


# Improving the Value of Conjoint Simulations

Adding variability to choice models can help shine a light on market behavior.

By Bryan Orme and Joel Huber

One of the reasons conjoint analysis has been so popular as a management decision tool has been the availability of a choice simulator. These simulators enable managers to perform "what if" questions about their market—estimating shares of choice under various assumptions about competition and their own offerings. Under the proper conditions, simulated shares of choice track closely with market shares. For example, simulators can predict the performance of a new offering, estimating both expected new sales as well as cannibalization from a company's existing brands. Alternatively, they can be used to estimate the direct and cross elasticity of price changes within a market. Simulators have also been useful as the logical guide to strategic simulations that prepare management for short- and long-term competitive responses. (See Additional Reading, pg. 20.)





## EXECUTIVE SUMMARY

Simulators developed from either conjoint or choice analysis estimate the market impact of different product/service configurations by predicting decisions for each individual and then adding up choices across simulated respondents. This article examines why these choice simulators provide a unique and robust vision of a market, particularly where strong differences exist among competitors. We further show that predictions of choice shares can be improved by adding two kinds of variability to the individual choice models.

Simulators begin with a simple choice model of individual customers. Aggregating these into expected shares for specific offerings, simulators offer a rich and robust vision of market behavior. Our goal here is to specify the properties of choice simulators that enable them to provide such a vision. After describing the components that make up a choice simulation, we define two critical properties, *differential impact* and *differential substitution*, that are important if simulators are to mimic market behavior. We then show that the correspondence to market behavior can be further improved by adding variability to the individual choice models. This tuning, through a process we call Randomized First Choice, adds variability to both partworths and to the overall profile values in a way that optimally matches simulators against experimental choice shares. We then show why this adjustment process is particularly useful in overcoming problems with aggregate choice models.

### WHAT IS A MARKET SIMULATION?

A good market simulator is like gathering a representative sample of a company's customers in one room to vote on product concepts and competitive scenarios. For each scenario, these customers are offered a choice among offerings, each of which is described by differing levels on a common set of attributes. The beauty of a simulation is that the choice models for each customer permit the analyst to estimate choice shares for a broad range of products in countless competitive contexts. The icing on the cake is that the simulated market audience never gets tired, asks for lunch breaks, or requires hourly payments.

How does a market simulator work? Suppose we have a way to predict how much customers will like different kinds of ice cream cones. Let's refer to those preferences as partworth utilities, and assume the following values for a given customer.

#### Partworth Utilities

Chocolate	0
Vanilla	30
Strawberry	40
\$0.60	50
\$0.80	25
\$1.00	0

We could use these utilities to predict the choice between a vanilla cone for \$0.80 or a strawberry cone for \$1.00 using the following additive rule:

$$\begin{aligned}\text{Vanilla (30 utilities) + \$0.80 (25 utilities)} &= 55 \text{ utilities} \\ \text{Strawberry (40 utilities) + \$1.00 (0 utilities)} &= 40 \text{ utilities}\end{aligned}$$

We'd predict that this particular customer would choose the vanilla cone. If we had data for 500 customers, we could count the number of times each of the two cones was preferred after summing their different partworth values, and compute a "share of preference," also referred to as a "share of choice."

#### Share of Choice

Vanilla @ \$0.80	300/500 = 0.60
Strawberry @ \$1.00	200/500 = 0.40

In this hypothetical market simulation, 60% of the customers preferred the vanilla and 40% the strawberry cone. This illustrates the most simple simulation approach, referred to as the first-choice model.

### WHY CONDUCT CHOICE SIMULATIONS?

Looking only at average preferences or at the partworth utilities can be misleading. It can mask important market forces caused by patterns of preference at the individual level. Marketers are generally not interested in averages, but in the targetable, idiosyncratic behavior of segments or individuals. For example, consider three respondents with the following partworths for color:

#### Partworths for Color

	Blue	Red	Yellow
Respondent A	50	40	10
Respondent B	0	65	75
Respondent C	40	30	20
Average	30	45	35

Relying only on average partworths, we would pronounce that red is the most preferred color, followed by yellow. However, if these colors were offered to each respondent, red would never be chosen under the first-choice model, yellow would be chosen once, and blue twice—the exact opposite of what aggregate utilities suggest. While this is a hypothetical example, it demonstrates that average utilities do not always tell the whole story. Many complex market effects like these can be discovered by conducting simulations. These examples help demonstrate the intuition behind the value of a simulator. The following sections articulate more precisely why simulators offer a unique and valuable image of the market dynamics.

### A MARKET STUDY

For illustrative purposes we will use a choice-based conjoint dataset of midsized televisions. We surveyed 352 respondents who

owned or were considering purchasing a mid-sized television set. The first part of the computer-based interview focused on attribute definitions (described in terms of benefits) for the six attributes included in the design. Sawtooth Software's tools helped create both the study and the analysis. The survey offered 27 choices among televisions respondents might purchase, and each choice involved five televisions described with six attributes: brand name (three levels), screen size (three levels), sound (three levels), picture-in-picture (available, not), channel blockout (available, not), and price (four levels). The choice set in Exhibit 1 illustrates the levels.

The first 18 choice tasks were combinations of attribute levels such as those shown in Exhibit 1. After completing these calibration tasks, respondents picked from an additional nine hold-out choice tasks, again including five alternatives.

### ELEMENTS OF MARKET SIMULATORS

Our study of television set preferences illustrates three building stages in a choice simulation:

1. Develop a preference model for each individual or homogeneous segment in the survey weighted to reflect the target population.
2. Estimate the relative utilities of the competitors in a choice set.
3. Apply the preference model to the competitive set to arrive at choice probabilities for each alternative and aggregate these to develop choice shares.

The first stage lets marketers predict for each person the value of a given choice alternative based on its characteristics. Historically, conjoint analysis has generated individual-level partworths, and the television data in this study are derived from hierarchical Bayes analysis of choice tasks. Exhibit 3 shows the partworths for a typical respondent. (See page 16.) The least desirable level within each attribute has been set at zero, so that other levels reflect differences from the least desirable level. For example, for this respondent JVC is the least liked brand, RCA being worth 18, and Sony 44 points more.

**EXHIBIT 2** Relative error: Mean absolute error predicting market share

	First Choice Rule	Product Variability	Attribute and Product Variability
Aggregate Logit	N/A	155%	121%
Latent Class	N/A	122%	113%
Hierarchical Bayes	116%	111%	107%

The second stage of a simulator takes these partworths and estimates utility values for potential market competitors. Generally, a simple additive model is used. For example, Exhibit 4 shows the predicted utility of three possible brands for this respondent. (See page 17.) The highest value is for the 28" Sony with surround sound for \$450, so the Sony would be the most likely alternative to be chosen.

Our primary focus is on the final stage—how to predict choice probabilities for those valuations. Although the Sony TV may be the most likely to be chosen, its probability of choice will be less than 1.0 thanks to fallible consumers. Thus, when estimating choice share we need to aggregate its unbiased probability. Our findings show that it generally will not be most accurate to assume that this person chooses the Sony simply because it has the highest utility. Paradoxically, to best predict choice shares across a group of respondents, it will be important to add variability to both the partworths given in Exhibit 3 and the total utilities for each alternative shown in Exhibit 4.

The prediction that the respondent in Exhibit 4 would choose the 28" Sony is called the first choice, or maximum utility rule. This rule is easy to understand and has a long history in conjoint analysis. Unfortunately, it often does not do well in practice. First-choice simulations can produce overly volatile share changes and demand functions. Because they lack the random factors that characterize actual purchase decisions, the

**EXHIBIT 1** Example choice set

JVC	RCA	JVC	Sony	JVC
25"	26"	25"	27"	25"
Stereo Sound	Surround Sound	Mono Sound	Surround Sound	Stereo Sound
Picture in Picture	Picture in Picture	No Picture in Picture	No Picture in Picture	No Picture in Picture
No Channel Blockout	Channel Blockout	No Channel Blockout	Channel Blockout	Channel Blockout
\$350	\$400	\$300	\$450	\$350

share estimates from first-choice simulations generate extreme probabilities, too close to zero or one

Randomized First Choice, as the name implies, begins with first choice and then adds appropriate variability. Most of the theory and mathematics behind the Randomized First Choice are well-known. Work by McFadden, Hausman, Wise, Revelt, and Train laid much of the groundwork upon which Randomized First Choice is built. (See Additional Reading.) However, to the best of our knowledge, those principles had not been synthesized into a generalized conjoint/choice market simulation model. Rather than use the partworths as measures of absolute preference, Randomized First Choice recognizes that a degree of error around these points by adding random variation to the partworths and also to the overall product utilities before computing share.

The two simple, intuitive types of variability are seen in Randomized First Choice adjustments. The first kind, attribute variability, occurs when a consumer is inconsistent in the relative weights or partworths applied to the attributes. For example, consider a TV purchaser who is price-sensitive on one retail visit but then values features on other visits. The second kind, product variability, occurs when a consumer simply chooses a different alternative on different choice occasions, typically through inconsistency in relative utilities of the alternatives. Although most simulators do not distinguish between these two forms of variability, our study found that they differ strongly in their treatment of competitor similarity. In particular, adding attribute variability preserves appropriate similarity relationships among alternatives while product variability clouds them. Additionally, product variability is needed to approximate fundamental choice inconsistency in the final choice. Thus, both kinds of variability are needed to appropriately model market behavior.

Randomized First Choice offers a general way to "tune" conjoint simulators to market behavior. Mathematically, it adds vari-

ation in the attribute values as well as variation in the final product valuation, as seen in Equation 1.

#### Equation 1

$$U_i = X_i (\beta + E_A) + E_P$$

where,

$U_i$  = Utility of product  $i$  for an individual or homogeneous segment at a moment in time

$X_i$  = Row vector of attribute scores for alternative  $i$

$\beta$  = Vector of partworths


$E_A$  = Variability added to the partworths (unique for each respondent, but held constant across alternatives)

$E_P$  = Variability added to product  $i$  (Gumbel-distributed, unique for each respondent and alternative)

The simulator then chooses the product that has the greatest utility in the set after the random components have been added. It generates each individual's choice probability by repeating this simulation numerous times per respondent with new random components each time.

Those familiar with logit will recognize that  $E_P$  is simply the error level in the logit model. The typical adjustment for scale in the logit model is mathematically equivalent to adjusting the variance of a Gumbel-distributed  $E_P$  in RFC simulations. The  $E_A$  term introduces corrections for product similarity and reflects taste variation as has been found in models by Hausman and Wise and in current work in mixed logit by Revelt and Train. (See Additional Reading.) In practice,  $E_A$  error works well in RFC simulations either Gumbel-distributed or distributed normally. The appropriate variance for both the  $E_A$  and the  $E_P$  error can be found by tuning the simulated shares to holdout choice shares or known market shares. The greater the added error variance, the less extreme the shares.

**EXHIBIT 3** Partworths for a typical respondent



It is not obvious why such a process would increase the predictive accuracy of simulators. The next section begins by detailing the desirable properties of any choice simulator. An experiment then follows that demonstrates the effectiveness of adding attribute and product variability, particularly when applied to aggregate and latent class segments, but also for individual choice models generated by hierarchical Bayes.

## TWO CRITICAL PROPERTIES OF MARKET SIMULATORS

What is so special about choice simulators that justifies making an added effort to implement them? We will examine the two properties of market simulators that enable them to track the complexity of market behavior. First, simulations display differential impact, the idea that marketing leverage is at a maximum when an alternative is near the threshold for choice within a choice set. Second, they exhibit differential substitution, through which new alternatives take disproportionate share from similar competitors. Because it may not be initially obvious why differential impact and differential substitution capture very important properties of market behavior, each of these is detailed below.

Differential impact is a central property of an effective choice simulator. With this, the impact of a marketing action depends on the extent that the alternative is near the purchase threshold. This point of maximum sensitivity occurs when the value of an alternative is close to that of the most valued alternatives in the set—when the customer is very close to choosing the company's offering. At that time, an incremental feature or benefit is most likely to win the business.

The differential impact implicit in a threshold model can best be understood by examining three cases reflecting different kinds of thresholds. First we present the linear probability model, which importantly defines the case of no threshold. Then we examine the other extreme, that of a first-choice model, which has the most extreme step-like threshold. Finally we consider the s-shaped model, reflected by choice models such as logit, probit, and Randomized First Choice, whose thresholds have been softened by the addition of variability.

If probability is a linear function of utility, then improving an attribute has the same effect on choice share regardless of how well it is liked. This linear probability model has many problems, the worst of which is a lack of differential impact. As shown in the first panel of Exhibit 5, adding, say, surround sound has the same share impact regardless of whether it is added to a high- or low-

end television. (See page 19.) By contrast, a threshold choice model specifies that the benefit from adding the feature mainly affects those consumers who are likely to change their behavior. The threshold model makes good sense—adding the feature does not affect a person who would have bought the brand anyway, nor does it influence customers who would never consider it. The differential impact brought about by a threshold model focuses managerial attention on the critical marginal customer and thereby avoids expensive actions that are less likely to alter market behavior.

The first-choice model is an extreme version of the threshold model. The first-choice model is mathematically equivalent to Equation 1 with no variability ( $\text{var}(E_p) = \text{var}(E_A) = 0$ ). As shown in the second panel of Exhibit 5, in the first-choice simulation the share of an alternative is zero until its utility surpasses those of others in the set. Once it passes that threshold, it receives 100%. The problem with the first-choice model is that it is patently false. We know that people do not make choices without variability. In studies of experimental choices, respondents given the same choice set (three-four alternatives, four-five attributes) choose a different alternative about 20% of the time. In our current study, with five alternatives and six attributes, respondents changed their choices from the same set 19% of the time.

Standard logit and probit models reflect a compromise between the first-choice and the linear models. Instead of the severe step function characterizing the first-choice model, the variability implicit in these models moderates the step into a smooth s-shaped function. They are equivalent to the model in Equation 1 with product variability added.

A little-understood benefit of a threshold model is that it can reflect complex patterns of interactions between, say, a feature and a particular brand simply through the simulation process. An interaction term specifies that a particular feature has a differential impact on particular brands. These interaction terms can be expressed in the utility function whereby the partworth of, say, price is different when Sony rather than JVC is the brand. While these interaction terms can and have been modeled as cross terms in the utility function, the same behavior can be represented by an aggregation of heterogeneous customers each following a threshold model. For example, consider a warranty  $\times$  price interaction. Such an interaction would arise if a warranty were more valuable when applied to low- over high-priced appliances. Suppose there are two segments, one valuing low price

**EXHIBIT 4** Utilities for attributes of three potential products

	Brand	Size	Sound	Blockout	Pict in Pict	Price	Total utility
<b>26" Sony, Surround Sound</b> \$450	44	46	104	0	0	0	<b>194</b>
<b>26" JVC, Stereo</b> \$300	0	0	82	0	0	98	<b>180</b>
<b>27" RCA, Block, P in P</b> \$400	18	25	0	42	60	36	<b>181</b>



and the other desiring high quality. Suppose further that those who value low price also value warranties, but those paying a high price feel there is less need for a warranty. In a simulation, adding a warranty to the high-priced brand will not cause the people to shift because it is unimportant to them. By contrast, adding it to the low-priced brand will shift those desiring low price. When these two segments are aggregated it may appear that the warranty affects the low-priced brand, whereas it really affects the segment valuing low prices. In that way heterogeneity arising out of a simulation can account for what would have otherwise been seen as an interaction between variables.

The heterogeneity account is also a more stable model than the curve-fitting exercise of the interaction cross terms. The operational problem with adding interaction terms to the utility function is that the numbers of such terms can grow uncontrollably large. Above we illustrated an example of price tiers, but many others are also possible. Consider all-too-common combinations of brand tiers where customers are simply not interested in certain brands, size tiers where a large size never passes the threshold for certain segments, and feature tiers, where certain groups are only interested in certain features. Modeling these with interaction terms in the utility function is complicated and can lead to overfitting or misspecification. The beauty of a simulator operating on individual threshold models is that it can approximate what appears to be a very complex pattern of interactive behavior solely through the aggregation of choices from simple main-effects individual models. (See Additional Reading.)

To summarize, differential impact is critical if we believe that impact on choice of, say, a new feature of a brand depends on utilities of competing brands. Brands close to the threshold will be far more sensitive than those that are far away. The threshold model within a simulator focuses managerial attention on those alternatives that are on the cusp, and in that way avoids putting managerial effort on alternatives that are already chosen, or would never be.

Differential substitution is a second benefit of an effective choice simulator. Its intuition follows from the idea that a new offering takes share disproportionately from similar ones. Differential substitution is particularly important because the dominant choice model, aggregate logit, displays no differential substitution. The logit assumption of proportionality implies that a new offering that gets, say, 20% of a market will take from each competitor in proportion to its initial share. Thus a brand with an initial 40% share loses eight percentage points ( $40\% \times 2$ ), and one with 10% share loses two percentage points ( $10\% \times 2$ ). Proportionality provides a naive estimate of substitution effects that can result in managerially distorted projections for the common case in markets where groups of brands are very similar and others are very different. For example, a product line extension can be expected to take proportionately most share from its sibling brands. Managers recognize this problem. Successful companies manage their portfolios with new brands that are strategically designed to maximize share taken from competitors while minimizing internal share losses. By contrast, proportionality

**Brands close to the threshold will be far more sensitive than those that are far away.**

glosses over these strategically important distinctions. Ignoring differential substitution could lead to the managerial nightmare of numerous line extensions whose cost to current brands is regularly underestimated.

An extreme, if instructive, example of differential substitution is the presence of a duplicate offering in the choice set. Economic theory often posits that a duplicate offering should take half the share of its twin, but none from its competitor. However, in practice this expectation is rarely met. If some consumers randomly pick a brand without deleting duplicates, then having a duplicate could increase total choice share. Indeed, the fight for shelf space is directed at capturing that random choice in the marketplace. To the extent that a duplicate brand increases the total share for that brand, we term the increase in total share from a duplicate share inflation. Clearly some share inflation is needed, but it is unclear how much. The  $E_A$  term in the Randomized First Choice model lets the analyst adjust the degree of penalty for product similarities that should operate in the model. Product alternatives defined by many of the same features will have correlated sums of  $E_A$  variability, which in turn leads these similar products to compete more closely within market simulations.

To summarize, in a typical market where strong differences exist among competitors, a simulator is needed to display expected differential substitution effects. The choices coming out of a simulator behave appropriately because similar alternatives have similar utility values. Because only one is chosen in each simulated choice, these alternatives take share from one another just as managers expect.

## TESTING THE VALUE OF RANDOMIZED FIRST CHOICE

In this example, we compared the ability of different simulators to predict nine holdout choice tasks. The holdout tasks were different from 18 calibration tasks in two respects. To test the share predictions of the different simulators, it was critical to have target sets for which shares of choice could be estimated. Respondents were randomly divided into four groups with approximately 90 in each group that would receive the same nine holdout choice tasks. Additionally, we designed the holdout choices to have some extremely similar alternatives. Four of the five alternatives in the holdout tasks were designed to be efficient choice sets that have statistically optimal utility and level balance. (See Additional Reading.) However, the fifth alternative duplicated another alternative in the set or duplicated all attributes except the two attributes judged least important. Normally such redundant alternatives would not be included because they add little to the discriminatory power of the choice set. However, we wanted very similar items to test the simulators' ability to predict share on highly similar alternatives, such as those populating most product lines.

We analyzed the choice data using three base methods for estimating respondent partworth utilities: aggregate logit, latent class, and hierarchical Bayes. The logic behind picking these three methods related to how much they accommodate individual differences. Aggregate logit is important because it reflects what happens when all respondents are analyzed as one homo-

geneous population. By contrast, latent class analysis seeks sets of latent segments (we used an eight-group solution) whose partworths best reflect the heterogeneity underlying the choices. (See Additional Reading.) Focusing even more on individual differences, hierarchical Bayes estimates different utilities for each respondent. It produces a posterior estimate of each individual's partworths that reflects a heterogeneous population prior conditioned on the particular choices each individual makes. (See Additional Reading.) Thus, hierarchical Bayes generates each individual's utility function, while latent class and aggregate logit typify popular ways to deal with markets as groups.

For each of these base models we examined the impact of three levels of variability within the Randomized First Choice framework. The first condition uses the first-choice rule that assumes respondents perfectly choose the highest valued alternative. It thus assumes no variability. The second condition adds the level of attribute variability that best predicts holdout choice shares. The third condition tunes both product and attribute variability to best predict the holdout choice shares. The mechanism of the tuning process is straightforward but tedious: we use a grid search of different levels of each type of variability until we find those that minimize the mean absolute error in predicting holdout choice shares.

To test the predictive validity of different models, we divided the 352 respondents into two matched samples of 176 respondents each. We again estimated the partworths under the different models independently within each of the two groups. We performed the grid search to optimize the variability terms that best predict the holdout choices for the first group, and used those same parameters for the second group to predict its holdout choices. We reported only the predictive accuracy for the second group.

## Results

Relative error measures the degree to which the different simulators predict the choice shares across all alternatives in the holdout tasks for the study. Exhibit 2 shows mean absolute error (MAE) predicting holdout stimuli as a percent of the test-retest MAE for repeated choice sets (3.5%). (See page 15.) For example, the 155% for aggregate logit indicates that adding product variability results in an error that is a little more than one and one-half times as great as 3.5% for the choice replication.

Exhibit 2 offers three important results. First, adding variability improves holdout share predictions. Regardless of the estimation technique, the error drops with additional variability. The second result refers to the performance of the different models. Comparing the rows shows that the more individualized methods do better than the more aggregate ones. The third result is that there is greater gain from adding variability to an aggregate level model than to the more individual-level models. Thus, aggregate logit gains the most, latent class is next, and hierarchical Bayes shows the smallest gains from adding variability.

We believe there is a simple reason for this last result. The individual-level models are not estimated with perfect accuracy, but have significant variation due to the noise in individual choices. Thus, when estimates from these models are put in a simulator they act as if variability has already been added to the partworths. However, in this case instead of coming from the Randomized First Choice process, attribute variability comes from the inherent variability in the estimation model.


## SUMMARY AND CONCLUSIONS

Our aim has been to demonstrate the usefulness of market simulations and to examine ways to build choice simulators that correctly reflect similarity effects. We began by showing how differential impact and differential enhancement help simulators model changes in choice in realistic and robust ways. Randomized First Choice is developed in the light of these principles by adjusting the simulator to achieve optimal levels of attribute and product variability.

This article has presented five key ideas about simulations:

1. Market simulations transform raw utility data into a managerially useful and appealing model: that of predicting market choice (share of preference) for different products. Under the proper conditions, shares of preference quite closely track with the idea of market share—something almost every marketer cares about.
2. Market simulations can capture idiosyncratic preferences occurring at the individual or group level. These “hidden” effects can have a significant impact on choice shares for products in market scenarios. Capturing such effects is

## EXHIBIT 5 Examples of linear, step, and S-shaped choice functions





important for accurately predicting choice share, especially when multiple product offerings are designed to appeal to unique segments of the market

- 3 Market simulations can reveal differential substitutability (cannibalism/cross-elasticity effects) among different brands or product features. If two similar brands are valued highly by the same respondents (have correlated preferences), these brands will tend to compete more closely. Product enhancements by one of these brands will result in the correlated brand losing more relative share than other less similar brands within the same simulation. Examining aggregate utilities can obscure these important relationships.
- 4 Market simulations can capture interaction effects between attributes. If the same respondents that strongly prefer the premium brand are also less price-sensitive than those who are more likely to gravitate toward a discount brand, simulations will reflect a lower price elasticity for the premium relative to the discount brand. A similar interaction effect can occur between many other types of attributes, such as model style and color.
- 5 Randomized First Choice provides a way to tune simulators by adding the appropriate level of attribute and product variability. Here it was used to tune simulated output to choice shares. However, it can be used more generally to adjust conjoint simulator results to known market shares. |

#### ADDITIONAL READING

Arora, Neeraj, Greg Allenby, and James L. Ginter (1998), "A Hierarchical Bayes Model of Primary and Secondary Demand," *Marketing Science*, 17, 29-44.

Green, Paul E. and Abba M. Krieger (1988), "Choice Rules and Sensitivity Analysis in Conjoint Simulators," *Journal of the Academy of Marketing Science*, 16 1, 114-127.

Green, Paul E. and V. Srinivasan (1990) "Conjoint Analysis in Marketing: New Developments with Implications for Research and Practice," *Journal of Marketing*, 54, 3-19.

Hausman and Wise (1978), "A Conditional Probit Model for Quantitative Choice: Discrete Decision Recognizing Interdependence and Heterogeneous Preferences," *Econometrica*, 43, 402-426.

Huber, Joel and Klaus Zwerina (1996), "The Importance of Utility Balance in Efficient Choice Designs," *Journal of Marketing Research*, 23, 307-317.

Kamakura, Wagner A. and Gary J. Russell (1989), "A Probabilistic Choice Model for Market Segmentation and Elasticity Structure," *Journal of Marketing Research*, 26, 339-390.

McFadden, D. (1973), "Conditional Logit Analysis of Qualitative Choice Behavior," in *Frontiers in Econometrics*, P. Zarembka ed. New York: Academic Press.


Orme, Bryan K. and Michael Heft (1999), "Predicting Actual Sales with CBC: How Capturing Heterogeneity Improves Results," Sawtooth Software Conference Proceedings.

Revelt, David and Kenneth Train (1998), "Mixed Logit with Repeated Choices: Household's Choices of Appliance Efficiency Level," *Review of Economics and Statistics*, (Forthcoming).

---

**Bryan Orme** is vice president of Sawtooth Software Inc. He plays an active role in new product development at Sawtooth Software and has published many articles on conjoint analysis and discrete choice techniques.

**Joel Huber** is professor at the Fuqua School of Business, Duke University. He has written numerous articles on the application of conjoint analysis to marketing problems.



# YOUR AD HERE

## Advertise in Marketing Research!

---

Contact **Richard Ballschmiede** at  
**1-800-AMA-1150 ext. 9076** for more information.