scale, and a literal logit transformation of these values is not possible. I have implemented a capability of estimating the best increment to add to the endpoints of the 0-100 interval before applying the logit transformation.

Description of the ATM Data

Thomas A. Wittenschlaeger and John A. Fiedler were kind enough to share with me the data used in their paper. The data pertain to suppliers of air traffic management (ATM) systems.

The study included 12 suppliers, but one company was unfamiliar to all but a few respondents. Because APM only asks for preference judgments involving brands familiar to each respondent, there was little preference information available for this brand—too little information to allow reliable estimation of its location in a product-market map based on preference data. This company was therefore not included in my analysis. Complete data for 11 companies and 14 attributes was available for 292 respondents, all of whom were included in the analyses described below.

Customer Region

The distribution of the 292 respondents among four regions of the world is shown in Table 4. I created the Other region by combining the small numbers of respondents from three regions: the rest of the Americas, Russia, CIS and Eastern Europe and Asia. Regions with few respondents were com-

Table 4. Region Segment Definitions

Label	Explanation
U	Canada, U.S.
W	Western Europe
M	Middle East and Africa
O	Other

bined so that every segment's position in the product-market map would be reliably estimated.

Company Familiarity

Every respondent also provided information about company familiarity using a 5-point scale. The 11 companies included in the analysis, and their average familiarity ratings, are shown in Table 5. (The excluded company had an average familiarity rating of only 1.34.)

Table 5. The Twelve Companies and Their Average Familiarity Ratings

Label	Average Rating
HG	3.61
RY	3.55
BO	3.44
LK	3.23
NR	2.95
SM	2.93
TM	2.89
NE	2.43
AL	2.37
CA	2.00
BD	1.99

Attribute Importance

We have complete ratings of attribute importance for all 14 attributes and all 292 respondents. The attribute definitions are provided in Table 6 along with the average importance rating for each attribute. These data will not be used to derive the product-market map, but will be referred to later when assessing the face validity of the map.

Company Perceptions

In addition, respondents provided partial information about their perceptions of the companies in terms of the attributes. Complete data for all companies and attributes would require $12 \times 14 = 168$ ratings from every respondent, which is too onerous a task and in any case ratings for unfamiliar brands or on unimportant attributes are likely to have little reliability. Therefore Sawtooth Software's APM collects perceptions only for those companies familiar to the respondent and only on those attributes of importance to him or her. Wittenschlaeger and Fiedler used these data to develop a product map. These data will be used here only to interpret the product-market map and assess its face validity.

Table 6. Definitions and Average Importance Ratings for 14 Attributes

Label	Explanation	Rating
ONT	Delivers on-time	4.69
ONB	Delivers on-budget	4.64
TRS	Is managed by a team I trust	4.60
LON	Provides long-term life cycle system support	4.53
GRW	Provides growth in functionality and capacity	4.53
ADV	Provides technically advanced solutions	4.46
ADP	Provides solutions that can adapt / accommodate to existing equipment	4.40
TUR	Offers turnkey solutions	4.13
COT	Maximizes use of commercial off-the-shelf products	4.12
MAN	Has installed many ATM systems	4.08
EXC	Offers products which exceed requirements	3.79
LWS	Offers the lowest price	3.64
LOC	Invests in local industry / economy	3.23
FIN	Is able to offer financing packages	3.01

ATM Preference Data

Finally, respondents indicate their relative preferences for each of several pairs of companies. The companies are simply identified by name in this task.

I fit the vector model to a logit transformation of the pairwise probabilities after adding an estimated increment of at least 5 to both ends of the 0-100 interval. (That is, the interval used to transform the probabilities was forced to be at least –5 to 105). This in effect forces the odds of choosing the less preferred alternative to be greater than or equal to 1:21.

The choice of 5 as a minimum size for the increment was somewhat arbitrary. An increment of 10 was found to be optimal when using the multivariate normal distribution to represent within-segment heterogeneity. However, replacing the multivariate normal distribution with the multivariate t improves model fit dramatically while requiring estimation of only one additional parameter. Estimating an unconstrained increment for the logit transformation and the degrees of freedom for the t distribution simultaneously for these data led to no increment for the logit but degrees of freedom for the t distribution so small that the estimation procedure became unstable. Constraining the increment to be at least 5 led to a larger estimate for degrees of freedom and stable estimates of the map. The resultant map is shown in Figure 7, where both axes are to the same scale.

Figure 7. A Two-Dimensional Map of Preferences

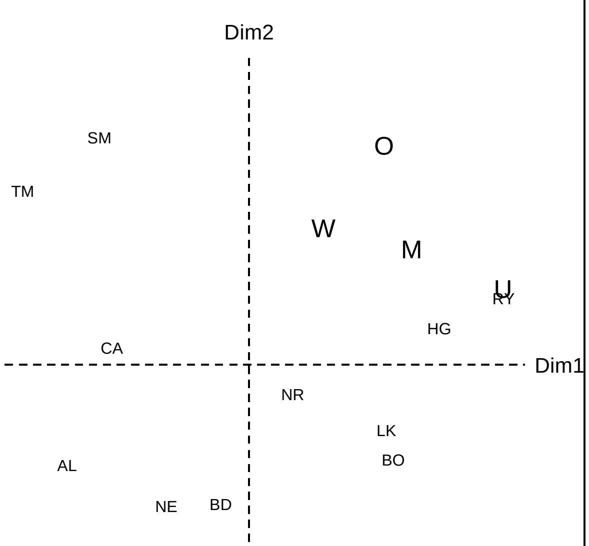

Interpreting the Map Using Brand Perceptions

While the map of Figure 7 shows both the companies and customer segments, it contains no information that allows us to interpret the map in terms of the 14 attributes included in the study (Table 6). Such an interpretation can be added to the map using a secondary analysis known as “property fitting.”

First, a table is created with as many rows as brands and as many columns as attributes. Each cell of this table shows the average rating received by the brand of that row on the attribute of that column. This newly created table is then related to the map one attribute at a time. Just how this is done is illustrated for the ADV attribute in Table 7. There we show the average ratings on ADV for the 11 brands together with the locations of these brands in the map. The relationship between ADV and the two other columns was then determined using regression analysis where the columns Dim1 and Dim2 are the independent variables in the regression.

Table 7. Relating ADV to the Map

	ADV	Dim1	Dim2
AL	3.28	-0.35	-0.18
BD	3.49	-0.05	-0.26
BO	3.69	0.29	-0.17
CA	3.43	-0.26	0.05
HG	3.81	0.38	0.09
LK	3.78	0.27	-0.11
NE	3.59	-0.16	-0.26
NR	3.66	0.09	-0.04
RY	3.91	0.50	0.15
SM	3.48	-0.29	0.46
TM	3.47	-0.43	0.36

The degree of success when performing this regression analysis separately for each of the 14 attributes in the study is shown in Table 8. Although each regression is based on only 11 observations, all but two of the regressions were statistically significant. The statistically insignificant regressions were for

FIN and LOC. Table 6 shows that these two attributes were rated by respondents as being least important of all, so these attributes should in fact be nearly irrelevant to brand preferences and unrelated to a map that explains these preferences.

Table 8. Relating Attributes to the Map

	R^2	P-value
ONT	0.82	0.00
ONB	0.62	0.02
TRS	0.95	0.00
LON	0.92	0.00
GRW	0.86	0.00
ADV	0.91	0.00
ADP	0.76	0.00
TUR	0.74	0.00
COT	0.60	0.02
MAN	0.58	0.03
EXC	0.88	0.00
LWS	0.65	0.01
LOC	0.39	0.14
FIN	0.15	0.53

The method of property fitting just described assumed that the attributes are linearly related to the benefits of the preference map. We have seen that attributes need not be linearly related to product benefits, but map interpretation is simplified when they are. I report here two analyses which indicate that the linear assumption is appropriate for relating attributes to benefits for these data.

The first check is to fit a quadratic regression of average brand perceptions to the benefit dimensions of the map. That is, three additional independent variables can be added to the regression shown in Table 7, corresponding to Dim1^2, Dim2^2 and Dim1*Dim2. These added terms failed to improve upon the vector model to a statistically significant extent for any of the 14 attributes. This is not surprising given that only 11 observations are available for each regression.

A second check makes use of additional information which is often collected by APM. Respondents provided perceptions of their ideal brand along with actual brands. A basis for deciding the appropriateness of a linear relationship between attributes and benefits is to compare the ratings given the ideal brand to the ratings given actual brands. If ratings for actual brands rarely straddle the ideal brand rating on an attribute, then this is a further indication that a linear relationship between the attribute and product benefits is appropriate.

I calculated for each respondent and attribute the ratings given the actual brands minus the rating given the ideal brand. Since the actual and ideal brands are rated on the same 5-point scale, the difference must always be an integer between 4 and –4. Over all respondents, brands and attributes only 5.7% of the differences were positive (4.3% by only one unit). This is not a lot of “straddling” of the ideal point.

Displaying the Relation of Attributes to Map Benefits

Portrayal of attributes in the map is the same as for segments and companies: the coefficients from the regression for each attribute provide its location in the map. Often attributes are shown as vectors radiating from the origin, but because the arrows can obscure other information, I simply plot the attributes as points. A display of the attributes as they relate to the dimensions of the map is shown in Figure 8. (The two attributes that coincide in Figure 8 are ONT and COT.)

The acronyms used for attributes are as shown in Table 6. Dim2 in the map distinguishes companies that “Have installed many ATM systems” (MAN) from “Offers the lowest price” (LWS). All other attributes are more closely associated with Dim1 in the positive direction. Attributes lying to the upper right seem to pertain to companies that best provide a customizable offering with substantial support. Examples are “Pro-

Figure 8. Relation of the Attributes to the Map

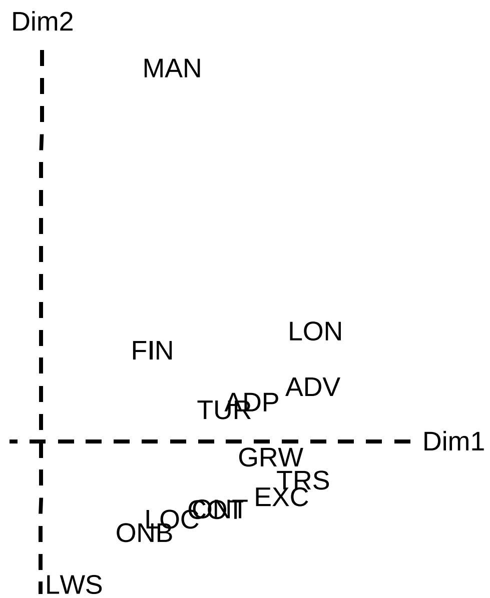

vides long-term life cycle support” (LON), “Provides technically advanced solutions” (ADV), “Provides solutions that can adapt / accommodate to existing equipment” (ADP), and “Offers turnkey solutions” (TUR). In contrast, attributes lying towards the lower right pertain to companies that offer a more standardized product, with the predictability that this allows: “Delivers on-budget” (ONB), “Delivers on-time” (ONT), and “Maximizes use of commercial off the shelf products” (COT). (Recall that LOC is not statistically significant.)

Putting it All Together: The Final Product-Market Map

Figure 7 and Figure 8 are combined into the single (busy!) product-market map shown as Figure 9. The distance of the attributes from the origin was reduced by the same fraction for all attributes and dimensions so that they would fit better into the display. The display is made somewhat easier to read by using the convention that the customer segments are denoted by single letters (cf. Table 4), the companies by two letters (cf. Table 5), and the attributes by three letters (cf. Table 6).

Concluding Comments

This paper has illustrated a method for estimating a product-market map from pairwise preferences for existing brands, such as is obtained using APM. Map interpretation was aided by regressing average brand perceptions onto the map.

The final map displays 4 customer segments, 14 attributes and 11 companies. In practice the map might be simplified for some purposes, perhaps by replacing the 14 attributes with descriptive labels of the map's dimensions that are based on these attributes.

Because the product-market map is based upon an analysis of customer preferences, it remains closely tied to these preferences in a quantitative sense and this property should be exploited. Simulators can be built using a spreadsheet software package to predict shares for all brands to help assess contemplated new or repositioned brands. An adequate discussion of the details on how to do this must await a separate paper.

Given the strategic value of the information provided by product-market maps, companies that make good use of this technology can expect to enjoy an important advantage over their competitors.

Figure 9. The Final Product-Market Map for the ATM Data