



Sawtooth Software

RESEARCH PAPER SERIES

Learning Effects in Preference Tasks: Choice-Based Versus Standard Conjoint

Joel Huber, Duke University
Dick R. Wittink, Cornell University
Richard M. Johnson, Sawtooth Software, Inc.
and
Richard Miller, Consumer Pulse
1992

LEARNING EFFECTS IN PREFERENCE TASKS: CHOICE-BASED VERSUS STANDARD CONJOINT

Joel Huber

Duke University

Dick R. Wittink

Cornell University

Richard M. Johnson

Sawtooth Software

Richard Miller

Consumer Pulse

Published in Sawtooth Software Conference Proceedings, 1992

INTRODUCTION

A current "hot topic" among conjoint users is the status of choice-based conjoint. Choice-based conjoint derives utilities directly from a series of choices among profiles. That is, instead of making judgments on the likelihood of choosing an alternative (as in full profile conjoint), or the degree of preference between a pair (as in ACA (ACA System for adaptive conjoint analysis, by Sawtooth Software)), respondents choose the best out of successive choice sets.

A major advantage of choice-based over standard conjoint is that its task more directly represents market behavior. After all, consumers do not generally rate alternatives in terms of preferences; they simply choose. Thus, one would reasonably expect choice-based conjoint to better reflect current demand. The disadvantage of choice-based conjoint is that it cannot estimate utilities for each respondent. Instead, choice-based conjoint requires the analyst to aggregate across respondents in order to derive stable coefficients, making it less effective than conventional conjoint for uncovering and defining segments.

In this study respondents complete a short choice-based task before and after ACA, enabling us to answer three questions. First, does simply taking ACA change attribute values? Second, are the results from ACA significantly different than from choice-based conjoint? Finally, can ACA be modified to approximate the choice-based results, thereby permitting us to achieve the "validity" of choices with the precision and individual-level analysis of ACA? The answer to all three questions is yes, but, perhaps the most interesting result is that the differences between ACA and choice-based conjoint follow a predictable pattern. ACA uncovers the microstructure of preferences: how a person would choose given sufficient information and time. By contrast, the choice results portray customers who are primarily motivated by brand name and price—an appropriate strategy assuming they have little time to make a decision. Thus the decision between the standard and choice-based conjoint depends on the kind of choice process that the company finds in the market. Elaborate searches are more likely to be reflected in the depth of processing captured by ACA, whereas choice-based conjoint is most appropriate in reflecting more immediate decision-making.

Impact of Conjoint on Attribute Importance

Consider first how partworths can be changed by the environment of a conjoint task. A conjoint task, generally, and ACA in particular, separates attributes so that they can be evaluated independently from the levels of other attributes. This separation occurs through two mechanisms. Respondents are often asked to assume that levels of other attributes are constant, and to focus on the particular ones displayed. Further, the attributes that are displayed are uncorrelated; the level of one attribute in a profile is not associated with the level of other attributes. Indeed, a necessary condition for an orthogonal array is that the probability of finding one attribute level remains unchanged regardless of other attribute levels in the profile. This independence enables the orthogonal array to efficiently derive partworths since the estimate for one attribute is not biased or contaminated by the others.

In contrast, attributes in the marketplace are anything but independent. Indeed, it is their dependence that enables people to make rational decisions with just a few attributes. Perhaps the best example of an attribute used to make a decision is that of brand name. Brand name provides substantial information about likely performance, expected benefits and problems in use, as well as the relative price of the item. For example, a person considering an IBM PC might justifiably expect very solid quality and service, good but not outstanding performance, and a price slightly higher than could be found through a knowledgeable search.

In the conjoint environment, such inferences are soon shown to be unreliable. There one finds IBM models whose prices and performance levels range unpredictably from very high to very low. In short, the brand name, which is a reasonable criterion in the marketplace, becomes less so in the orthogonalized world of a conjoint exercise. This lessened value of brand as an indicator of quality leads to the prediction that names, such as IBM, will come to have less importance following a conjoint exercise.

Next consider price. High price conveys two pieces of information. A high price conveys negative information about the sacrifice of having to pay more, while at the same time, conveying positive information about high quality. Thus, there is the sacrifice aspect of high price, which is almost always aversive, and the inferential aspect, which is almost always desired. Conjoint, however, tends to drive out this inferential mechanism since prices are, by design, unassociated with quality levels found. To the extent that the positive inferential mechanism attached to price is minimized, the importance of the sacrifice aspect of price should increase. Thus, the expected effect of conjoint on price is opposite to that predicted for brand name. Price should increase in importance (become more aversive) as the association from high price to high quality is diminished in a conjoint task.

Summarizing, we hypothesize that the process of conjoint decouples attribute associations. That is, by showing attributes mixed with other attributes in unexpected combinations, respondents learn that inferences from one attribute to others are ineffective. The conjoint process thus teaches respondents to evaluate attributes individually and not by their inferred levels on other attributes. Thus a "name" brand will become less important if one cannot infer other good qualities from it, and a good (low) price should become more important because one does not infer other negative qualities on the basis of that price.

Most other attributes, like weight, size, and performance can be expected to act like price. That is, they are negatively correlated with each other in the environment. People will generally expect that a better level in one requires a sacrifice in the others. For example, decreasing the weight of a laptop will generally exact a penalty in terms of size, performance, or price. To the extent that a conjoint task decouples associations among negatively correlated attributes, then these attributes should become more important. This thinking leads to the hypothesis that a conjoint exercise makes brand name less important and most other attributes more important.

Below we describe a study that evaluates the impact of taking a conjoint task on attribute importance. Choice-based conjoint is used to measure changes in attribute weights and consistency due to an intervening ACA task. Then we consider the issue of the degree to which ACA matches these choice tasks and how it might be possible to convert from one to the other.

A Study to Examine the Impact of a Conjoint Task on Attribute Values

To evaluate the effect of conjoint on attribute values, we need a way to measure these values that is minimally invasive. Choice-based conjoint provides a fine way to measure relative tradeoffs at the group level and can be accomplished in a very short period of time. The technique we use is randomized choice-based conjoint. With this procedure, each respondent makes choices from a small number of sets whose attributes are randomly drawn from the set of all manipulated attribute levels. A multinomial logit model then derives partworth utility scores that predict these choices. The choice-based attribute partworths are analogous to standard partworths except that they predict the probability of choice rather than a preference rating.

The study (Wittink *et al.*, in this volume) involves the selection of laptop computers conducted by IntelliQuest. ACA is used to evaluate six attributes of laptop computers: brand name, price, relative speed, battery time, total weight and exterior size. Each of these attributes has four levels, with ranges shown in the first column of Table 1. The conjoint choices, given just before and after the ACA exercise, involve two pairs and two triples. For each respondent, profiles are randomly generated from the 64 possible profile combinations. These same choices are then repeated after ACA. Repeating the choice sets gives us a measure of consistency within respondents, while making them differ randomly across respondents permits us to make stable estimates of aggregate utilities and to test for any possible interactions.

For ease of exposition here, the four-level attributes are linearized to have proportional differences but to range between 0 and 1. Thus, for example, the four weights (5, 6, 7, 8) are normalized (to 0, .33, .66, 1). This normalization means that a coefficient for an attribute measures the value of its range. Brand name could not be linearized so its four levels are represented by the contrast between that brand and Librex.

The utilities shown in Table 1 reflect the raw utilities from the logit analysis divided by the consistency measures in the bottom row. We modified the raw utilities so that the various attribute coefficients would be comparable across methods. Understanding the reason for this transformation requires a brief digression into the nature of multinomial logit.

Multinomial logit defines the error level as a constant, so that greater fit is represented by larger coefficients. Those who use logit may have seen evidence of this property in the disquieting finding that, as one increases fit by adding more parameters, the absolute values of the original coefficients tend to rise. Logit can be usefully contrasted with linear regression in this regard since regression coefficients remain unbiased as one adds noise. Thus, if one adds random error to a dependent variable in regression, the coefficients remain quite constant, although the standard errors around those coefficients increase. With logit, since the error term is fixed, adding noise to the choices makes the coefficients smaller, while keeping the standard errors the same.

This property of logit in which the coefficients change, rather than the error term, hinders comparisons of coefficients across different data sets. For example, if the coefficient for price is different before and after ACA, it is unclear whether that is due to greater noise after conjoint or to structural differences in relative utility value. Table 1 separates the consistency (or noise) component from the relative utility component. The partworths are scaled by a linear dilation through the consistency factor so that the utility coefficients for the second choices are as close as possible to those for the first choices.

An examination of the rescaled utilities for the choice set before and after conjoint supports the hypothesis given earlier. Brand name becomes less important as indicated by the lower adjusted utility values for the three brand dummy variables. By contrast, after ACA, price, performance and battery time become more important. The changes in calculator weight and exterior size are not statistically significant. Overall, these results show that after a conjoint exercise, brand name becomes relatively less important and other attributes more important, as predicted by the idea that conjoint reduces associational links among attributes.

A second result is that the respondents are approximately 40% less consistent after ACA than before. This evidence of fatigue is unexpected, but may be due to the fact that the second choice set came at the end of a relatively long survey that contained a substantial section about computer usage in addition to the conjoint exercise.

Table 1 Utilities of Laptop Attributes For First and Second Choice Sets			
Attribute:	First choices: Before ACA	Second choices: After ACA	ACA predicted choices
IBM vs Librex	1.23 (.11)	1.02 (.16)	0.97* (.09)
Toshiba vs Librex	1.13 (.11)	0.74* (.16)	0.77* (.09)
Dell vs Librex	0.92 (.11)	0.73 (.16)	0.77 (.09)
Price \$2300 to \$3600	1.61 (.08)	2.08* (.12)	1.48 (.08)
Performance: 20% better to 10% worse	1.19 (.10)	1.41 (.15)	1.35 (.09)
Battery time 4 to 2 hours	0.63 (.07)	0.83* (.11)	1.01* (.07)
Total Weight 5 to 8 lbs	0.38 (.08)	0.21 (.12)	0.68* (.07)
Exterior size 8"x 10" to 10"x 12"	0.07 (.07)	-.05 (.12)	0.39* (.06)
Consistency	1	0.6*	1.4*

*Difference between weight given and first choice weight is significant at $p < .05$

How Well Does ACA Reflect Choice-based Conjoint Results?

We have shown that simply participating in ACA affects subsequent choices. Another important question to ask is the extent to which ACA's predicted choices approximate those found in the choice-based survey. We use the ACA utilities to predict each respondent's choice using the maximum utility rule: always choose the object with the greatest utility. A subsequent logit analysis then predicts these simulated choices as a function of the attributes. This analysis, shown in the third column of Table 1, reveals two interesting results.

First, ACA has much less variability than either the first or second choice, being 40% more consistent than the first choice battery and more than twice as consistent as the second choices. This makes sense; ACA's choices come from an analysis of a logically consistent model based on about 20 judgments. The average of these is likely to be more consistent than four individual choices. Second, we see that brand name and price are less important in ACA relative to the functional attributes. This decreased importance of brand name may be due to the impact of ACA in reducing the association between the brand name and the attributes -- the same account as given for why brand name became less important after ACA. The lowering of price importance found in the ACA compared with choices requires a different explanation, since limiting price associations should increase the importance of price, rather than decrease

it, as found. It may be that the focus on trading off the benefits of the alternatives, as found in ACA, underplays price. Consider the clever salesman who is able to steer customers to a higher priced product by focusing on the performance attributes. However, when one goes to choose, the critical attributes become brand name and price. In other words, focus is put on a few global attributes, rather than trying to balance a large number of attributes.

Which Measure is Correct?

The previous results raise an intriguing and important question about the usefulness of standard versus choice-based conjoint. We show that error levels and relative weights change depending on the method used to collect the data. Two results are clear. First, the fact that attribute associations are decoupled by standard conjoint leads to less reliance on brand name and to more reliance on functional attributes. Second, choice appears to additionally differ from ACA in putting more weight on price. But which measure is correct? Which measure provides a better estimate of what will happen in the market? Clearly, this study cannot answer that question, since there is no ultimate criterion by which to test the utilities found. We can however, usefully speculate on contexts in which one measure or the other may be more useful.

The big difference between ACA and choice-based conjoint is in the depth of processing. The four choices take the average respondent about two minutes, contrasting with ACA's time of 15-20 minutes. One must question how representative those choices are to those made in the marketplace. Consider first the category under study, laptops. Purchases of laptops are generally not made in anything like 30 seconds; people spend significant time discussing a wide range of features. Thus, the lack of attention to the "less important attributes" such as unit weight and exterior size may not be replicated in market choice. Further, market decisions are often made on the basis of recommendations from magazines or from people who have studied the options closely. Thus, a company following the advice of choice-based conjoint might well underestimate the importance of minor attributes and the depth of processing that occurs in market decisions. In contrast, ACA's depth may approximate these decisions better.

For consumer goods, a different argument leads the same recommendation to use standard over choice-based conjoint. On any given purchase, consumers are typically very insensitive to anything other than brand name and price. Consumers learn about package goods through usage and through interacting with other users. A product that does not work well will

eventually lose share in the marketplace. In a sense, the lower weight for brand names in ACA is consistent with the idea that ultimately it is features that people buy, not brand names. A brand name comes to have value over time because it is associated with good features. Choice-based conjoint appears to be representing this immediate response to the market offering and thus should better reflect current market share. However, it would be very risky, if not foolhardy, for a company to reduce features because short-run profits are maximized. Eventually, customers will learn, and that pattern of learning is likely to be more closely reflected in ACA's utilities than in choice-based ones.

Another way to express the foregoing is to say that ACA provides a reasonable normative model for decisions. It approximates customers' needs as expressed through careful tradeoffs. Thus, if a company really believes it is in its long run best interest to give customers what they need, ACA appears to give a better characterization of those needs than choice-based conjoint.

Can ACA be Adjusted to Predict Choice?

Suppose one's goal is to approximate the choice-based results, but desires the additional accuracy of ACA. That is, Table 1 indicates that choice-based conjoint puts greater weight behind brand names and price relative to the other attributes. A logit analysis predicting choice on ACA's part-utilities can be used to determine this optimal reweighting.

Table 2
Weights for ACA's Part-Utilities to Predict Choices In the First Choice Set

Attribute:	Coefficient	(Standard error)
Brand Name	1.45*	(.09)
Price	1.30*	(.07)
Performance	0.95	(.07)
Battery time	0.75*	(.08)
Weight	0.64*	(.10)
Size	0.37*	(.11)

* Coefficient is significantly different from equal weight (all 1.0), p<0.05.

The part-utility of each attribute is the utility assigned to the particular level of each attribute. For example, the part-utility for IBM might be .5 while the part-utility for 20% greater performance might be .7, with similar kinds of numbers for the other four attributes. Table 2 shows how the weighting assumed by ACA should be modified to best predict choice.

It turns out that the revised model predicts choices very well. The replication consistency of the first and the second choices is 71%. That is, if one uses the replicated choices to predict the first choices, one would be right in 71% of the time. However, with unadjusted ACA, the

hit rate increases to 74%. Further, the reweighted ACA from Table 2 increases the hit rate to 78%. All of these differences in hit rates are statistically significant.

Thus, the ACA model of choices predicts choices better than choices themselves! Some of this result comes from the use of the erratic second choices as the basis for prediction, but much comes from the stability in the ACA utilities, compared to the noisiness of actual choice. The result does imply that one can augment the benefit of choice-based conjoint by combining it with ACA's output.

Discussion

These results tell us more about choice-based conjoint than ACA, although they are relevant to both. First, we find evidence that choice-based conjoint is susceptible to an intervening task. If respondents are tired, it may appear in the consistency of the choice-based conjoint. Second, the act of taking ACA also has the effect of increasing the relative importance of functional attributes, while decreasing the importance of brand name as a surrogate for those attributes. These results are consistent with the idea that a conjoint task generally lessens the impact of associations among attributes. Third, choice-based conjoint reveals a more simplified decision rule than does ACA, with greater emphasis on brand name and prices. Finally, we show that choice-based conjoint can be predicted quite well by modifying the relative weights given utility components for each attribute to better reflect choice behavior.

Choice-based conjoint has become increasingly popular lately, largely because its task appears closer to actual choice in the marketplace. Our results here are relevant to that perception. Choice-based conjoint appears to tap immediate decisions, where the attributes are clearly displayed for the decision-maker. Most decisions for laptops, in particular, and durables in general, are not made quickly. Further, even if individual consumers may not trade off features, they read magazines and speak to people who do. Thus, there is good reason to believe that the high emphasis on brand and price found in choice-based conjoint may not accurately track long-term market response to feature improvements.

The implications for the design of conjoint studies is inescapable—collect both standard and choice-based conjoint. Standard conjoint provides a detailed image of the microstructure of customer needs, one that permits segmentation of each respondent and reveals tradeoffs for attributes that might be otherwise ignored. Choice-based conjoint, for its part, portrays a more accurate picture of short-term customer response to brand name and price. The relative weighting of choice-based and standard conjoint depends on the kinds of projections the marketing researcher needs to make. To the extent that the perceptions of features, brand names and prices are expected to remain stable in the future, then the choice-based analysis is quite appropriate. However, many strategies in the marketplace involve altering features and relative prices—upsetting the correlational structure among the attributes. In that competitive environment, the effect of conjoint of breaking down the associations among attributes may be the best way to mimic the effects of competing companies in an active market.

Finally, the ability to predict choice-based conjoint with ACA provides a way to have the advantages of both methods, which we may call choice-adjusted ACA. In this procedure, Ci3 (Ci3 System for computer interviewing, by Sawtooth Software), can be used to build the same

kind of short randomized choice task as employed in this study. Then ACA's partworth utilities are adjusted to predict these choices and provide input to the choice simulator. Choice-adjusted ACA permits aggregate choices to be predicted without losing information on each individual respondent. That is, one can segment respondents by their choice-adjusted utilities. By contrast, generally with choice-based conjoint there is not enough information from the 16 or so choices from each respondent to provide satisfactory stability for grouping individuals. Thus, ACA is needed even if one does not accept the argument that ACA's partworths are the appropriate ones for a customer-centered firm to use. It is needed to enable choice-based conjoint to better do its job.