Sawtooth Software

RESEARCH PAPER SERIES

Multistage Conjoint Methods to Measure Price Sensitivity

Jon Pinnell
IntelliQuest, Inc.
June 1994

530 W. Fir St.
Sequim, WA 98382
(360) 681-2300
www.sawtoothsoftware.com

Multistage Conjoint Methods to Measure Price Sensitivity

Jon Pinnell
Director of Marketing Sciences
IntelliQuest, Inc.
Austin, TX

For presentation at the 1994
Advanced Research Techniques Forum

June, 1994
Beaver Creek, Colorado

The author wishes to thank Bryan Orme
for conducting many of the
calculations in this paper.

Multistage Conjoint Methods to Measure Price Sensitivity

Introduction

Over the last decade, conjoint analysis has become widely accepted for research involving product design, as well as measuring the market's sensitivity to pricing changes. Conjoint analysis has as its objective the measurement of part-worth utilities for product features, typically at the level of the individual respondent, and then the estimation of preference for "configured" products through simulations.

The accuracy of these conjoint simulations has been a topic of debate almost as long as conjoint has been around (See Wittink and Walsh, 1988, for a good review).

In practice, many conjoint applications must deal with many attributes, more than can be presented effectively in full profiles. One such form of conjoint analysis that has received much attention, and has been the topic of much debate, is Sawtooth Software's Adaptive Conjoint Analysis (ACA). (For a review of studies concerning the validity of ACA, see Johnson, 1991.)

Some researchers using ACA have found that shares of preference for expensive products are over-stated, and that those shares drop too slowly when simulated prices are increased. Put another way, the importance of price is understated. This issue is particularly acute in conjoint studies with many non-price attributes and when price is included as just one more attribute.

Several hypotheses exist as to why this condition occurs. One hypothesis is a potential lack of independence between some of the non-price attributes in the study. That is, conjoint methods that present less than full profiles to the respondent require the respondent to keep an "all other things equal" mind set, and to remember that the concepts are identical on all omitted attributes. But this may be extremely difficult for respondents. For instance, quality, performance, and reliability could all be included in one study, but might not be seen as independent by the respondent. A product represented as higher in quality might also be seen as likely to be higher in performance and reliability. Price, in contrast, might not be represented by multiple attributes. This could result in "multiple counting" of some attributes, but not price.

A second hypothesis is that data collection techniques which force respondents to pay attention to all attributes may make the importance of all attributes more similar. This would tend to lessen the importance of the more important attributes, of which price is probably one.

The purpose of this paper is not to explore these alternative explanations of the effect. Rather, this paper seeks to understand ways to deal with the phenomenon, given that it occurs.

It is interesting to consider ways this problem can manifest itself. Four signs can indicate a problem.

- The importance of price being too small, but this is too subjective to be useful

- The dollar values for product feature differences being too large, which is again too subjective

- The market sensitivity being too small for price changes, which starts to become more concrete, and finally...

- The severe overestimation of shares of preference for products at the high-end of the product line, which is clearly the most telling.

The simulated shares of preference can be compared either to market share of existing products or, exercising slightly more control, to choices the same respondents made in the survey. Consider the two following hypothetical distributions of shares of preference estimated from a conjoint simulator and shares of choices from the same respondents.

Figure 1

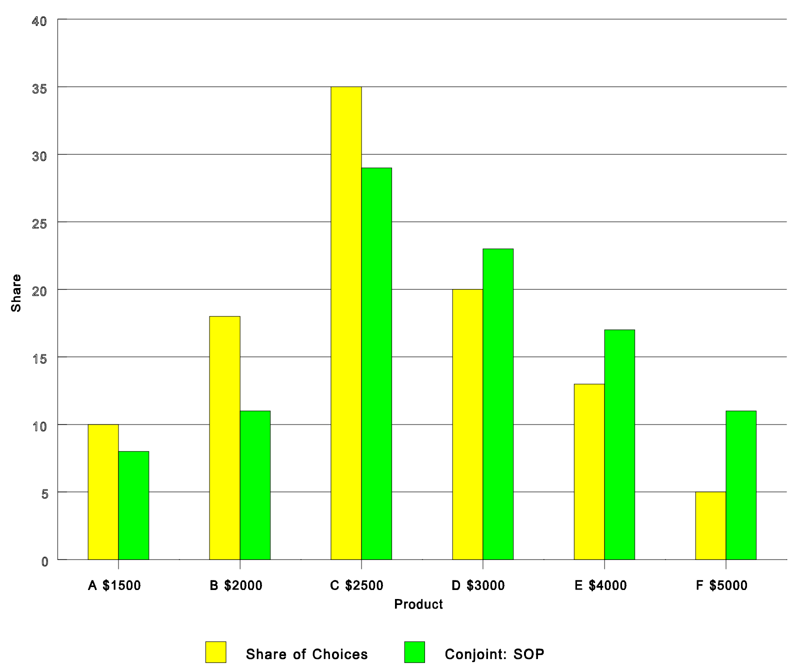

Notice that they have the same mean absolute deviation of estimated share, but in the first distribution (Figure 1) all of the positive errors are above the average price and all of the negative are below the average price. That is different from the second distribution (Figure 2) in which the errors appear to be random with respect to price.

Figure 2

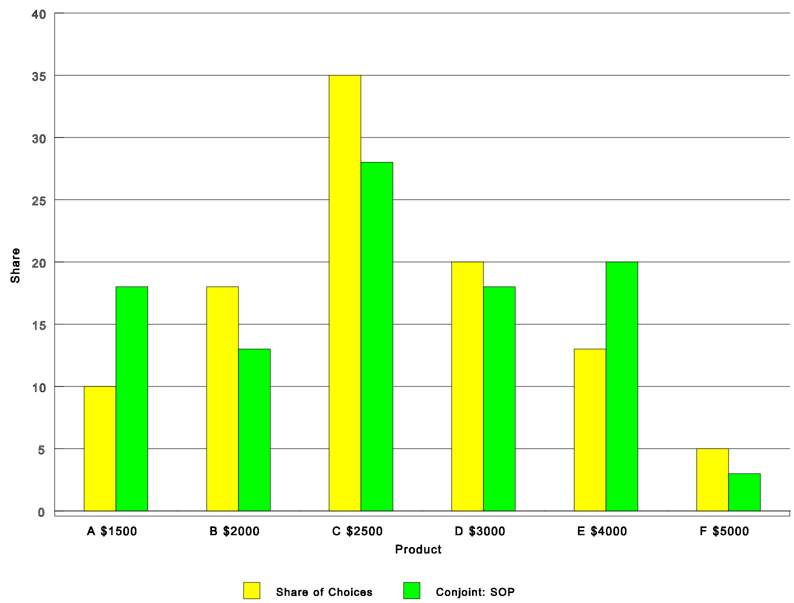

The first point I would like to make is that it is a good idea to include holdout choices in every conjoint study in addition to the primary conjoint task. The findings from those questions can be used not only to create distributions like these to test for overstatement of share of high priced products, but also to measure consistency at the level of the individual respondent.

The second point I would like to make is to consider "dual conjoint" as a remedy to the problem of mis-estimation of the effect of price.

One approach to solving this problem, and providing a more accurate determination of the importance of price, has been dual conjoint analysis. In dual conjoint, each respondent participates in two conjoint studies. One conjoint study is used to examine a possibly large number of product attributes and features, and a second conjoint study with the same respondents concentrates on just price, a performance variable, and brand. The first design is used to quantify all of the relevant attributes and features other than price; the second measures price sensitivity in combination with only major variables such as brand and significant performance differences. Typically, the second study is used only to solve for the effect of price, relative to other attributes. Because the second study is used to solve for very few parameters, less information is required per respondent.

There are several ways to implement dual conjoint. There are necessarily two conjoint studies for each person, and the method which is best for the first study might not be the best for the second. As stated above, the problem is particularly acute in the presence of many (10 or more) attributes. When there are many attributes, and computer administration is possible, hybrid methods, like ACA, are preferred for the first study (Green and Srinivasan, 1990). Two converging fields of conjoint analysis make the question of which method is best for the second administration particularly interesting. One option is to use a second ratings-based conjoint method as the second method, or dual. Alternatively, one could use a choice-based conjoint method.

IntelliQuest has recently completed a large methodological study to compare alternative implementations of dual conjoint. The study compared ACA alone with two dual methods, both of which used ACA as the first stage, and which used either ratings-based conjoint or choice-based conjoint as the second stage. Methods were compared using three different criteria.

First, each method was measured in terms of its ability to produce accurate predictions of individual respondent choices. The choice tasks used for this purpose were administered twice in the interview, once at the beginning and once at the end so that we could measure the reliability of respondent choices themselves.

Second, each method was used to calculate the monetary difference in utility between the respondent's most preferred brand and second most preferred brand.

Finally, aggregate shares of preference were compared using a preference simulation model.

Research Design

To test the alternative data collection methodologies, IntelliQuest screened and mailed disk-by-mail surveys to 1,505 personal computer purchase influencers in business settings. During the pre-qualifying stage, all respondents indicated they would be involved in a computer purchase in the next three months. Of the 1,505 respondents who were mailed survey diskettes, 715 returned completed surveys, representing a 48% response rate. Of the 715 returned surveys, 63 were excluded as not qualified or internally inconsistent, leaving a total of 652 qualified respondents.

The disk-by-mail survey included a similar ACA module for all respondents. In addition, respondents were randomly assigned to two groups. One group received a ratings-based conjoint while the other group received discrete choice, or choice-based conjoint. The ACA module included the following 11 attributes:

Brand
Microprocessor
Performance
Monitor
Storage
Model (desktop or tower)
Channel
Service
Warranty
Upgrade ability
Price

Table 1 indicates the levels studied for each attribute. The Brand attribute included five brands that were customized for each respondent. Respondents were each asked to indicate the brand they were most likely to purchase, the brand they were second most likely to purchase, third most likely and so on. The brand attribute in conjoint was customized based on each individual's most likely brand, second most likely, third most likely, fifth most likely, and tenth most likely. For the cases in which respondents did not rank order ten brands, reasonable brands were imputed. If a respondent ranked only two brands, a high share brand that had not been mentioned was imputed as the third place brand, a medium share brand was imputed as the fifth place brand, and a low share brand was imputed as the tenth place brand.

Respondents were divided into high and low price tiers based on the prices they expected to pay for the personal computer they were going to purchase. The two groups evaluated the following prices:

Low Price Tier	High Price Tier
$1,500	$2,700
$1,800	$3,200
$2,200	$3,800
$2,500	$4,300

Frequently, dual conjoint studies include price and brand only in the second task. For this design, we included price and brand in both the first and second conjoint task. This design allows the evaluation of ACA alone versus a dual design. Since the utilities for ACA were estimated using a main effects design, the inclusion of price and brand should have no impact on the utility estimates of the other parameters.

The ratings-based and choice-based conjoint designs each included only Price, Brand, Performance, Warranty, Service, and Channel. These six attributes were present in every task; that is, both the ratings-based and choice-based conjoints used full profiles of the subset of attributes. The ratings-based presented paired comparisons of concepts while the choice-based presented triples with a none option. The ratings-based conjoint involved 14 pairs and all respondents saw the same concepts (except for the customized brands and price). The choice-based conjoint used randomized designs. Each respondent provided first and second choices to eight choice tasks. Price and brand were customized in the dual stage in the same way as in the ACA conjoint.

Empirical Results

Predictive Ability

Methods will be compared initially by their ability to predict respondent choices. In addition to the ACA and the ratings- or choice-based conjoints, each respondent also provided first and second choices for four hold-out discrete choice tasks. The holdouts were presented to each respondent twice, once at the beginning of the interview and once at the end. The holdout tasks were comprised of three concepts defined on the same six attributes as the second ratings-based or choice-based conjoint task. The pre holdouts were identical to the post. They also included customized brand and price for each respondent, so each respondent's concepts were unique. By collecting first and second choice among three alternatives for each task, three pairwise statements can be inferred. They are:

> The first choice is more preferred than the second choice
> The first choice is more preferred than the non-chosen concept
> The second choice is more preferred than the non-chosen concept

Since each respondent had four holdout tasks, each containing three concepts, each respondent made 12 implied preference judgements involving pairs of concepts. Those were used to evaluate the consistency of each respondent's preferences, as well as the predictive capabilities of each alternative methodology.

Each conjoint method produced a different set of utilities that may be used to predict which concept each respondent should prefer in each pair. We assessed the success of each conjoint method by noting the average number of successful predictions for each respondent, as well as the percentage of pairwise predictions that were correct.

Success in predicting preferences depends partly on the consistency of what is being predicted: if the respondents answer at random, then we would not expect any conjoint method to predict those answers very well. We assessed the consistency of preferences for which the respondent answered consistently in both administrations of the holdout tasks.

The average number of consistent pairwise preferences per respondent was 10.125, for a percentage of 84.38. It can be shown (Wittink and Johnson, 1991) that the maximum expected success in prediction for any conjoint method given this level of reliability is 91%.

Findings

ACA Alone

Using each respondent's utilities from ACA, we are able to predict each respondent's preference for each concept in the hold-out tasks. By simply summing the part-worth utilities for each task and identifying the one with the largest sum, we can predict a first choice. Likewise, we can predict a second choice. We can then identify how many of the twelve inferred pairwise statements of preference ACA correctly predicts.

ACA correctly predicts 9.944 or 82.87% of the 12 pairwise comparisons. This number is slightly lower than what would be obtained from predicting the post holdouts from the pre-holdouts (10.125).

We can test the significance of this difference by comparing, for each individual, the number correctly predicted out of 12 for the pre-holdouts and for ACA. The difference in prediction success, at the level of the individual, has a mean of 0.182, and a t-ratio of 2.36.

One additional point should be made concerning ACA. The holdouts measured six attributes, while the ACA interview measured 11, so part of the resolving power of ACA was wasted in measuring predictive ability of these holdout choices.

As outlined previously, the goal of the study was to compare various implementations of dual conjoint. ACA will serve as the first part of the two dual designs, but will also serve as benchmark in addition to the pre-holdouts.

Given that, the first dual execution to be tested was ACA plus ratings-based conjoint.

<u>ACA plus Ratings-based</u>

The full profile ratings-based conjoint is designed as a three attribute study with brand, price, and a composite performance attribute. The composite performance attribute is created by combining levels of several other attributes. For instance, in this study performance, warranty, service, and channel were combined to form a new attribute (the performance bundle). Using this design, each respondent provides enough information that we can calculate the partworth utilities for all levels of each of the three attributes for each individual. We can then calculate the utility of the performance bundle by summing ACA utilities and, comparing that to the utility of the performance bundle in the second design, we can develop a scaling factor by which these two conjoint studies can be combined. In this way, we can use all of the ACA utilities except price, which is estimated and combined from the ratings-based conjoint. After performing these calculations and manipulations, we have a new set of utilities, which are identical to the ACA set except for price. Comparing to ACA alone, ACA plus ratings-based correctly predicted an average of 0.35% fewer of the 12 pairwise relationships. This difference however is not significant.

Using the individual utilities from the full-profile ratings-based conjoint, it is also possible to bridge *both* price and brand into the other ACA utilities. Conducting these calculations and manipulations we have another set of utilities for each person. These can also be evaluated relative to ACA alone in terms of their ability to predict the pairwise relationships from the hold-out choices. This method improves upon ACA's prediction by an average of 0.72% of the 12 pairwise relationships. This compares favorably to ACA plus ratings-based bridging price only (-0.35%) and pre-holdouts (+0.78%). This method, however, does not represent a significant improvement over ACA alone or ACA plus ratings-based conjoint bridging price only, nor is it significantly worse than pre-holdouts.

ACA plus Choice-based

The choice-based study can also be used to estimate individual level utilities. Unlike the ratings-based method, however, each person doesn't provide enough information to allow us to estimate utilities for each level of each attribute for each person. Just like the ratings-based conjoint, we create a regression equation which equates all of the features in the concept to the response given (either a choice or a rating). Then, instead of solving for the features, we can substitute the ACA utilities for everything besides price, leaving only price and a scaling factor to be estimated. Price can be dummy coded in which a unique coefficient will be found for all but one of the levels of price. Alternatively, we can solve for a single price parameter, either assuming price to be linear, or following some curvi-linear distribution, like the log of price. We will consider the effects coded and the linear case.

First, we will consider using effects coding to estimate price. Again, comparing to ACA alone, this dual method correctly predicts 0.72% fewer of the 12 pairwise statements. This difference is not significant. Since we used a split sample between ratings based and choice based, we cannot conduct tests of significance at the individual level between choice-based and ratings-based.

With the choice-based method, we can estimate utilities for price, not as n-1 separate regression coefficients, but by assuming price to be linear as one parameter. This approach has the benefit of using fewer degrees of freedom, as well as slightly coercing price to behave correctly.

As conjectured, the linear model of price improves the predictive ability of the choice-based conjoint method. The linear model develops utilities that correctly predict 0.5% fewer of the 12 pairwise statements than ACA alone, and 0.22% more than ACA plus choice based with price effects coded.

Weighted ACA

In addition to the utilities derived from the three methods outlined above, two additional methods were tested. Both additional methods involve re-weighting the price utilities as determined by ACA. This has the benefit of allowing ACA to determine the shape of the price utilities, while adjusting their height by including additional information. Since, as mentioned above, we included price in the ACA study as well as the dual, we were able to specify a regression equation using the ACA utilities to predict choices from the dual component. Solving that equation, we are able to identify a single weight that, when applied to all respondents, best fits the choices from the dual study. Additionally, by conducting the regression analysis for each person individually, we are able to identify the best weight for each person individually.

The idea of weighting utilities is not widely practiced, but a novel and potentially powerful approach (for instance, see Huber, 1992).

By weighting ACA's price at the level of the individual, such that each person has a unique weight, we were unable to improve on ACA's predictive ability. In fact, the weights for some people actually decreased performance at the individual level because of reversals (negative coefficients).

If, however, we solve for a single weight to be applied to all individuals' price utilities, we are able to improve ACA's performance. The aggregate weighted ACA correctly predicts an additional 0.235 % out of 12 of the pairwise preference comparisons. Although this improvement seems relatively small, its importance will be demonstrated in the next two sections.

Conclusion

All of these methods do well at predicting hold-out choices. ACA plus ratings-based conjoint appeared to be as, or slightly more, successful than ACA alone, and ACA plus choice-based conjoint appeared to be slightly less successful than ACA alone, although neither difference was significant. These results are summarized in Table 2.

Dollar Value of Utility Differences

The second criterion to test was the dollar value associated with a specified difference in utilities. In this study, the brands included in the conjoint sections were customized for each respondent individually, based on their most likely brand. Given that, we can easily compare the price difference required to equalize the utility difference between each person's most likely and second most likely brands.

If a respondent's price utilities are linear with dollars, then utilities can be expressed in a dollar metric just by scaling them so that utility differences are equal to price differences. Our price attribute had several levels, so it was not likely that any respondent's price utilities would be exactly linear with actual dollars of price. We therefore used just two middle price levels to estimate a scaling constant for each respondent, choosing the constant that made each respondent's difference in utilities for those two levels equal to the difference in those prices.

An interesting point, and seemingly a paradox, is that if a method understates the importance of price, it will overstate the dollar values of other features. Put another way, the more important price is, the smaller that dollar value of features will be. Although that may initially seem counter-intuitive, let's think about a hypothetical example with two respondents' utilities.

	Respondent A	Respondent B
Most likely brand	1.0	1.0
Second most likely brand	0.0	0.0
$1,000	1.0	2.0
$2,000	0.0	0.0

For both respondents, the difference in brands is worth one utile. For respondent A, saving $1,000 is also worth one utile, while for respondent B saving $1,000 is worth 2 utiles. According to the standard conjoint practice, we would say that price is more important to respondent B than respondent A.

We convert each respondent's utilities to a dollar metric by scaling them so that the difference in utiles for price is equal to the difference in price levels, or $1,000. As a result, respondent A places a value for most likely brand over second most likely brand of $1,000. On the other hand, respondent B places a value of only $500 dollars for the same difference, because price was more important to this respondent.

Findings

While the findings for the predictive ability (above) of each method were very similar, the findings for dollar per utile are very divergent, neatly dividing the methods into two classes. With utilities derived from unadjusted ACA, the dollar value to equalize the difference in preference between the most likely and second most likely brand is approximately $200.

With utilities derived from aggregate-weighted ACA, ACA plus ratings-based conjoint, AND ACA plus choice-based conjoint, the dollar value needed to equalize the difference is approximately $100. These findings are summarized in Table 3.

While the differences in predictive ability (as stated above) probably aren't large enough to be meaningful, this halving of "required" price is very meaningful.

Conclusion

Although I am not sure what the test of truth should be with this criterion, the fact that three methods produce the same finding gives great weight to that solution as being more accurate. The lower price is also more heavily supported by managerial opinion.

There are a couple of other very interesting conclusions to be made based on this finding. Although the dual approach including choice-based conjoint did least well at predicting individual choices, it doesn't seem to give biased answers. Rather, it seems to provide too little information (or too much noise) to allow effective estimation of individual utilities, even for a single price attribute.

<u>Shares of Preference</u>

Based on the aggregate level findings concerning dollars, it would also be interesting to predict shares of preference at the aggregate level. Each of the methods discussed above provides a unique set of utilities that can be used to calculate predicted shares. These shares can then be compared to the actual percent of time each of the holdout concepts was chosen. In this controlled situation, the first choice model is probably a better simulation model than a probabilistic model (Elrod, 1989). Using a first choice model, we can compare each of the methods in terms of total absolute deviation of the estimated shares.

As a benchmark, we can use the pre-holdouts to produce the first measure of deviation. The total absolute deviation of predicted shares using pre-holdouts is 0.45. Unadjusted ACA produces a TAD of 1.71.

ACA plus ratings-based conjoint improves on ACA with a TAD of 0.88, if only price is bridged, and 0.82 if price and brand are bridged.

ACA plus choice based conjoint also provides an improvement over ACA. The TAD using effects coding is 0.74, and solving for linear price utilities, the TAD is 0.81.

ACA plus either ratings or choice based conjoint reduce the error from ACA alone by about half.

The very interesting finding, though, is when we use the weighted ACA utilities. When we solve for a weight for each person individually, we have a TAD of 0.85, on par with either of the dual methods. However, when we use an aggregate level weight, we more than halve the TAD again. Using ACA with an aggregate level weight, the TAD is under 0.35, even better than pre-holdouts.

Findings
As with the dollar value per utility difference, the findings from this section are quite clear. Weighting ACA with a single aggregate level weight works best. ACA weighted individually, or combined with ratings-based or choice-based conjoint all do about equally well, and about half as well as the aggregate weighted ACA. ACA alone does about half as well as ACA plus ratings-based or ACA plus choice-based.

The findings from this section suggest again that ACA correctly measures the shape of the utilities, but not the height. Also, the findings suggest that getting the height correct is more important than the shape, but getting both correct provides the best model. The results for the various methods are summarized in Table 4.

Conclusions
The three criteria outlined each produced rank orders of the effectiveness of the various methods. The three criteria each produced slightly different rank orders. The meaning of these rank orders, fortunately, is fairly easy to interpret. Overall, the dual methods did better than ACA alone, but weighted ACA tended to outperform even the dual methods. However, like a dual method, the

weighted ACA does require additional information for each respondent, most likely a choice-based conjoint exercise. This approach has the benefit of utilizing ACA to measure the shape of price utilities, and choice to measure the height.

The choice-based approach could have other benefits as well. Aggregate level analysis would allow the researcher to investigate interactions involving price. The choices could also be used as a validation of the respondent's consistency of the first conjoint task.

Discussion

It strikes me that there are a number of key points to take away from these analyses.

It is not evident from this limited example if these problems are specific to ACA, or if all conjoint methods can be improved through similar adjustments.

It is not clear whether all attributes are impacted like price, or if price is a unique attribute whose importance is understated because it is only counted once. While the analysis from this data set suggests that price is unique, that finding must be confirmed on other data sets.

It is interesting that the ratings-based dual does better than the choice-based dual. Two reasons come to mind. First, choices contain less information in each observation than ratings. While choices tell you which is preferred, ratings tell you by how much it is preferred. Second, and less obvious, is that the convention of dual conjoint frequently involves the second ratings-based conjoint having three attributes: price, brand, and a bundle. We, at least, tend to create the bundled or composite attributes to have levels that are clearly monotonic. Basically, create at least three products, one relatively poor, one okay product, and one very good product. We did follow that convention for the ratings-based dual, so those respondents saw a simple performance attribute in which Performance, Warranty, Service, and Channel all varied simultaneously. However, in the choice-based dual, we elected to present more complex performance information in which these four components varied independently. In retrospect, it would have been preferable for the purposes of this comparison to use the same attribute combinations in the ratings-based and choice-based duals. The simpler presentation in the ratings-based task may have led to its slight superiority in predicting holdouts.

Summary

To summarize, researchers have found that when price is included as just one of many conjoint attributes, it tends to receive too little importance. We compared ACA alone with two Dual conjoint methods, one of which used ACA together with a separate ratings-based conjoint module to estimate price utilities, and another which used ACA together with a separate choice-based conjoint module to estimate price utilities.

All methods work quite well, and nearly equally so, in predicting individual holdout choices.

The ratings-based dual method was slightly better than ACA, and the choice-based method was slightly worse, though the differences were not significant. The choice-based dual method apparently suffered from the fact that choices provide too little information for reliable estimation of individual price utilities.

Although all methods had approximately the same accuracy of choice predictions, they made different predictions. Compared to ACA alone, the other methods attributed twice as much importance to price.

We also examined the accuracy of predictions of aggregate shares of preference for the holdout concepts. We found that both of the dual methods provided aggregate predictions with about half as much error as those of ACA alone. However, the best aggregate predictions were obtained by using utilities from ACA alone, but rescaling price utilities to have approximately twice their natural importance.

We conclude that if price is just one of many ACA attributes, special action should be taken to ensure that it receives proper importance. For greatest accuracy of individual choice predictions, we suggest a dual approach involving ratings-based conjoint. For greatest accuracy in predicting aggregate shares of preference, we recommend a choice-based dual approach. An additional possibility, and I believe the key finding from the study, is to use price utilities from ACA, but after adjusting them by a scale factor common to all respondents that best predicts individual choices.

References

Elrod, Terry and S. Krishna Kumar (1988), "Bias in the First Choice Rule For Predicting Share," in *Sawtooth Software Conference Proceedings,* Ketchum, ID: Sawtooth Software.

Green, Paul and V. Srnivasan (1990), "Conjoint Analysis in Marketing: New Developments With Implications for Research and Practice," *Journal of Marketing,* (October), 3-19.

Huber, Joel (1992), "Learning Effects in Preference Tasks: Choice-Based Versus Standard Conjoint," in *Sawtooth Software Conference Proceedings,* Ketchum, ID: Sawtooth Software.

Johnson, Richard M. (1991), "Comments on Studies Dealing With ACA Validity and Accuracy, With Suggestions For Future Research," Working Paper, Sawtooth Software.

Wittink, Dick and Richard M. Johnson (1991), "Estimating the Agreement Between Choices Among Discrete Objects and Conjoint-Ratings-Based Predictions After Correcting for Attenuation," Working Paper, Cornell University.

Wittink, Dick and John Walsh (1988), "Conjoint Analysis: Its Reliability, Validity, and Usefulness," in *Sawtooth Software Conference Proceedings,* Ketchum, ID: Sawtooth Software.

Table 1

Brand	Customized based on each respondent's most, second, third, fifth and tenth most likely brands
Microprocessor	Intel Microprocessor AMD Cyrix
Performance	Below average performance Average performance Above average performance
Monitor	Average monitor (VGA 640 X 480) Above average monitor (SVGA 1023 X 768) Far above average monitor (XVGA 1280 X 1024)
Storage	Average storage (80 MB hard drive) Above average storage (120 MB hard drive) Far above average storage (240 MB hard drive)
Model	Small desktop model Large desktop model Tower model
Channel	Order over the telephone Obtain from a retail store Obtain from a sales person at your site
Service	Ship back to manufacturer for service Service at local dealer On-site service
Warranty	90 day warranty 1 year warranty 5 year warranty
Upgrade ability	Processor can not be upgraded Can upgrade to a faster processor in the future Can upgrade to a faster processor now
Price	Customized based on expected expenditures $1500/2700 $1800/3200 $2200/3800 $2500/4300

Table 2
Predictive Ability

Predicting post hold-out choices

ACA alone:	82.87%
ACA plus ratings-based (price only):	-0.35% relative to ACA alone
ACA plus ratings-based (price and brand):	+0.72% relative to ACA alone
ACA plus choice-based (effects coded):	-0.72% relative to ACA alone
ACA plus choice-based (linear):	-0.50% relative to ACA alone
Weighted ACA:	+0.235% relative to ACA alone
Pre hold-outs predicting post hold-outs:	84.38%

Table 3

Approximate Dollar Value of Utility Differences

ACA alone:	\$200
ACA plus ratings-based:	\$100
ACA plus choice-based:	\$100
Weighted ACA:	\$100

Table 4
Total Absolute Deviation - Aggregate Shares of Preference

ACA alone:	1.71
ACA plus ratings-based:	0.88/0.82
ACA plus choice-based:	0.74/0.81
Weighted ACA:	0.35