Conjoint Analysis: Marketing Engineering Technical Note[1]

Table of Contents

Introduction

Conjoint analysis is a technique for measuring, analyzing, and predicting customers' responses to new products and to new features of existing products. It enables companies to decompose customers' preferences for products and services (typically provided as descriptions or visual images) into "part-worth" (or utilities) associated with each level of each attribute of the product. They can then recombine the part-worths to predict customers' preferences for any possible combination of attribute levels, and the likely market share or revenue that a new product is likely to achieve when introduced into a market in which other competing products may already be available. They can also use conjoint analysis to determine the optimal product concept or to identify market segments that value a particular product concept highly.

---

[1] This technical note is a supplement to Chapter 6 of Principles of Marketing Engineering, by Gary L. Lilien, Arvind Rangaswamy, and Arnaud De Bruyn (2007). © (All rights reserved) Gary L. Lilien, Arvind Rangaswamy, and Arnaud De Bruyn. Not to be re-produced without permission.

There is a vast and growing literature on conjoint analysis. Here we provide an overview of the core analytical aspects of "full profile" conjoint analysis, and we briefly mention other approaches, such as Hierarchical Bayes conjoint analysis. The interested reader can explore the referenced articles for further details and enhancements.

Conjoint Analysis for Product Design

For commodity products with a single attribute (e.g., price), it is not that difficult to come up with a rank-order of the available products in terms of that single attribute. For example, everyone should prefer a lower-priced product, when all products are identical. However, how do we generate such preference orders if products have multiple attributes (conjoined attributes), as most products do? This is the original issue that motivated the development of conjoint measurement. Since then, conjoint measurement has gradually evolved into a comprehensive approach, called conjoint analysis, for measuring and understanding consumer preferences, and using the resulting measures for simulating market reactions to potential new products. There are three stages in a typical conjoint study: (1) Design of a data collecting instrument, (2) Collecting data from consumers, and (3) Analyzing the data and simulating market response. Here, we focus on the first and third stages, namely, design and simulations

Designing a conjoint study: Conjoint Analysis starts with the premise that a product category could be described as a set of attributes. For example, pizzas could be considered to have the following attributes: size, brand, type of crust, topping, amount of cheese, type of sauce, price, etc. Every pizza could then be described as a combination of levels of those product attributes; for example, large Papa John's thick crust pepperoni pizza with extra cheese and tomato sauce, priced at $12.95.

The objective of the design stage is to specify a set of product bundles for which we obtain customers' overall evaluations, in such a way that those evaluations could then be decomposed into the part-worth value that each customer attaches to each level of each attribute. To develop such a design is not a simple task. For example, if there are 6 attributes, each with four possible levels, then we could create 4^6 (= 4096) different products. It is not reasonable to ask each customer to evaluate all of those bundles. Instead, in this case, if a

customer rates as few as 25 product bundles, that is sufficient for estimating the part-worths for the attribute levels. One way to select the product bundles is to ensure that they satisfy an orthogonality constraint. This means that across the selected product bundles, each level of an attribute combines in roughly the same proportion with the levels of other attributes. In other words, if we select any two attributes A and B, then the probability of finding the attribute level B_i in a product bundle is the same irrespective of the particular attribute level of A_j found in that product bundle. One of the common methods of finding such orthogonal combinations is through the "Addelman" designs (Addelman, 1962). Many software packages can create such designs automatically. Knowledgeable users can create alternative designs that account for both "main effects" and interactions, or create non-orthogonal designs that are still efficient in obtaining information from the study participants (e.g., adaptive conjoint analysis). According to Wittink and Cattin (1989), commercial applications of conjoint analysis used a median number of 16 product profiles for obtaining respondent evaluations.

The number of independent parameters to be estimated is equal to

$$\left\{ \sum_{i=1}^{N} (n_i - 1) \right\} - 1 \qquad (1)$$

where N is the number of attributes and n_i is the number of levels of attribute i. For each product attribute we can arbitrarily set the lowest utility value (say, equal to zero). We can also arbitrarily set the maximum total utility from any product (say, equal to 100).

In some circumstances orthogonal designs can result in unrealistic products, such as when respondents perceive some of the attributes used in the study to be correlated—automobile horsepower (hp) and gas mileage (mpg) typically have a high negative correlation, but orthogonal designs could result in hypothetical products that combine high hp with unrealistically high levels of mpg. If a product is unrealistic in an orthogonal combination, there are several possible remedies: (1) We can combine the attributes and develop a new set of levels for the combined attribute. (For example, hp and mpg might be combined into a "performance" attribute with high performance associated with high hp and low mpg, and low performance associated with low hp and high mpg.) (2) We can replace unrealistic

products by substituting other combinations (perhaps generated randomly, but not duplicating the retained combinations). While this approach compromises orthogonality, it will rarely affect the estimated utility functions significantly if we replace only a few bundles (say, less than five percent). (3) We can select other orthogonal combinations (although this remedy requires special expertise).

An additional consideration in developing a suitable design is the exact nature of the data collection instrument. There are several options here, including, (1) asking respondents to sort and rank-order a set of cards, each containing a description of a product bundle, (2) asking respondents to rate each product bundle, say on a scale of 0 to 100, to reflect their likelihood of buying that product, (3) presenting respondents a sequence of product bundles, two at a time, and asking them to assign 100 points between them, and (4) offering respondents a sequence of sets of product bundles and asking them to choose one product from each set. Each of these data collection options has an associated set of costs and benefits.

Additional details about the foregoing aspects of conjoint analysis are available in a number of published sources, including Green, Krieger and Wind (2001) and Hauser and Rao (2003).

Using conjoint data for market simulations: Depending on the exact nature of the data collected, there are various options for analyzing the data and creating a part-worth function for each respondent. The simplest approach is to use dummy variable regression with ratings or rank-order data (data collection options 1 and 2 listed in the previous section).

$$R_{ij} = \sum_{k=1}^{K} \sum_{m=1}^{M_k} a_{ikm} x_{jkm} + \varepsilon_{ij}, \tag{2}$$

where

j = a particular product or concept included in the study design;

R_{ij} = the ratings provided by respondent i for product j; (Alternatively, the rankings could be reversed so that higher numbers represent stronger preference, and then used as if they are similar to interval-scaled ratings);

a_{ikm} = part-worth associated with the mth level (m=1, 2, 3, ..., M_k) of the

4

kth

attribute;

M_k = number of levels of attribute k;

K = number of attributes;

x_{jkm} = dummy variables that take on the value 1 if the mth level of the

kth

attribute is present in product j and the value 0 otherwise; and

ε_{ij} = error terms, assumed to be normal distribution with zero mean

and variance equal to σ^2 for all i and j.

To facilitate interpretation, the a_{ikm}'s obtained from regression can be rescaled so that the least preferred level of each attribute is set to zero and the maximum preferred product combination is set to 100, producing results that are more easily interpreted. Letting $\tilde{a}_{ikm}$'s denote the estimated (rescaled) part-worths, the utility u_{ij} of a product j to customer i is equal to

$$u_{ij} = \sum_{k=1}^{K} \sum_{m=1}^{M_k} \tilde{a}_{ikm} x_{jkm}, \tag{3}$$

Note that product j can be any product that can be designed using the attributes and levels in the study, including those that were not included in the estimation of the part-worths in Eq. (2).

A major reason for the wide use of conjoint analysis is that once part-worths ($\tilde{a}_{ikm}$'s) are estimated from a representative sample of respondents, it is easy to assess the likely success of a new product concept under various simulated market conditions. One might ask: What market share would a proposed new product achieve in a market with several specific existing competitors? To answer this question we first specify all existing products as combinations of levels of the set of attributes under study. If more than one competing product has identical attribute levels, we need to include only one representative in the simulation.

Transforming preferences to choices: To complete the simulation design we must specify a choice rule to transform part-worths into the product choices that customers are most likely to make. The three most common choice rules are maximum utility, share of utility, and logit.

Maximum utility rule: Under this rule we assume that each customer chooses from the available alternatives the product that provides the highest utility value, including a new product concept under consideration. This choice rule is most appropriate for high-involvement purchases such as cars, VCR's, and other durables that customers purchase infrequently.

We can compute the market share for a product by counting the number of customers for whom that product offers the highest utility and dividing this figure by the number of customers in the study. In computing overall market shares it may sometimes be necessary to weight each customer's probability of purchasing each alternative by the relative volume of purchases that the customer makes in the product category:

$$m_j = \frac{\sum_{i=1}^{I} w_i p_{ij}}{\sum_{j=1}^{J} \sum_{i=1}^{I} w_i p_{ij}}, \qquad (4)$$

where

I = number of customers participating in the study;

J = the number of product alternatives available for the customer to choose from, including the new product concept;

m_j = market share of product j;

w_i = the relative volume of purchases made by customer i, with the average volume across all customers indexed to the value 1; and

p_{ij} = proportion of purchases that customer i makes of product j (or equivalently, the probability that customer i will choose product j on a single purchase occasion).

Share of utility rule: This rule is based on the notion that the higher the utility of a product to a customer, the greater the probability that he or she will choose that product. Thus each product gets a share of a customer's purchases in proportion to its share of the customer's preferences:

$$p_{ij} = \frac{u_{ij}}{\sum_j u_{ij}} \text{ for } j \text{ in the set of products J,}\tag{5}$$

where u_{ij} is the estimated utility of product j to customer i.

We then obtain the market share for product i by averaging p_{ij} across customers (weighting as in Eq. (4) if necessary). This choice rule is particularly relevant for low-involvement, frequently purchased products, such as consumer packaged goods.

This choice rule is widely applied in conjoint studies and often provides good estimates of market shares. However, as Luce (1959) notes, this rule requires that utilities be expressed as ratio-scaled numbers, such as those obtained from constant-sum scales where the customer allocates a fixed number of points (say, 100) among alternatives. Unfortunately, data from most conjoint studies do not satisfy this requirement.

•*Logit choice rule*: This rule is similar to the share-of-utility rule, except that the underlying theoretical rationale is different. To apply the share-of-utility model, we assume that the utility functions are basically accurate— but an element of randomness occurs in translating utilities into choice. In applying the logit choice rule we assume that the computed utility values are mean realizations of a random process, so that the brand with the maximum utility varies randomly, say from one purchase situation to the next. The choice rule then gives the proportion of times that product j will have the maximum utility:

$$p_{ij} = \frac{e^{u_{ij}}}{\sum_j e^{u_{ij}}} \text{ for } j \text{ in the set of products J,}\tag{6}$$

Both the share-of-utility and the traditional logit rules share a questionable property known as IIA (independence from irrelevant alternatives). The choice probabilities from any subset of alternatives depend only on the alternatives included in the set and are independent

of any alternatives not included. This property implies that if, for example, you prefer light beers to regular beers, then adding a new regular beer (an irrelevant alternative) to your choice set would nevertheless lower your probability of choosing a light beer, a counterintuitive result.

How should we select among these choice rules? The maximum utility rule (also called the first choice rule) is simple and elegant, and choices predicted by using this rule are not affected by positive linear transformations to the utility function. This rule is particularly relevant for high-ticket items and in product categories where customers are highly involved in the purchase decisions. However, this rule predicts more extreme market shares, i.e., it has a tendency to produce market shares closer to zero and one than the other choice rules. Also, it is less robust—small changes in utility values of products can drastically change their market shares. On the other hand, the market share predictions made by the share-of-preference and logit choice rules are sensitive to the scale range on which utility is measured. The market share predictions of the share-of-utility rule will change if one adds a constant value to the computed utility of each product, but they are unaltered if all utility values are multiplied by a constant. Market share predictions of the logit choice rule are not altered if one adds a constant to the utilities, but will change if one multiplies all utilities by a constant. Thus, each rule has its advantages and limitations.

One way to choose among these three rules is this: First, for each rule, compute the predicted market shares of just the existing products. Then use the choice rule that produces market shares that are closest (in the sense of least squares) to the actual market shares of these products (this assumes that we are using a representative sample of customers for the study). This approach can be formalized using yet another choice rule, called the alpha rule, proposed by Green and Krieger (1993). We describe this rule next.

Alpha rule: This rule is a weighted combination of the maximum utility rule and the share-of-preference rule, where the weight is chosen to ensure that the market shares computed in the simulation are as close as possible to the actual market shares of the existing products. Specifically, we choose an alpha (a) in the following formula to maximally recover the observed market shares of existing products:

$$p_{ij} = \frac{u_{ij}^{\alpha}}{\sum\limits_{j} u_{ij}^{\alpha}} \qquad\qquad (7)$$

To determine the best value of a, we minimize the "entropy" representing the extent of departures of computed markets shares of existing products from their actual observations:

$$Entropy = \sum_{j} m_{j} \ln\left(\frac{m_{j}}{\hat{m}_{j}(\alpha)}\right) \qquad\qquad (8)$$

where j is an index to represent an existing product, m_j is the actual market share for product j, and $\hat{m}_j$ is the computed market share of product j for any given α. More details about this procedure can be found in Green and Krieger (1993). Orme and Huber (2000) have recently proposed another rule called *randomized first choice* that shows promise as an alternative to the four choice rules that we have outlined. Conceptually, it combines the ideas inherent in the logit choice rule and the alpha rule.

Additional Considerations in Simulations

Computing contribution instead of market share: Products that deliver high market shares need not necessarily result in high profitability for the company. Market share computations do not take into account the costs of manufacturing each product profile. A simple (and rough) way to measure incremental contribution of a product bundle (price – unit variable costs) is to first define a base product bundle and its contribution margin. Next, we can specify the incremental variable costs (positive or negative) for each attribute level, compared to the attribute level of the base product bundle. Finally, we can set the revenue index potential for the base product at 100 and measure the revenue index of every other product with respect to this base level. We show this below with a numerical example:

Unit contribution of base product		Market share (as per any selected choice rule)		Normalization factor	
	×		×		= 100

Suppose the unit contribution of the base product is $2 and the market share

as per the maximum utility rule is 25 percent, then the normalization factor is two. Now, if the incremental contribution of another product bundle is $1 as compared to the base product, and its computed market share per the maximum utility rule is 40 percent, then its revenue index is ($2 - $1)$\times$ 40 $\times$ 2 = 80. Note that any additional fixed costs have to be added separately, outside the model. Note also that this computation ignores the price potential (i.e., what the market might be willing to pay) for all product bundles, except the base product. These simplifications suggest that the base product has to be selected carefully to ensure that interpretations of the revenue index are both meaningful and appropriate for the given context.

Segmenting customers based on their preferences: In our discussion so far, we have focused on individual-level conjoint analysis and predicted market behavior by aggregating individual customer choices. That is, we implicitly grouped all customers into one segment. An important question is this: How should we conduct market simulations if the market consists of distinct groups of customers? There are several options:

1. *Post-hoc segmentation:* We could use traditional cluster analysis of the part-worth data to identify segments of customers with differing preference structures (see technical note on segmentation and targeting at www.mktgeng.com).

2. *Latent class segmentation:* This approach is most useful when we don't know (or cannot make a good guess about) the number of segments (see technical note on segmentation and targeting at www.mktgeng.com); however, this approach requires larger sample sizes than traditional cluster analysis.

3. *Hierarchical Bayes:* We assume that customers come from one or more populations (i.e., from a finite mixture of populations), but that customers in each population have different part-worth functions determined according to a specified distribution (e.g., multivariate normal). We can then estimate the "posterior point estimates" of each respondent's part-worth function, conditional on that respondent belonging to a given population (segment). These part-worth estimates may be used in simulations to determine the expected market share for any specified product in any segment, and across the overall population. Hierarchical Bayes methods are computer-intensive, but software

packages are now available to simplify their practical applications.

In a comprehensive simulation study comparing latent class method with Hierarchical Bayes, Andrews, Ansari, and Currim (2002) find that both have comparable out-of-sample prediction accuracy, although the Hierarchical Bayes had somewhat better with the in-sample data.

Choice-based conjoint analysis: Another approach to conjoint analysis is choice-based conjoint, originally proposed by Louviere and Woodworth (1983). Here, customers are presented with several sets of product profiles and are asked to choose the product in each set that they prefer the most. Exhibit 1 shows an example of a choice task with four profiles in a set. Thus, we directly measure customer choices, rather than measuring their preferences, and then converting the preferences to choices by using choice rules. The basic premise is that when we measure customer choices, the research process more closely resembles what people will actually do in the marketplace. The sets of profiles presented to customers are carefully selected according to experimental design criteria. The resulting "choice data" can then be analyzed using the multinomial logit model (see technical note on Choice Modeling at www.mktgeng.com). Several studies comparing the relative performance of the traditional "full-profile" conjoint and choice-based conjoint seem to indicate that both models tend to predict equally well (see Elrod, Louviere, and Krishnakumar 1992). The choice-based conjoint has the advantage of offering statistical tests of attribute weights and market shares. However, it is an aggregate model that does not offer direct measures of utility functions at the individual level, making it difficult to incorporate segmentation analyses of the part-worth data.

Hierarchical Bayes approaches can also be used with choice-based conjoint analysis. This approach enables us to ask subsets of questions from a larger sample of respondents (thereby reducing respondent fatigue) but also, at the same time, obtain useful posterior point estimates of the entire part-worth function for each respondent. Further details about these methods can be found in Allenby and Ginter (1995) and Johnson (2000).

Which Checking Account Would You Choose?

Services included with all four accounts	• ATM card with free unlimited use of your bank's ATM machines • Unlimited free check writing with no per-check charges • Unlimited free access to automated account information over telephone			
Monthly fee and minimum balance required	**Account #1** No minimum balance required. Pay a monthly fee of $10 if your account balance drops below $0.	**Account #2** Maintain a $500 minimum checking balance to avoid a $10 monthly fee.	**Account #3** Maintain a $1,500 minimum checking balance to avoid a $10 monthly fee.	**Account #4** Maintain a combined balance of at least $15,000 in your checking and savings accounts to avoid a $25 monthly fee.
Competitive interest paid on checking accounts	No	No	Yes	Yes
Use of other banks' ATM machines	$1.50 for each ATM visit	$1.00 for each ATM visit	$1.00 for each ATM visit	Free
Priority access telephone number	No	No	No	Yes
Choose the one account you prefer	❑	❑	❑	❑
❑ None of these				

EXHIBIT 1

A typical choice-based conjoint task for respondents. Note that respondents can opt not to choose any of the four profiles presented. Source: Cohen 1997.

Summary

Conjoint analysis is a customer preference measurement and analysis technique that is widely used in several areas of marketing. It is particularly useful for designing new products that are likely to perform well in the marketplace, and for determining the "optimal" changes to make to existing products (e.g., changes to price, changes to product features) to improve their market performance. There are three basic stages in a conjoint analysis: (1) designing the study, (2) getting data from customers, and (3) conducting simulations to determine good product designs. We described some of the technical aspects that are relevant for the first and third stages. Conjoint analysis is an active area of research, and new methods and enhancements are continuing to be developed, primarily because of advances in modeling and computing technologies.

References

(This is not the complete set of references. Others need to be added from our master list).

Addelman, Sidney, 1962, "Orthogonal Main Effects Plans for Asymmetrical Factorial Experiments," *Technometrics* 4**,** 21-46.

Allenby, Greg M., and Ginter, James L., 1995, "Using extremes to design products and segment markets, *Journal of Marketing Research*, Vol. 32, No. 4 (November), pp. 392-403.

Andrews, Rick L., Ansari, Asim, and Currim, Imran S. , 2002, "Hierarchical Bayes Versus Finite Mixture Conjoint Analysis Models: A Comparison of Fit, Prediction, and Partworth Recovery," *Journal of Marketing Research*, Vol. XXXIV (February), 87-98.

Elrod, Terry; Louviere, Jordan J.; and Krishnakumar, S. Davey, 1992, "An empirical comparison of ratings-based and choice-based conjoint models," *Journal of Marketing Research*, Vol. 24 (August), pp. 368-377.

Green, Paul E., and Krieger, Abba M., 1993, "Conjoint analysis with product-positioning analysis," in *Handbook of MS/OR in Marketing*, Jehoshua Eliashberg and Gary Lilien, eds., Elsevier Publishing, The Netherlands, pp. 467-515.

Green, Paul E.; Krieger, Abba M.; and Wind, Yoram (Jerry), 2001, "Thirty years of conjoint analysis: Reflections and prospects," *Interfaces*, Vol. 31, No. 3, Part 2 of 2 (May-June), pp. S56-S73.

Hauser, John R. and Rao , Vithala R., 2003, "Conjoint Analysis, Related Modeling, and Applications," in Market Research and Modeling: Progress and Prospects: A Tribute to Paul E. Green (International Series in Quantitative Marketing), Yoram (Jerry) Wind and Paul E. Green (Editors), ISQM, International Series in Quantitative Marketing, Kluwer Academic Publishers.

Johnson, Richard M., 2000, "Understanding HB: An Intuitive Approach," Sawtooth Software Research Paper Series, www.sawtoothsoftware.com.

Louviere, Jordan J., and Woodworth, George, 1983, "Design and analysis of consumer choice or allocation experiments: A method based on aggregate data," *Journal of Marketing Research*, Vol. 20 (November), pp. 350-367.

Luce, R. Duncan, 1959, *Individual Choice Behavior*, John Wiley and Sons, New York.

Orme, Bryan, and Huber, Joel, 2000, "Improving the value of conjoint simulations*, Marketing Research,* Vol. 12, No. 4, pp. 12-20.

Wittink, Dick and Cattin, Philipe, 1989, "Commercial Use of Conjoint Analysis: An Update," *Journal of Marketing*, 53 (July), 91–96.