Adjusting Choice Models to Better Predict Market Behavior

GREG ALLENBY *   allenby_1@cob.osu.edu
Ohio State University

GERALDINE FENNELL*
Consultant, Dublin, Ireland

JOEL HUBER*
Duke University

THOMAS EAGLE
Eagle Analytics

TIM GILBRIDE
Notre Dame University

DAN HORSKY
University of Rochester

JAEHWAN KIM
Korea University

PETER LENK
University of Michigan

RICH JOHNSON
Sawtooth Software

ELIE OFEK
Harvard University

BRYAN ORME
Sawtooth Software

THOMAS OTTER
Ohio State University

JOAN WALKER
Boston University

Abstract

The emergence of Bayesian methodology has facilitated respondent-level conjoint models, and deriving utilities
from choice experiments has become very popular among those modeling product line decisions or new product

* Co-chairs. Author order is alphabetical.

introductions. This review begins with a paradox of why experimental choices should mirror market behavior despite clear differences in content, structure and motivation. It then addresses ways to design the choice tasks so that they are more likely to reflect market choices. Finally, it examines ways to model the results of the choice experiments to better mirror both underlying decision processes and potential market choices.

Keywords: Bayesian analysis, extended model of behavior, motivating conditions

Adjusting Choice Models to Better Predict Market Behavior

The last five years have seen a revolution in experimental choice models. While choice experiments have been around since Louviere and Woodworth (1983), their use for product line additions and optimization has recently been greatly expanded. These choice experiments have different names, such as "choice-based conjoint" or "stated preference models," but they all derive preference functions from successive choices among options described along defined product attributes.

The revolution occurred when the availability of hierarchical Bayes and similar methodologies permitted the estimation of reliable respondent-level utility models from 10–20 hypothetical choices. Before that development, heterogeneity had to be represented in the parameters of an aggregate choice model, typically represented as interactions between the parameters reflecting the alternatives and the characteristics of the respondents. Specifying these interactions could result in unmanageably large models. For example, crossing 20 parameters against 20 customer characteristics produces 400 cross terms, which themselves might fall prey to overfitting. With the new approach, one very large model is replaced by a large number of simpler main effects models, one for each respondent.

The availability of reliable respondent-level models greatly improves the ability of choice models to inform product formulation, product line and pricing decisions. In particular, segmentation, targeting, positioning and exploring the impact of alternative scenarios on a brand's market share are generally easier to implement and more accurate with respondent-level parameters.

In this paper, we review what we know about choice experiments and suggest issues that would benefit from additional research. Our review and analysis is not meant to be exhaustive, but instead reflects the diverse views of the practitioner and academic members of our task force at the Colorado Choice Symposium. Our first theme is to point to differences between choice tasks and market behavior and ask how the choice task should be structured to bring the two together. Our second theme is that, while new ways to handle heterogeneity have moved us from complex to simple model specifications, behavioral theories behind a respondent's encoding, evaluating and decoding stimuli are rich in detail that may not be reflected by the simple linear models. Human behavior is complex, both within choice tasks and in the marketplace, offering opportunities for modifying tasks and research designs to mimic that behavior and for developing the statistical models to represent it.

How Choice Experiments are Different from Market Choice

From the perspective of collecting market data, experimental choices have a number of attractive properties, particularly when compared with alternative tasks. First, it is choices that occur in the marketplace, not ratings or rankings. Second, because respondents make choices every day, they appear ready to make hypothetical choices about almost anything. Finally, assuming that the repeated choices in a choice experiment correspond to those in the marketplace, choices can be directly related to market share. This correspondence provides both theoretical and intuitive justification for the use of choice experiments.

Despite overt similarities, the repeated choice task is different in fundamental ways from marketplace choice (DeSarbo and Green, 1984). These differences define a challenge that provides the focus of this paper: how to adjust the choice task and its analysis to best predict marketplace behavior. Four principal differences are listed below.

1. *Information in choice experiments is clearer and more complete.* Compared with marketplace behavior, information about choices is presented in ways that facilitate a rational choice among the alternatives. In conjoint studies, characteristics are arrayed in a grid cleanly displaying comparable information about each option. By contrast, in the marketplace, attribute information about options may be hard to find and often is not comparable across alternatives.

2. *The choice set is determined by the experimenter rather than the consumer.* By providing a researcher-determined choice set, a conjoint exercise fails to reproduce conditions that prospects experience that predispose them to spend their resources in the product category under study. Such conditions can be expected to determine the attributes that they value in the product category, and to guide their choice among brand offerings. Thus, presenting choice sets designed by the researcher fails to conform to the process by which people select options to consider and then choose from that set.

3. *Differential motivation to process that information.* In contrast to a marketplace choice, conjoint respondents may have little motivation to make the best choice. This lack of motivation can lead to simplification of the choice task. Opposing this motivational difference, the clarity of the information provided and its relative simplicity in the experimental choices may lead to more careful information processing in the hypothetical choice. Worse, there may be learning effects and fatigue effects peculiar to experimental settings. In all, we cannot say that choices are more or less error-prone in experimental versus marketplace contexts. However, it is reasonable to conclude that the *ability* to be accurate is greater in conjoint, whereas the *motivation* to be so is greater in the marketplace.

4. *The experimental choice context is more abstract.* Market choice is anchored in a particular context, taking place in a particular store, at a specific time, carrying a wealth of feelings and beliefs associated with particular uses. By necessity, choice experiments can only ask respondents to recall or imagine such contexts of purchase and use.

Simply put, experimental choices do not look much like those in the marketplace. They differ with respect to the options available, the type and accessibility of information presented on those options, and the ability and motivation to make a good choice. Given these palpable differences, the appropriate questions involve why it is reasonable to project from an experimental scenario to the marketplace and how we can improve these projections. Before answering these questions, it is useful to summarize what we know about the behavior of respondents in repeated choice experiments.

- Time for a given choice initially takes an average of 40 seconds, but soon asymptotes to around 12 seconds per choice, providing evidence of both learning and substantial simplification (Johnson and Orme, 1996).
- The relative importance of price information compared with brand name information doubles as one makes successive choices, suggesting that focal attention shifts over time (Johnson and Orme, 1996).
- The correspondence between choices and holdout choices increases with progress through the choice tasks (Huber et al., 1993). More simply, later choices are more reliable than earlier choices.
- Successive experimental choices are relatively context-free. That is, they are largely independent of external reference points and are unaffected by the patterns of dominance or compromise within the set (Huber, 2004). Accordingly, they are unlikely to directly apply to marketplace choices, which are domain and context dependent.

Thus, there is evidence of an increase in efficiency and a change in focus as respondents progress through successive experimental choice tasks. Respondents are learning both what their preferences are for this set of options, and how to express them more quickly and reliably. Clearly, making judgments faster and more reliably entails simplification. This simplification can take two forms (Huber et al., 2002). First, there is *attribute truncation,* in which the information from a number of less important attributes is simply ignored. Second, there is *level focus*, whereby attention is focused on particular levels within attributes. A common form of level focus involves the tendency to screen out alternatives with undesirable attribute levels.

The shift in importance from brand names to prices demonstrates that this learning is not just truncation to simplify the task, but can involve justifiable refocusing. In such a process, respondents initially choose familiar brands, but as they learn that desired characteristics and expected prices are not reliably associated with those brands, their brand loyalty decreases while their price sensitivity increases. Recently, models have been proposed that accommodate unobserved changes in part-worths for individual respondents (Otter et al., 2004).

Further, we know that people simplify choices in the marketplace. Most people in supermarkets take only a small amount of time, and many times are not even aware of competitive offerings or even the price paid (Dickson and Sawyer, 1990). The appropriate question, then, is not whether choice experiments are overtly similar to actual choices; they are clearly not. The question is whether the simplification in conjoint mirrors the simplification in the marketplace. Are the attributes ignored in conjoint the same ones? Is the loss aversion

displayed in conjoint also revealed in the marketplace? Is the sensitivity to attributes comparable?

The foregoing goal of matching the simplification of the experiment with that of the marketplace would seem to argue for choice experiments that mimic it as much as possible. Indeed, developments in computer imaging have enabled the use of virtually designed images of automobiles instead of descriptions, and the use of virtual reality in the selection of beer from a cooler. Taken to its logical extreme, a market-mimicking strategy would lead to virtual stores where the market impact can be experimentally tested. However, we believe that quite apart from its cost in time and complexity, there are good reasons to avoid such simplistic mimicking of what happens in the marketplace.

As discussed above, the attributes and the alternatives in conjoint studies are clearly displayed. Thus, the result of conjoint is an estimate of the impact on choice *given that the attributes and competitive set are known to both the respondent and the researcher.* By contrast, the short-term impact of a change in the regular price, package size or featured ingredient for a product in a store environment can be minimal. People are most likely to purchase what they purchased last, simply because they do not notice the change. Thus, actual markets often experience a strong time lag before changes in offerings are realized into changes in sales. Conjoint can be viewed as a way to predict how people are likely to react once new information about brands and competitors becomes known. Seen that way, there is an advantage to the clear layouts of conjoint choices and the fact that the choice set is unambiguously determined. It results in a model of choice given that the consumer can access a set of well defined competitive conditions (Louviere et al., 2003).

Improving the Choice Experiment Task

We make three proposals about the structure of the choice task. First, attribute ranges and levels should be set at approximately the level of the market to be simulated; second, the task should deliberately be made simpler, and finally, the correlational structure should be more orthogonal than is found in the market. Each of these recommendations is detailed below.

Ranges and Levels Should Match Those of the Market Being Simulated

The range of levels presented in each attribute has a large impact on generated utilities, a fact that is well known, but little discussed. Because of this susceptibility of the outcome to range, we recommend setting the range of each attribute to equal or slightly exceed (Green and Srinivasan, 1978) those in the simulated market. This strategy establishes appropriate design parameters for choice experiments and avoids indeterminacy from range effects.

A choice study should also mimic the average levels of the attributes found in the market to be simulated. For example, if one is interested in knowing the impact of a general rise in an attribute level, such as the impact of a new feature or the impact of all prices rising, then these differences should be present in the choice experiment. While we believe that the distribution of attribute levels should match the relevant market as much as possible,

this matching strategy is less effective with respect to two other characteristics of the choice task, information load and attribute correlations. As we next propose, these often should not correspond to the market being simulated.

Limit the Complexity of the Choice Task

Keeping the choice task simple is consistent with the idea that choice experiments reflect what people would choose if given clearly displayed attribute levels that can be reliably processed. Generally, one should be careful about giving people more than 20 pieces of new information in a choice set. This implies about three alternatives if there are six attributes, or 20 alternatives if there is only price differing. Given that conjoint choices typically take only 12 seconds, it would be difficult for respondents to make reasonable decisions with more than 20 bits of information. The resulting task with 20 pieces of information still is complex enough to reveal both simplification and a focus on the important attribute levels.

One way to make the choices manageable while collecting information on larger numbers of attributes is to ask respondents to make choices on partial rather than full profiles. For example, a partial profile study would get choices across eight attributes by displaying a subgroup of three at each choice occasion. Research has shown that these partial profiles can result in better predictions to holdout choices than full profiles (Chrzan and Patterson, 1999).

Be Cautious About Correlated Attributes

The second area where mimicking may be problematic is with respect to the correlations across attributes. In efficient markets, important attributes often are negatively correlated, meaning that an improvement in one attribute is associated with undesirable levels of other attributes. In conjoint, such correlations can be simulated by conditional coding of attributes in which, for example, the more durable alternatives get higher prices. The problem arises when respondents simplify their choice by focusing on one correlated attribute while ignoring the other. In that case, the estimate of the omitted attribute will be biased. These distorted coefficients are less likely to occur with orthogonal designs because the coefficients are statistically independent of each other. For example, if price is noticed and durability is ignored, price will get the appropriate negative coefficient and durability will be non-significant, mirroring the underlying process.

Notice that the "distorted" results from conditional designs may appropriately mirror the short-run behavior of a market. In the short-run, backwards coefficients are common—an increase in price with no improvement in quality generally results in greater dollar sales. Still, while markets may be slow to respond, they are not irrational. As consumers gradually learn that one offering is dominated by another in the market, its sales will appropriately decrease. The important conclusion here is that orthogonal designs do a better job of representing prospects' behavior that has evolved over time. Thus, if the goal of conjoint is to reflect immediate responses by consumers, as in a shelf/price study, then the correlations should remain. However, if the goal is to predict prospects' reactions given time to understand and adjust, then an orthogonal design is generally preferred.

Choice Experimental Design

There are three primary sources for experimental choice designs. Designs for attributes with two levels have been developed by Street and Burgess (2004). More generally, Warren Kuhfeld (2004) at SAS has developed a series of programs that generate designs using a combination of tabled designs and search routines. These can accommodate designs of almost any dimensionality and produce utility-balanced designs (Huber and Zwerina, 1996). At the other end of the spectrum, Sawtooth Software has programs that generate random designs. These designs are not completely random, but are designed to draw from each level and each pair of levels in an approximately balanced fashion. They are nearly as efficient (e.g., 95%) as optimal designs for large designs (e.g., 100 choice sets), and have the advantage of not depending on a particular functional form.

A few researchers have attempted to build optimal choice designs (Johnson et al., 2003; Tobia et al., 2004) for individual respondents. Several forms are possible: some ask general questions about the preference order of levels within attributes, while others use the information from past choices to determine the next ones. While these methods have worked in simulations, their performance with real people has been disappointing. This is clearly an area where both theoretical and practical research need to converge to develop truly adaptive and efficient choice designs.

Improving Choice Analysis

The previous section has considered ways the task can be changed to increase the validity of choice experiments. In this section we consider ways in which improvements in predicting marketplace behavior can be achieved by improving the models used to analyze choice data. We discuss four approaches to improving the model: (i) using a model that better reflects actual decision making; (ii) parameterizing the model to improve the summary of important information; (iii) judiciously employing error terms to improve the precision of the estimates; and (iv) model fusion. We discuss each recommendation in detail below.

Using Models that Reflect Actual Behavior

The availability of hierarchical Bayes and similar methodologies has facilitated the estimation of choice models for each respondent. Researchers have found that, once heterogeneity is taken into account, a simple linear model without interaction terms is often adequate. For example, early research using weekly scanner data often contained multiple price coefficients (e.g., regular price, deal price, lagged price, reference price, etc.) that were found to be unnecessary when analyzing household-level data (Bell and Lattin, 2000).

However, there are many applications where a simple linear model fails to reflect the underlying behavioral process and leads to poor predictions. This occurs, for example, when consumers purchase multiple varieties of an offering (e.g., multiple flavors) because of the effects of satiation (Kim et al., 2002). Multiple purchases can be viewed as interior solutions in an economic choice model, solutions that are incompatible with the corner

solutions of a linear utility specification. A non-linear specification is required to avoid overprediction of volume from temporary price reductions in the face of variety seeking behavior and satiation (Kim et al., 2004).

Linear compensatory choice models also do not address simplifying choice heuristics such as attribute truncation and level focus. Real choice tasks and experimental choice tasks are characterized by consumers who focus on a subset of product attributes (Gilbride et al., 2004) or use the levels of certain attributes to eliminate alternatives from the choice set (Gilbride and Allenby, 2004). The presence of either simplifying heuristic results in an abrupt change to the choice probability that is difficult to estimate because the resultant likelihood surface is not differentiable. Bayesian methods, including data augmentation, offer a viable method for dealing with the irregularity of the likelihood surface, avoiding the need to rely on standard linear models that assume poor performance on one attribute can be compensated by strong performance on another. Models that incorporate variable selection and screening rules result in improved predictions for consumers who place zero value on particular attributes and choice alternatives, conditions that frequently exist in the marketplace.

At a more fundamental level, choice models currently lack variables that reflect an extended model of behavior that begins with the domain-specific motivating conditions prospects face in the context of their lives, for which they seek product attributes and bene-fits, generate brands to consider, and select brands to buy and use (see Fennell, 1988, 1997). In general, such extended models of behavior can point to independent variables capable of improving marketplace predictions of choice models.

Parameterizing Models to Estimate Key Information

Researchers make marketplace predictions with specific policy implications in mind. These implications often involve a non-linear function of a subset of the model parameters. For example, conjoint studies are often conducted to help determine the price prospects will pay for adding a feature to an offering. The standard parameterization estimates the amount consumers are willing to pay for a change in an attribute-level as the ratio of two coefficients—the part-worth of the attribute-level change divided by the price coefficient. The use of a diffuse prior for the price coefficient with support near zero implies a willingness-to-pay distribution with undefined moments. In connection with weakly informative data, such a prior typically yields posterior estimates of willingness-to-pay that lack face validity.

An alternative parameterization of the likelihood function, which directly estimates policy variables such as willingness-to-pay (see Train, 2003), can provide a solution to the problem of an unbounded posterior distribution. While a re-parameterization of the likelihood does not affect the information contained in the likelihood function (Zehna, 1966), it does affect Bayesian analyses where prior distributions are introduced, particularly when there is limited information per unit of analysis, a condition present in nearly all marketing analyses. Sonnier et al. (2004) discuss the conditions that favor either parameterization in terms of fit to the data and make the point that estimating willingness-to-pay directly is always preferable if willingness-to-pay is of ultimate interest to the user. The general idea of choosing a

parameterization that allows for direct estimation of focal parameters is likely to apply in many contexts.

While the estimation of model parameters at the respondent level can have clear advantages when it comes to reflecting actual behavior, an aggregation challenge may arise when the researcher seeks to present summary measures for sensitivity analyses regarding the reaction of the market as a whole (for example, the change in profit expected to the firm when a product attribute is improved and the price is re-set optimally). Though the choice simulator can be used repeatedly in many cases to yield numerical results, this procedure can be cumbersome and fail to provide insight into which estimated parameters are driving the final results. Ofek and Srinivasan (2002) offer a closed form solution to answer "what if" type questions that clarifies the mechanisms whereby changes in offerings translate into changes in share or profitability.

Judicious Use of Error Terms

Marketplace behavior often involves choice among many alternative product offerings. Choice alternatives can number in the hundreds for both durables (e.g., automobiles) and non-durables (e.g., grocery items). In these instances, model error terms, if associated with each unique alternative, will have the effect of producing predicted choice probabilities that insufficiently reflect differential substitution among these alternatives. One approach to improving marketplace predictions is to reduce the number of independent error terms in a model. This approach makes sense in cases where the unobservable components of multiple offerings that give rise to the error term come from the same latent factors. For example, packaged goods offerings are often available in multiple sizes. If the stochastic portion of marginal utility comes from unobserved factors that affect brand selection, then one approach to reducing the number of error terms in a model would be to assign the same error realization to all offerings containing the same brand (see Allenby et al., 2004).

An alternative approach is to substitute (*post hoc*) the kind of error used during the simulation phase using an approach termed Randomized First Choice. Rather than using the draws for each respondent from the posterior distribution of the parameters, one can avoid using that large data file by collapsing to the average of beta for each respondent. Then, one can add IID normal error to each part worth and simulate choice for each respondent hundreds or thousands of times, following the maximum utility rule. Empirical research suggests the accuracy of the simulations under Randomized First Choice is better than using the full information contained in the draws (Orme and Baker, 2000).

Model Fusion

Information for improving marketplace predictions is often available in multiple datasets. For example, both stated and revealed preferences provide information about the utility of an offering. One approach to combining or integrating this information is to include data from one source as a covariate in a model of the other (see Horsky et al., 2004a, 2004b). Alternatively, information across datasets can be combined by forming a joint likelihood

function with common parameters. Such approaches have served as a discussion topic in a previous Choice Symposium (Ben-Akiva et al., 1994), and have been shown to lead to improvements in the algebraic signs of the coefficients, as well as the precision of the estimates. Moreover, such a modeling framework facilitates tests of the stability of preferences (β) and can aid in correcting specification errors. Stated preference data may require corrections for various response biases, while revealed preference data often require information that controls for contextual effects, e.g., whether a purchase is for personal or business use.

Generically, let the marketplace model be denoted $Y_{MP} = f(x_{MP}, z | \alpha_{MP})$. where Y_{MP} are marketplace choices, f is a known functional form, x_{MP} are variables that are specific to the marketplace, z are other variables, and α_{MP} are the model's parameters. Similarly, let the model for the laboratory be $Y_{LAB} = g(x_{LAB}, z | \beta_{LAB})$, where Y_{LAB} are choices in the laboratory, g is a known functional form, x_{LAB} are variables that are specific to the experimental setting, and β_{LAB} are the model's parameters. The z variables, for example price or package size, are common in both equations. The functional forms for the laboratory and marketplace may be different, but should be motivated by the same behavioral model, such as random utility theory. Currently, respondent-level models are independently estimated in both marketplace and laboratory data. Meta-analysis studies, such as Renken et al. (2004) relate aggregate estimates from the marketplace α_{MP} and laboratory β_{LAB} to each other.

To calibrate choice models to the marketplace, it is necessary to collect choice-based conjoint from people who are also part of a purchase panel. These respondents have underlying parameters or latent variables ξ, such as price sensitivity, that drive the parameters in the marketplace and laboratory: $\alpha_{MP} = h_{MP}(\xi | \theta_{MP})$ and $\beta_{LAB} = h_{LAB}(\xi | \theta_{LAB})$, where h_{MP} and h_{LAB} are functions of ξ that depend on unknown parameters. The functions, h_{MP} and h_{LAB} may or may not have the same functional form. The objective of the calibration study would be to uncover the relationships between the choice parameters from the marketplace and laboratory and the common parameters.

The calibration could be extended to a number of different product categories and time periods to search for consistent relationships between ξ, α_{MP} and β_{LAB}. A methodology for performing this multi-category analysis would be hierarchical Bayes, where the category-level parameters are assumed to arise from a distribution of heterogeneity. This model could then be used to predict the relationship for categories not in the analysis (see Lenk and Rao, 1990 for an example of this method in new product adoption, and Lenk 1992 for coupon redemptions).

Concluding Remarks

Bayesian methods have been successfully applied in marketing for over ten years, yielding better predictions of marketplace behavior, with larger and more complex choice models than previously possible. It is now common, for example, for practitioners to estimate random-effect choice models with 50 or more dimensions. In addition, innovations in data collection that better correspond to the differences between laboratory and marketplace environments, and choice models that are more reflective of a respondent's cognitive process, are continuously being developed and reported in the marketing literature. This trend will

likely continue in the future, with the prediction of marketplace behavior offering a rich domain for theoretical and applied research.

References

Allenby, Greg M., Thomas S. Shively, Sha Yang, and Mark J. Garratt. (2004). "A Choice Model for Packaged Goods: Dealing with Discrete Quantities and Quantity Discounts," *Marketing Science* 23, 95–108.

Bell, David R. and James M. Lattin. (2000). "Looking for Loss Aversion in Scanner Panel Data: The Confounding Effect of Price Response Heterogeneity," *Marketing Science* 19, 185–200.

Ben-Akiva, Moshe, Mark Bradley, Taka Morikawa, Julian Benjamin, Thomas Novak, Harmen Oppewal, and Vithala Rao. (1994). "Combining Revealed and Stated Preferences Data," *Marketing Letters* 5, 335–350.

Chrzan, Keith and Michael Patterson. (1999). "Comparing the Ability of Full and Partial Profile Choice Experiments to Predict Holdout Choices," Paper presented at AMA's ART Forum, Santa Fe, NM, 1999.

DeSarbo, Wayne S. and Paul E. Green. (1984). "Concepts, Theory, and Techniques, Choice-Constrained Conjoint Analysis," *Decision Science* 15, 291–323.

Dickson, Peter R. and Alan G. Sawyer. (1990). "The Price Knowledge and Search of Supermarket Shoppers," *Journal of Marketing* 54, 42–53.

Fennell, Geraldine. (1997). "Value and Values: Relevance to Advertising." In L. Kahle and L. Chiagouris (eds.), *Values, Lifestyles, and Psychographics*. Hillsdale, NJ: Erlbaum.

Fennell, Geraldine. (1988). "Action as Counterchange: Identifying Antecedents of the Domain and Goal of Action," In L. Alwitt (ed.), *Proceedings 95th Annual Convention of the American Psychological Association*. Washington, D.C.: American Psychological Association.

Gilbride, Timothy J., Greg M. Allenby, and Jeff Brazell. (2004). "Models of Heterogeneous Variable Selection," Working paper, Ohio State University.

Gilbride, Timothy J. and Greg M. Allenby. (2004). "A Choice Model with Conjunctive, Disjunctive, and Compensatory Screening Rules," *Marketing Science* 23, 391–406.

Green, Paul E. and V. Srinivasan. (1978). "Conjoint Analysis in Consumer Research: Issues and Outlook," *Journal of Consumer Research* 5, 103–123.

Horsky, Dan and Paul Nelson. (1992). "New Brand Positioning and Pricing in an Oligopolistic Market," *Marketing Science* 11, 133–153.

Horsky, Dan, Paul Nelson, and Steven S. Posovac. (2004). "Stating Preference for the Ethereal but Choosing the Concrete: How the Tangibility of Attributes Affects Attribute Weighting in Value Elicitation and Choice," *Journal of Consumer Psychology* 14, 132–140.

Horsky, Dan, Paul Nelson, and Sanjog Misra. (2004). "Observed and Unobserved Preference Heterogeneity in Brand Choice Models," working paper, Simon School of Business University of Rochester.

Huber, Joel. (2004). "The Distorting Impact of Dominance and Compromise in Choice Based Conjoint," Paper presented at AMA's ART Forum, Whistler, Canada, 2004.

Huber, Joel, Dick R. Wittink, John A. Fiedler, and Richard Miller. (1993). "The Effectiveness of Alternative Preference Elicitation Procedures in Predicting Choice," *Journal of Marketing Research* 30, 105–114.

Huber, Joel and Klaus Zwerina. (1996). "The Importance of Utility Balance in Efficient Choice Designs," *Journal of Marketing Research* 33, 307–317.

Huber, Joel, Dan Ariely, and Greg Fischer. (2002). "Expressing Preferences in a Principal-Agent Task: A Comparison of Choice, Rating and Matching," *Organizational Behavior and Human Decision Processes* 87, 66–90.

Johnson, Rich and Bryan Orme. (1996). "How Many Questions Should You Ask in Choice-Based Conjoint Studies," Technical Report, Sawtooth Software: available at http://www.sawtoothsoftware.com/technicaldownloads.shtml#howmany.

Johnson, Rich, Joel Huber, and Lynd Bacon. (2003). "Adaptive Choice-Based Conjoint". In *Proceedings of the 10th Sawtooth Software Conference*, April 2003, San Antonio, Texas.

Kim, Jaehwan, Greg M. Allenby, and Peter E. Rossi. (2004). "Volumetric Conjoint Analysis," working paper, Ohio State University.

Kim, Jaehwan, Greg M. Allenby, and Peter E. Rossi. (2002). "Modeling Consumer Demand for Variety," *Marketing Science* 21, 229–250.

Kuhfeld, Warren. (2004). "Experimental Designs, Efficiency, Coding and Choice Designs," Technical Report TS-722C, SAS: available at http://support.sas.com/techsup/tnote/tnote_stat.html#market

Lenk, Peter. (1992). "Hierarchical Bayes Forecasts of Multinomial Dirichlet Data Applied to Coupon Redemptions," *Journal of Forecasting* 11, 603–619.

Lenk, Peter and Ambar Rao. (1990). "New Models from Old: Forecasting Product Adoption by Hierarchical Bayes Procedures," *Marketing Science* 9, 42–53.

Louviere, Jordan and George Woodworth. (1983). "Design and Analysis of Simulated Consumer Choice Allocation Experiments: An Approach Based on Aggregate Data," *Journal of Marketing Research* 20, 350–367.

Louviere, Jordan L., Deborah J. Street, and Leonie Burgess. (2003). "A 20+ Years Retrospective on Choice Experiments". In Yoram Wind and Paul E. Green (eds.), *Market Research and Methodology: Progress and Prospects*, New York: Kluwer Academic Press, pp. 201–214.

Ofek, Elie and V. Srinivasan. (2002). "How Much Does the Market Value an Improvement in a Product Attribute?" *Marketing Science* 21, 398–411.

Orme, Bryan and Gary Baker. (2000). "Comparing Hierarchical Bayes Draws and Randomized First Choice for Conjoint Simulations". In *Proceedings of the 8th Sawtooth Software Conference*, March 2000, Hilton Head Island, South Carolina.

Otter, Thomas, Sylvia Frühwirth-Schnatter, and Regina Tüchler. (2004). "Unobserved Preference Changes in Conjoint Analysis," Working paper, Ohio State University.

Renken, Tim, Greg Rogers, and Joel Huber. (2004). "A Comparison of Conjoint and Scanner Data-Based Price Elasticity Estimates," Paper presented at AMA's ART Forum, Whistler, Canada, 2004.

Severin, Valerie C., Leonie Burgess, Jordan Louviere, and Deborah J. Street. (2004). "Comparing Statistical Efficiency and Respondent Efficiency in Choice Experiments," working paper, Department of Mathematical Sciences, University of Technology, Sydney, Australia.

Sonnier, Garrett, Andrew Ainslie, and Thomas Otter. (2004). "The Effects of Different Parameterizations on Heterogeneous Choice Models," working paper, UCLA

Street, Deborah J. and Leonie Burgess. (2004). "Optimal and Near-optimal Pairs for the Estimation of Effects in 2-level Choice Experiments," *Journal of Statistical Planning and Inference* 118, 185–199.

Tobia, Olivier, John R. Hauser, and Duncan Simester (2004), "Polyhedral Methods for Adaptive Choice-based Conjoint Analysis," *Journal of Marketing Research* 41, 116–131.

Train, Kenneth. (2003). *Discrete Choice Methods with Simulation*, Cambridge: Cambridge University Press.

Zehna, Peter W. (1966). "Invariance of Maximum Likelihood Estimators," *Annals of Mathematical Statistics* 37, 744.